Pollution release tied to invariant manifolds: A case study for the coast of Florida

Francois Lekien[a,b,*], Chad Coulliette[a,c], Arthur J. Mariano[d], Edward H. Ryan[d], Lynn K. Shay[d], George Haller[e], Jerry Marsden[a]

[a] *Control and Dynamical Systems, California Institute of Technology, CA, USA*
[b] *Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544, USA*
[c] *Environmental and Engineering Science, California Institute of Technology, CA, USA*
[d] *Rosenstiel School of Marine and Atmospheric Science, University of Miami, FL, USA*
[e] *Department of Mechanical Engineering, Massachusetts Institute of Technology, MA, USA*

Abstract

High-resolution ocean velocity data has become readily available since the introduction of very high frequency (VHF) radar technology. The vast amount of data generated so far, however, remains largely unused in environmental prediction. In this paper, we use VHF data of the Florida coastline to locate Lagrangian coherent structures (LCS) hidden in ocean surface currents. Such structures govern the spread of organic contaminants and passive drifters that stay confined to the ocean surface. We use the Lagrangian structures in a real-time pollution release scheme that reduces the effect of industrial contamination on the coastal environment.
© 2005 Elsevier B.V. All rights reserved.

1. Introduction

The release of pollution in coastal areas [1–3] may have a dramatic impact on local ecosystems, especially if the pollution recirculates near the coast rather than leaving for the open ocean. Due to the sensitive dependence of Lagrangian particle motion on initial conditions, identical parcels of fluid released at the same time but from two slightly different locations may produce vastly different contaminant distributions. It is commonly believed that this concept of sensitivity to initial conditions only applies to trajectories starting at the same time but at slightly different locations. However, time-dependent vector fields generate flows on the extended phase space (space-time) and particle starting at the same location but at slightly different time are also subject to sensitive dependence on their

* Corresponding author.
 E-mail address: lekien@princeton.edu (F. Lekien).

initial conditions (cf. Section 7). Except for degenerate cases, trajectories with high sensitivity to the initial release *position* are also highly sensitive to the initial release *time*. As a result, one needs a detailed understanding of the nonlinear dynamics of surface currents to devise pollution schemes that result in favorable outcomes.

Such a detailed understanding may be gained from very high frequency (VHF) radar technology (see Section 2), which produces well resolved real-time surface velocities at select coastal locations. In this paper, we show how such velocity data can be used to uncover hidden invariant manifolds in the Lagrangian particle dynamics of the ocean surface. We then use the location of numerically computed Lagrangian coherent structures (as defined in [4]) to devise an automated pollution release scheme that minimizes the harmful recirculation of contaminants near the coastline.

In our analysis, we use a combination of accurate surface current measurement [5,6] and recent developments in nonlinear dynamical systems theory [7]. In contrast to earlier approaches to pollution control in simple flow models [8–12], we rely on real-time velocity data collected by coastal Doppler radar systems. Discounting measurement errors (typically of the order of 5 cm s^{-1}), we consider the VHF velocity data shown in Fig. 1 to be a faithful discretized representation of the actual surface velocity field.

We locate time-dependent coherent structures in the VHF velocity data by employing the Direct Lyapunov exponent (DLE) algorithm proposed by Haller [7]. Local maximizing curves or ridges [4] of the DLE field are time-dependent manifolds governing mixing and transport in the fluid flow. These structures are known to be quite robust with respect to measurement and processing errors [13]. In other words, the measurement error on the velocity field (VHF radar data) does *not* affect much the geometry and position of the structures presented here. A formal proof is given in [13] but a simpler, more intuitive justification can easily be derived from the definition given in Section 5. The DLE field is the derivative of the flow map with respect to the initial conditions. As a result, measurement errors in the velocity field are *averaged along trajectories* before it affects the DLE field. For unbiased observation stations and long trajectories spanning a large domain, the averaged Lagrangian error becomes

very small with respect to the corresponding Eulerian error.

Other methods for extracting invariant manifolds from geophysical data include the advection of judiciously chosen material lines [14–16], relative and absolute dispersion [17], finite-size Lyapunov exponents [18], and the use of a Lagrangian version of the Okubo-Weiss criterion [19]. Reviews of the different techniques can be found in [7,20]. In this work, we favor Lagrangian structures computed using Direct Lyapunov exponent field for their robustness to Eulerian errors.

Analyzing the location of the extracted coherent structure relative to a hypothetic pollution release spot, we identify the main frequency components of the manifold's own motion. We then use these components to predict the short-term motion of the structure, with an emphasis on predicting its environmentally friendly positions. Our automated pollution scheme uses the environmentally friendly time windows to release contaminants. As we show by simulations, this scheme achieve a significant reduction in recirculating pollutants.

The organization of this paper is as follows. The experimental setting used to measure the velocity along the coast of Florida is described in Section 2. This data is used in Sections 3 and 4 to integrate numerically parcels of pollution. Based on the definition of Lagrangian Coherent Structure given in Section 5, a Lagrangian barrier is computed for the domain of interest in Section 6. Section 7 uses the barrier to minimize the effect of the pollution in coastal Florida and a final algorithm is given in Section 8.

2. VHF radar data along the coast of Florida

The use of radio frequencies to measure ocean surface currents has received attention in recent coastal oceanographic experiments, using high frequency (HF) radar techniques [21,22]. Recent surface current observations from ocean surface current radar (OSCR) using the very high frequency (VHF) mode reveal complex surface current patterns near the coast of Florida [5,6].

The OSCR VHF system was deployed for the Southern Florida Ocean Measurement Center (SFOMC) four-dimensional current experiment from June 25 to August 25, 1999. Radio waves are backscattered from
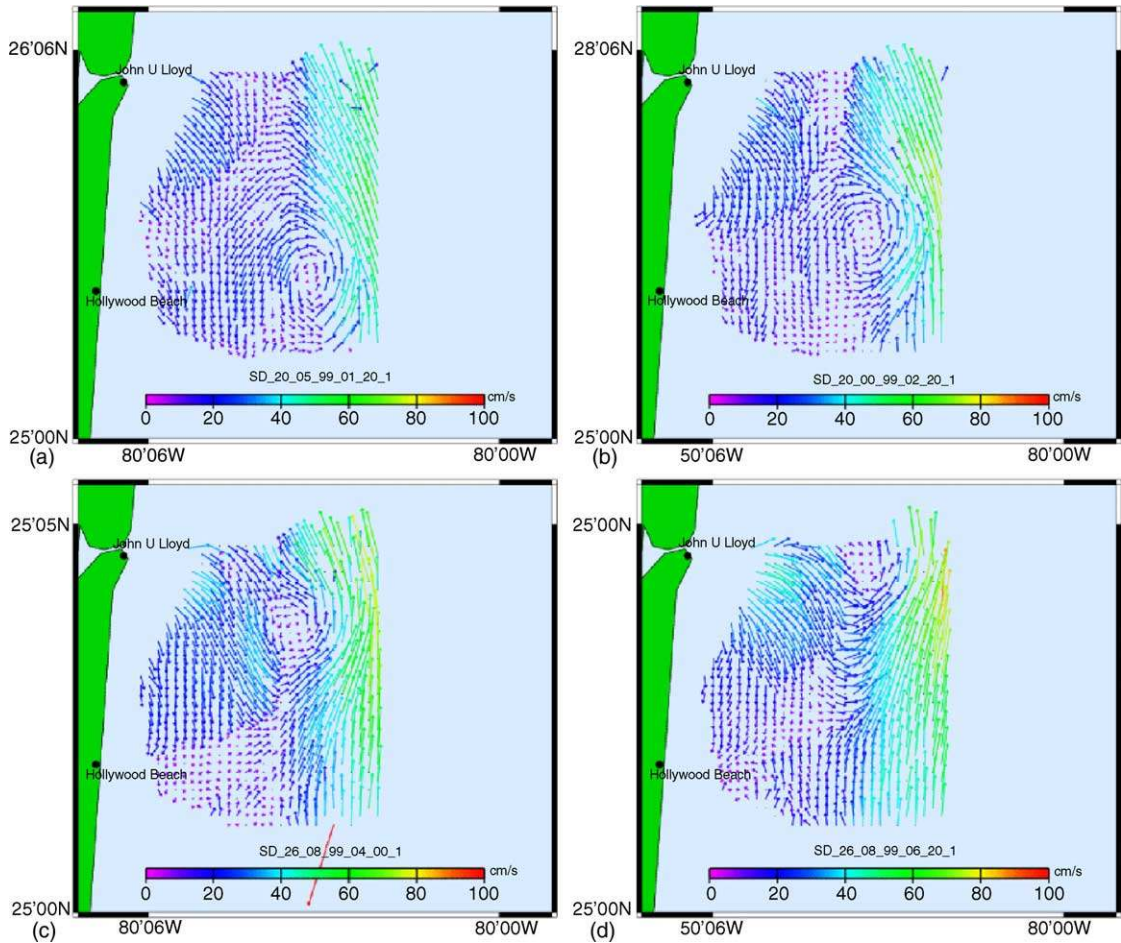
Fig. 1. Surface velocity maps obtained by VHF radar along the southeast coast of Florida, near Fort Lauderdale, during the SFOMC 4D Current Experiment on June 26, 1999: (a) 01:20 GMT; (b) 02:20 GMT; (c) 04:00 GMT; (d) 05:20 GMT. The sequence shows a northward propagating submesoscale vortex. The translation speed of the vortex is about 30 cm s^{-1} and its horizontal scale is 2–3 km [5].

the moving ocean surface by surface waves of one-half of the incident radar wavelength.

This Bragg scattering effect [23] results in two discrete peaks in the Doppler spectrum. In the absence of surface current, spectral peaks are symmetric and their frequencies are offset from the origin by an amount proportional to the surface wave phase speed and the radar wavelength. If there is an underlying surface current, Bragg peaks in the Doppler spectrum are displaced by the radial component of current along the radar's look direction. Using two radar stations sequentially transmitting radio waves resolves the two-dimensional velocity vector [5,6]. The resultant

data set represents coastal ocean surface currents mapped over a 7 km × 8.5 km domain at 20 min intervals with a horizontal resolution of 250 m at 700 grid points. The radars, transmitting at 50 MHz, were located in John Lloyd State Park, Dania Beach, Florida (Master) and an oceanfront site in Hollywood Beach, Florida (Slave), which are separated by 7 km.

Given the narrow shelf off Ft. Lauderdale and the strong Florida Current that intrudes onto the shelf on one-to-three day time scales [6,24], the VHF radar domain is ideally located for purposes of examining a wide spectrum of coastal and oceanographic processes. At times, the speed of the Florida Current exceeds

$2 \, \text{m s}^{-1}$ just 5 km offshore [6]. The average ambient relative vorticity is 4f, where f is the local Coriolis parameter, and maximum relative vorticity exceeds 10f [24]. The dominant period in the velocity data increases from ten hours near-shore to 5 days offshore. There is also significant energy at 27 h, the inertial period.

