3
Lagrangian Coherent Structures

Shawn C. Shadden

3.1
Introduction

Mounting evidence suggests that fluid advection can effectively be studied by considering special material surfaces, which are referred to here as Lagrangian coherent structures (LCSs).[1)] What makes these material surfaces special is their distinguished attracting or repelling nature. Notably, LCS are often locally the most strongly attracting or repelling material surfaces in the flow, and as such have a strong influence on the flow topology. In fact, by understanding their evolution, one can often reveal *mechanisms* that underly complex laminar, and even turbulent, fluid transport in conspicuous detail.

Conceptually, LCS can be approached from the dynamical systems perspective, or from a more physically based fluid mechanics perspective. Starting from the former, we note that fluid advection is described by the equation

$$\dot{\mathbf{x}}(\mathbf{x}_0, t_0, t) = \mathbf{u}(\mathbf{x}, t) \tag{3.1}$$

where $\mathbf{u}(\mathbf{x}, t)$ is the velocity field of a fluid and $\mathbf{x}(\mathbf{x}_0, t_0, t)$ describes the motion (trajectory) of a fluid element, or equivalently a *material point*,[2)] starting at position $\mathbf{x}_0$ at time t_0; we assume nominally volume-preserving flow, $\nabla \cdot \mathbf{u} \approx 0$. Since the motion of fluid is, generally speaking, chaotic, revealing salient flow features helps us understand how the flow is organized. Indeed, the study of *coherent* or *organizing* structures in fluid mechanics has surely been of interest for as long as we have

1) Herein we use LCS to abbreviate both singular and plural forms.

2) The concept of a fluid element is an idealization – but one that is overwhelmingly used in modeling fluid mechanics, for example, the Navier-Stokes equation and resulting solution are based on this assumption. Resulting flow data is typically not "amended" by diffusion for LCS computations, as LCS targets advection, which indeed is the primary mode of transport over the length and time scales typically considered. For many applications though, interest lies in the advection of matter that is approximately transported by the fluid (e.g., bubbles, aerosols, suspensions, and emulsions); in such cases, the dynamics (Eq. (3.1)) can be appropriately augmented [5], or LCS of the fluid can be seen as approximate LCS for the advected matter, and under appropriate conditions, their relevancy can be established rigorously [6, 7], see also [8,9].

Transport and Mixing in Laminar Flows: From Microfluidics to Oceanic Currents,
First Edition. Edited by Roman Grigoriev.

studied fluid motion. In more recent history, the field of dynamical systems has assisted in making this effort precise. Generically, Eq. (3.1) is a nonlinear ordinary differential equation, which is the focus of study in dynamical systems theory. Hyperbolic fixed points and their associated stable and unstable manifolds[3),4)] are known to be organizing structures in dynamical systems, at least in the more mathematically tractable problems traditionally considered in that field [1], as described below. Briefly, in steady 2D flow, stable and unstable manifolds of hyperbolic fixed points behave as separatrices, partitioning regions with similar dynamics and providing the skeletal structure of the flow topology. Steady 2D flow is exceptional though, because it is integrable. Conversely, time-dependent and/or 3D flow is typically not integrable, and fluid motion generically is chaotic, making understanding the flow topology exponentially harder. The simplest unsteadiness that can be added to a system is a periodic perturbation, yet this simple change in the velocity field often translates to complex change in the fluid motion, that is, chaos. In such systems, periodic trajectories can be seen as fixed points using a stroboscopic view (Poincare map) and the stable and unstable manifolds of these fixed points can reveal a template organizing the fluid's chaotic motion. Therefore, computing these manifolds is central to understanding transport in periodic systems, and indeed this aspect of chaotic advection has become well-developed, see, for example, [2–4].

In practical applications, fluid flows span a spectrum of unsteadiness. As complexity increases, the motivation for identifying organizing structures becomes increasingly compelling, yet applying invariant manifold theory in complex flow problems is challenging, for example, it is difficult to even *define* these concepts satisfactorily for temporally aperiodic flows, which is well-documented [10–12]. Hence, the definitions for, and relevancy of, stable and unstable manifold theory in systems with general time dependence are not obvious, and research in this area continues to mature. At least in the sense presented here, these manifolds have fallen under the broader category LCS – a terminology adopted from Haller and Yuan [13]. This terminology helps broaden our description, and categorization, of organizing structures from the confines of the traditional definitions of stable and unstable manifolds, while maintaining a practical methodology for revealing "effective" stable and unstable manifolds in flows with arbitrary time dependence.

From the fluid mechanics perspective, fluid has a fascinatingly complex behavior and understanding the nature of this complexity is one of the great challenges in fluid mechanics. Growth in technology has enabled tremendous computational and empirical capabilities to derive velocity field data in wide-ranging applications. However, obtaining data is certainly not enough, we must be able to properly

3) A manifold can be viewed herein as a smooth surface, and it is tacitly assumed that a surface is meant herein to mean hypersurface (i.e., for 2D systems, surface is synonymous with curve).

4) For 2D systems, eigenvalues can be purely imaginary and, thus, the velocity gradient "nonhyperbolic." In 3D, since complex eigenvalues occur in conjugate pairs, this requires one, or all three, eigenvalue to remain zero. While this is certainly possible, it is presumably less common over finite time intervals in complex flows (in the interior at least, since no-slip boundaries are clearly nonhyperbolic).

assimilate the complex information encoded in that data. Comprehending the motion of the fluid is paramount in studying fluid mechanics but often precise understandings are elusive due to the inherent complexity of fluid flow. As described in Ref. [14], "human temporal perception struggles with untangling chaotic trajectories in a turbulent flow." Thus, we are tempted to more tractable descriptions, based on instantaneous rate-of-change information, such as the velocity field or related Eulerian measures. However, the relevancy of these fields, or coherent structures derived from them, to the actual fluid motion is dubious since time dependence can "wreak havoc" on conclusions drawn from instantaneous information. Furthermore, time series of instantaneous fields frequently mislead our perception from the reality of the integrated behavior of the fluid motion over time. These troubles are well-documented, but perhaps not widely appreciated. On the other hand, Lagrangian-based measures that explicitly track the fluid motion are less common, and in many cases *ad hoc* or qualitative. The advantage of the methods presented herein is the ability to systematically and succinctly encode key Lagrangian information into a single field from which we can define structures that remain relevant in space and time, are coordinate-frame invariant, and perhaps most importantly convey information regarding fundamental *mechanisms* of fluid transport.

The structures we compute are termed *Lagrangian* because they are defined from the fluid motion, as opposed to an instantaneous Eulerian snapshot, and because they are themselves material surfaces advected by the flow. They are also termed *coherent* because they have distinguished stability compared to nearby material surfaces, and consequentially, LCS can often be identified with familiar coherent flow features. For example, the coherent patterns traced out by visual markers (e.g., dye) we see emerge when visualizing fluid flow are often manifestations of underlying LCS. The Lagrangian nature of LCS makes them "transport barriers;" however, it is their distinguished stability, which often translates into separatrix behavior, that substantiates their importance. In this chapter, we review and summarize the main concepts underlying this approach to studying fluid advection, and how these structures are commonly computed in practice and related practical concerns. For underlying mathematical details, or specific application-oriented computational results, we have attempted to provide appropriate references.

3.2 Background

Let us revisit the nature of hyperbolic fixed points in steady flow for developing intuition regarding LCS and to make ideas introduced in the section above more concrete. Consider the steady vortex dipole shown in Figure 3.1a – a pedagogical application that we will build upon. Vortices are considered the building blocks for turbulence, and thus understanding their structure is fundamental to fluid mechanics. In the vortex dipole, there are two hyperbolic saddle-type fixed points. Their stable/unstable manifolds are the trajectories that asymptote to these points in forward/backward time. In general, stable and unstable manifolds act as separatrices,

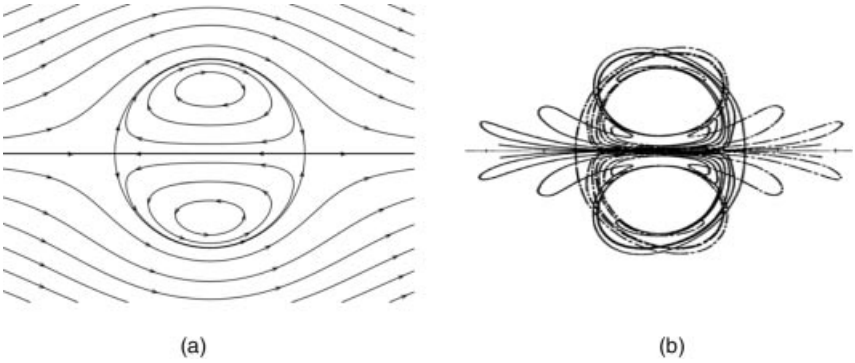

Figure 3.1 Panel (a): Streamlines for a steady vortex dipole. Two hyperbolic saddle-type stagnation points are located upstream and downstream. Their stable and unstable manifolds form so-called *heteroclinic* connections between the saddle points, and behave as *separatrices* partitioning the flow into regions of distinct dynamics–a common property of stable and unstable manifolds. Panel (b): Periodically perturbed vortex dipole. Heteroclinic tangle formed by transverse intersection of stable and unstable manifolds when a perturbation is applied to the steady vortex dipole, reproduced from [15] with permission.

partitioning regions of different dynamics. This behavior is exemplified here as these manifolds define the vortex boundaries, separating recirculating fluid from non-recirculating fluid. In general, hyperbolic fixed points, which are required to be saddle-type due to incompressibility of fluid, form permanent hubs, drawing in fluid and expelling the fluid in diverging directions. Namely, the saddle points in the vortex diapole pull in the fluid along the stable manifolds and expel the fluid in diverging directions along the unstable manifolds. Restating, the fixed points are hubs directing the flow, and the invariant manifolds delineate how fluid will be directed by those hubs. This behavior is central to the importance of stable/unstable manifolds, and as it turns out, this behavior remains relevant in complex unsteady flows; however, the role of fixed points is often replaced by appropriately behaved moving trajectories.

The clean partitioning of dynamically distinct regions in the domain by stable and unstable manifolds, and lack of mixing, as seen in Figure 3.1a, occurs primarily in 2D time-independent problems. The geometrical structure of the transport topology is usually much more complicated when systems become time-dependent or 3D. When steady 2D systems, as in Figure 3.1a, are subjected to periodic perturbations, the fixed points typically perturb to periodic orbits, and their associated stable and unstable manifolds transform from neat heteroclinic connections to complex heteroclinic tangles, as illustrated in Figure 3.1b, leading to the theory of lobe dynamics, see, for example, [15, 16]. This transverse intersecting of manifolds is one of the hallmarks of chaos. If one visualizes the chaotic dynamics along with these manifolds, the dynamics are revealed to be, in fact, rather orderly and predictable. A great deal of developments have been made in understanding chaotic dynamics in periodically, and quasi-periodically, perturbed systems using invariant manifold theory, see, for

example, [2–4]. Notably, such information has proved very valuable in the understanding, and practical design, of efficient fluid mixing.

Several difficulties exist for computing stable and unstable manifolds in nonperiodic systems. First, it is not immediately clear what one should compute the invariant manifolds of. For steady flow, hyperbolic fixed points and for periodic flow, hyperbolic periodic trajectories are the *persistent hubs* organizing the flow topology. However, fixed or periodic trajectories are not commonly encountered in most practical applications due to the general unsteadiness of most flows. Furthermore, trajectories that are hyperbolic at one instant might quickly change their stability – something that is not a problem when working with fixed points in steady and periodic systems. Next, the very definition of stable/unstable manifolds becomes unwieldily in the "real world." Defining stable and unstable manifolds requires asymptotic limits [1, 17]. However, such limits are irrelevant for practical applications since (1) the fluid's motion is known only from an inherently finite data set, for example, from a computational fluid dynamics simulation or empirical measurement, and (2) since trajectories can gain or loose hyperbolicity over time and we desire to understand inherently transient phenomena. Amending invariant manifold theory as needed to rigorously and unambiguously establish and compute stable and unstable manifolds of hyperbolic trajectories in nonperiodic, finite-time systems takes care (see, e.g., [11, 18, 19]). Herein, we take a "less rigorous" approach that seeks to uncover such manifolds, and perhaps more general organizing structures, using hyperbolicity measures, such as the finite-time Lyapunov exponent (FTLE) distribution, as described in Section 3.3. Numerous and wide-ranging applications have demonstrated, indeed, that structures analogous to stable and unstable manifolds are effectively determined in nonperiodic, and even turbulent, flow, and in fact this approach has revealed that such structures are seemingly ubiquitous in fluid mechanics. LCS have garnered much attention because they often provide a more precise (and temporally relevant) depiction of common fluid mechanics constructs (e.g., vortices, flow separation, and stirring) than prevailing methods of flow characterization, making them very useful indeed. In the context of mixing, while lobe dynamics generally becomes less common in nonperiodic systems, LCS continue to control the stretching, folding, and alignment mechanisms underlying kinematic mixing.

