

Optimal Pollution Mitigation in Monterey Bay Based on Coastal Radar Data and Nonlinear Dynamics

Chad Coulliette

Control & Dynamical Systems and Environmental Science & Engineering,
California Institute of Technology

Francois Lekien*

Mechanical and Aerospace Engineering, Princeton University

Jeffrey Paduan

Naval Postgraduate School

George Haller

Mechanical Engineering, Massachusetts Institute of Technology

Jerry Marsden

Control & Dynamical Systems, California Institute of Technology

May 28, 2007

*Correspondence to: 220 Engineering Quadrangle, Princeton University, Princeton, NJ 08544,
1-609-258-0303, lekien@princeton.edu.

Abstract

High-frequency (HF) radar technology produces detailed velocity maps near the surface of estuaries and bays. The use of velocity data in environmental prediction, nonetheless, remains unexplored. In this paper, we uncover a striking flow structure in coastal radar observations of Monterey Bay, along the California coastline. This complex structure governs the spread of organic contaminants, such as agricultural run-off which is a typical source of pollution in the bay. We show that a HF radar-based pollution release scheme using this flow structure reduces the impact of pollution on the coastal environment in the bay. We predict the motion of the Lagrangian flow structures from finite-time Lyapunov exponents of the coastal HF velocity data. From this prediction, we obtain optimal release times, at which pollution leaves the bay most efficiently.

Introduction

Pollution in coastal areas may impact the local ecosystem dramatically if the pollutants recirculate near the coast rather than leaving for the open ocean, where they are dispersed and then safely absorbed (1, 2). This article shows that accurate current measurements and dynamical systems theory can help in designing timed pollution release with the desirable outcome. Inspired by previous investigations (3), we consider a holding tank where pollutants can be temporarily stored and released at a later time.

The focus of our study is the Elkhorn Slough and the Duke Energy Moss Landing power plant, both of which are located near the Moss Landing Harbor in Monterey Bay (see Fig. 1). The Elkhorn Slough is a regular source of organic con-

taminants such as dichlorodiphenyl-trichloroethane (DDTs) and polychlorinated biphenyl (PCBs) from agricultural run-off, phthalic acid esters (PAEs) from plasticizer manufacturing, insecticidal sprays, wetting agents and repellents, and polycyclic aromatic hydrocarbons (PAHs) from the combustion of natural fossil fuels (1, 2). In addition, the Moss Landing power plant is a source of thermal pollution, which exhausts through a pipe that extends 200 meters into Monterey Bay (see Fig. 2).

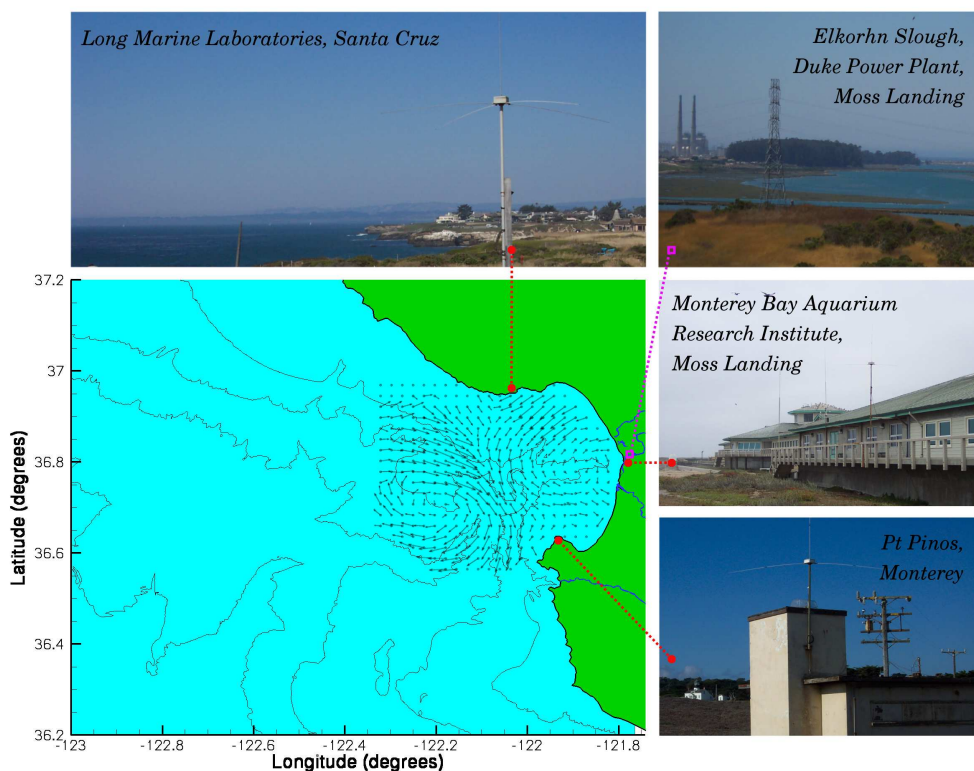


Figure 1: Locations of three CODAR SeaSonde HF radar systems around Monterey Bay. Top, bottom and right photographs show the HF radar antenna at Santa Cruz, Point Pinos, and Moss Landing that were used to measure the current data, respectively. Also shown are the footprint at 08:00 GMT on August 8th, 2000 (4–6) and bottom topography contours at various depths.