During this experiment, surface current observations (Fig. 2) revealed Florida Current intrusions over the shelf break, wavelike structures along the inshore edge of the current and numerous submesoscale vortices. One example started at 01:20 GMT on July 20, 1999 (Fig. 1) when a submesoscale vortex was located along the southern part of the VHF-radar domain just inshore of the Florida Current. Surface currents within the vortex ranged from 20–30 cm s^{-1} at a diameter of about 1–1.25 km from the vortex's center. The vortex's northward displacement of about 6 km occurred over a 5-h period. See [5] for a more detailed analysis of this vortex.

3. Interpolation and velocity field

The discrete data set containing the radar measurement does not constitute a velocity field, a vector function giving the velocity at each point and at each time. Experimental data must be interpolated between grid points and the resulting interpolating function is used as a velocity field. Since the smoothness of the flow depends directly on the smoothness of the velocity field and the methods used in this paper involve the first derivative of the flow, we require that the velocity field used is differentiable in space and time and that its derivatives are continuous in space and time, which is usually denoted by $\mathbf{v} \in C^1$, or more explicitly:

$$\mathbf{v} = \begin{pmatrix} u(x, y, t) \\ v(x, y, t) \end{pmatrix} \in C^1(\Omega \times \mathbb{R} \to \mathbb{R}^2), \qquad (1)$$

where $\Omega \subseteq \mathbb{R}^2$ is the spatial domain of interest. Data is provided at the vertices of a regular mesh (x_i, y_j), at discrete times t_k. We use the tricubic local interpolation scheme described in [25] that represents each component of the velocity as a piecewise third order polynomial function. The polynomials are C^∞ inside each cell and the interpolated function is globally C^1 (see proof

in [25]). In addition, the functions $\partial^2 u / \partial x^2$, $\partial^2 u / \partial t^2$, $\partial^2 u / \partial t^2$ $\partial^3 u / \partial x \partial y \partial t$ and the corresponding derivatives for v are also continuous. As a result, the local tricubic interpolator described here is truly C^1 in space and time. Uniqueness, existence and C^1 smoothness of solutions for the system is therefore guaranteed (see [26], for example).

Notice that the methods described in this paper are independent on the model (HF radar data and tricubic interpolation). For instance, modal analysis [27,28] or empirical modes [29] can be used to interpolate the experimental data. A smooth model velocity field can alternatively be obtained from high resolution ocean modeling [30–32]. Recently, data assimilation was extended to use Lagrangian data (floats and drifters), broadening the range of geophysical flow available for Lagrangian studies [33,34]. The Lagrangian coherent structures presented in this paper are robust to errors and variations in the velocity field [13].

4. Numerical experiments

The complexity of the flow resulting from the integration of the interpolated HF radar data becomes evident from tracking different realizations of a fluid parcel–a model for a spreading contaminant–released at the same time, but at a slightly different location. The results for two such numerical experiments are shown in Figs. 3 and 4. The complete animation, along with others, are available from http://www.lekien.com/~francois/papers/rsmas. The analysis uses two parcels of particles launched at 09:45 GMT on July 10, 1999. Using the available high-resolution VHF velocity data, the fluid particles are advected using a 4th order Runge-Kutta-Fehlberg algorithm (RKF45) combined with 3rd order tricubic interpolation in both space and time. The motivation for using such a complex interpolator is that the resulting velocity field is C^1 in extended phase space [25].

Note that large concentrations of the black contaminant remain for relatively long times near the coast, whereas the white parcel exits the domain quickly to the north and are advected into the open ocean. The latter scenario is highly desirable, because it minimizes the impact of the contaminant on coastal waters, by causing it to be safely dispersed into the open ocean.
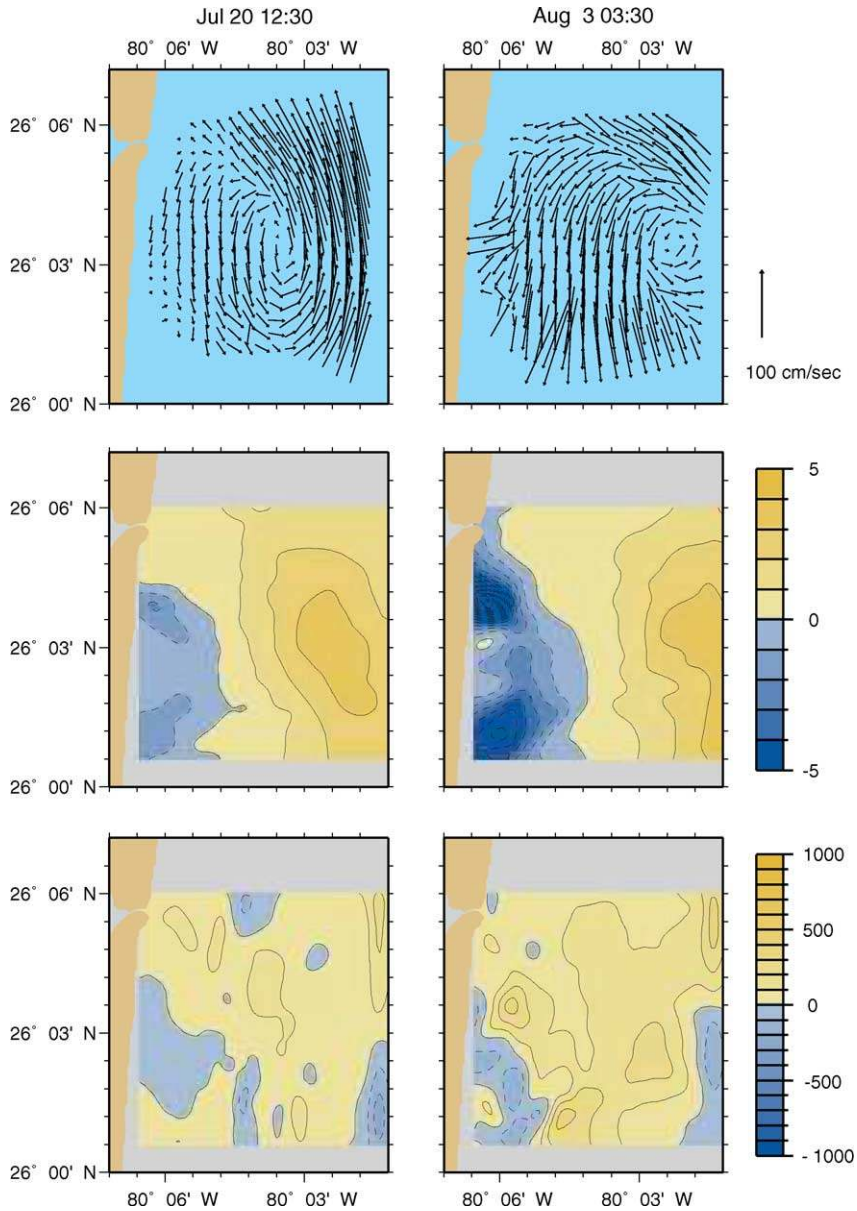
Fig. 2. (Top panel) The velocity pattern obtained by HF radar along the coast of Florida, near Fort Lauderdale, for two different time periods, July 20 and August 3, 1999. A submesoscale vortex is evident in each velocity map. (Middle panel) The corresponding normalized (by the local Coriolis parameter) relative vorticity anomaly fields. The mean vorticity, of order $4f$, was removed from each estimate to reveal the anomalies. Large positive relative vorticity values are associated with the vortices that are elongated due to the velocity shear of the Florida Current. Large negative values are found in the vicinity of a near-shore topographic step. (Bottom Panel). The corresponding horizontal divergence fields are calculated from spline fits to the velocity data.
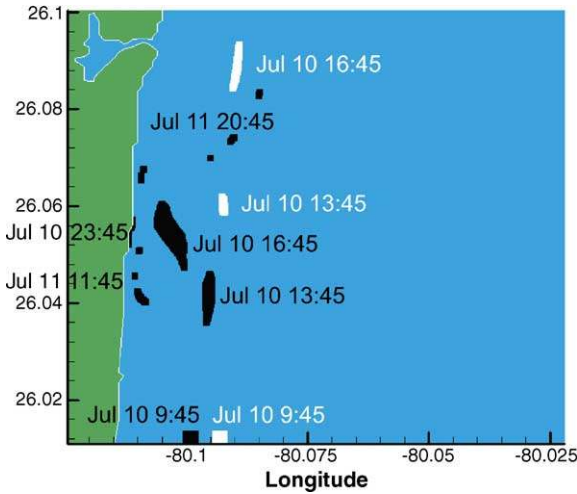
Fig. 3. Two parcels of contaminant released at exactly the same time, but at slightly different initial locations on July 10, 1999 at 09:45 GMT. The white parcels leave the domain quickly as they are advected by the northward flowing Florida Current. The black parcels re-circulate near the coast for more than 36 h. *Animation available at* http://www.lekien.com/~francois/papers/rsmas.

This observation provides an opportunity to understand and predict differing evolution patterns of a fluid parcel, depending on its initial location and time of release. Such patterns are known to be delineated by repelling material lines or finite-time stable manifolds [35–38], which we compute in the next section.

Notice that the domain studied in Fig. 4 is not completely closed by a coastline. Particles can leave the domain through the open-boundary (southern, western and northern edges). Once particles leave the domain, no information is available to follow their track and they are disregarded. One might wonder if our conclusion holds in the real world where the white parcel can possibly be pushed back into the domain of Fig. 4 by unknown currents. This would invalidate the reasoning above and the proposed contaminant control scheme. Such a problem is inherent to open-boundaries and must be studied from two different approaches:

- **Open-boundary flow**. An estimate on how likely parcels can leave the domain and re-enter it *in a reasonable amount of time* can be derived from a study of the flow on a section of the open-boundary.

In the case of the white parcel of Fig. 4 (or more generally, particles leaving the domain through the northern open-boundary), one can notice that the velocity vectors are almost always pointing *toward the outside* of the domain. During the numerical experiment (July 10, 1999 → July 12, 1999), less than 0.1% of the velocity vectors on the northern edge were indicating inflow. Moreover, such vectors were of relatively small magnitude and localized near the shoreline at the beginning of the experiment. It is therefore very unlikely that particles forming the white parcel will re-enter the domain through the northern edge. Re-entrant particles must travel a long distance outside the domain and re-enter through the western or southern open-boundary. After such a long travel time, diffusion and dispersion have destroyed the parcel of contaminant significantly and the resulting very low concentrations of white contaminant can be ignored. To generalize this notion, we computed the inflow and the outflow through the northern segment of the open-boundary. For a given segment of open-boundary $\partial\Omega$ the outflow is defined as

$$J_{\text{out}} = \int \text{d}t \int_{\partial\Omega} U(\hat{\mathbf{n}} \cdot \mathbf{v}) \text{d}l, \tag{2}$$

where the unit normal vector $\hat{\mathbf{n}}$ is pointing outside the domain and

$$U(z) = \begin{cases} 0 & \text{if } z \leq 0 \\ z & \text{if } z > 0 \end{cases}. \tag{3}$$

Similarly, we define the inflow as

$$J_{\text{in}} = \int \text{d}t \int_{\partial\Omega} U(-\hat{\mathbf{n}} \cdot \mathbf{v}) \text{d}l. \tag{4}$$

For the months of July and August 1999, we have $J_{\text{in}}/(J_{\text{in}} + J_{\text{out}}) < 0.01\%$. This justifies the fact that particles leaving the northern edge of the domain do not typically re-enter the domain through this edge. If enough diffusion of the contaminant and dispersion due to sub-scale processes occur, the white parcel can safely be disregarded when it leaves the domain.