We note that the term LCS is used in various communities to mean different things. An important distinction is that LCS is used here to describe coherent *surfaces*, and not coherent *regions* in the flow. Nonetheless, often due to their separatrix nature, such LCS end up identifying boundaries to regions of coherent dynamics, as stable/unstable manifolds often do (cf. Figure 3.1, see also Refs. [20, 21] for comparison with boundaries to almost-invariant regions). In theory, LCS encompasses the notion of stable and unstable manifolds of hyperbolic trajectories, and in practice, as often computed, typically reveal these manifolds in systems where they are known to exist. However, it is not clear that we require that LCS be stable/unstable manifolds, at least in a traditional sense. Indeed, it is perhaps advantageous to think of an LCS, defined by Haller [22], more broadly as a locally strongest repelling or attracting material surface.

3.3 Global Approach

Considering LCS as the most repelling or attracting material surfaces we could ask what is the best way to compute them. Testing the expansion or contraction about every material surface in the flow is hopeless, since there is an infinite possibility of surfaces to test and there is no good way to discretize the set of all possible surfaces. Instead, we take a more indirect approach. We discretize the fluid domain with a dense grid of material points, measure the Lagrangian expansion rate (roughly "hyperbolicity") about each material point, plot the spatial distribution, and extract surfaces that maximize the measure. This is essentially how *repelling* LCS are most commonly computed. *Attracting* LCS are analogously computed by reversing time, as expansion in backward time implies contraction in forward time. Since the computation is so similar, we often do not distinguish the direction of integration in discussions of methodology. One may contrast this "global approach" to how stable/unstable manifolds are traditionally computed, whereby the computation begins with the identification of a *specific* hyperbolic trajectory and stable and unstable manifolds are grown from the stable and unstable subspaces of that trajectory. The global approach proceeds independent of any specific hyperbolic trajectories, whereby the distribution of an appropriate hyperbolicity measure generally reveals *all* influential finite-time hyperbolic structures (e.g., the relevant hyperbolic trajectories and their associated stable and unstable manifolds).

One should keep in mind that when we speak of attracting and repelling surfaces, we do so in the context of nominally volume-preserving fluid flows. Thus, material points about a repelling LCS tend to compress in the tangential direction and expand in the normal direction. Likewise, material points about an attracting LCS will tend to expand along the LCS and contract in the normal direction toward the structure.

The global approach was built on the observation that material points straddling stable/unstable manifolds typically separate faster in forward/backward time than pairs of points not straddling such manifolds. Alternatively, as previously stated, LCS can more generally be defined as material surfaces with a most repelling or attractting property, without regard to them being stable or unstable manifolds of something.[5)] This strategy for computing LCS owes to many people. Bowman wrote a paper [23] that proposed the use of a finite-strain (FS) field for locating stable and unstable manifolds. As summarized by Jones and Winkler [24], Bowman's "idea is to infer the manifold geometry by considering the stretching associated with the hyperbolicity of these structures." Bowman's paper captured the essence of how LCS are often computed nowadays (this paper was rejected from publication yet has been numerously acknowledged). As noted by Bowman, the computation of the FS field is similar

5) Stable/unstable manifolds do not always produce ridges in the FTLE field, see, for example, [34, 40, 41]. This has mainly been demonstrated in simple flows, as noted in [41], that lack sufficient spatial heterogeneity or with "uniform" chaos. That is, while stable and unstable manifolds usually act as separatrices, they are not necessarily the most repelling or attracting surfaces. Recently, Haller [22] introduced the notion of "weak" LCS and how to "constrain" the computation to distinguish the invariant manifolds.

to the computation of the FTLE field, which had been used in 1991 by Pierrehumbert [25] and in 1993 by Pierrehumbert and Yang [26] to reveal "structure" in atmospheric flows. Doerner *et al.* [27] published a paper that argued that local maxima of the FTLE field coincide with stable manifolds of hyperbolic fixed points (focussing on steady systems). Winkler [28] noted that such measures are closely related to measures of relative dispersion (RD), especially in practice, which, for example, von Hardenberg *et al.* [29] also utilized, along with FTLE fields, as a means to locate invariant manifold-type structures in geophysical flows. More precise quantification of the FTLE approach was described by Haller [30, 31], followed by Shadden *et al.* [32] and Lekien *et al.* [33]. Haller and Yuan [13] also proposed the hyperbolicity time (HT) measure that was more specific in its measure of hyperbolicity than existing "finite-stretching" measures, and Haller's papers [13, 30, 34, 35] developed a more rigorous foundation to LCS compuation than previously available; indeed, terminology and basic ideas we discuss here were significantly influenced by these fundamental works. The similar, but distinct, finite-size Lyapunov exponent (FSLE), developed by Aurell *et al.* [36], which measures time to separate a set distance instead of the FTLEs measure of distance separated in a set time, was also proposed [37–39] as a suitable measure in practice.

3.3.1 FTLE

So far, the spatial distribution of the FTLE appears to be an effective way to detect LCS computationally; in fact, it is often considered the "de facto" method. For illustration purposes, the application of this approach to a more practical version of the vortical flow considered above is shown in Figure 3.2. Briefly, planar velocity data was measured surrounding a mechanically generated vortex ring using particle image velocimetry (PIV). The velocity data was used to compute trajectory data, and subsequently FTLE fields in both forward and backward time (FTLE computation is described in more detail in Section 3.4). Snapshots of the forward and backward FTLE fields, and the instantaneous PIV velocity field, are plotted at one time instant in Figure 3.2. Visually, the FTLE fields reveal distinct curves of high FTLE; these curves

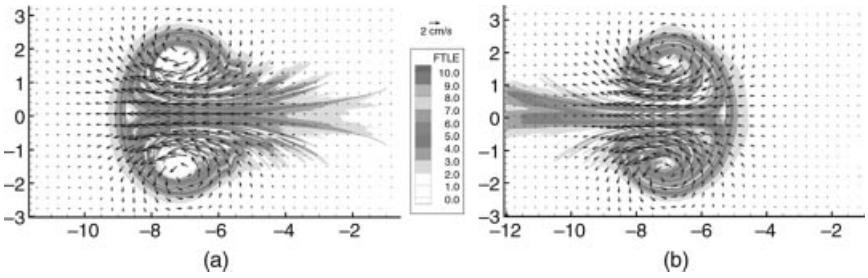

Figure 3.2 Backward and forward time FTLE fields, panels (a) and (b) respectively, for empirical vortex ring reveal attracting/repelling LCS that define vortex boundary and control transport and mixing. For further details see [42, 43].

are LCS analogous to the stable and unstable manifolds shown in Figure 3.1b (unlike the computation of the manifolds in Figure 3.1b, the LCS computation proceeds regardless of the time dependence of $\boldsymbol{u}(\boldsymbol{x}, t)$, making it applicable in more general settings). Together, these LCS reveal the *precise* vortex boundary, and precisely how fluid is entrained/detrained to/from the vortex ring as it forms and propagates; highly relevant information that is notably absent when employing traditional analyses techniques (e.g., velocity field, vorticity, streamlines, and Q-criterion).

Lyapunov exponents, in general, have long been used to determine the predicability, or sensitivity to initial conditions in dynamical systems. The FTLE measures the separation rate between initially close trajectories. Roughly, if two material points are initially separated by a small distance $|\xi_0|$ at time t_0, then the separation at some later time will be

$$|\xi_t| \approx e^{\Lambda(t-t_0)}|\xi_0| \tag{3.2}$$

or equivalently

$$\frac{|\xi_t|}{|\xi_0|} \approx e^{\Lambda(t-t_0)} \tag{3.3}$$

where Λ denotes the FTLE, which is a function of space, time and integration length. To see this we first note that the precise separation between two arbitrary material points starting at time t_0 from positions $\boldsymbol{x}_0$ and $\boldsymbol{x}_0 + \xi_0$ is given by

$$\xi_t = F_{t_0}^t(x_0 + \xi_0) - F_{t_0}^t(\boldsymbol{x}_0) \tag{3.4}$$

where $F_{t_0}^t : \boldsymbol{x}_0 \mapsto \boldsymbol{x}(\boldsymbol{x}_0, t_0, t)$ denotes the flow map. By expanding the function $F_{t_0}^t(\boldsymbol{x}_0 + \xi_0)$ in ξ_0 in Eq. (3.4), the distance between the two points over time is

$$|\xi_t| = \left|\nabla F_{t_0}^t(\boldsymbol{x}_0) \cdot \boldsymbol{x}_0\right| + \mathcal{O}(|\xi_0|^2) \tag{3.5}$$

which, by definition of the Euclidean norm, can be rewritten as

$$|\xi_t| = \sqrt{\boldsymbol{e}^\top \cdot \nabla F_{t_0}^t(\boldsymbol{x}_0)^\top \cdot \nabla F_{t_0}^t(\boldsymbol{x}_0) \cdot \boldsymbol{e}}\, |\xi_0| + \mathcal{O}(|\xi_0|^2) \tag{3.6}$$

where $\boldsymbol{e} = \xi_0/|\xi_0|$ and $\top$ denotes transpose. The stretching factor, or coefficient of expansion, between the two points becomes

$$\lim_{|\xi_0| \to 0} \frac{|\xi_t|}{|\xi_0|} = \sqrt{\boldsymbol{e}^\top \cdot \nabla F_{t_0}^t(\boldsymbol{x}_0)^\top \cdot \nabla F_{t_0}^t(\boldsymbol{x}_0) \cdot \boldsymbol{e}} \tag{3.7}$$

when $|\xi_0|$ becomes small. Letting $\lambda^i(\boldsymbol{x}_0, t_0, t)$ denote the ith eigenvalue, numbered in decreasing order, of the right Cauchy–Green strain tensor

$$C(\boldsymbol{x}_0, t_0, t) = \nabla F_{t_0}^t(\boldsymbol{x}_0)^\top \cdot \nabla F_{t_0}^t(\boldsymbol{x}_0) \tag{3.8}$$

and $\boldsymbol{e}^i(\mathbf{x}_0, t_0, t)$ the associate eigenvectors, the coefficients of expansion in the eigenvector directions, that is, $\xi_0 = |\xi_0|\mathbf{e}^i$, are given by

$$\lim_{|\xi_0|\to 0}\frac{|\xi_t|}{|\xi_0|}=\sqrt{\lambda^i(\xi_0,t_0,t)} \tag{3.9}$$

Note, $\sqrt{\lambda^i}=\sigma^i$ is the singular value of the deformation gradient $\nabla F_{t_0}^t(\boldsymbol{x}_0)$. The FTLEs are defined by taking the natural logarithm of the coefficients of expansion, and dividing by the length of time the expansion occurred over (*integration time*) to get the average logarithmic expansion rate

$$\Lambda^i(\boldsymbol{x}_0,t_0,t)=\frac{1}{|t-t_0|}\ln\sqrt{\lambda^i(\boldsymbol{x}_0,t_0,t)}=\frac{1}{t-t_0}\ln\sigma^i(\boldsymbol{x}_0,t_0,t) \tag{3.10}$$

The largest FTLE, Λ^1 is often referred to without distinction as the FTLE, Λ.