In contrast to earlier approaches to timed pollution release from holding tanks (3, 7–10), we avoid the use of simplified models and target measured ocean data directly. This strategy accommodates constantly changing flow conditions, an essential requirement for any pollution control algorithm of practical use. Another novel feature of our study is the use of finite-time dynamical systems methods (11–13) for the analysis of HF radar data. The recent interest in the development and application of such methods stems from the realization that stirring in mesoscale geophysical flows is governed by coherent structures of finite lifespan (13–16).



Figure 2: Aerial view of the Elkhorn Slough and the Duke power plant (right circle on the photograph). The plant exhausts warm, desalinated water through a pipe that extends 200 meters off the beach. The circle on the left of the photograph indicates the outlet of the pipe and the plume.

The presence of coherent features in geophysical flow data prevents the ap-

plication of homogeneous and isotropic turbulence theory (11) while the temporal irregularity and spatial complexity of such data renders the classic techniques of chaotic advection inapplicable (17–20). Most coastal flows fall into this intermediate regime. They are too energetic to be modeled as steady or periodic, but there is insufficient energy to reach a state where homogeneous turbulent diffusion would be adequate for understanding transport. This quasi-turbulent regime is chaotic, and thus extremely sensitive to initial conditions. Only a small change in the initial position or the release time of some material can considerably affect its trajectory. This is why a dynamical systems approach to transport is often needed for coastal flows.

In this paper, we use the radar measurements to identify the Lagrangian Coherent Structures (LCSs), which govern chaotic stirring of any Lagrangian particles. Specifically, we use Lyapunov exponents to find a highly convoluted LCS that repels nearby fluid parcels and, hence, acts as a barrier between two different types of motion: recirculation and escape from the bay. Recent work (13) shows that the flux across the LCS is negligible for the lifetime of the structure. Release of pollutants on one side of this moving fluid structure will result in sustained recirculation of the contaminant in the bay. If, however, pollution is released on the other side of the repelling material line, then the contamination will quickly clear from coastal regions and head towards the open ocean. Clearly, the latter scenario is more desirable. We propose an algorithm that uses real-time HF radar data to predict release times leading to the desired pollution behavior. A similar approach should work for optimizing the release of pollution into the atmosphere, rivers, lakes, or other waterways where sufficiently accurate wind or current data is available, and the release of pollution can be contained until an appropriate release time.

Higher frequencies are typically necessary for smaller regions (21).

High-frequency radar measurements

Our analysis makes use of high frequency (HF) radar technology (4–6), which is now able to resolve time-dependent Eulerian flow features in coastal surface currents. Such an HF radar installation has been operating in Monterey Bay since 1994 (6). In our study, we use data from this installation, acquired by the three HF radar antennas (shown in Fig 1.), binned every hour on a horizontal uniform grid with 1 km by 1 km intervals. An example of an HF radar footprint of the bay at 05:00 GMT, August 12, 2000 is shown in Fig. 1.

The surface current patterns in Monterey Bay are part of a dynamic upwelling system dominated by along shore wind forcing. The counterclockwise circulation pattern shown in Fig. 1, including the strong jet-like flow from north to south across the mouth of the Bay, is representative of the currents under strong, upwelling-favorable (from the northwest) wind conditions. Such winds are common, particularly during the summer months. However, periods of three to five days of upwelling favorable winds are generally followed by a shorter period of weak or reversed winds known as relaxation periods. During relaxation periods, the surface currents are generally weaker and less organized and they often exhibit a narrow band of south-to-north flow across the mouth of Monterey Bay.

To connect with the vast literature on dynamical systems, notice that the availability of measured velocities in the bay removes the need for a model based on partial differential equations. If the position of a fluid particle in Monterey Bay is

referred to as a vector \mathbf{x} , it obeys the ordinary differential equation

$$\dot{\mathbf{x}} = \mathbf{v}(\mathbf{x}, t) . \quad (1)$$

where $\mathbf{v}(\mathbf{x}, t)$ is the velocity at time t and position \mathbf{x} . The form of Eq. (1) is a generic time-dependent dynamical system (11) and demonstrates the connection between a measured velocity field and the vast literature on dynamical systems techniques. Rather than modeling, we are demonstrating a method for analyzing Lagrangian trajectories computed from any velocity field: measured, modeled or assimilated. In this article, the velocity field, $\mathbf{v}(\mathbf{x}, t)$ is provided by the high-frequency (HF) radar measurements of near-surface currents in Monterey Bay.