- **Diffusion and dispersion**. The argument above only holds if there is a reasonable difference between the motion of the tracers and the pure
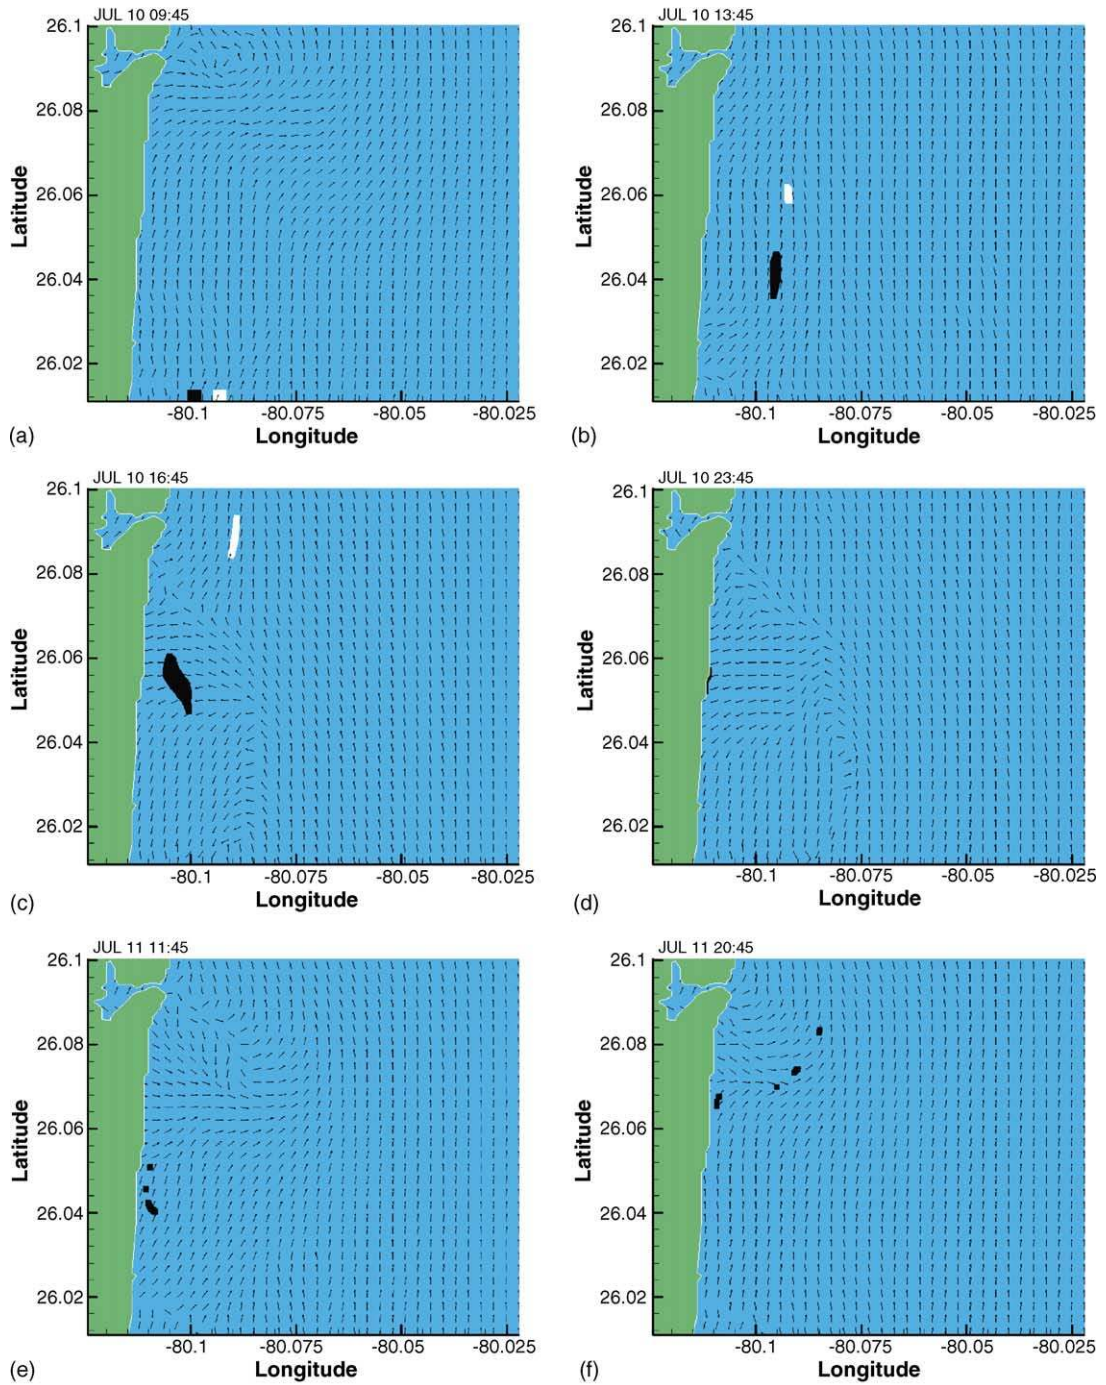
Fig. 4. A time sequence of the motion of two parcels released almost at the same position on July 10, 1999 at 09:45 GMT. The interpolated velocity from the radar and the position of the parcels is shown for (a) July 10 09:45 GMT, (b) July 10 13:45 GMT, (c) July 10 16:45 GMT, (d) July 10 23:45 GMT, (e) July 11 11:45 GMT and (d) July 11 20:45. *Animation available at* http://www.lekien.com/~francois/papers/rsmas.

Lagrangian model. In reality, contaminant are not advected exactly as fluid particles and their motion is best described by an advection-diffusion equation [39,40], for example. The parcels in Fig. 4 only represent the motion of contaminant for short integration time. For longer integration times, such as particles leaving through the northern edge and re-entering the domain through the western or southern edge, the advective terms destroy this localized representation and smooth the distribution of contaminants. In addition to diffusion, a correct model of the motion of such contaminant must include statistical processes to model sub-grid processes that are not resolved in the VHF radar data [41,42]. For small trajectories, the idealized Lagrangian parcel is a reasonable approximation of the concentration of contaminant [43]. For long trajectories, diffusion and dispersion tend to average the concentrations [44]. The objective of the contaminant control scheme in this paper is to *avoid peaks of large concentration*. As a result, we are only concerned with the small time effect of the release (e.g., before parcels could re-enter the domain). In the long term, diffusion and dispersion play in favor of our objective.

A recently developed nonlinear technique, called the direct Lyapunov exponent (DLE) analysis [7], identifies repelling or attracting material lines associated to Lagrangian coherent structures (LCS) in velocity data by means of curves experiencing maximum material stretching. Measured HF radar data allow the computation of such structures and the identification of environmentally friendly release spots.

5. Lagrangian coherent structures

The study of transport and mixing in fluid flows can be highly simplified by the use of finite-time invariant manifolds or Lagrangian coherent structures [14,16,45–47]. This work does not involve discussing the existence or uniqueness of hyperbolic trajectories. A preliminary discussion on this topic can be found in [4]. In this paper we use the Direct Lyapunov Exponent algorithm [7] to compute Lagrangian structures and divide the domain into regions of qualitatively different dynamics.

5.1. Direct Lyapunov exponents

The DLE algorithm starts with the computation of the flow map, the map that takes an initial fluid particle position $\mathbf{x}_0$ at time t_0 to its later position $\mathbf{x}(t, \mathbf{x}_0)$ at time t. To perform this analysis, a uniform grid of 200×200 particles are launched at time t_0. Each particle is advected using a 4th order Runge-Kutta-Fehlberg algorithm and a 3rd order interpolation for $t - t_0 = 25\,\text{h}$. An analysis of the influence of the integration time $t - t_0$ in the computation of DLE maps can be found in [48].

These particle trajectories are used to approximate the flow map, which associates current positions to positions at time t_0. The coastline is modeled as a free-slip boundary. Particles that cross the open boundaries of the domain on the northern, eastern and southern edges are disregarded. These numerical algorithms have been compiled into a software package, MANGEN, that is available from the authors upon request [49].

The spatial gradient of the flow map ($\partial \mathbf{x}/\partial \mathbf{x}_0$) is a 2×2 matrix and can be computed by using using finite differences in the grid of trajectories [37]. The direct Lyapunov exponent $\sigma_t(\mathbf{x}_0, t_0)$ is defined as the normalization of the largest singular value of the spatial gradient of the flow map [37]. More specifically,

$$\sigma_t(\mathbf{x}_0, t_0) = \frac{\ln \lambda_{\max}\left(\left[\frac{\partial \mathbf{x}(t, \mathbf{x}_0)}{\partial \mathbf{x}_0} \right]^\top \left[\frac{\partial \mathbf{x}(t, \mathbf{x}_0)}{\partial \mathbf{x}_0} \right] \right)}{2|t - t_0|}, \qquad (5)$$

with the superscript $^\top$ referring to the transpose of a matrix. Notice that the denominator of Eq. (5) normalizes the Lyapunov exponent. The logarithm of the largest singular value of the gradient of the flow map has a linear dependence in the integration time $(t - t_0)$. Dividing the result by $(t - t_0)$ allows for a more constant value of σ_t when $(t - t_0)$ changes. This is particularly important for domains with an open-boundary where the computation of certain trajectories must be stopped when they leave the domain. Such trajectories are disregarded when they exit the domain and the Lyapunov exponent must be computed with a smaller integration time. As noted above, such smaller integration time do not influence the results presented here. Contaminants only follow pure Lagrangian advection as a first approximation. In the long term, diffusion and dispersion will naturally distribute the contaminant at a very low concentration. Our goal is to study the short-term

peaks of concentration, hence Direct Lyapunov exponents with short integration times.

5.2. Lagrangian ridges

Repelling material lines are maximizing ridges of the scalar field $\sigma_t(\mathbf{x}_0, t_0)$ [7,13]. By ridge of a scalar field, we mean a gradient line of this field that has maximum curvature in the orthogonal direction [4]. More specifically, a ridge is a C^1 curve $\mathbf{c}(s)$, $s \in]a, b[$, satisfying the following conditions:

(1) $\mathbf{c}(s)$ is parallel to $\nabla \sigma_t(\mathbf{c}(s), t_0)$,
(2) $\dfrac{d\mathbf{c}(s)}{ds} \neq 0$,
(3) $\Sigma(\mathbf{n}, \mathbf{n}) = \min_{\|\mathbf{u}\|=1} \Sigma(\mathbf{u}, \mathbf{u})$, where $\mathbf{n}$ is a unit vector normal to the curve $\mathbf{c}(s)$ and $\Sigma = \partial^2\sigma_t/\partial\mathbf{x}^2$ is the second derivative of σ thought of as a bilinear form evaluated at the point $\mathbf{c}(s)$.