The eigenvectors $\boldsymbol{e}^i$ form an orthonormal basis. Nearly all generic perturbations ξ_0 will have a component in the $\boldsymbol{e}^1$ direction, and growth of the perturbation will eventually be dominated by growth of that component since it has the largest eigenvalue. This observation, along with the proceeding derivations, justifies Eq. (3.3).

We last note that geometrical features of the fields $\lambda(\boldsymbol{x}_0,t_0,t)$, $\sigma(\boldsymbol{x}_0,t_0,t)$, or $\Lambda(\boldsymbol{x}_0,t_0,t)$ are roughly equivalent since the mappings between these measures (Eq. (3.10)) are monotonic. For example, local extrema of one field will correspond to local extrema of the other. However, because of the often exponential divergence of particles straddling LCS, the logarithmic-based FTLE field can produce crisper structures. For most purposes, these fields, and *ad hoc* measures introduced above, can be used interchangeably, or with minor modification, in the discussions that follow. We focus on FTLE because it has been most widely used, and more rigorous results exist for this method [22, 31–33, 35].

3.3.2 FTLE Ridges

LCS appear as local maximizing surfaces of FTLE. In some ways, it is striking that when one plots a suitable scalar measure of hyperbolicity, such as the FTLE, one can clearly reveal an otherwise hidden template of transport. That is, why should we expect that if we plot the spatial distribution of the FTLE, it will form well-defined surfaces and not be more diffuse? Why should we expect these surfaces to be influential separatrices, or for them to be transported by the flow? Perhaps equipped with the right perspective such behavior seems sensible and justifiable, but it is not obvious such coherence will persist in the complexity of turbulent flow, especially when employing *averaged* stretching statistics such as the FTLE. Clearly, material surfaces with distinguished stability persist in all sorts of applications (cf. Section 3.6) and are readily detectable using suitable measures such as the FTLE.

Shadden *et al.* [32] formalized the definition of LCS as ridges of the FTLE field, an idea introduced in [31], and that definition was extended to higher dimensional systems in Lekien *et al.* [33]. These papers helped to make this approach precise, and

consequently to help better derive properties of "ridges" of FTLE and to form a basis for computational strategies. In [32], we defined a ridge as essentially a curve on which the FTLE is locally maximized in the transverse direction, leading to a "second-derivative ridge" definition that required that a ridge was a curve (more generally, hypersurface in high dimensions) where the first derivative of the FTLE field must be zero in the normal direction, $\boldsymbol{n}$, and the second derivative of the field must be negative and minimum in the normal direction. Furthermore, we desired hyperbolic hypersurfaces where only the largest Lyapunov exponent be positive, with the rough idea being that there should be attraction within the surface and repulsion normal to it. These conditions for 3D systems can be stated mathematically as

1) $\nabla\Lambda^1 \cdot \boldsymbol{n} = 0$
2) $\langle \boldsymbol{n}, \nabla^2\Lambda^1\boldsymbol{n}\rangle = \mu_{\min}\left(\nabla^2\Lambda^1\right) < 0$
3) $\Lambda^1 > 0 > \Lambda^2, \Lambda^3$

where $\mu_{\min}(\nabla^2\Lambda^1)$ is the minimum eigenvalue of $\nabla^2\Lambda^1$.

As the computation of FTLE/LCS matured, deficiencies were noted, see, for example, [32, 34, 40, 41] and discussion in section below. To address these issues, Haller [22] more recently derived necessary and sufficient conditions for LCS in terms of invariants of the strain tensor. Recalling that the FTLE is derived from the largest eigenvalue of the finite-time strain tensor, this approach can be seen as taking a step back to the more complete stretching information encoded in $C(\boldsymbol{x}_0, t_0, t)$. A primary motivation was to ensure that LCS are locally the *most normally* repelling structures (as motivated below). Those conditions can subsequently be used to refine the FTLE ridge definition; namely, sufficient conditions that can be checked are that $\boldsymbol{e}_1 = \boldsymbol{n}$ and $|\nabla\boldsymbol{e}_1 \cdot \boldsymbol{e}_1| \leq 1$, and

$$\langle \boldsymbol{n}, \nabla^2\lambda^1\boldsymbol{n}\rangle < -2\lambda_1|\nabla\boldsymbol{e}_1 \cdot \boldsymbol{e}_1|\left(\frac{\lambda_1}{\lambda_2} + \frac{\lambda_1}{\lambda_3}\right)$$

note this last condition is in terms of λ_is, which can be rewritten in terms of Λ_is via Eq. (3.10). Under these conditions, condition 3 above can be relaxed to $\Lambda^2 \neq \Lambda^1 > 0$. An important condition for normal hyperbolicity is $\boldsymbol{e}_1 = \boldsymbol{n}$; interestingly, it turns out that $\boldsymbol{e}_1 \to \boldsymbol{n}$ for persistent LCS, that is, the direction of dominant expansion/contraction converges to the normal direction.

The above mathematical criteria can be taken as a working definition of LCS, but we note that LCS are physical objects in the flow, and as such exist independent of any metric used to compute them. Nonetheless, mathematically precise definitions, as given above, are often necessary, if higher level definitions are not amenable to rigor or algorithmic implementation.

3.3.3 Nature of Stretching

The study of stretching in fluid flow has a long history, which we cannot possibly overview here. However, we do want to point out some pertinent comments

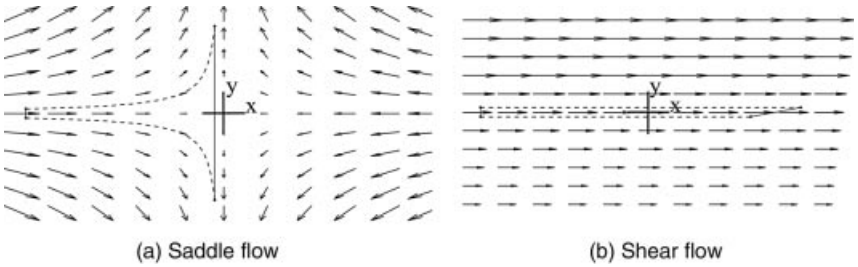

Figure 3.3 Stretching about a material line can be due to normal expansion (a) or tangential expansion (b). Certain stretching diagnostics used for LCS detection (e.g., FTLE) can obscure this distinction.

regarding stretching and hyperbolicity. A common physical interpretation of FTLEs regards the deformation of a spherical fluid element. Based on the arguments above, an *infinitesimal* sphere of fluid will deform into an ellipsoid when advected by the flow. The principal axes of the ellipsoid align with the eigenvectors of the strain tensor (Eq. (3.8)), and each principal axis is equal to the original radius of the sphere scaled by the corresponding coefficient of expansion. Thus, the FTLEs represent the average logarithmic expansion rates of each principal axis. Note for volume-preserving flow, the sum of Lyapunov exponents must be zero.

Hyperbolicity is often synonymous with its manifestation, stretching, in fluid mechanics. In a broad sense, a trajectory is hyperbolic if infinitesimal perturbations to that trajectory expand or contract over time. How this should be interpreted in finite time is not obvious. Trajectories can change behavior over time, and long-term limits cannot be computed, or are perhaps irrelevant since we may be interested in transient behavior. The physical interpretation in the proceeding paragraph motivates defining finite-time hyperbolicity in terms of FTLEs.[6)] For example, Wiggins [44] defines a trajectory as hyperbolic over a finite time interval if none of its FTLEs is zero. Correspondingly, a case could be made that the semimajor axis, or equivalently the largest FTLE, can be viewed as a *measure* of hyperbolicity for that time interval since nearly all perturbations approach this expansion rate (Eq. (3.3)).

This interpretation of hyperbolicity is a bit generic though, because almost all trajectories are hyperbolic under such a definition. Also, while the FTLE, or more generically any time-averaged Lagrangian measure of stretching, will indicate levels of "hyperbolicity" for a trajectory, such measures can obscure the *nature* of the stretching. A revealing juxtaposition is shown in Figure 3.3. In both panels (a) and (b), the x-axis is a material surface. In panel (a), we see growth of a line element straddling the axis in the normal direction for a flow field topologically equivalent to flow near a hyperbolic trajectory. Alternatively, in panel (b), we see

6) Hyperbolicity has also been phrased in terms of "exponential dichotomy" [45].

growth of two points straddling the x-axis in the tangential direction for a shear flow.[7] In Figure 3.3a, the x-axis is a stable manifold and the general stretching behavior observed is heuristically what one typically associates with an LCS. However, locations of high shear certainly can cause high values of FTLE, and surfaces of maximal shearing rate may lead to ridges in the FTLE (or similar) field. This paradox was noted in [22, 31], with the shear profile shown in Figure 3.3b taken from [22]

$$\begin{aligned} \dot{x} &= 2 + \tanh(y) \\ \dot{y} &= 0 \end{aligned} \tag{3.11}$$

For this flow, the shear rate is locally maximized at $y = 0$ and thus the FTLE field would admit a ridge along the x-axis. However, it would seemingly be awkward to refer to the x-axis in Figure 3.3b as an LCS since it does not act as a stable manifold, or more generally a separatrix.

One can exclude such maximal shear surfaces if one is more choosy in defining a hyperbolic surface. The importance of normally hyperbolic invariant manifolds has been noted [17] and similar ideas have been employed for identifying LCS with *normally* hyperbolic material surfaces. Normal hyperbolicity can be defined in multiple ways in finite time. For example, as presented in Haller and Yuan [13], one could require that normal perturbations to the surface *continuously* expand/contract in the normal direction over the complete time interval of interest, leading to, for example, the HT approach to LCS detection [13, 30]. More recently, Haller [22] required that at the end of the time interval of interest, a perturbation that started normal to the surface must have expanded/contracted *more* in the direction normal to the surface than in any other direction; this condition is likely easier to verify in practice. In any sense, the example shown in Figure 3.3b is certainly not normally hyperbolic since there is no expansion of the line element in the direction normal to the material surface (x-axis).

Inevitably in most practical applications, stretching along a hyperbolic material surface will have tangential and normal components, placing its behavior somewhere between the ends of the spectrum presented in Figure 3.3, with the additional complication that the stretching characteristics can change along the surface in both space and time. Refinements to our views of hyperbolicity offer alternatives for us to refine how we ultimately define a maximally stretching surface as an LCS. At this point, only a few works have distinguished different categories of maximal stretching surfaces in practical applications. Preliminary results indicate that most structures are normally hyperbolic [46]; however, other types of structures are certainly possible [47]. More experience is needed at this time to understand the implications

7) In the saddle flow, one may notice that the growth is exponential in time, whereas the growth in the shear flow is linear in time. From these examples, one might expect that for sufficient integration time, normally hyperbolic structures become more pronounced than shear structures when plotting FTLE. While this may tend to happen, it is not always the case. Furthermore, distinction of "exponential growth" for finite-time, nonperiodic systems is not straightforward; for example, the shear flow has nonzero FTLE – more generally, any type of growth in finite time can be bounded by an exponent.

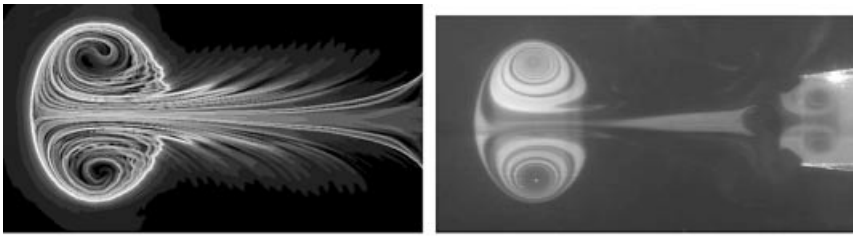

Figure 3.4 Comparison of backward time FTLE and attracting LCS (panel a) with a florescence visualization (panel b) by Paul S. Krueger and John O. Dabiri. The PIV data used to compute FTLE and the florescence visualization come from similar experimental setups, but different experiments and different times during each experiment. Nonetheless, structures revealed in both fields are similar, as florescence aligns with attracting LCS. The role of repelling LCS is less obvious from the flow visualization.

of such distinctions in practical flows. This could perhaps lead to deeper understandings on the nature of fluid, and help refine how we approach studying it.