Since the velocity data is measured, there is some measurement error, as well as vectors that could not be resolved in some areas or at some times. Various techniques such as Open-boundary Modal Analysis (OMA) (22) are available for filtering, interpolating and extrapolating this data. Ocean modeling and data assimilation schemes have also proved to be an adequate source of dynamical systems (11, 23, 24).

We chose to use HF radar data *without any filtering, interpolation or extrapolation*. The objective of this work is to extract and use the coherent structures from the data, without any possible correlation with a filtering method or a model. Once the existence of a flow structure has been established for unfiltered data, modal techniques can be used to increase the smoothness of the measurements and structures.

The HF radar data gives the velocity field $\mathbf{v}(\mathbf{x}, t)$ but we are concerned with making deductions based on the resulting *flow* $\mathbf{x}(t; t_0, \mathbf{x}_0)$, i.e., the solution of

Eq. (1) that satisfies $\mathbf{x}(t_0) = \mathbf{x}_0$, where t_0 and \mathbf{x}_0 are the initial time and position of the trajectory.

The temporal complexity of the currents becomes evident from tracking different evolutions of a fluid parcel (a model for a blob of contaminant) released at the same precise location, but at slightly different times. We show the results of two such numerical experiments in Fig. 3.




Figure 3: Evolution of two parcels of contaminants released from the same position near Moss Landing at 22:00 GMT, August 6, 2000 (black) and 09:00 GMT, August 7, 2000 (white), plotted together with the snapshot of surface currents observed at 08:00 GMT on 8 August 2000. The motion of the two parcels is shown through daily snapshots over eight days. Note that the black parcel remains in the bay, while the white parcel departs from the bay.

Using available HF velocity data, we advected the fluid particles using a 4th order Runge-Kutta algorithm combined with tricubic interpolation in space and time. These particle trajectories are used to approximate the flow map, which associates initial positions, \mathbf{x}_0 , to final positions, \mathbf{x} . These numerical algorithms have been compiled into a software package called ManGen¹.

Figure 3 shows that one contaminant parcel remains in the bay, whereas the other parcel exits the bay and moves immediately towards the open ocean. The latter scenario (the white parcel on Fig. 3) is highly desirable, because it minimizes the impact of the contaminant on coastal waters, by causing it to be safely dispersed in the open ocean. This observation inspires us to understand and predict different evolution patterns of the same fluid parcel, depending on its initial location and time of release.

Lagrangian Coherent Structures

To understand the evolution of fluid parcels, we use a geometric description of mixing from nonlinear dynamical systems theory. Autonomous and time-periodic fluid flows have long been known to produce chaotic advection (25), i.e., irregular stirring of fluid parcels. Instrumental in this stirring are stable and unstable manifolds of hyperbolic fluid trajectories (26). These structures are material curves formed by fluid trajectories that converge to (resp. diverge from) a hyperbolic trajectory. For near-incompressible flows, the convergence within a stable manifold causes the manifold itself to repel nearby fluid parcels. As a result, stable manifolds act as repelling material lines that send fluid parcels on their two sides to different spa-

¹<http://www.lekien.com/~francois/software/mangen>

tial regions. Conversely, unstable manifolds act as attracting material lines, targets along which fluid parcels spread out and form striations. We refer to attracting and repelling material lines jointly as hyperbolic material lines.

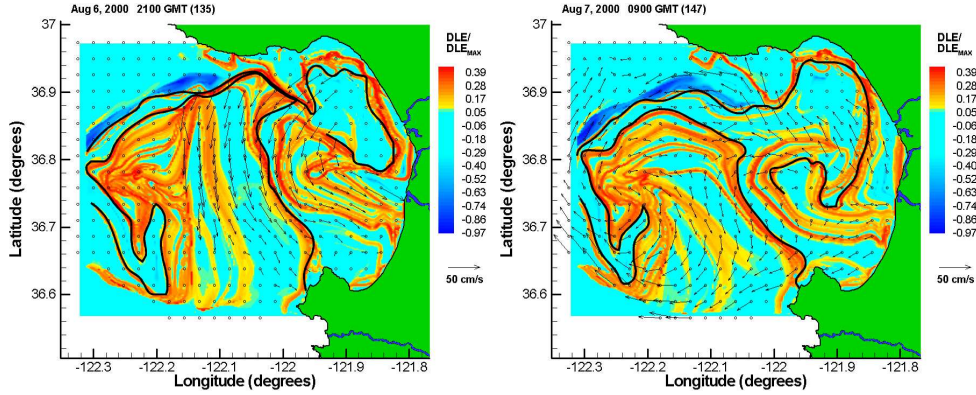


Figure 4: Distribution of $\mathcal{L}_t(t_0, \mathbf{x}_0)$ in Monterey Bay at time $t_0 = 21:00$ GMT, August 8, 2000 (left panel) and at time $t_0 = 09:00$ GMT Aug 7, 2000 (right panel). Superimposed on these plot are the dominant stable LCS as indicated by the ridges of $\mathcal{L}_t(t_0, \mathbf{x}_0)$ (black curves).