Notice that the ridges defined here are not necessarily contour level sets of σ_t. They correspond to "water-dividing lines" of the field. The ridges of the Lyapunov exponent field as defined in this section are not necessarily Lagrangian and do not strictly behave as material lines. They are instantaneous screenshots of the future Lagrangian dynamics. However, recent studies [4] have shown that the Lyapunov exponent σ_t is Lagrangian for large integration times $(t - t_0)$ (in the sense that the change of its value along trajectories varies as the inverse of $(t - t_0)$). In the definition above, a Lagrangian structure is a gradient line of σ that has maximum curvature in the transverse direction. For typical coastal flows such as the one studied here, such a ridge has a relatively constant value of σ_t and is, in good approximation, a Lagrangian line of the flow (see [4] for proofs and numerical results).

Relative dispersion and finite-size Lyapunov exponents have been successfully used to identify chaotic patterns in fluid flows [50,51]. In small domains such as the one depicted in this paper, the velocity field is typically too small for the diffusion coefficients to give any relevant information about transport in mixing. Recent extensions [52,53] allow for the extraction of information for a particular length scale or time scale. Such methods were verified experimentally in [54]. In this work, we seek precise Lagrangian lines rather than chaotic regions or patterns. Direct Lyapunov exponents

and Lagrangian coherent structures [4,37] provide such a framework. A fundamental difference is that there is no need for estimating relevant length scales. The algorithm described above will extract both small and large-scale structures. Some small-scale structures might have high exponents. However, they do not shadow the large-scale structures of interest. Lagrangian structures are identified as ridges [4] and the intensity of the Lyapunov exponents is much less relevant than the curvature along the ridge, for example. Once extracted, the ridges differ by their *length* and can be classified or eliminated based on their length, rather than intensity.

The same procedure performed backward in time (i.e., for $t < t_0$) renders attracting material lines at t_0 as ridges of $\sigma_t(\mathbf{x}_0, t_0)$. These curves are not apparent to naked-eye observations and are not easily deduced from velocity field plots, yet they govern global mixing patterns in the fluid [55]. Such Lagrangian structures in measured ocean data have previously been inaccessible due to lack of an efficient extraction methods and coarse resolution of the observations.

6. Data analysis

Direct Lyapunov Exponents are used to analyze the Lagrangian trajectories in the VHF radar domain. In particular, we want to be able to define pollution barriers and pathways near the southeast coast of Florida.

Selected frames of the contour level sets of the maximum Lyapunov exponents are shown in Fig. 5. During the experiment, the plot reveals a strong stable Lagrangian structure attached to the coast near Fort Lauderdale, propagating to the southeast. This structure acts as a quasi Lagrangian barrier between the coastal recirculating zone (southwest of the LCS) and the Florida Current (northeast of the same LCS). Recent work [4] shows that the flux of particles crossing the Lagrangian structure is negligible. As a result, the LCS also acts as a material barrier. Superimposed on Fig. 5 are the two parcels used in Figs. 3 and 4. The average lifetime of the Lagrangian barrier is about 48 h and parcels typically travels from South to North in less than 12 h, so we expect that the barrier has a significant effect on the flow.

Analysis of Fig. 3 reveals that any particle northeast of the barrier (white parcel) is flushed out of the domain
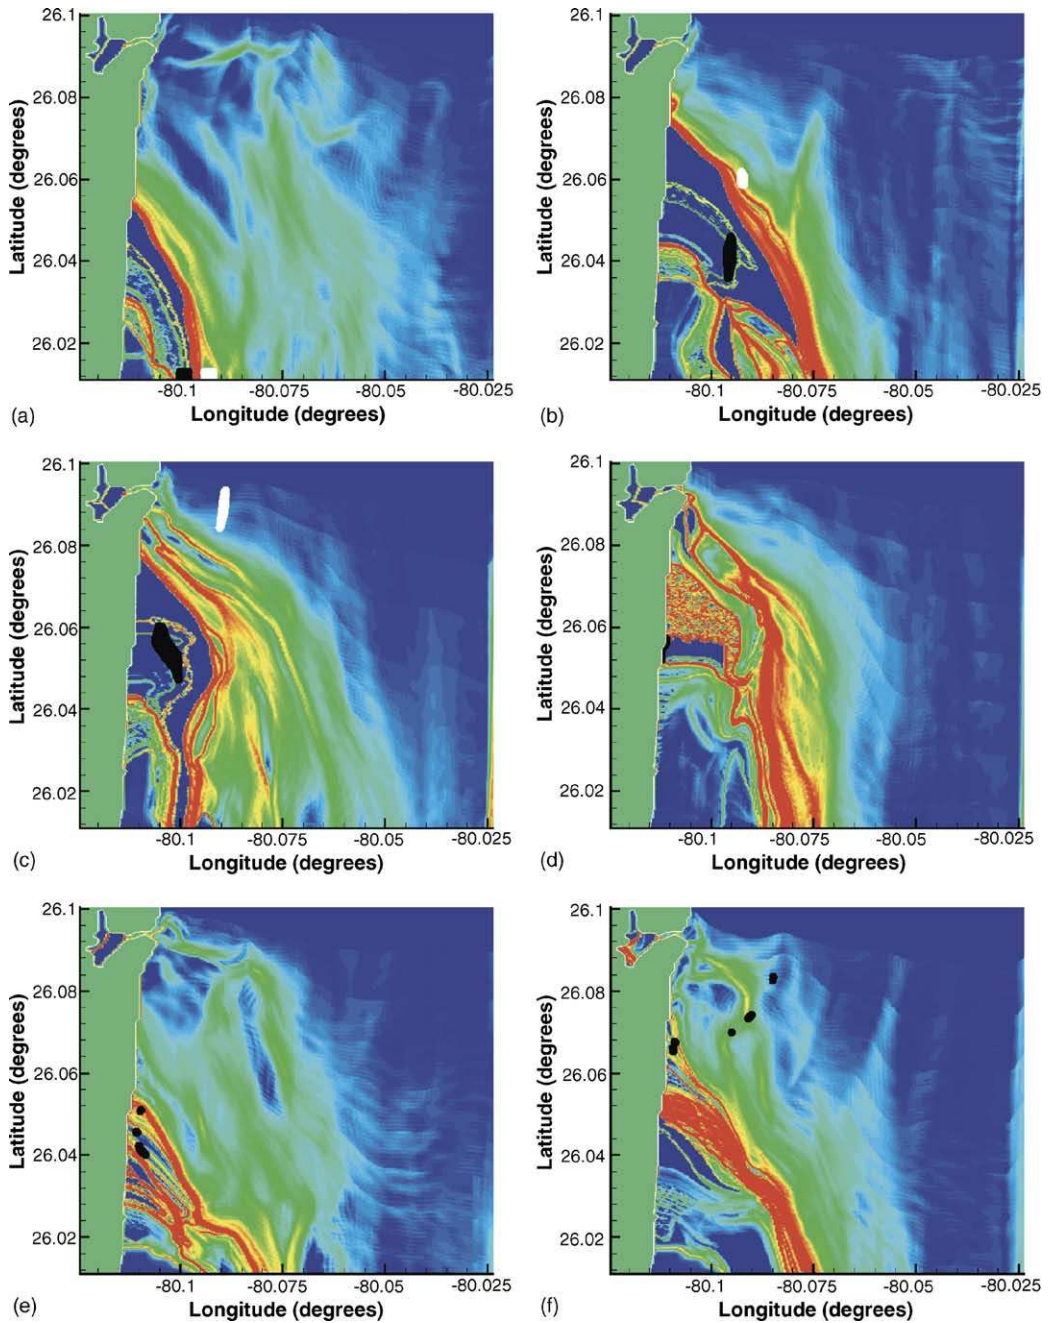
Fig. 5. Level Sets of the maximal Direct Lyapunov Exponents σ along the coast of Florida on (a) July 10, 1999 09:45 GMT, (b) July 10 13:45 GMT, (c) July 10 16:45 GMT, (d) July 10 23:45 GMT, (e) July 11 11:45 GMT and (d) July 11 20:45. The simulation shows repelling material lines attached to the coast near Fort Lauderdale. Superimposed on each figure panel are the respective positions of the two parcels from Fig. 4. Every particle North of the coherent structure flows through the northern open-boundary. It is non-optimal to release contaminants below the branch of the manifold because it will remain between the coast and the manifold for a long time. *Animation available at:* http://www.lekien.com/~francois/papers/rsmas.
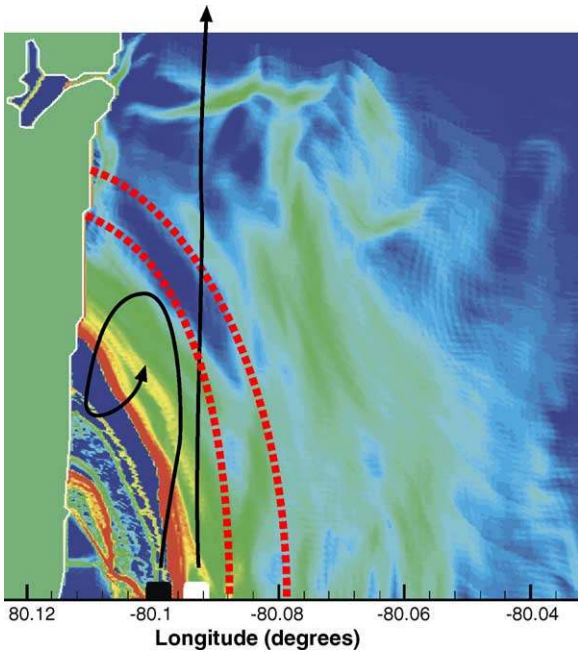
Fig. 6. Level Sets of the maximal Direct Lyapunov Exponents σ along the coast of Florida on July 15, 1999 at 9:45 GMT. The dashed line represents future positions of the repelling coherent structure. The parcel North of the manifold flows through the northern open-boundary (white parcel). It is dangerous to release contaminants below the branch of the coherent structure because they will persist between the coast and the structure for a long time (black parcel).

in only a few hours. In contrast, parcels starting south-west of the barrier (black parcel) typically re-circulate several times near the Florida coast before they finally rejoin the current. Interestingly, such behavior is not obvious from a simple observation of the velocity foot-prints, which is typical of fast varying time-dependent flows. In this mathematical framework, the surface currents are not necessarily influencing particle paths directly, but the currents influence the Lagrangian structures, such as causing transport barriers and pathways, and the Lagrangian structures act directly on particle paths (Fig. 6).