We note that any notion of hyperbolicity must be associated to a specific time interval, exemplifying one of the main idiosyncrasies of finite-time analysis. Thus, our view on the stability of a surface in general depends on our horizon of knowledge. For example, perturbations to a surface could be marked by mostly tangential stretching during one time interval and switch to primarily normal expansion using a slightly longer horizon, or a surface could be marked by expansion using one integration time and then attraction using another. Likewise, a ridge in the FTLE field may only appear as a ridge using a set range of integration times.

Finally, we point out that attracting LCS will generally be distinguishable from advection of material points since fluid is attracted to and along these surfaces. This point connects to arguably the most intuitive method of understanding flow – the visualization of some visual marker that follows the fluid, "tracers." The utility of flow visualization is not to track individual particles but to use tracers to reveal coherent features that reveal how the flow is organized. We now know that such coherent patterns commonly observed are manifestations of attracting LCS (cf. Figure 3.4). However, we point out that repelling LCS are essentially hidden from flow visualization, even though these structures play a fundamental role in transport. Therefore, the computation of attracting and repelling LCS both broadens (by revealing a new class of hidden coherent structures) and deepens (by making the organizing surfaces precise) our understanding of coherent structures in fluid flow.

3.3.4 Objectivity

A desirable property of the global approach is that it is independent of coordinate frame, that is, objective. This cannot be said of most common flow diagnostics.

A quantity is called objective if it remains invariant under translations and rotations, that is, coordinate transformations $\boldsymbol{x} \mapsto \boldsymbol{y}$ of the form

$$\mathbf{y} = Q(t)\,\mathbf{x}(\mathbf{x_0}, t_0, t) + \mathbf{b}(t) \tag{3.12}$$

where $Q(t)$ is a time-dependent rotation matrix and $\mathbf{b}(t)$ represents a time-dependent translation. As before, let

$$\xi_t = \boldsymbol{x}(\boldsymbol{x}_0 + \xi_0, t_0, t) - \boldsymbol{x}(\boldsymbol{x}_0, t_0, t) \tag{3.13}$$

Applying a general transformation of coordinates according to Eqs. (3.12) and (3.13) becomes

$$\zeta_t \doteq Q(t)\,\boldsymbol{x}(\boldsymbol{x}_0 + \xi_0, t_0, t) + b(t) - Q(t)\,\boldsymbol{x}(\boldsymbol{x}_0, t_0, t) - b(t) \tag{3.14}$$

$$\zeta_t = Q(t)\xi_t \tag{3.15}$$

Note, $Q(t)$ is an isometry, so $|\zeta_t| = |\xi_t|$. Therefore, *relative* stretching measures such as Λ remain invariant under coordinate transformation.

3.4 Computational Strategy

LCS are effectively revealed by computation of the FTLE field. The basic strategy for computing the FTLE field is straightforward. Essentially, a grid of material points is advected and the deformation of the grid is used to compute the distribution of FTLE. Notably, to compute the FTLE *field*, we need trajectory information over a *grid* of points. Out of necessity, trajectory information is usually derived from postprocessing velocity field data obtained by computational fluid dynamics or empirical measurement. While computationally intensive, this approach is highly flexible, allowing for variations in the distribution and resolution of trajectory information, which can ultimately influence the accuracy and performance of LCS identification. The use of trajectory information independent of velocity field information is possible, and indeed desirable since it avoids introducing modeling and computational errors, but rarely a viable option.

The standard algorithm starts with the initialization of a grid of material points at time t_0. For simplicity, we assume these points form a structured grid $\boldsymbol{x}_{ijk}(t_0)$. This simplifies gradient computation and does not require *a priori* knowledge of the flow topology, which makes the discussion most broadly applicable. The initial locations of these points, as opposed to the final locations, represent the locations at which FTLE will be computed – the FTLE grid. The material points are integrated by the velocity field for some finite time interval, $T = t - t_0$,

$$\mathbf{x}_{ijk}(t) = \mathbf{x}_{ijk}(t_0) + \int_{t_0}^{t} \boldsymbol{u}(\mathbf{x}_{ijk}(\tau), \tau)\mathrm{d}\tau \tag{3.16}$$

using numerical integration. Note, for measures such as FSLE or HT, the integration time, T, of each trajectory depends on the stretching characteristics, whereas measures such as the FTLE, RD, or FS nominally use a fixed horizon for each material point. Such integration requires the velocity field to be interpolated in space and time. Higher-order integration and interpolation schemes ensure greater accuracy and smoothness of the computed results, respectively. One may assume that since we have set out on a numerically daunting task of measuring exponential growth, and computing quantities based on first and higher order derivatives of the flow, that the integration and interpolation schemes are critical. Even though precise trajectory computation is difficult in chaotic flow, LCS identification is typically reliable (cf. Section 3.5). Notably, integration and interpolation errors do not typically produce spurious structures, nor will they cause influential LCS to be undetected. Of course variations in identified LCS location, smoothness, and so on will depend on the numerical method, but variations are typically small if care is taken to ensure the integration and interpolation schemes used are reasonable in the light of the flow conditions. Efficient integration on unstructured grids was described in [48].

Once the final positions, $\mathbf{x}_{ijk}(t)$, are computed, the deformation gradient can be evaluated at each point in the structured grid $\mathbf{x}_{ijk}(t_0)$ by finite differencing, for example, using central differencing

$$\nabla F_{t_0}^{t}\big|_{(x_{ijk}(t_0),y_{ijk}(t_0),z_{ijk}(t_0))} \approx \begin{bmatrix} \frac{x_{(i+1)jk}(t)-x_{(i-1)jk}(t)}{x_{(i+1)jk}(t_0)-x_{(i-1)jk}(t_0)} & \frac{x_{i(j+1)k}(t)-x_{i(j-1)k}(t)}{y_{i(j+1)k}(t_0)-y_{i(j-1)k}(t_0)} & \frac{x_{ij(k+1)}(t)-x_{ij(k-1)}(t)}{z_{ij(k+1)}(t_0)-z_{ij(k-1)}(t_0)} \\ \frac{y_{(i+1)jk}(t)-y_{(i-1)jk}(t)}{x_{(i+1)jk}(t_0)-x_{(i-1)jk}(t_0)} & \frac{y_{i(j+1)k}(t)-y_{i(j-1)k}(t)}{y_{i(j+1)k}(t_0)-y_{i(j-1)k}(t_0)} & \frac{y_{ij(k+1)}(t)-y_{ij(k-1)}(t)}{z_{ij(k+1)}(t_0)-z_{ij(k-1)}(t_0)} \\ \frac{z_{(i+1)jk}(t)-z_{(i-1)jk}(t)}{x_{(i+1)jk}(t_0)-x_{(i-1)jk}(t_0)} & \frac{z_{i(j+1)k}(t)-z_{i(j-1)k}(t)}{y_{i(j+1)k}(t_0)-y_{i(j-1)k}(t_0)} & \frac{z_{ij(k+1)}(t)-z_{ij(k-1)}(t)}{z_{ij(k+1)}(t_0)-z_{ij(k-1)}(t_0)} \end{bmatrix}. \tag{3.17}$$

Once the deformation gradient is computed, straightforward evaluation of its largest singular value, and subsequently the FTLE (or similar), at each location $(x_{ijk}(t_0), y_{ijk}(t_0), z_{ijk}(t_0))$ is possible. This procedure provides the FTLE field and thus reveals LCS at the time instant t_0. For unsteady flow, the FTLE field and locations of LCS generally change over time and the above procedure can be repeated for a range of times t_0 to provide a time series of FTLE fields, and consequently, a time history of the LCS movements. Note that to compute a time series of FTLE fields, a new grid of material points is "released" at a series of release times t_0 and integrated from t_0 to $t_0 + T$.

A few notes are in order. First, we recall that the linearized flow map, that is, deformation gradient, can be derived from solving the linearized dynamical system

$$\frac{\mathrm{d}}{\mathrm{d}t}\nabla F_{t_0}^{t}(\boldsymbol{x}_0) = \nabla \boldsymbol{u}(\boldsymbol{x}(t),t)\cdot\nabla F_{t_0}^{t}(\boldsymbol{x}_0) \tag{3.18}$$

ergo perturbations can be in theory integrated by the linearized flow. Nonetheless, we typically integrate the nonlinear flow and then use finite differencing to approximate the deformation gradient. In practice, the two methods are not the same due to the finite differences involved. Second, the computation of trajectories is highly parallel since particles are assumed to not interact with each other or the flow. Therefore, one can readily exploit parallel architectures for trajectory computation by partitioning in space (or time when computing a time series of FTLE fields), see, for example, [49, 50]. Third, in the computation of a time series of FTLE fields, the integration interval for a particular release will typically overlap with the integration time interval of later, or previously, released grids. Therefore, it would be advantageous if each release could "use the same flow map information." This is typically not possible since material points in each release will not occupy the same position at the same time; however, Brunton and Rowley [51] demonstrate that under suitable approximation, an efficient flow map interpolation strategy can be devised. Finally, several variations of the above procedure are possible and in many cases desirable; the steps outlined here represent the rudimentary components.

3.4.1 Grid-Based Computation

Here we point out a few practical comments regarding the use of the FTLE grid. First, the computation of the deformation gradient described above is equivalent to considering the stretching of three orthogonal line elements, in the x-, y-, and z-directions, with each column representing the stretching from each line element. From Eq. (3.3) and the discussion that followed, we might expect that stretching of each element becomes similar. Indeed, the FTLE could be approximated at each point by considering the stretching of a single line element. This, in fact, has been done, see, for example, [25, 26, 29]. However, the evaluation of the full-deformation gradient does not incur substantial computational expense, and can provide more accurate results and additional information that could be useful for LCS identification (cf. Section 3.2). Furthermore, several related stretching measures can be derived from the deformation gradient, and using elements that span all linearly independent directions, we are more likely to "straddle" all LCS (cf. discussion below), than when using a single line element at each grid point.

In Section 3.1, we showed that the FTLE defines the average logarithmic expansion rate about a trajectory for *infinitesimal* perturbations. Likewise, the entries in Eq. (3.17) should be defined from infinitesimal elements to determine the deformation gradient. In the computation of FTLE, the neighbors in the FTLE grid itself are typically used for finite differencing. Using this strategy, the grid should, in theory, be highly resolved. For most applications, the resolution needed would, in fact, be cost prohibitive. The exponential growth we seek to reveal quickly leads to "overflow" as the integration proceeds (cf. Winkler [28], Section 5.4 for estimates) and we would need to integrate more material points than is often practically possible to remain under theoretical limits. Alternatively, at each grid point, one can consider the advection of local perturbations that are much smaller in magnitude than the FTLE

grid spacing. In theory, these perturbations could be made small enough so that arbitrary accuracy in the FTLE computation is achieved. Similarly, one can rescale the perturbation as it grows too large – a traditional approach in computing Lyapunov exponents. However, the actual value of the FTLE we compute is not our primary motivation – the computation of LCS is our primary motivation; thankfully so, since in most cases, computing FTLE by finite differencing of the FTLE mesh does not accurately capture the theoretical FTLE values (cf. Section 2.3 of the work by Tabor and Klapper [52], and references therein, for discussions on the stretching of finite elements).

It is in fact the "straddling" of these manifolds that appears more important (which is likely why "less rigorous" stretching measures often reveal LCS); for example, refer back to the FTLE field shown in Figure 3.2. Notice that the value of FTLE can quickly decrease away from the LCS. Therefore, even if we perfectly computed the FTLE values over a coarse grid using infinitesimal perturbations, we may not likely see any ridges in the FTLE field since generally we would not expect the grid points to lie on, or sufficiently close to, the LCS. However, by differencing the computational grid as outlined above, LCS that lie between grid points will be detected, even for relatively coarse meshes. Consequently, if one is interested in knowing the approximate location of the LCS, then a coarse grid can be used to obtain the approximate LCS location, and then the FTLE grid can be adaptively refined near LCS to iteratively improve the location estimates. We originally discussed these ideas in [53]. Sadlo *et al.* [54, 55] implemented a grid adaptation method that targeted FTLE ridges; however, they used a more targeted approach that, instead of refining results from a coarse FTLE grid, used other *a priori* knowledge of the flow as a starting point for refinement. Similarly, Garth *et al.* [56] implemented an approach similar in spirit where grid refinement was based on accuracy of flow map determination, rather than FTLE ridges (note, though, the flow map is, in a sense, most sensitive near LCS).