Recent progress in nonlinear dynamical systems has extended the above geometric picture to velocity fields with general time dependence, such as the surface velocity field of Monterey Bay. Families of hyperbolic material lines continue to organize finite-time mixing in such flows, even when the flow becomes quasi-turbulent (12). Several numerical algorithms and theoretical criteria have been proposed to identify hyperbolic material lines in general velocity data sets (11, 13–15, 27–30). Here we compute finite-time Lyapunov exponents (27). We start with a grid of initial particle positions \mathbf{x}_0 distributed across the domain at time t_0 . These are mapped to a later position $\mathbf{x}(t; t_0, \mathbf{x}_0)$ at time t_0 . We begin by computing the

Cauchy-Green strain tensor at time t ,

$$\mathbf{C}_t(\mathbf{x}_0, t_0) = \left[\frac{\partial \mathbf{x}(t; t_0, \mathbf{x}_0)}{\partial \mathbf{x}_0} \right]^\top \left[\frac{\partial \mathbf{x}(t; t_0, \mathbf{x}_0)}{\partial \mathbf{x}_0} \right], \quad (2)$$

where the superscript \top refers to the transpose of a matrix. We compute the eigenvalues of $\mathbf{C}_t(\mathbf{x}_0, t_0)$ at some time t , long before or after the reference time t_0 . The largest eigenvalue $\sigma_t(\mathbf{x}_0, t_0)$ of \mathbf{C}_t is also the largest singular value of the flow map and, for an infinitesimal grid spacing, typically behaves as $\sigma_t(t_0, \mathbf{x}_0) \sim e^{2\gamma(t-t_0)}$. As shown in (13), the coefficient γ approximates the rate of stretching about the trajectory $\mathbf{x}(t; t_0, \mathbf{x}_0)$. As a result, we define,

$$\gamma_t(\mathbf{x}_0, t_0) = \frac{1}{t - t_0} \ln \sqrt{\sigma_t(\mathbf{x}_0, t_0)}, \quad (3)$$

as our “stretching coefficient”. To compare results at different times t_0 , we define

$$\mathcal{L}_t(\mathbf{x}_0, t_0) = \frac{\gamma_t(\mathbf{x}_0, t_0)}{\max_{\mathbf{x}_0} \{\gamma_t(\mathbf{x}_0, t_0)\}}, \quad (4)$$

as the normalized finite-time Lyapunov exponent. We are interested in local maximizing curves or “ridges” of the scalar field $\mathcal{L}_t(t_0, \mathbf{x}_0)$ because they represent repelling material lines (14, 31). By ridges $\mathbf{c}(s, t)$, where $s \in]a, b[$, we mean a gradient curve of \mathcal{L} that minimizes the (negative) curvature of \mathcal{L} in the direction orthogonal to the ridge at each time t . More precisely, the ridge is a smooth curve $\mathbf{c}(s, t)$ that satisfies,

$$\frac{\partial \mathbf{c}}{\partial s} \times \nabla \mathcal{L} = 0. \quad (5)$$

where the cross product ensures that $\mathbf{c}'(s, t)$ is parallel to $\nabla \mathcal{L}$. This first condition

implies that the ridge is one of the many gradient curves of \mathcal{L} . We select isolated, distinguished gradient curves by requiring also

$$\mathbf{n}^\top \frac{\partial^2 \mathcal{L}}{\partial \mathbf{x}^2} \mathbf{n} = \min_{\mathbf{u} \neq \mathbf{0}} \frac{\mathbf{u}^\top \frac{\partial^2 \mathcal{L}}{\partial \mathbf{x}^2} \mathbf{u}}{\|\mathbf{u}\|^2} \quad (6)$$

where \mathbf{n} is the unit normal vector to the ridge at point $\mathbf{c}(s, t)$. The second condition states that, among all the gradient curves of \mathcal{L} , the ridge is the one that maximizes the curvature in the normal direction.

A more extensive description of the ridges of \mathcal{L} and their properties can be found in (13). If the Lyapunov exponent \mathcal{L} is viewed as the altitude, a ridge corresponds to a continental divide in the landscape. Any particle sitting on top of a ridge can fall on either side if there is a perturbation. The valleys on each sides of a ridge correspond to regions of qualitatively different dynamics. Particles on the ridges are *sensitive to initial conditions* because, depending on the direction of the initial perturbation, they can easily fall in different valleys. The same procedure performed backward in time (i.e., for $t < t_0$) would render attracting material lines at t_0 as ridges of $\mathcal{L}_t(t_0, \mathbf{x}_0)$.

The ridges of $\mathcal{L}_t(t_0, \mathbf{x}_0)$ divide the flow into regions of qualitatively different Lagrangian behavior (13, 32). Particles trapped inside the same loop of a LCS behave similarly and can be assimilated to a coherent mass of fluid. For this reason, the LCS provide a simple and geometric way to investigate the underlying velocity field and its action of particle trajectories. In this paper, we study the relationship between a symbolic fate (recirculating into the bay or escaping to the ocean) and the initial position with respect to a LCS .