7. Minimizing the effect of pollution

The location of the base of the Lagrangian coherent structure (along the coastline) can be used as a criterion to minimize the effect of coastal pollution. We

will refer to the intersection of the coastline and the La-grangian structure as the barrier point. For the region and time period analyzed here, factories and sewage discharge pipelines along the coast should not release anything if the barrier point is located *North* of them. Optimizing the release site by moving the source of pollution is unrealistic. Optimizing the release times is a much more implementable operation. The Direct Lyapunov exponents typically give the influence of the initial position $\mathbf{x}_0$ on the final position $\mathbf{x}(t; t_0, \mathbf{x}_0)$ of the trajectory that started at $\mathbf{x}_0$. However, a change in initial time $t_0 \rightarrow t_0 + \delta t_0$ can be interpreted as a change of initial position $\mathbf{x}_0 \rightarrow \mathbf{x}_0 - \mathbf{v}(\mathbf{x}_0, t_0)\delta t_0$ and we have

$$\frac{\partial \mathbf{x}(t; t_0, \mathbf{x}_0)}{\partial t_0} = -\frac{\partial \mathbf{x}(t; t_0, \mathbf{x}_0)}{\partial \mathbf{x}_0}\mathbf{v}(t_0, \mathbf{x}_0), \quad (6)$$

so sensitivity with respect to the initial position can always be interpreted as sensitivity with respect to the initial condition *in the direction of the initial velocity vector* $\mathbf{v}(\mathbf{x}_0, t_0)$ (see Fig. 7). Moreover, in this case, we used a slip boundary condition and the velocity field is always tangent to the coastline. In other words, for any $\mathbf{x}_0$ on the coastline, $\mathbf{v}(\mathbf{x}_0, t_0)$ is tangent to the coastline, so optimizing the *release site along the coastline* is equivalent to optimizing the *release time*. More precisely, we have,

$$\frac{\partial \mathbf{x}(t; t_0, \mathbf{x}_0(s))}{\partial t_0} = \mp\frac{\partial \mathbf{x}(t; t_0, \mathbf{x}_0(s))}{\partial s}\|\mathbf{v}(t_0, \mathbf{x}_0(s))\|, \quad (7)$$

where s is the arc-length along the coastline and the sign is determined by the projection of the velocity on the tangent vector $\partial \mathbf{x}/\partial s$. Eq. (7) shows that sensitivity
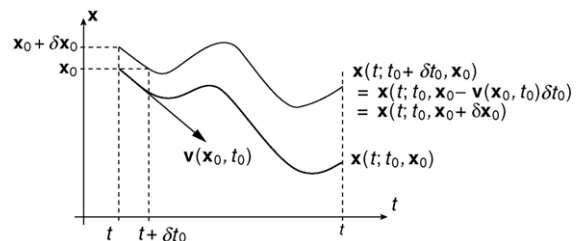

Fig. 7. Sensitivity to initial conditions must be studied in terms of spatial perturbations $\delta \mathbf{x}_0$ *and* temporal perturbations δt_0 of the initial condition $(\mathbf{x}_0, t_0)$. However, the law of motion $\dot{\mathbf{x}} = \mathbf{v}(\mathbf{x}, t)$ gives directly $\mathbf{x}(t; t_0 + \delta t_0, \mathbf{x}_0) = \mathbf{x}(t; t_0, \mathbf{x}_0 - \mathbf{v}(\mathbf{x}_0, t_0)\,\delta t_0)$ and sensitivity to initial time can be derived from the dependence in $\mathbf{x}_0$.
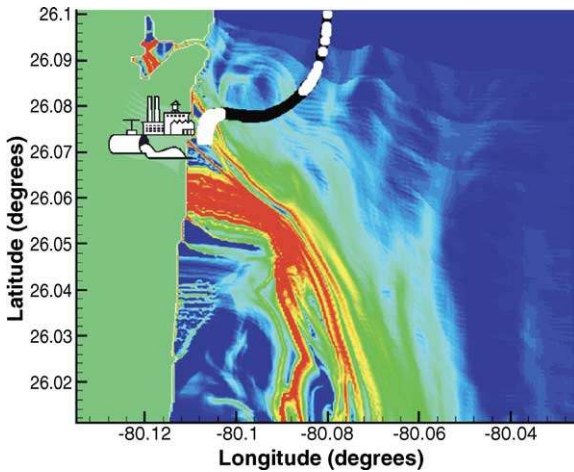
Fig. 8. Imaginary source of pollution along the Florida coast. The black spots are the resulting contaminants from a factory releasing at a constant rate. Superimposed on this figure are the white spots of a factory releasing only during environmental friendly time windows. *Animations available at* http://www.lekien.com/~francois/papers/rsmas.
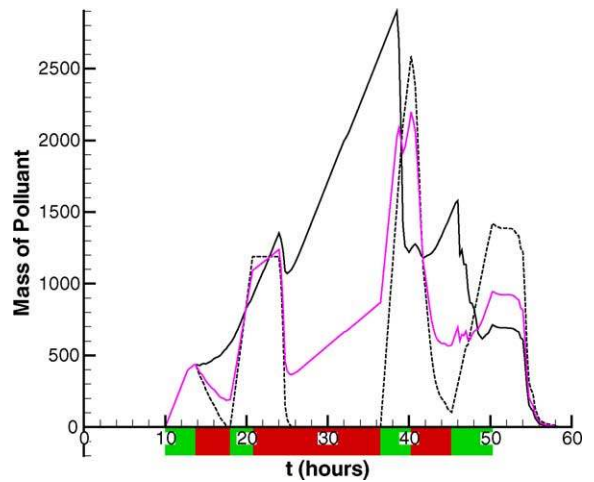

Fig. 9. Three different release strategies for pollutants at a factory on the coast. The black line shows the mass of pollutant in the coastal area for a uniform release in time. The dashed line is for releasing pollutants at times determined by the DLE analysis. The purple line is a strategy that is based on the DLE analysis and the criterion of minimizing peak values.

in terms of the release time is equivalent to sensitivity in terms of release position along the coastline and we can decide to keep the pollution source fixed and modify the release schedule, which is much easier to implement in practice.

To illustrate how an efficient pollution release algorithm can be set up, a fixed imaginary source of pollution is placed along the coastline. Using the DLE plots of Fig. 5, we identify zones of (green) favorable release[1] and (red) dangerous release.[2]

To minimize the effect of coastal pollution, we propose using a holding tank that stores contaminants during dangerous release times. The tank stores pollution during the half-period of the barrier point oscillation, during which contaminants should not be released. The contents of the tank are released once the barrier point moves South of the source of pollution as shown in Fig. 8.

The black spots on Fig. 8 are the trace of the pollution of a factory releasing at a constant rate. Superimposed on this figure is the white trace of contaminants released only during time windows determined

using our DLE algorithm. The total mass of contaminant in the coastal area for the two modes is shown on Fig. 9.

The two sources of pollution release the exact same mass of pollutant in the ocean. However, by obtaining information from the DLE results, the white factory (dashed curve) is able to reduce the effect of the pollution in the shallow coastal area by a factor of three.

Also shown on Fig. 9 are the results of a third numerical experiment. In many cases, the damage to the environment is a function of the maximum concentration of contaminant. From this viewpoint, the algorithm (releasing nothing during "red" zones and as much as possible during "green" zones) does not seem to be efficient. The peak of maximum concentration for the white factory has only decreased by a small amount. The results reveal that a long "red" zone can lead to the accumulation of large amount of pollution in the tank. If such a zone is followed by a short "green" zone, the large content of the tank is released quickly and creates a peak in the concentration. To set up a more elaborate algorithm, we release a minimum flux of pollution into the ocean independently of the type of zone (red or green). We define a new degree of freedom α, the per-

---

[1] When the structure is below the position of the factory.

[2] When the structure is above the factory.

centage of contaminant produced that will be released during "red" zones. Here, the solid curve of Fig. 9 corresponds to $\alpha = 100\%$, the dashed curve to $\alpha = 0\%$. The purple curve of Fig. 9 corresponds to $\alpha = 33\%$ (i.e., 1/3 of the pollution is released at all time). Fig. 9 shows that a significant reduction of the peak of maximum concentration can be obtained using an appropriate partial release during zones that are marked dangerous by the DLE algorithm.

8. An automated pollution control algorithm

The purpose of this section is to show that it is possible to implement a pollution release scheme based on the motion of a LCS in real time. If, at anytime, the DLE field was available, a simple decision system based on the position of the LCS at the present time should be able to select the beginning and the end of favorable release zones (see previous section). Unfortunately, if the velocity is known up to the present time, the DLE field can only be computed up to a certain time in the past. As we approach the present time, DLE can only be computed over a short amount of time and does not converge towards the expected Lagrangian structure. Instead, the plot reveals the maximum eigenvalues of the linearized flow. We determined empirically that 8 h at least were necessary to obtain a correct picture of the LCS. Our algorithm needs to be able to predict the position of the Lagrangian structure at least 8 h in the future. Notice that we do not attempt to predict the Eulerian velocity field. We propose to extrapolate the position of the Lagrangian barrier up to the present time.

8.1. Prediction algorithm

Several predictive methods could be used. One might think of using a model of the area, use it to predict the velocity field and compute the DLE field based on predicted velocities. However, the values of the DLE field are sequences in time and can also be pre-

dicted directly. We elected a simple prediction scheme. Since we are only interested in the latitude of the moving barrier, for each time, the latitude of the barrier is computed up to the latest time in the past where DLE can be computed (i.e. 8 h before the present time). This forms a time sequence that we would like to predict for at least 8 h. We refer to this operation as a *short time Lagrangian prediction* because the velocity (Eulerian field) is *never* interpolated in time during this process.

The spectrum of each sequence has been computed using the last 50 h and a few examples are shown in Fig. 10. A complete animation of the spectrums computed at each time can be found at http://www.lekien.com/~francois/papers/rsmas. We identified the different components of the oscillations of the Lagrangian barrier as the frequency at the maximum of each peak in the spectrum. Fig. 11 shows the significant frequencies that were identified at each time step. Notice that there is very little tidal influence on the spectrum of Fig. 11. These spectrums are based on the motion of Lagrangian particles along the coastline. Tidal oscillations influence the flow mostly in the direction orthogonal to the coastline for such a linear coast. Notice that, even for more complicated costlines, non-tidal dominant frequencies usually appear in the motion of Lagrangian structures (see [49] for example). The weak amplitude of the tidal frequencies in this example is a direct consequence of the fact that the Lagrangian motion studied is parallel to the coastline. Simulations using longer time sequences (e.g. Fig. 12) reveal that the computed frequencies are not sensitive to the length of the time sequence used. Within an appropriate range (50–200 h), using more data does not necessarily provide more accurate predictions.