If the grid spacing is too large, then neighboring grid points can span multiple LCS, or portions of the same LCS, and results become uncorrelated with any particular LCS. More generally, as the grid spacing increases, the FTLE field becomes uncorrelated with the actual flow topology. As grid spacing decreases, LCS detection is improved; however, computational cost grows exponentially. There is a common misconception that the FTLE grid need not be more resolved than the velocity grid used to drive the computations. However, there should be a distinction between Lagrangian and Eulerian information. Certainly, two points starting inside the same element of the velocity mesh are not expected to have the same dynamics. Indeed, for example, it is certainly possible for more than one LCS to lie between velocity grid points. Generally, the FTLE mesh is more highly resolved than the velocity mesh.

3.4.2
Integration Time

The Lagrangian perspective is of central importance to the study of LCS, which bases our analysis on the behavior of the fluid over time. When computing FTLE, or similar, an integration time should be chosen that is long enough, so the dominant flow

features have a chance to emerge, but not so long that the final positions used to compute the deformation gradient are uncorrelated to the flow features we are trying to expose back at the initial time the material points were released. The choice of integration time used to compute FTLE is often *ad hoc*, and in many applications, this is little concern. In general, LCS can shrink, grow, appear, and disappear with changes in integration time. Typically, LCS are robust and long-lived (compared to Eulerian time scales), and thus variations of the FTLE topography with integration time are gradual. Specifically, the following generalizations often hold: (1) As integration time is initially increased, FTLE distribution tends to sharpen along LCS and the exposed length of LCS tends to grow; the growth often results from the FTLE computation tracking how the influence of hyperbolic hubs extend out into the domain. (2) The exposed location of a particular LCS in the FTLE field is typically not sensitive to changes in integration time (except perhaps for relatively short integration times when LCS are not well-defined). That is, holding t_0 fixed in the FTLE computation, the identified location of LCS should remain fixed if the LCS is a material surface.

Consider again Figure 3.3a. The particles appear to exponentially diverge near the hyperbolic fixed point. The further away the particles are from the fixed point along the stable manifold, the longer it takes for them to reach, and be significantly stretched apart by, that hub. Considering LCS as a stable manifold of a hyperbolic trajectory, stretching at a particular section of the LCS may not become pronounced until one considers a sufficiently long time horizon. For example, Figure 3.5 shows the variation of the FTLE with integration time for the double-gyre system considered in [32] given by the stream function

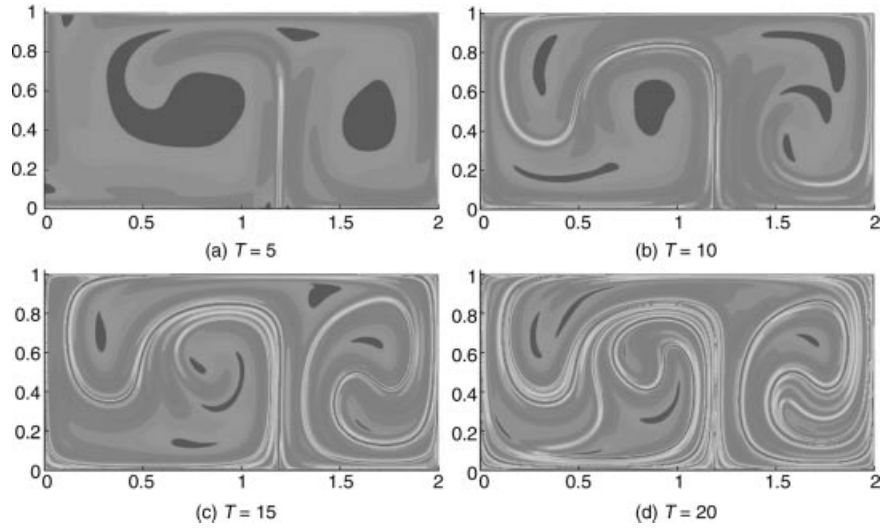

Figure 3.5 The double-gyre FTLE field at $t = 0$ for $A = 0.1$, $\omega = 2\pi$, and $= 0.25$ for increasing integration times T.

$$\gamma(x, y, t) = A \sin(\pi f(x, t)) \sin(\pi y) \tag{3.19}$$

where

$$\begin{aligned} f(x, t) &= a(t)x^2 + b(t)x \\ a(t) &= \varepsilon \sin(\omega t) \\ b(t) &= 1 - 2\varepsilon \sin(\omega t) \end{aligned} \tag{3.20}$$

over the domain $[0, 2] \times [0, 1]$. There is a stable manifold attached to the bottom boundary and as the integration time is increased, more of the manifold is revealed in the FTLE plot.

One may notice in Figure 3.5d that a segment of the LCS close to the hyperbolic trajectory has become faint, which is not a rendering artifact. This may seem paradoxical; however, material points that were released from this location have started to reach a condition where their final locations are "uncorrelated" with the local flow topology. A higher grid resolution is needed for the averaged growth rate over this integration time to be representative of the local flow topology at this location. Note that this occurs mainly for systems for which material points are likely to revisit the same region; for example, one does not readily observe this phenomenon for the vortex example (Figure 3.2), which is why we choose instead to show the variation of T for the double gyre flow.

Structured approaches to determining adequate integration times have been proposed. Winkler [28] used ideas from Boffetta *et al.* [57] to predict appropriate integration times for various examples in computing RD to restrict stretching to the "exponential regime," thereby limiting the influence of longer term diffusive behavior. In the double-gyre flow, the hyperbolic trajectory on the bottom boundary persists for all time. In most practical applications, hyperbolic trajectories may exist only over finite time intervals, and similarly for their stable and unstable manifolds. Therefore, as the integration time for the FTLE becomes large, it is possible that LCS existing over various subsets of time begin to "cloud" the FTLE field, and subsequently our interpretation of the flow. Lekien and Leonard [58] provided a structured approach to filter out appropriate time scales from the dynamics to determine an appropriate integration time. As mentioned toward the end of Section 3.2, the direction of dominant expansion/contraction converges to the normal direction for a normally hyperbolic LCS; in fact this is a necessary condition [22]. Therefore, the time needed for this alignment to converge (up to some ε) can be taken as a lower bound for the integration time needed to identify normally hyperbolic LCS.

From the practical view, the velocity data is only known over a finite interval of time and, for nonperiodic flow, the integration of trajectories used to compute FTLE cannot proceed beyond the timespan of the data set. Furthermore, unless all boundaries of the spatial domain of the data set are invariant, material points used to compute the FTLE field are often flushed from the domain before the desired integration time is reached, unless the open boundaries are sufficiently far from the region one is computing FTLE. Quite frequently, practical limits are imposed on the

integration time length before dynamics-based limits become relevant, especially when computation and measurement of the velocity data is performed with Eulerian features in mind. To summarize, integration time is usually determined by grid spacing (computational cost), time scale of the dynamics, and/or availability of data.

In the FTLE computation, material points that leave the domain before the desired integration time can be treated in a number of ways. Often it is not possible to extrapolate the velocity field outside the bounds of the data set without a high degree of arbitrariness. For material points that exit early, it is common to compute the FTLE for that point, and due to the finite-differencing strategy (cf. Eq. (3.17)), at neighboring points in the FTLE grid.[8)] A drawback is that various parts of the domain can end up using different integration times to compute their FTLE value, and notably, the stretching rate (logarithmic or otherwise) is typically not constant over a trajectory in finite time. Therefore, this strategy can result in spurious structures in the FTLE field, for example, marking boundaries to portions of fluid that have exited the domain. Also, in the author's experience, applications where material points are likely to exit during FTLE computation, the $1/|T|$ scaling in the definition of the FTLE can often amplify the FTLE values of points that leave within a small integration time and it is often undesirable to include this scaling. Conversely, in applications where points do not leave the domain, the $1/|T|$ scaling is unnecessary, as it only shifts all values uniformly.

Along with spurious LCS that can result from open domains, limited data sets can lead to dominant coherent structures failing to emerge before particles exit the domain. In such cases extrapolation is often desirable, even if arbitrary. For example, in applications of studying LCS in heart flow problems, we have often extended the velocity field information by creating linear fields outside of the domain of data, for example, to represent blood pumped out of the heart. Tang *et al.* [59] developed a similar approach of extending atmospheric flow data using linear velocity fields to study clear air turbulence over the Hong Kong airport, and introduced a smoothing to ensure C^1 continuity of the fields. The advantage of linear fields is that they maintain a homogeneous stretching rate outside the domain. Nonetheless, LCS revealed in the original domain of data can be subject to the extrapolated velocity field in nonlinear ways. For example, for neighboring points in the FTLE grid, one point can be advected outside of the domain into the linear field, whereas the other point may remain inside the domain for a significantly longer time. The subsequent stretching is due to a combination of the linear flow with the original nonlinear flow, making the combination nonlinear and the stretching rate nonconstant. This is more of a concern in applications with open and closed boundaries, as opposed to applications with strictly open boundaries. Also, the linear field outside the data domain may look nothing like the actual velocity field outside the data domain. However, the use of linear fields tends to minimize the introduction of additional arbitrary stretching

8) When implemented as such, this excludes "cases when neighboring trajectories all leave the domain earlier than the trajectory considered" [59], since the FTLE would already have been computed for that trajectory. Note that this approach does not stop integration prematurely (as that would cascade through the grid), it just computes FTLE prematurely.

since it holds stretching rates constant once elements leave the domain, or as put by Tang *et al*. [59] tends to "lock-in and enhance" existing FTLE ridges.

As mentioned previously, positive integration times reveal repelling LCS in the FTLE field and negative integration times reveal attracting LCS in the FTLE field. Expansion is nominally easier to compute than contraction; therefore, we reverse time to effectively measure contraction as expansion. In theory however, when we advect a grid of material points, both types of information are encoded. Indeed, Haller and Sapsis [60] recently showed that when the integration time is positive, troughs of the *smallest* FTLE field reveal attracting LCS (or equivalently ridges of the negative of the smallest FTLE). The advantage of this approach is that only a single numerical run is needed, instead of separate numerical runs in forward and backward time. Note, however, that the smallest FTLE field is plotted at the *final* locations of the material points. Therefore, the resulting field is defined over a highly distorted grid, which often must be interpolated to a regular grid for analysis.

3.4.3
LCS Extraction

As presented here, the computation of FTLE fields is often seen as a means to an end; ultimately, we desire to locate LCS to better understand the flow topology. FTLE provides a measure that enables relatively easy and robust determination of LCS location *and* strength. We use FTLE precisely because it highlights these structures; therefore, one could argue that the more we can highlight, or filter out, LCS from the background FTLE field the better. Initial applications of the framework focused on demonstrating the existence and implications of these structures, which indeed was critical to this method gaining acceptance and wider use. Now that the subject has begun to mature, LCS are becoming part of, or have the potential to become part of, broader computational objectives, for example, decision making, design, transport calculations, and optimization. In this light, it is often necessary to have the ability to effectively parameterize or extract LCS from FTLE plots. Finally, LCS filtering or extraction is necessary for effectively visualizing these structures in 3D. Not surprisingly how to best extract LCS is closely related to determining what is the best way to allocate the FTLE computational costs to target LCS.