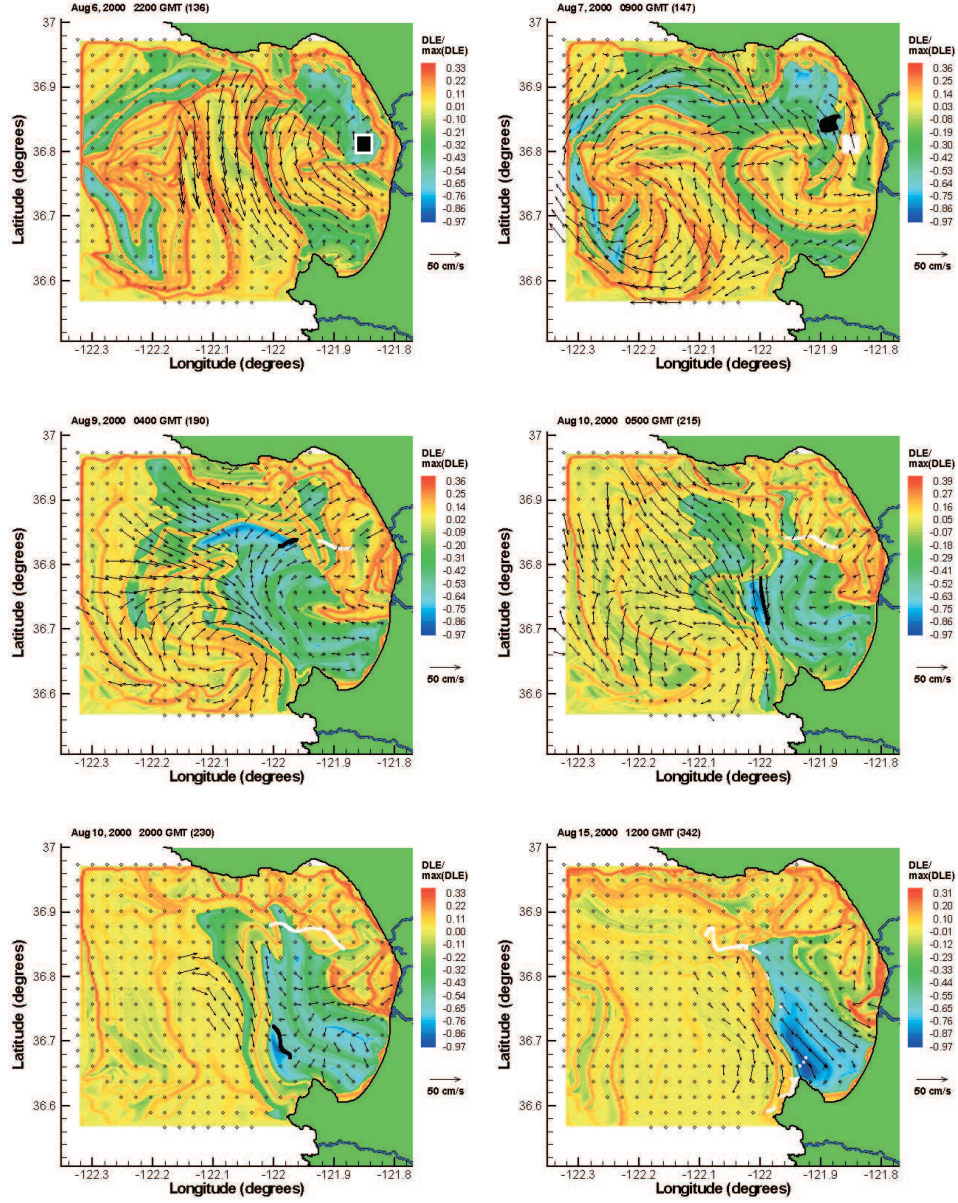


Figure 5: Two parcels of contaminants released from the same position near Moss Landing at 22:00 GMT, August 6, 2000 and at 09:00 GMT, August 7, 2000. The black arrows show instantaneous surface velocities captured by the HF radars. The ridges of the \mathcal{L}_t field reveal the hidden Lagrangian structure of the bay at the same time instants.

Analysis of HF radar data

We have performed the above analysis on a grid of fluid particles launched at 06:00 GMT on August 8, 2000. Using available HF velocity data, we advected the fluid particles for 200 hours, used their positions to approximate the flow map, and then numerically differentiated the flow map with respect to the initial positions of the particles. In this computation, we used a 4th order Runge-Kutta algorithm combined with 3rd order polynomial interpolation in space and time (33). We considered the coastline a free-slip boundary, and disregarded particles that crossed the linear fluid boundaries of the domain on the northern, southern and western edges. A sample result of such a computation is shown in Fig. 4, where the scalar distribution $\mathcal{L}_t(t_0, \mathbf{x}_0)$ is calculated over the initial grid \mathbf{x}_0 .