Based on the relevant frequencies T_i, we use a very simple predictive algorithm based on a Fourier decomposition where the motion of the Lagrangian barrier is represented as a finite sequence of Fourier modes

$$y = A_0 + \sum_{i=1}^{N} A_i \cos\left(\frac{2\pi}{T_i}t\right) + B_i \sin\left(\frac{2\pi}{T_i}t\right), \quad (8)$$

and, the optimal model is obtained when one minimizes the error

$$\epsilon = \sqrt{\sum_k \left(y_k - A_0 - \sum_{i=1}^{N} A_i \cos\left(\frac{2\pi}{T_i}t_k\right) - B_i \sin\left(\frac{2\pi}{T_i}t_k\right)\right)^2}, \quad (9)$$
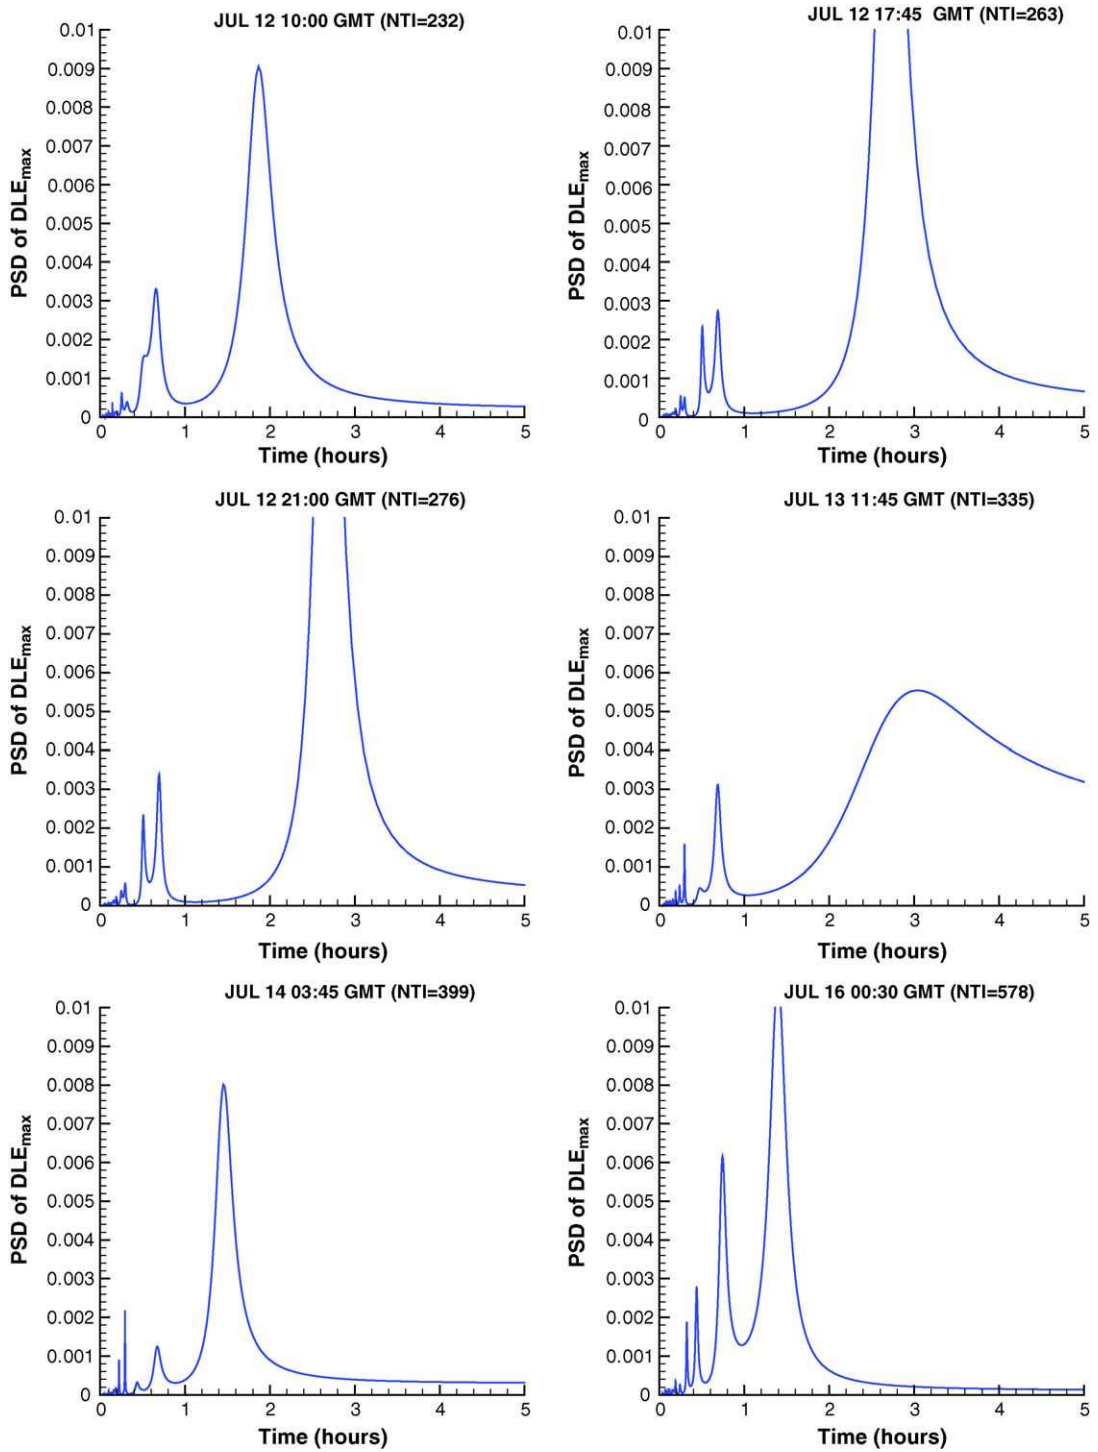
Fig. 10. Entropy spectrum of the time sequence of the latitude of the Lagrangian barrier computed using 50 h of data at selected time steps.
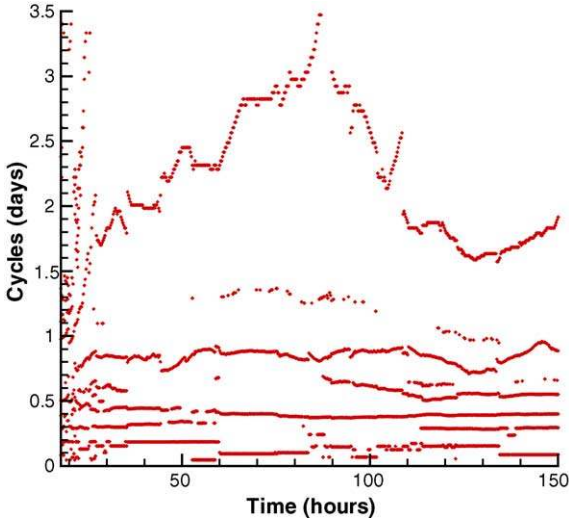
Fig. 11. Spectral peaks in the time sequence of the latitude of the Lagrangian barrier computed using 50 h of data.

where y_k are the K positions of the Lagrangian barrier measured in the last 50 h (at times t_k). The error is minimum when
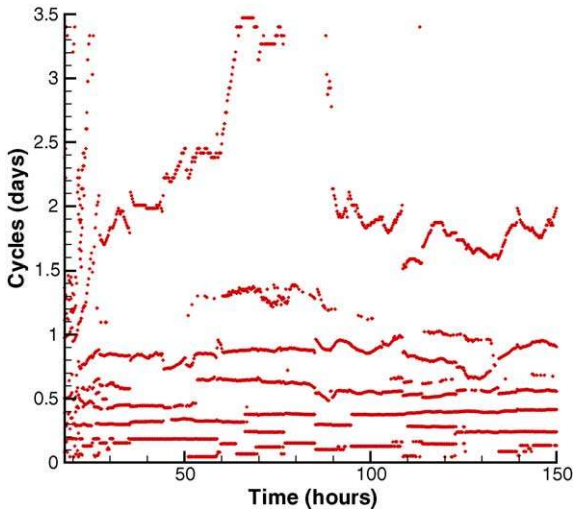
$$\frac{\partial \epsilon}{\partial A_i} = 0, \tag{10}$$

Fig. 12. Spectral peaks of the time sequence of the latitude of the Lagrangian barrier computed using 75 h of data.

and

$$\frac{\partial \epsilon}{\partial B_i} = 0, \tag{11}$$

for all i, which leads to the linear system in $(2K + 1)$ unknowns[3]

$$(Q_{ij}) \begin{pmatrix} \vdots \\ A_j \\ \vdots \\ B_j \\ \vdots \end{pmatrix} = \begin{pmatrix} \vdots \\ \sum_k y_k \cos \frac{2\pi t_k}{T_j} \\ \vdots \\ \sum_k y_k \sin \frac{2\pi t_k}{T_j} \\ \vdots \end{pmatrix}, \tag{12}$$

where

$$Q_{ij} = \sum_k f_i(t_k) f_j(t_k), \tag{13}$$

and

$$f_i(t) = \begin{cases} 1 & \text{if } i = 1 \\ \cos \frac{2\pi t}{T_{i-1}} & \text{if } i = 2, 3, \ldots, K+1 \\ \sin \frac{2\pi t}{T_{i-K-1}} & \text{if } i = K+2, \\ & K++3, \ldots, 2K+1 \end{cases} \tag{14}$$

In this paper, the linear system given by Eq. (12) was solved using the *GNU Scientific Library (GSL)* for maximum 13 modes and each mode is a dominant frequency automatically extracted from the spectrums computed at each time slice. The system is ill-conditioned when some frequencies are close to each other (e.g., if $T_i \approx T_{i'}$, the determinant of the linear system is close to zero, in which case the coefficients A_i, B_i, $A_{i'}$ and $B_{i'}$ can grow arbitrarily large). To avoid this problem, groups of close peaks, defined as relative distances less than 1%, are ignored and combined into a single resonant frequency. Moreover, we allow the algorithm to select the optimal number of modes for each time step. The linear system is solved with the jth most significant modes ($3 \leq j \leq 13$) and the best result is kept.

At each time step, once the number K of modes is established and the coefficients A_k and B_k are computed, the sequence in Eq. (8) is used to extrapolate the

---

[3] The $(2K+1)$ unknowns are $A_0, A_1, A_2, \cdots A_K$ and $B_1, B_2, \ldots B_K$.

position of the barrier during the past 8 h where DLE was unavailable as well as in the future. Fig. 13 shows some of these results for selected time slices.

Sometimes the prediction can be very accurate for many hours in the future, such as on July 12 10:00 GMT or July 14 03:45 GMT. However, large errors can occur even at the present time such as July 12 17:45 GMT. On July 12 21:00 GMT and July 13 11:45 GMT, the system automatically switched to a low-mode analysis because a better fit of the past data could not be obtained with more modes. Results can be significantly improved such as on July 13 11:45 or have a strong divergence such as July 12 21:00 GMT. This phenomena has been observed several times and is a consequence of the presence of close frequencies in the spectrum. Removing lower modes usually provides a sufficient solution.

While no attempts were made to improve the prediction method, the objective was to study the possible advantages of a pollution algorithm. Thus, it sometimes used erroneous data as input, even though it may have been obvious not to be able to predict the position of the barrier accurately.