For 2D systems, we often *visually* filter out LCS by mapping colors (which could include intensities of a particular color) to FTLE values. In such cases, ridges are visually assessed from the color (e.g., Figures 3.2 and 3.5) or intensity (e.g., Figures 3.4 and 3.6) mapping. In general, when dealing with FTLE values themselves, we are temped to associate LCS as sets of globally high FTLE (e.g., threshold out low FTLE values); however, FTLE ridges are defined in terms of spatial derivatives of the FTLE field (Section 3.2) not the actual FTLE values per se. In 3D, visual identification of LCS from unfiltered color or intensity mapping is not possible. One must either extract the LCS surfaces, or use volume rendering to visualize the LCS in the FTLE plots (or reduce to 2D by considering sections). In regards to volume rendering, the primary goal is essentially to skip as much of the volume as possible while retaining the subvolumes containing the LCS. In the most basic sense, volume rendering can be

achieved by mapping the FTLE to color and opacity scales, whereby FTLE becomes transparent as it decreases (thresholding can be seen as a naive implementation). This works well for visualizing the most well-defined LCS. At least in 2D, opacity mappings can also be helpful when showing multiple fields (e.g., superimposing forward and backward time FTLE fields); for example, see work of Garth *et al.* [56] for a systematic procedure that couples FTLE plots with a "texture-based" representation [61] of the flow and leverages GPU to accelerate FTLE computation.

We have previously discussed (e.g., [32, 48, 53]) how LCS can be extracted based on the definition in Section 3.2. Using previous notation, let $\mathbf{v}_{\min}$ be the unit eigenvector associated $\mu_{\min}$, which must be normal to the LCS. Along an LCS $\nabla\Lambda \cdot \mathbf{v}_{\min} = 0$ and therefore potential LCS can be extracted from zero level sets of the field $\alpha = \nabla\Lambda \cdot \mathbf{v}_{\min}$. There is a sign ambiguity with computing $\mathbf{v}_{\min}$ and therefore the sign must be chosen in a locally consistent manner – this is critical for proper ridge extraction. This method has been applied in various applications [48, 53, 62], often when specific LCS were used for transport rate computations, to various degrees of success. A difficulty with this approach is that smooth second derivatives are notoriously difficult to compute from discrete data, and the computation of $\nabla^2\Lambda$ is no exception. Generally, this approach distinguishes *all* ridges of FTLE, notably many of which our eyes would pass over in the FTLE field. Therefore, one looses the relative strength information encoded in the FTLE plot and significant contextual information. Hence, LCS extraction is perhaps more critical when driving additional computations where an exact parameterization is needed than for visual appeal. It is usually desirable to filter spurious ridges or weak LCS by checking additional conditions, such as setting a minimum threshold for Λ or maximum threshold for $\mu_{\min}$; the first indicating strength of hyperbolicity and the second how well-defined the ridge is, which can be tied to Lagrangian [32, 33] and robustness [22] properties. The application of ridge extraction to the mechanically generated vortex is show in Figure 3.6. Unfortunately, this filtering can cause portions of otherwise well-defined LCS to be skipped where the criteria locally fail (presumably due to computational anomalies) as also seen in the figure. A similar ridge extraction approach was adopted by Sadlo *et al.* [54, 55] who also noted troubles with the practical application of this method in noisy data sets and offered some insights on addressing these problems.

Mathur *et al.* [46] introduced a clever way to distinguish FTLE ridges. Essentially, once the FTLE field is computed, material points are seeded near LCS and advected by the FTLE gradient instead of the flow. In this way, the gradient of the FTLE field $\nabla\Lambda$ pushes the points up onto the ridge. Certain criteria check once each point approximately reaches the top of the ridge and the computation for the point is stopped. In other work, Lipinski and Mohseni [63] proposed a ridge-tracking algorithm for efficient LCS computation in 2D. The basic idea of this method is to target the FTLE computation along LCS, and in doing so provide tight estimates of the LCS locations. The approach starts by computing FTLE along lines that intersect the domain. Local maxima of FTLE along these lines are assumed to correspond to intersections with LCS. These locations, and estimates of the local orientation of the

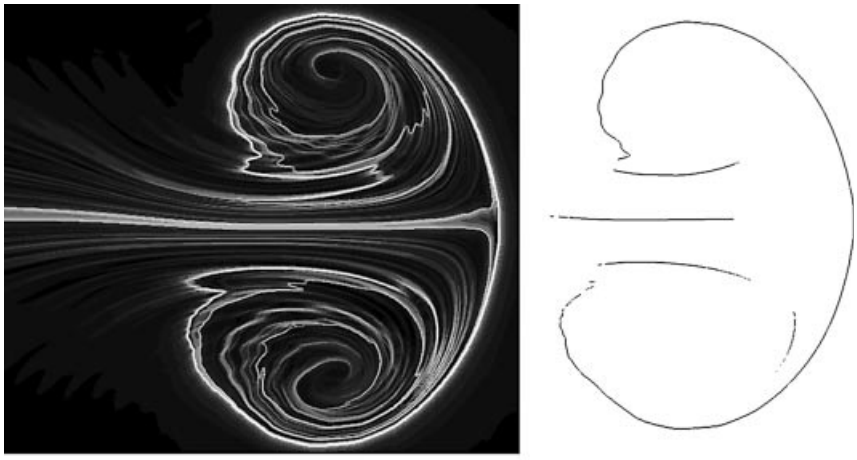

Figure 3.6 Results of ridge extraction from forward time FTLE field of mechanically generated vortex ring. Ridges with $\Lambda < 70\%$ of maximum, and $\mu_{\min} > -4800$ are not shown.

LCS, are used to allocate FTLE mesh points only in the directions of the LCS. FTLE is then computed at the newly seeded points to further the location and orientation estimate, and the process continues sequentially to grow ("track") the location of the LCS in both spatial directions.

We note that in most cases, the extraction of LCS from FTLE fields has been directed at 2D problems, or in 3D problems by considering 2D sections. However, the need for extraction is arguably most compelling in 3D. We also note that since LCS are material surfaces, it has been observed that if one is able to extract LCS at one time instant, the time evolution of the LCS can be obtained by advecting it as a material surface (e.g., this is leveraged in [55, 63, 64]). However, because LCS can shrink, grow, appear, and disappear, it is unclear how broadly useful this is in practice. Greatest utility may be for attracting LCS [64]. That is, high FTLE values indicate where particles are *most* sensitive to changes in initial conditions; therefore, errors in the position estimate for repelling LCS will quickly grow, whereas for attracting LCS they will shrink. More notably though, because of the typical normal expansion about repelling LCS, fluid along the surface tends to compress, decreasing the scope of LCS revealed over time. Presumably, one may consider the advection of a repelling LCS *backward* in time to efficiently display past locations of repelling LCS, or advect attracting LCS as material surfaces forward in time to display future locations. However, in real-time applications, future positions of the repelling LCS are most sought after since these structures require future knowledge to obtain. That is, if data is known up to the real time t and an integration time of $|T|$ is used for FTLE computation, repelling LCS can only be computed up to at past time $t_0 = t - |T|$, using a traditional approach. Hence, for a realtime estimate, the material surface must be advected forward in time, subject to the prior mentioned difficulties.

3.5 Robustness

The computation of LCS is typically derived from trajectory information obtained by processing Eulerian velocity data. The velocity data often contains inherent modeling or measurement errors. Furthermore, since the velocity data is discrete, the integration process must fill the gaps in space and time with an interpolated field, which introduces additional deviations from the true field. Even if these errors are small, they can accumulate during integration and lead to significant errors in the final position estimates used to compute the deformation gradient. Thus, one may wonder if the location, or even existence, of a particular LCS holds for the true system, or just the "model" system? This is a particularly difficult question to answer since one typically does not know the exact nature of the errors involved. One can approach studying this problem in a number of ways. Theoretically, one can choose a certain class of errors and derive rigorous estimates for the robustness of LCS. Empirically, one can systematically introduce known differences between the true and model fields and quantify changes in the results, or compare results from realizations with unknown but distinct errors. Finally, one can compare results of LCS computations with independent measurements indicating LCS location from the true system. All of these have been done. In general, it has been shown that LCS are surprisingly robust considering the integrated error influence.

In Ref. [31], Haller showed that if the model velocity field is close in a time-weighted norm to the true velocity field, then a hyperbolic material line of the model system implies a hyperbolic material line in the true system. The weighted norm is such that it permits relatively large errors as long as they are local in time – those that do not accumulate to large errors when integrated over time. Haller showed, as one might otherwise assume, that strongly repelling/attracting and/or long-lived LCS are the most robust, which more recently [22] was translated into conditions on FTLE ridges, which essentially imply that the ridge should be well-defined. Olcay *et al.* [65] investigated how velocity field resolution and random errors affect LCS. In their approach, they degraded a vortical flow field by locally smoothing the flow and sampling it on spatially sparser grids, subsampling the data in time, and adding Gaussian noise. They found that mean LCS locations were robust to these variations in the velocity data, but locally the LCS could significantly deviate from its baseline location. In all cases, the existence of the LCS was robust. Shadden *et al.* [62] compared LCS computed from radar measurements of ocean surface currents with GPS-tracked surface buoy data. They showed that even though the highly noisy radar data led to errors in trajectory computation that approach 20–30 km over the integration interval used to compute FTLE, the error in location of the LCS appeared to be on the order of only a few kilometers. More recently, Harrison and Glatzmaiera [66], in studying the California Current system, compared LCS computed from a mesoscale eddy field generated from Fourier filtering satellite altimetry observations and from the full observational satellite altimetry. They showed that the LCS found were relatively insensitive to both sparse spatial and temporal resolution, and to the velocity filtering.

3.6 Applications

The computation of LCS has proven to be effective in wide range of fluid mechanics problems. Applications of this method have ranged from microfluidic to geophysical and from laminar to turbulent. In addition to this wide scope, the recent growth in popularity of this approach is also likely due in large part to this method being relatively straightforward to implement and hence highly accessible. Here, we briefly list some representative publications describing the computation of LCS to study practical fluid transport problems. In many of these publications, unique perspectives on LCS computations, as well as their relevance to understanding transport processes in the respective application, are presented.

Transport in the ocean and atmosphere has wide-ranging importance. Historically, the computation of stable and unstable manifolds has seen preference in this area, and not surprising the earliest computations of LCS using FTLE fields focused on these application areas [23, 25, 26, 29]. More recent studies of atmospheric mixing include the works of Refs. [47, 67, 68] and studies of the polar vortices include the works of Refs. [69–71]. In regards to ocean mixing, publications include Refs. [62, 66, 72–79]. Beron-Vera, Olascoaga *et al.* [80–86] have provided a number of fruitful applications of this method to marine science. Tang and Peacock explored the implications of the approach for locating internal ocean waves [87]; several groups have used this approach to studying transport associated with hurricanes/cyclones [8, 14, 88, 89] and to study transport of ocean [90] and atmospheric [91] contaminants.

The use of LCS has also been used extensively to study vortex rings [42, 92–94], as well as hairpin vortices in turbulent channel flow [95], and vortex shedding [96–99]. In aeronautics, LCS has been used to study clear air turbulence [100] and flow separation [32, 101]. In the realm of biofluids, LCS applications include studies of respiration [102], animal swimming [5, 103–110], and cardiovascular fluid mechanics [48, 111–113]. Finally, several studies have used LCS to understand the mechanics of fluid mixing devices [46, 53, 114].

3.7 Conclusions

LCS are core surfaces that organize fluid transport. LCS computation arose as a means to compute effective stable and unstable manifolds of hyperbolic trajectories in temporally aperiodic systems, which was motivated by previous discoveries that such manifolds organized transport in temporally periodic flows. Indeed, in periodic and aperiodic systems alike, these structures often govern the stretching, folding, and alignment mechanisms that control kinematic mixing. As the application of this framework branched out to increasingly complex problems, LCS have become more broadly associated with distinguished material surfaces. Notably, while the entirety of a fluid is composed of a continuum of material surfaces, those with locally the strongest stability (attracting or repelling) are often influential in dominating

advection patterns. The computation of LCS has enabled the ability to better perceive and appreciate this influence. To date, LCS are effectively identified by plotting Lagrangian-based measures of hyperbolicity, such as the FTLE. We have shown how such computations proceed in practice, and a number of theoretical and practical concerns associated with this approach. Wide-ranging applications have shown the utility of LCS toward enabling deeper understanding of transport-related phenomena over a wide range of spatial scales and temporal variability. Continued theoretical, computational, and application-oriented developments in this regard will hopefully further enable greater comprehension of the nature of unsteady fluid transport and its implications.