In agreement with the above general discussion, local maximizing curves, or ridges, on this plot form repelling material lines that act as moving barriers to transport. Note the highly convoluted maximizing curve that attaches to the southern coastline of the bay near Point Pinos (34), as seen in both panels of Fig. 4. The black curve can be viewed as a stable LCS — a curve of fluid particles converging to an attachment point moving back and forth along the coast in the vicinity of Point Pinos. This stable LCS divides the bay into two regions of different parcel behavior. Fluid mechanicians might recognize the black curve as a streakline which originates at the upwelling source in the center of the bay.

Fluid parcels on one side of the stable LCS will recirculate in the bay after they pass by the coastal attachment point. Parcels on the other side of the LCS exit to the open ocean after passing by the attachment point. This is the reason underlying the different parcel behaviors in Fig. 3: the same release location fall on different

sides of the stable LCS on August 6 and August 7. Figure 5 illustrates this point by superimposing the instantaneous positions of the stable LCS on snapshots of parcel positions. Recall that the behavior of the white parcel is highly desirable for the evolution of pollutants.

Optimal pollutant release times

The Elkhorn Slough and the Duke Energy Moss Landing Power Plant are both located near the Moss Landing Harbor, which is on the eastern shore of Monterey Bay. Both contribute to pollutants entering Monterey Bay.

An important consequence of the above analysis is *the existence of time intervals where released contaminants have either a high or low impact on the environment*. Our objective is to show that a pollution control algorithm based on a nonlinear dynamic analysis with Lyapunov exponents can achieve a significant *reduction in the impact of a contaminant in a coastal area, without reducing the total amount of contaminants released*.

To facilitate the discussion, we consider an exhaust pipe similar to that mentioned in Fig. 2, which carries pollution (e.g., chemical, thermal) from Moss Landing and the Elkhorn Slough to an offshore release site shown in Fig. 6.

Although building a pipeline is not necessary for our method, it is necessary to have some control over the release time and location of the pollutants. Thus to expedite our explanation we will imagine a pipeline which carries the contaminants from the Moss Landing area to an offshore release site at the same location that the black and white parcels were released. This hypothetical pipeline and release location are shown in Fig. 6.

For any given time, we consider a portion of the previously discussed LCS as it ascends along the coastline of the bay from Moss Landing, meandering past Santa Cruz. The meandering of the LCS causes it to intersect the axis of the pipeline in several points. These intersection points can be counted by following the LCS, starting from its coastal attachment point. We refer to the first intersection point as $\mathcal{L}_t^{peak}(t_0)$.

The end of the pipe is at the same location as the release site for the white and black parcels featured in Fig. 3. Figure 6 also shows the instantaneous intersection of the stable LCS (revealed by a ridge of the \mathcal{L}_t field) and the axis of the hypothetical pipeline.

The motion of the intersection point along the axis of the hypothetical pipeline is complicated, which is evident from the time history of the intersection location in Fig. 7. Superimposed on this plot are the release times and release location of the white and black parcels of Fig. 3. Recall that the reason for their different future behaviors is the difference in their initial position relative to the curve of Fig. 7. In particular, the white parcel exits the bay quickly because it enters the flow when the point of release lies between the \mathcal{L}_t^{peak} (curve of Fig. 7) and the edge of the pipeline (horizontal line on Fig. 7).

Notice that Fig. 7 proves the existence of time intervals where pollutants are quickly advected outside the bay. The objective of our pollution release algorithm is to maximize pollutant release during these time windows and to store pollutants in a tank outside these intervals.




Figure 6: A hypothetical pipeline carries contaminants to be released in the bay from the Moss Landing area. Also shown is the instantaneous intersection point of a peak in the LCS field (i.e., \mathcal{L}_t^{peak}) and the axis of the pipeline.

Real-time Coastal Pollution Management

Based on the analysis in the previous section, it is tempting to think that the intersection curve in Fig. 7 predicts directly times of pollution release that will lead to a quick exit from the bay. Why not simply release pollution when the curve indicating the \mathcal{L}_t^{peak} is well above the horizontal line marking the outlet of the pipeline? As in the case of the white parcel, such a release would certainly guarantee that the contaminant is initially east of the stable LCS and hence leaves the bay quickly as it will approach the Monterey Peninsula west of the separation line near Point




Figure 7: Oscillations of \mathcal{L}_t^{peak} along the axis of the pipeline. The zero reference time corresponds to 07:00 GMT, August 1, 2000. The horizontal line marks the location of the outlet of the pipeline. The black and white squares represent the release time and release longitude of the parcels featured in Fig. 3.

Pinos.