Our goal is to show that even in the presence of errors in the prediction of the position of the Lagrangian barrier, a significant reduction of the concentration of pollution can be achieved by the algorithm.

To determine favorable release time intervals, only the relative position of the predicted barrier with respect to the latitude of the release site is important. The release site is indicated by a the horizontal line on Fig. 13 and a favorable release interval occurs when the latitude of the barrier point is below the release site. To assess the quality of the prediction algorithm, the predicted curve was used at each time to decide whether or not the system was in a favorable zone. The answer to this question was compared at each time with the correct answer based on the actual time sequence of the DLE ridge.[4] The predicted curve agrees with the actual curve in 78% of the cases. In this experiment, the actual curve only stays below the release site about 40% of the time. The algorithm was used to simulate with different release sites, and all of them had a success rate ranging from 70 to 80%.

_____

[4] The "actual" DLE ridge is computed at the end of the experiment, thus with all data available.

8.2. Decision algorithm

To account for the high error rate on the raw decision algorithm (favorable release zones), the factory or another release center is expected to simulate a more complex decision algorithm to decide which portion of the pollutant is to be released and which is to be stored in the tank. The data from the 4 last predictions (last hour) is used to increase the accuracy. A persistent low impact release time interval occurs when at least 75% of the last 4 points indicate that it is a good release interval. A persistent high impact zone corresponds to at least 75% of the points detecting an unfavorable interval. In other cases, the system is in an undetermined state.

Based on the type of zone that the factory thinks it is in, the release flux or position of the valve to the tank was computed as follows:

- If the system is in a *low impact* zone, the production of the factory is released in the ocean. If not empty, the tank is emptied at maximum speed;
- If the system is in an *high impact* state, the production of the factory is moved to the tank. If the tank is full, the production is released in the ocean but the tank is not emptied; and,
- If the system is in an *undecidable* state, the action depends on the amount of liquid stored in the tank. If the tank is less than half full, the production is stored in the tank. If the tank is more than half full, the pollution produced is emptied in the ocean but the tank is not to be emptied.

The action taken in the undecidable case results from the following observation: if we are uncertain that the release will have a low impact on the environment, it is better not to take any risk and keep the pollutants produced in the tank.

However, if the tank is more than half full, it is a much higher risk to store more pollutant because a long high impact zone could follow. In this case, the system may be forced to release pollution during the unfavorable zone, simply because the tank has reached its maximum capacity. When the tank is more than half full, it is less risky to release the pollution when there is at least a 50% chance that the release will be favorable. In addition to the framework presented above, we learned in the previous section that a small percentage of the
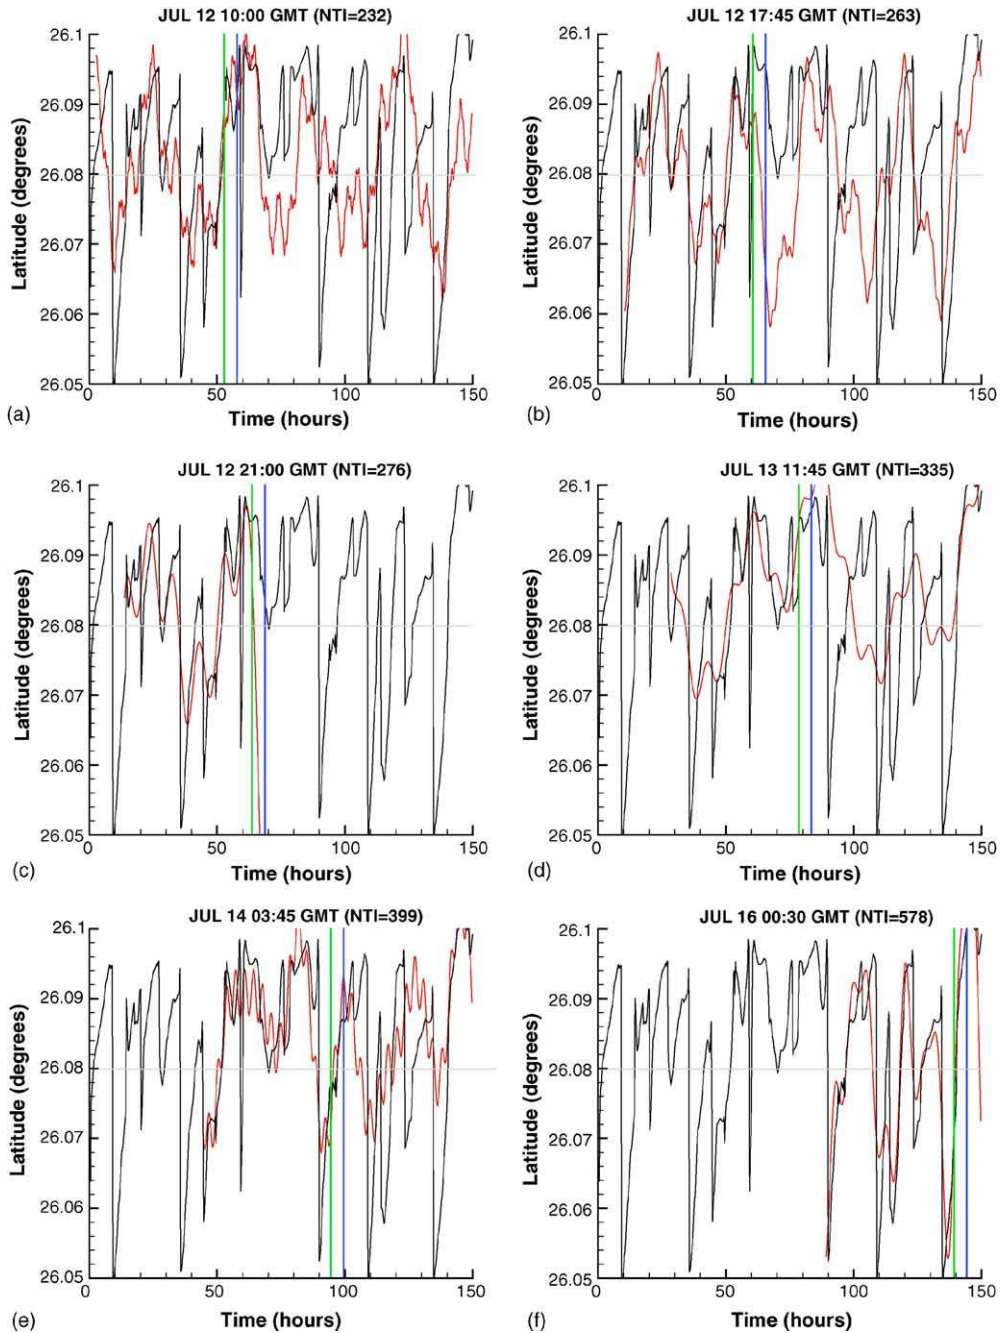
Fig. 13. Predictive motion of the intersection of the coherent structure with the coastline. The prediction is shown for selected times and uses the last 50 h of data. The black curve represents the actual time sequence that can be computed at the end of the time of the experiment. The blue vertical line in each figure represents the actual time and the green vertical line represents the time up to which DLE can be computed. In other words, the position of the Lagrangian structure is known up to the green vertical line. The red curve shows the Fourier sequence that was used to predict the position of the LCS in the future. It has been fitted to the measured data up to the green line and may or may not make a good prediction during the 8 h preceding the current time and in the future.
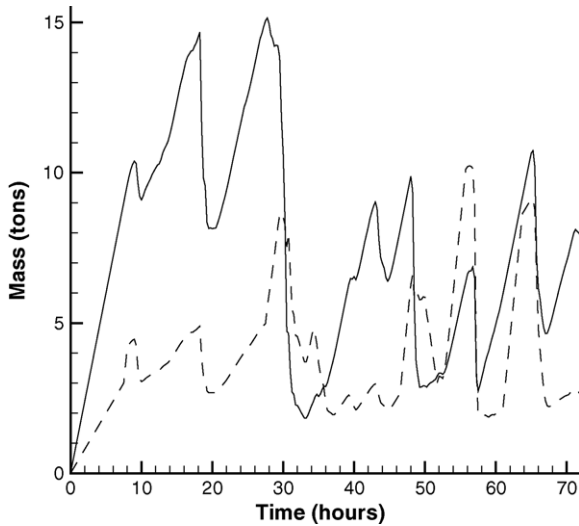
Fig. 14. Concentration of pollutant. The solid curve shows the concentration of pollutants in the coastal area of a source releasing at a constant rate. The dashed curve shows the effect of a source of pollution releasing the same amount of pollutant, but using the DLE algorithm to reduce its effect on the coastal environment.

pollution has to be released in the ocean at all times to achieve effective reduction of the peak of maximum concentration.

The result of such an experiment is presented on Fig. 14. The factory was producing 1.2 tonnes of pollutant per hour and at least 0.4 tonnes h^{-1} were to be dumped in the ocean ($\alpha = 33\%$). The tank had a maximum capacity of 30 tonnes and the maximum release rate from the tank was set to 3.6 tonnes h^{-1}.

9. Conclusions

We have shown the existence of a set of repelling material lines in the VHF-radar derived surface current fields acquired along the Hollywood, Florida coast in July 2003. We have also shown how these material lines can be used to minimize the effect of coastal pollution by determining and predicting optimal release times. This approach can be used for simulating trajectories of buoyant contaminants or the trajectories of nearly Lagrangian tracers.[5] The data source can be

VHF radar data or any other current data source, such as data-assimilated ocean models that approximate the near-surface velocity field to some reasonable level of accuracy. The advantage of using ocean models is that the velocity provided is 3D+1, and thus we can explore the Lagrangian structures that develop at various depths. We have shown that a real-time experimental realization of our pollution release is possible and can efficiently reduce the impact of a polluting source in a coastal area without reducing productivity. Applications of the algorithm are limited to regions where an experimental setting (such as VHF radars) or an accurate model of the flow is available. The cost of such an installation is usually high but can be shared by several sources of pollution in a nearby area.

The algorithm is only efficient when an effective barrier exists in the domain,[6] attaches to the coastline and oscillates about the pollution sources. In any other case, the algorithm presented here does not reduce the effect of the pollutant better than a constant release scheme. Future improvements include the combination of this algorithm with other methods that are not based on Lagrangian structures to combine the beneficial effects of each method. A major effort is currently furnished to adapt these methods to three-dimensional spaces and capture three-dimensional effects, such as upwelling and downwelling in coastal areas.

---

[5] By *nearly* Lagrangian tracers, we refer to particles following almost exactly the currents. In this case, the Lagrangian structures

computed in this paper and the suggested algorithm apply to their transport.

[6] Coastal flows may not have any barriers for certain periods of time or at all.

References

[1] F.G. Prahl, E. Crecellus, R. Carpenter, Polycyclic aromatic hydrocarbons in Washington coastal sediments: an evaluation of atmospheric and riverine routes of introduction, Environ. Sci. Technol. 18 (1984) 687–693.