Acknowledgments

This chapter has benefited from the keen insight of George Haller, who provided valuable recommendations, and more generally helped pave the way for LCS computation.

References

1 Wiggins, S. (1994) *Normally Hyperbolic Invariant Manifolds in Dynamical Systems*, Springer, New York.
2 Ottino, J.M. (1989) The kinematics of mixing: Stretching, chaos, and transport, in *Cambridge Texts in Applied Mathematics*, Cambridge University Press, Cambridge; New York, J.M. Ottino. ill.; 24 cm.
3 Wiggins, S. (1992) *Chaotic Transport in Dynamical Systems*, Springer, New York.
4 Aref, H. (2002) The development of chaotic advection. *Physics of Fluids*, **14** (4), 1315–1325.
5 Peng, J. and Dabiri, J.O. (2009) Transport of inertial particles by Lagrangian coherent structures: Application to predator-prey interaction in jellyfish feeding. *Journal of Fluid Mechanics*, **623**, 75–84.
6 Haller, G. and Sapsis, T. (2008) Where do inertial particles go in fluid flows? *Physica D*, **237**, 573–583.
7 Sapsis, T. and Haller, G. (2010) Clustering criterion for inertial particles in two-dimensional time-periodic and three-dimensional steady flows. *CHAOS*, **20** (1), 017515.
8 Sapsis, T. and Haller, G. (2009) Inertial particle dynamics in a hurricane. *Journal of the Atmospheric Sciences*, **66** (8), 2481–2492.
9 Tallapragada, P. and Ross, S.D. (2008) Particle segregation by Stokes number for small neutrally buoyant spheres in a fluid. *Physical Review E*, **78** (3), Part 2).
10 Miller, P.D., Jones, C.K.R.T., Rogerson, A.M., and Pratt, L.J. (1997) Quantifying transport in numerically generated velocity fields. *Physica D*, **110** (1–2), 105–122.
11 Haller, G. and Poje, A.C. (1998) Finite time transport in aperiodic flows. *Physica D*, **119** (3–4), 352–380.
12 Ide, K., Small, D., and Wiggins, S. (2002) Distinguished hyperbolic trajectories in time-dependent fluid flows: Analytical and computational approach for velocity fields defined as data sets. *Nonlinear Processes in Geophysics*, **9**, 237–263.
13 Haller, G. and Yuan, G. (2000) Lagrangian coherent structures and mixing in two-dimensional turbulence. *Physica D*, **147** (3–4), 352–370.
14 du Toit, P. and Marsden, J.E. (2010) Horseshoes in hurricanes. *Journal of*

Fixed Point Theory and Applications, **7**, 351–384.

15 Romkedar, V., Leonard, A., and Wiggins, S. (1990) An analytical study of transport, mixing and chaos in an unsteady vortical flow. *Journal of Fluid Mechanics*, **214**, 347–394.

16 Romkedar, V. (1990) Transport rates of a class of 2-dimensional maps and flows. *Physica D*, **43** (2–3), 229–268.

17 Fenichel, N. (1971) Persistence and smoothness of invariant manifolds for flows. *Indiana University Mathematics Journal*, **21**, 193–225.

18 Mancho, A.M., Small, D., Wiggins, S., and Ide, K. (2003) Computation of stable and unstable manifolds of hyperbolic trajectories in two-dimensional, aperiodically time-dependent vector fields. *Physica D*, **182** (3–4), 188–222.

19 Duc, L.H. and Siegmund, S. (2008) Hyperbolicity and invariant manifolds for planar nonautonomous systems on finite time intervals. *International Journal of Bifurcation and Chaos*, **18** (3), 641–674.

20 Froyland, G., Padberg, K., England, M.H., and Treguier, A.M. (2007) Detection of coherent oceanic structures via transfer operators. *Physical Review Letters*, **98**, 224503.

21 Froyland, G. and Padberg, K. (2009) Almost invariant sets and invariant manifold-connecting probabilistic and geometric descriptions of coherent structures in flows. *Physica D*, **238**, 1507–1523.

22 Haller, G. (2011) A variational theory of hyperbolic Lagrangian coherent structures. *Physica D*, **240**, 574–598.

23 Bowman, K. (1999) Manifold geometry and mixing in observed atmospheric flows, http://geotest.tamu.edu/userfiles/213/manifold_geometry.pdf.

24 Jones, C.K.R.T. and Winkler, S. (2002) Invariant manifolds and lagrangian dynamics in the ocean and atmosphere, in *Handbook of Dynamical Systems*, vol. 2 (ed. B. Fiedler), Elsevier, Amsterdam, pp. 55–92.

25 Pierrehumbert, R.T. (1991) Large-scale horizontal mixing in planetary-atmospheres. *Physics of Fluids A*, **3** (5), 1250–1260.

26 Pierrehumbert, R.T. and Yang, H. (1993) Global chaotic mixing on isentropic surfaces. *Journal of the Atmospheric Sciences*, **50** (15), 2462–2480.

27 Doerner, R., Hubinger, B., Martienssen, W., Grossmann, S., and Thomae, S. (1999) Stable manifolds and predictability of dynamical systems. *Chaos Solitons & Fractals*, **10** (11), 1759–1782.

28 Winkler, S. (2001) Lagrangian dynamics in geophysical fluid flows. Ph.D., Brown University.

29 von Hardenberg, J., Fraedrich, K., Lunkeit, F., and Provenzale, A. (2000) Transient chaotic mixing during a baroclinic life cycle. *CHAOS*, **10** (1), 122–134.

30 Haller, G. (2001) Distinguished material surfaces and coherent structures in three-dimensional fluid flows. *Physica D*, **149** (4), 248–277.

31 Haller, G. (2002) Lagrangian coherent structures from approximate velocity data. *Physics of Fluids*, **14** (6), 1851–1861.

32 Shadden, S.C., Lekien, F., and Marsden, J.E. (2005) Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows. *Physica D*, **212** (3–4), 271–304.

33 Lekien, F., Shadden, S.C., and Marsden, J.E. (2007) Lagrangian coherent structures in *n*-dimensional systems. *Journal of Mathematical Physics*, **48** (6), 065404.

34 Haller, G. (2000) Finding finite-time invariant manifolds in two-dimensional velocity fields. *CHAOS*, **10** (1), 99–108.

35 Haller, G. (2001) Lagrangian structures and the rate of strain in a partition of two-dimensional turbulence. *Physics of Fluids A*, **13**, 3368–3385.

36 Aurell, E., Boffetta, G., Crisanti, A., Paladin, G., and Vulpiani, A. (1997) Predictability in the large: An extension of the concept of Lyapunov exponent. *Journal of Physics A – Mathematical and General*, **30** (1), 1–26.

37 Boffetta, G., Lacorata, G., Redaelli, G., and Vulpiani, A. (2001) Detecting barriers to transport: A review of different techniques. *Physica D*, **159**, 58–70.

38 Joseph, B. and Legras, B. (2002) Relation between kinematic boundaries, stirring, and barriers for the antarctic polar vortex. *Journal of the Atmospheric Sciences*, **59** (7), 1198–1212.

39 Koh, T.Y. and Legras, B. (2002) Hyperbolic lines and the stratospheric polar vortex. *CHAOS*, **12** (2), 382–394.

40 Lekien, F. (2003) Time-dependent dyanmical systems and geophysical flows Ph.D., California Institute of Technology.

41 Branicki, M. and Wiggins, S. (2010) Finite-time Lagrangian transport analysis: stable and unstable manifolds of hyperbolic trajectories and finite-time Lyapunov exponents. *Nonlinear Processes in Geophysics*, **17** (1), 1–36.

42 Shadden, S.C., Dabiri, J.O., and Marsden, J.E. (2006) Lagrangian analysis of fluid transport in empirical vortex ring flows. *Physics of Fluids*, **18** (4), 047105.

43 Shadden, S.C., Katija, K., Rosenfeld, M., Marsden, J.E., and Dabiri, J.O. (2007) Transport and stirring induced by vortex formation. *Journal of Fluid Mechanics*, **593**, 315–331.

44 Wiggins, S. (2005) The dynamical systems approach to Lagrangian transport in oceanic flows. *Annual Review of Fluid Mechanics*, **37**, 295–328.

45 Coppel, W.A. (1978) *Dichotomies in Stability Theory*, Springer, Berlin; New York.

46 Mathur, M., Haller, G., Peacock, T., Ruppert-Felsot, J.E., and Swinney, H.L. (2007) Uncovering the Lagrangian skeleton of turbulence. *Physical Review Letters*, **98** (14), 144502.

47 Beron-Vera, F.J., Olascoaga, M.J., Brown, M.G., Kocak, H., and Rypina, I.I. (2010) Invariant-tori-like Lagrangian coherent structures in geophysical flows. *CHAOS*, **20** (1), 017514.

48 Shadden, S.C., Astorino, M., and Gerbeau, J.-F. (2010) Computational analysis of an aortic valve jet with Lagrangian coherent structures. *CHAOS*, **20** (1), 017512.

49 Jimenez, R. (2011) CUDA LCS Software http://www.its.caltech.edu/raymondj/LCS/.

50 du Toit, P., Newman, http://www.cds.caltech.edu/pdutoit/Philip_Du_Toit/Software.html (Last access Nov. 12, 2010).

51 Brunton, S.L. and Rowley, C.W. (2010) Fast computation of finite-time Lyapunov exponent fields for unsteady flows. *CHAOS*, **20** (1), 017503.

52 Tabor, M. and Klapper, I. (1994) Stretching and alignment in chaotic and turbulent flows. *Chaos, Solitons & Fractals*, **4**, 1031–1055.

53 Shadden, S.C. (2006) A dynamical systems approach to unsteady systems Ph.D., California Institute of Technology.

54 Sadlo, F. and Peikert, R. (2007) efficient visualization of lagrangian coherent structures by filtered AMR ridge extraction. *IEEE Transactions on Visualization and Computer Graphics*, **13** (5), 1456–1463.

55 Sadlo, F., Rigazzi, A., and Peikert, R. (2010) Time-dependent visualization of lagrangian coherent structures by grid advection, in *Topological Methods in Data Analysis and Visualization* (eds. V. Pascucci, X. Tricoche, H. Hagen, and J. Tierny), Springer, Berlin, pp. 151–165.

56 Garth, C., Gerhardt, F., Tricoche, X., and Hagen, H. (2007) Efficient computation and visualization of coherent structures in fluid flow applications. *IEEE Transactions on Visualization and Computer Graphics*, **13** (6), 1464–1471.

57 Boffetta, G., Celani, A., Cencini, M., Lacorata, G., and Vulpiani, A. (2000) Nonasymptotic properties of transport and mixing. *CHAOS*, **10** (1), 50–60.

58 Lekien, F. and Leonard, N. (2004) *Dynamically consistent lagrangian coherent structures*, in *8th Experimental Chaos Conference, volume 742 of AIP Conference Proceedings* (eds. S. Boccaletti, B.J. Gluckman, J. Kurths, L.M. Pecora, R. Meucci, and O. Yordanov), Florence, Italy, pp. 132–139.

59 Tang, W., Chan, P.W., and Haller, G. (2010) Accurate extraction of Lagrangian coherent structures over finite domains with application to flight data analysis over Hong Kong International Airport. *CHAOS*, **20** (1), 017502.

60 Haller, G. and Sapsis, T. (2011) Lagrangian coherent structures and the smallest finite-time Lyapunov exponent. *CHAOS* (preprint).

61 Li, G.S., Tricoche, X., and Hansen, C.D. (2006) GPUFLIC: Interactive and accurate dense visualization of unsteady flows, in *Eurographics/ IEEE-VGTC Symposium on Visualization* (eds. T. Ertl, K. Joy, and B. Santo), The Eurographics Association, pp. 29–34.