The above method is flawed for practical applications, because any point of the \mathcal{L}_t^{peak} curve in Fig. 7 is constructed from future velocity data over the next 200 hours. In other words, to predict when and where to release pollution on Monday, we would need knowledge of the currents in the bay up until approximately Tuesday of the following week. Such future data is clearly unavailable at the time when a decision has to be made. Trying to predict the velocity field in the bay for more than 3 days might be unrealistic, or at least very difficult, because of the spatial and temporal complexity of the flow. Instead, we propose a focused Lagrangian

prediction.

As a first step, we modify our calculation of $\mathcal{L}_t(t_0, \mathbf{x}_0)$. We fix $t = 22:00$ GMT, Aug 6, 2000 as today, or the “present time”, when we would like to make our prediction. For any earlier time t_0 , we calculate the peak of the \mathcal{L}_t ridge; this means that the future window in our computation is gradually shrinking to zero as t_0 approaches the present time t . As expected, this results in a gradual (albeit surprisingly slow) growth of error between the actual \mathcal{L}_t^{peak} (computed with a constant 200 hour future window) and the real-time \mathcal{L}_t^{peak} (computed with a shrinking future window). The actual and the real-time \mathcal{L}_t^{peak} locations, as functions of time, are plotted in Fig. 8.

The real-time $\mathcal{L}_t(t_0, \mathbf{x}_0)$ peak curve approximates the actual (200 hour) curve with an error less than, approximately, 10% up to 8 hours before the “present time.” During the last 8 hours, the error on the predicted \mathcal{L}_t^{peak} becomes prohibitive. Note that the inserts in Fig. 8 show slices of the \mathcal{L}_t contours along the axis of the pipeline at $t - t_0 = 20$ hours and $t - t_0 = 100$ hours. We observe from inserts that the position of the barrier is identified by a sharp ridge. It is best to identify the ridge not by its maximum magnitude, but by the gradient in a direction that is approximately orthogonal to the \mathcal{L}_t ridge. The \mathcal{L}_t ridge intersects the axis of our pipeline in a nearly orthogonal direction, so we will use the axis of the pipeline to examine the gradient of the \mathcal{L}_t ridge. In Fig. 9, we examine the maximum value of the ridge and the gradient of the \mathcal{L}_t ridge as a function of the time used to compute the \mathcal{L}_t ridge. Note that the maximum value of the \mathcal{L}_t ridge shown in the left panel of Fig. 9 displays no behavior which indicates a clear choice for the computational time needed to evolve the \mathcal{L}_t contours. However, the gradient of the \mathcal{L}_t is more useful. During the first 8 hours, the longitudinal component of




Figure 8: Oscillation of the \mathcal{L}_t^{peak} along the axis of the pipeline over a 150 hour period, from 07:00 GMT, August 1, 2000 to 22:00 GMT, Aug 6, 2000. The green curve is the real-time curve based on information up to the “present time” (computed with a *shrinking* time window), with the \mathcal{L}_t^{peak} located from the gradient of a numerical maximization along the pipe axis. The red curve is the actual \mathcal{L}_t^{peak} location (computed with a *constant* 200-hour time window). The inserts show a slice of the \mathcal{L}_t contours along the axis of the pipeline at $t - t_0 = 20$ hours and $t - t_0 = 100$ hours.

the \mathcal{L}_t gradient increases linearly with the time used to compute the \mathcal{L}_t contours. After the first 8 hours, the magnitude of the longitudinal component of the gradient begins to oscillate due to nonlinear effects. Thus the minimum integration time which provides a well-defined \mathcal{L}_t ridge is approximately 8 hours, which matches the previous qualitative observation. The magnitude of the longitudinal component of the gradient may still increase after 8 hours, but its growth is no longer linear in time, and thus additional computational time is not as beneficial after the first 8

hours. Consequently, we need to stop our real-time \mathcal{L}_t calculation about eight hours *before* the “present time” to take advantage of the steep increase in the gradient during that time.




Figure 9: The left panel shows the relative maximum of $\mathcal{L}_t(t_0, \mathbf{x}_0)$ as a function of computational time. The right panel shows the maximum longitudinal component of the gradient of \mathcal{L}_t . During the first 8 hours the gradient continues to grow in magnitude, thus making the \mathcal{L}_t ridge more pronounced and identifiable. After 8 hours the magnitude of longitudinal component of the gradient oscillates.

As a second step, we identify the main frequency components of the real-time \mathcal{L}_t^{peak} curve over the shortened time interval $[t_0, t - 8 \text{ hours}]$. Shown in Fig. 10, the power spectrum density of the real-time \mathcal{L}_t^{peak} curve highlights seven dominant frequency components, with the importance of each frequency determined by the area under the corresponding peak in the spectrum. Surprisingly, the most influential component in this particular time interval is not the tidal oscillation (with a period of 24 hours) or any of its harmonics, but rather a component with a period of 8.6 days. As it was already obvious from Fig. 7, this means that, during the 22 days observed, the \mathcal{L}_t^{peak} stays on one side of the outlet for about 4.3 days before it

crosses to the other side. Note that the 8.6 day period was computed from 22 days of data. It is consistent with the major wind reversals observed during this data collection period, but will most likely change based on seasonal changes in winds, currents outside the bay, and other factors driving flow in the bay.