[2] D.W. Rice, C.P. Seltenrich, R.B. Spies, M.L. Keller, Seasonal and annual distribution of organic contaminants in marine sediments from the elkhorn slough, moss landing harbor and near-shore Monterey Bay, California, Environ. Pollut. 82 (1993) 79–91.

[3] K. Verschueren, Handbook of Environmental Data on Organic Chemicals, Van Nostrand Reinhold Co, New York, 1983.

[4] S. Shadden, F. Lekien, J. Marsden, Definition and properties of lagrangian coherent structures: Mixing and transport in two-dimensional aperiodic flows, Physica D (2005) submitted.

[5] L.K. Shay, T.M. Cook, B.K. Haus, J. Martinez, H. Peters, A.J. Mariano, P.E. An, S. Smith, A. Soloviev, R. Weisberg, M. Luther, VHF radar detects oceanic submesoscale vortex along the Florida Coast, EOS, Trans., Am. Geophys. Union 81 (19) (2000) 209–213.

[6] L.K. Shay, T.M. Cook, H. Peters, A.J. Mariano, R. Weisberg, P.E. An, A. Soloviev, M. Luther, Very high frequency radar mapping of the surface currents, IEEE J. Oceanogr. Eng. 27 (2002) 155–169.

[7] G. Haller, Distinguished material surfaces and coherent structures in 3D fluid flows, Physica D 149 (2001) 248–277.

[8] T. Webb, R.B. Tomlinson, Design procedures for effluent discharge into estuaries during ebb tide, J. Environ. Eng. 118 (1992) 338–362.

[9] R. Smith, Optimal use of holding tanks to avoid pollution in narrow estuaries, IMA J. Appl. Math. 51 (1993) 187–199.

[10] J.H. Bikangaga, V. Nussehi, Application of computer modeling techniques to the determination of optimal effluent discharge policies in tidal water systems, Water Res. 29 (1995) 2367–2375.

[11] R.T. Giles, Optimal strategies for discharging pollutants into narrow estuaries, Water Res. 29 (1995) 563–569.

[12] R. Smith, Using small holding tanks to reduce pollution in narrow estuaries, J. Hydraul. Eng. 124 (1998) 117–122.

[13] G. Haller, Lagrangian coherent structures from approximate velocity data, Phys. Fluids A 14 (2002) 1851–1861.

[14] H. Ridderinkhof, J.T.F. Zimmerman, Chaotic stirring in a tidal system, Science 258 (1992) 1107–1111.

[15] P.D. Miller, C.K.R.T. Jones, A.M. Rogerson, L.J. Pratt, Quantifying transport in numerically generated velocity fields, Physica D 110 (1997) 1–18.

[16] C. Coulliette, S. Wiggins, Intergyre transport in a wind-driven, quasigeostrophic double gyre: An application of lobe dynamics, Nonlin. Proc. Geophys. 8 (2001) 69–94.

[17] A. Provenzale, Transport by coherent barotropic vortices, Ann. Rev. Fluids Mech. 31 (1999) 55–93.

[18] T.Y. Koh, B. Legras, Hyperbolic lines and the stratospheric polar vortex, Chaos 12 (2) (2002) 382–394.

[19] C.G. Farnetani, H. Samuel, Lagrangian structures and stirring in the earth's mantle, Earth Planet. Sci. Lett. 206 (3–4) (2003) 335–348.

[20] G. Boffetta, G. Lacorata, G. Radaelli, A. Vulpiani, Detecting barriers to transport: a review of different techniques, Physica D 159 (1–2) (2001) 58–70.

[21] D. Prandle, The fine-structure of nearshore tidal and residual circulations revealed by HF radar surface current measurements, J. Phys. Oceanogr. 17 (1987) 231–245.

[22] H.C. Shay, L.K. Graber, D.B. Ross, R.D. Chapman, Mesoscale ocean surface current structure etected by HF Radar, J. Atmos. Ocean. Tech. 12 (1995) 881–900.

[23] R.H. Stewart, J.W. Joy, Hf radio measurements of surface currents, Deep-Sea Res. 21 (1974) 1039–1049.

[24] H. Peters, L.K. Shay, A.J. Mariano, T.M. Cook, Current variability on a narrow shelf with large ambient vorticity, J. Geophys. Res. Oceans 107 (C8) (2002) art. no. 3087.

[25] F. Lekien, J. Marsden, Tricubic interpolation in three dimensions, Int. J. Numer. Methods Eng. 63 (3) (2005) 455–471.

[26] M.W. Hirsch, S. Smale, Differential Equations Dynamical Systems and Linear ALgebra, Academic Press, New York, 1974.

[27] B.L. Lipphardt Jr., A.D. Kirwan Jr., C.E. Grosch, J.K. Lewis, J.D. Paduan, Blending HF radar and model velocities in Monterey Bay through normal mode analysis, J. Geophys. Res. Oceans 105 (C2) 425–450 (2000) 3.

[28] F. Lekien, C. Coulliette, R. Bank, J. Marsden, Open-boundary modal analysis: Interpolation extrapolation and filtering. J. Geophys. Res. Oceans. (109), C12004, doi:10.1029/2004JC002323, 2004.

[29] C.W. Rowley, T. Colonius, R.M. Murray, Model reduction for compressible flows using pod and galerkin projection, Physica D 189 (1–2) (2004) 115–129.

[30] P.E. Lermusiaux, Data assimilation via error subspace statistical estimation. Part II: Middle Atlantic Bight shelfbreak front simulations and ESSE validation, Month. Weath. Rev. 127 (7) (1999) 1408–1432.

[31] P.F.J. Lermusiaux, A.R. Robinson, Data assimilation via error subspace statistical estimation. Part I: Theory and schemes, Month. Weath. Rev. 127 (7) (1999) 1385–1407.

[32] I. Shulman, C.R. Wu, J.K. Lewis, J.D. Paduan, L.K. Rosenfeld, J.C. Kindle, S.R. Ramp, C.A. Collins, High resolution modeling and data assimilation in the Monterey Bay area, Contin. Shelf Res. 22 (8) (2002) 1129–1151.

[33] A. Molcard, L.I. Piterbarg, A. Griffa, T.M. Ozgokmen, A.J. Mariano, Assimilation of drifter observations for the reconstruction of the Eulerian circulation field, J. Geophys. Res. Oceans 108 (C3) (2003) Art. No. 3056.

[34] A. Molcard, A. Griffa, T.M. Ozgokmen, Lagrangian data assimilation in multilayer primitive equation ocean models, J. Atmos. Ocean. Technol. 22 (1) (2005) 70–83.

[35] B. Sandstede, S. Balasuriya, C.K.R.T. Jones, P. Miller, Melnikov theory for finite-time vector fields, Nonlinearity 13 (4) (2000) 1357–1377.

[36] G. Lapeyre, B.L. Hua, B. Legras, Comments on 'Finding finite-time invariant manifolds in two-dimensional velocity fields', Chaos 11 (2001) 427.

[37] G. Haller, Lagrangian structures and the rate of strain in a partition of two-dimensional turbulence, Phys. Fluids A 13 (2001) 3368–3385.

[38] G.C. Yuan, L.J. Pratt, C.K.R.T. Jones, Barrier destruction and lagrangian predictability at depth in a meandering jet, Dynam. Atmos. Oceans 35 (1) (2002) 41–61.

[39] I. De Vries, R.N.M. Duin, J.C.H. Peeters, F.J. Los, M. Bokhorst, R.W.P.M. Laane, Patterns and trends in nutrients and phytoplankton in dutch coastal waters: Comparison of time-series analysis ecological model simulation and mesocosm experiments, ICES J. Marine Sci. 55 (4) (1998) 620–634.

[40] M. Drago, B. Cescon, L. Iovenitti, A three-dimensional numerical model for eutrophication and pollutant transport, Ecol. Model. 145 (1) (2001) 17–34.

[41] I. Shulman, S.H.D. Haddock, D.J. McGillicuddy, J.D. Paduan, W.P. Bissett, Numerical modeling of bioluminescence distributions in the coastal ocean, J. Atmos. Ocean. Technol. 20 (7) (2003) 1060–1068.

[42] J.D. Paduan, I. Shulman, HF radar data assimilation in the Monterey Bay area, J. Geophys. Res. Oceans 109 (C7) (2004) Art. No. C07S09.

[43] E. Boss, L. Thompson, Lagrangian and tracer evolution in the vicinity of an unstable jet, J. Phys. Oceanogr. 29 (2) (1999) 288–303.

[44] M.J. Tremblay, J.W. Loder, F.E. Werner, C.E. Naimie, F.H. Page, M.M. Sinclair, Drift of sea scallop larvae placopecten magellanicus on Georges Bank – A model study of the roles of mean advection larval behavior and larval origin, Deep Sea Res. Part II – Topic. Stud. Oceanogr. 41 (1) (1994) 7–49.

[45] A.C. Poje, G. Haller, Geometry of cross-stream mixing in a double- gyre ocean model, J. Phys. Oceanogr. 29 (1999) 1649–1665.

[46] A.C. Poje, G. Haller, I. Mezic, The geometry and statistics of mixing in aperiodic flows, Phys. Fluids 11 (10) (1999) 2963–2968.

[47] G. Haller, G. Yuan, Lagrangian coherent structures and mixing in two-dimensional turbulence, Physica D 147 (2000) 352–370.

[48] F. Lekien, N. Leonard, Dynamically consistent lagrangian coherent structures, in: Experimental Chaos, vol. 742, American Institute of Physics, March 2004, pp. 132–139.

[49] F. Lekien, C. Coulliette, G. Haller, J. Paduan, J. Marsden, Optimal pollution release in Monterey Bay based on nonlinear analysis of coastal radar data, Environ. Sci. Technol. (2005) submitted.

[50] G. Boffetta, A. Celani, A. Crisanti, A. Vulpiani, Relative dispersion in fully developed turbulence: Lagrangian statistics in synthetic flows, Europhys. Lett. 46 (2) (1999) 177–182.

[51] G. Boffetta, M. Cencini, S. Espa, G. Querzoli, Experimental evidence of chaotic advection in a convective flow, Europhys. Lett. 48 (6) (1999) 629–633.

[52] V. Artale, G. Boffetta, A. Celani, M. Cencini, A. Vulpiani, Dispersion of passive tracers in closed basins: Beyond the diffusion coefficient, Phys. Fluids 9 (11) (1997) 3162–3171.

[53] G. Boffetta, A. Celani, M. Cencini, G. Lacorata, A. Vulpiani, Nonasymptotic properties of transport and mixing, Chaos 10 (1) (2000) 50–60.

[54] M.C. Jullien, Dispersion of passive tracers in the direct enstrophy cascade: Experimental observations, Phys. Fluids 15 (8) (2003) 2228–2237.

[55] G.A. Voth, G. Haller, J.P. Gollub, Experimental measurements of stretching fields in fluid mixing. Phys. Rev. Letters 88 (25), (2002), 254501.1–254501.4.