62 Shadden, S.C., Lekien, F., Paduan, J.D., Chavez, F.P., and Marsden, J.E. (2009) The correlation between surface drifters and coherent structures based on high-frequency radar data in Monterey Bay. *Deep-Sea Research Part Ii-Topical Studies in Oceanography*, **56** (3–5), 161–172.

63 Lipinski, D. and Mohseni, K. (2010) A ridge tracking algorithm and error estimate for efficient computation of Lagrangian coherent structures. *CHAOS*, **20** (1), 017504.

64 Ferstl, F., Buerger, K., Theisel, H., and Westermann, R. (2010) Interactive separating streak surfaces. *IEEE Transactions on Visualization and Computer Graphics*, **16** (6), 1569–1577.

65 Olcay, A.B., Pottebaum, T.S., and Krueger, P.S. (2010) Sensitivity of Lagrangian coherent structure identification to flow field resolution and random errors. *CHAOS*, **20** (1)

66 Harrison, C.S. and Glatzmaier, G.A. (2011) Lagrangian coherent structures in the California current system - sensitivities and limitations. *Geophysical and Astrophysical Fluid Dynamics*, doi: 10.1080/03091929.2010.532793.

67 Duran-Matute, M. and Velasco Fuentes, O.U. (2008) Passage of a barotropic vortex through a gap. *Journal of Physical Oceanography*, **38** (12), 2817–2831.

68 Resplandy, L., Levy, M., d'Ovidio, F., and Merlivat, L. (2009) Impact of submesoscale variability in estimating the air-sea CO_2 exchange: Results from a model study of the POMME experiment. *Global Biogeochemical Cycles*, **23**, GB1017.

69 Rypina, I.I., Brown, M.G., and Beron-Vera, F.J. (2007) On the lagrangian dynamics of atmospheric zonal jets and the permeability of the stratospheric polar vortex. *Journal of the Atmospheric Sciences*, **64** (10), 3595–3610.

70 Lekien, F. and Ross, S.D. (2010) The computation of finite-time Lyapunov exponents on unstructured meshes and for non-Euclidean manifolds. *CHAOS*, **20** (1), 017505.

71 de la Camara, A., Mechoso, C.R., Ide, K., Walterscheid, R., and Schubert, G. (2010) Polar night vortex breakdown and large-scale stirring in the southern stratosphere. *Climate Dynamics*, **35** (6), 965–975.

72 Lekien, F. and Coulliette, C. (2007) Chaotic stirring in quasi-turbulent flows. *Philosophical Transactions of the Royal Society A – Mathematical Physical and Engineering Sciences*, **365** (1861), 3061–3084.

73 d'Ovidio, F., Isern-Fontanet, J., Lopez, C., Hernandez-Garcia, E., and Garcia-Ladon, E. (2009) Comparison between Eulerian diagnostics and finite-size Lyapunov exponents computed from altimetry in the Algerian basin. *Deep-Sea Research Part I-Oceanographic Research Papers*, **56** (1), 15–31.

74 Rypina, I.I., Brown, M.G., and Kocak, H. (2009) Transport in an idealized three-gyre system with application to the adriatic sea. *Journal of Physical Oceanography*, **39** (3), 675–690.

75 Gildor, H., Fredj, E., Steinbuck, J., and Monismith, S. (2009) Evidence for submesoscale barriers to horizontal mixing in the ocean from current measurements and aerial photographs. *Journal of Physical Oceanography*, **39** (8), 1975–1983.

76 Rossi, V., Lopez, C., Hernandez-Garcia, E., Sudre, J., Garcon, V., and Morel, Y. (2009) Surface mixing and biological activity in the four Eastern Boundary Upwelling Systems. *Nonlinear Processes in Geophysics*, **16** (4), 557–568.

77 Branicki, M. and Malek-Madani, R. (2010) Lagrangian structure of flows in the Chesapeake Bay: challenges and

perspectives on the analysis of estuarine flows. *Nonlinear Processes in Geophysics*, **17** (2), 149–168.

78 Carlson, D.F., Fredj, E., Gildor, H., and Rom-Kedar, V. (2010) Deducing an upper bound to the horizontal eddy diffusivity using a stochastic Lagrangian model. *Environmental Fluid Mechanics*, **10** (5), 499–520.

79 Rypina, I.I., Pratt, L.J., Pullen, J., Levin, J., and Gordon, A.L. (2010) Chaotic advection in an archipelago. *Journal of Physical Oceanography*, **40** (9), 1988–2006.

80 Beron-Vera, F.J., Olascoaga, M.J., and Goni, G.J. (2010) Surface ocean mixing inferred from different multisatellite altimetry measurements. *Journal of Physical Oceanography*, **40** (11), 2466–2480.

81 Beron-Vera, F.J. (2010) Mixing by low- and high-resolution surface geostrophic currents. *Journal of Geophysical Research-Oceans*, **115**.

82 Reniers, A.J.H.M., MacMahan, J.H., Beron-Vera, F.J., and Olascoaga, M.J. (2010) Rip-current pulses tied to Lagrangian coherent structures. *Geophysical Research Letters*, **37**, L05605.

83 Beron-Vera, F.J. and Olascoaga, M.J. (2009) An assessment of the importance of chaotic stirring and turbulent mixing on the West Florida Shelf. *Journal of Physical Oceanography*, **39** (7), 1743–1755.

84 Olascoaga, M.J., Beron-Vera, F.J., Brand, L.E., and Kocak, H. (2008) Tracing the early development of harmful algal blooms on the West Florida Shelf with the aid of Lagrangian coherent structures. *Journal of Geophysical Research-Oceans*, **113** (C12), C12014.

85 Beron-Vera, F.J., Olascoaga, M.J., and Goni, G.J. (2008) Oceanic mesoscale eddies as revealed by Lagrangian coherent structures. *Geophysical Research Letters*, **35** (12), L12603.

86 Olascoaga, M.J., Rypina, I.I., Brown, M.G., Beron-Vera, F.J., Kocak, H., Brand, L.E., Halliwell, G.R., and Shay, L.K. (2006) Persistent transport barrier on the West Florida Shelf. *Geophysical Research Letters*, **33** (22), L22603.

87 Tang, W. and Peacock, T. (2010) Lagrangian coherent structures and internal wave attractors. *CHAOS*, **20** (1), 017508.

88 Rutherford, B., Dangelmayr, G., Persing, J., Schubert, W.H., and Montgomery, M.T. (2010) Advective mixing in a nondivergent barotropic hurricane model. *Atmospheric Chemistry and Physics*, **10** (2), 475–497.

89 Rutherford, B., Dangelmayr, G., Persing, J., Kirby, M., and Montgomery, M.T. (2010) Lagrangian mixing in an axisymmetric hurricane model. *Atmospheric Chemistry and Physics*, **10** (14), 6777–6791.

90 Coulliette, C., Lekien, F., Paduan, J.D., Haller, G., and Marsden, J.E. (2007) Optimal pollution mitigation in monterey bay based on coastal radar data and nonlinear dynamics. *Environmental Science & Technology*, **41** (18), 6562–6572.

91 Tang, W., Haller, G., Baik, J.-J., and Ryu, Y.-H. (2009) Locating an atmospheric contamination source using slow manifolds. *Physics of Fluids*, **21** (4), 043302.

92 Shadden, S.C., Katija, K., Rosenfeld, M., Marsden, J.E., and Dabiri, J.O. (2007) Transport and stirring induced by vortex formation. *Journal of Fluid Mechanics*, **593**, 315–331.

93 Olcay, A.B. and Krueger, P.S. (2008) Measurement of ambient fluid entrainment during vortex ring formation. *Experiments in Fluids*, **44**, 235–247.

94 Olcay, A.B. and Krueger, P.S. (2008) Measurement of ambient fluid entrainment during laminar vortex ring formation. *Experiments in Fluids*, **44** (2), 235–247.

95 Green, M.A., Rowley, C.W., and Haller, G. (2007) Detection of Lagrangian coherent structures in three-dimensional turbulence. *Journal of Fluid Mechanics*, **572**, 111–120.

96 Lipinski, D., Cardwell, B., and Mohseni, K. (2008) A Lagrangian analysis of a two-dimensional airfoil with vortex shedding. *Journal of Physics A-*

Mathematical and Theoretical, **41** (34), 344011.

97 Eldredge, J.D. and Chong, K. (2010) Fluid transport and coherent structures of translating and flapping wings. *CHAOS*, **20** (1), 017509.

98 Yang, A., Jia, L., and Yin, X. (2010) Experimental study of a starting vortex ring generated by a thin circular disk. *Journal of Bionic Engineering*, **7** (Suppl S), S103–S108.

99 Cardwell, B.M. and Mohseni, K. (2008) Vortex shedding over a two-dimensional airfoil: Where the particles come from. *AIAA Journal*, **46** (3), 545–547.

100 Tang, W., Mathur, M., Haller, G., Hahn, D.C., and Ruggiero, F.H. (2010) Lagrangian coherent structures near a subtropical jet stream. *Journal of the Atmospheric Sciences*, **67** (7), 2307–2319.

101 Ruiz, T., Boree, J., Tran, T., Sicot, T.C., and Brizzi, L.E. (2010) Finite time Lagrangian analysis of an unsteady separation induced by a near wall wake. *Physics of Fluids*, **22** (7), 075103.

102 Lukens, S., Yang, X., and Fauci, L. (2010) Using Lagrangian coherent structures to analyze fluid mixing by cilia. *CHAOS*, **20** (1), 017511.

103 Peng, J., Dabiri, J.O., Madden, P.G., and Lauder, G.V. (2007) Non-invasive measurement of instantaneous forces during aquatic locomotion: A case study of the bluegill sunfish pectoral fin. *Journal of Experimental Biology*, **210** (4), 685–698.

104 Peng, J. and Dabiri, J.O. (2007) A potential-flow, deformable-body model for fluid–structure interactions with compact vorticity: Application to animal swimming measurements. *Experiments in Fluids*, **43** (5), 655–664.

105 Franco, E., Pekarek, D.N., Peng, J., and Dabiri, J.O. (2007) Geometry of unsteady fluid transport during fluid–structure interactions. *Journal of Fluid Mechanics*, **589**, 125–145.

106 Peng, J. and Dabiri, J.O. (2008) An overview of a Lagrangian method for analysis of animal wake dynamics. *Journal of Experimental Biology*, **211** (2), 280–287.

107 Peng, J. and Dabiri, J.O. (2008) The 'upstream wake' of swimming and flying animals and its correlation with propulsive efficiency. *Journal of Experimental Biology*, **211** (16), 2669–2677.

108 Lipinski, D. and Mohseni, K. (2009) Flow structures and fluid transport for the hydromedusae Sarsia tubulosa and Aequorea victoria. *Journal of Experimental Biology*, **212** (15), 2436–2447.

109 Wilson, M.M., Peng, J., Dabiri, J.O., and Eldredge, J.D. (2009) Lagrangian coherent structures in low Reynolds number swimming. *Journal of Physics-Condensed Matter*, **21** (20), 204105.

110 Green, M.A., Rowley, C.W., and Smits, A.J. (2010) Using hyperbolic Lagrangian coherent structures to investigate vortices in bioinspired fluid flows. *CHAOS*, **20** (1), 017510.

111 Shadden, S.C. and Taylor, C.A. (2008) Characterization of coherent structures in the cardiovascular system. *Annals of Biomedical Engineering*, **36** (7), 1152–1162.

112 Xu, Z., Chen, N., Shadden, S.C., Marsden, J.E., Kamocka, M.M., Rosen, E.D., and Alber, M. (2009) Study of blood flow impact on growth of thrombi using a multiscale model. *Soft Matter*, **5** (4), 769–779.

113 Vetel, J., Garon, A., and Pelletier, D. (2009) Lagrangian coherent structures in the human carotid artery bifurcation. *Experiments in Fluids*, **46** (6), 1067–1079.

114 Santitissadeekorn, N., Bohl, D., and Bollt, E.M. (2009) Analysis and modeling of an experimental device by finite-time lyapunov exponent method. *International Journal of Bifurcation and Chaos*, **19** (3), 993–1006.