To complete our prediction procedure, we used all the significant frequencies of the spectrum of this curve to predict the location of \mathcal{L}_t^{peak} along the axis of the pipeline into the near future. The amplitudes and phases of the prediction curve are determined by minimizing the norm of the difference (i.e., the integral of the squared difference) between fitted and real-time \mathcal{L}_t values. The left panel of Fig. 11 shows the predicted \mathcal{L}_t^{peak} together with the actual and the real-time locations of \mathcal{L}_t^{peak} . Note how faithfully the predicted curve reproduces the main features of the actual \mathcal{L}_t^{peak} oscillations.

In particular, the left panel of Fig. 11 predicts that releasing contaminants from the pipeline between 3 hours and 110 hours from the present time (22:00 GMT, Aug 6, 2000) will cause most of the pollution to exit Monterey Bay without recirculation. On the other hand, pollution released after 110 hours is not expected to leave the bay immediately due to the excursion of the actual \mathcal{L}_t^{peak} curve into longitudes on the coastal side of the pipe outlet. In this case, the algorithm should wait for about 3 hours and prepare to release pollutants and empty the holding tank for a period of about 107 hours. Not only does the algorithm predicts wheter or not to release pollution, but it also provides an estimate of the length of the discharge period and, hence, also set the rate at which the tank should be emptied.

On the left panel of Fig. 11, the next predicted “red” period starts 110 hours from the present time, while the actual red period turned out to start 118 hours after the present time. This means that the error in predicting the end of the release




Figure 10: Power spectrum density of the real-time \mathcal{L}_t^{peak} oscillations shown in Fig. 7. The dominant wavelength is 8.6 days and the spikes at 48 hours and 4 days indicate harmonics associated with the 24-hour tidal oscillation. The importance of each frequency is proportional to the area below the corresponding spike.

interval was approximately 8 hours with an horizon of 4 days.

To illustrate the efficacy of the above pollution release scheme, we repeated the same prediction procedure for a different “present time”, $t = 20:00$ GMT, Aug 17, 2000. The left panel of Fig. 11 shows that similar performances are achieved. In this case, the algorithm correctly predicts that the LCS is too far east and that




Figure 11: Actual, real-time, and predicted \mathcal{L}_t^{peak} location along the axis of the pipeline. The horizontal line marks the location of the outlet of the pipe. The color bar indicates the periods of desirable releases (green) and the periods to avoid (red). Each panel corresponds to a different “present time.”

pollutants should be redirected to the holding tank. It also predicts that the next “green” interval is 60 hours from the present time.

It is worth noticing that, in the second case, the period of the dominant mode in the \mathcal{L}_t^{peak} oscillation was 9.26 days. The difference in oscillation wavelength (8.6 days for $t_0 = 22:00$ GMT, Aug 6, 2000 and 9.26 days for $t_0 = 20:00$ GMT, Aug 17, 2000) is to be expected since this flow is highly time-dependent. This is evidence that a static analysis of the flow will never be sufficient to make predictions about Lagrangian transport in Monterey Bay. A nonlinear analysis of real time current measurements such as that described in this article is necessary.

More generally, the prediction method described above determines environmentally friendly future time windows. These windows last for about 100 hours, over which most of the pollution released from the pipeline will advect towards the open ocean. We marked the bottom of Fig. 11 with green bars for time periods that

result in the pollution exiting the bay and with red bars for release times that cause the pollution to remain within the bay.

Discussion

From the simulations and predictions presented in the previous section, the following general principles emerge:

- For best performance, the holding tank must be able to hold contaminants produced over approximately 5 days. Using such a tank, we can wait, if needed, for the entire disadvantageous half-period of the main \mathcal{L}_t^{peak} mode to pass.
- Previous work on optimal pollution release has focused on releasing the pollution at high tide or some constant time shift from high tide (9), but the methods used in such studies only hold for simplified models of coastal flows. When using actual current data as we are here, it can be seen that using such a release scheme for pollutants in Monterey Bay would not give optimal results. The complicated flow patterns in Monterey Bay, although influenced by tidal fluctuations, have their bay-scale retention characteristics set by the longer period fluctuations associated with the coastal wind forcing.
- The influence of the length of the pipeline reveals the chaotic nature of the flow in Monterey Bay. Selecting a longer pipeline will raise the horizontal line (outlet position) in Fig. 11, which in turn leads to shorter time windows for optimal release. This is the opposite of what we would expect a longer pipeline to do, that is, cause the pollution to exit the bay sooner. Shorter

pipelines do not, however, necessarily lead to quicker clearance, because pollutants would fall between secondary peaks of \mathcal{L}_t (visible in Fig. 6) and the coastline, thus requiring more revolutions around the bay before exiting.

These principles demonstrate the importance of a nonlinear analysis of the velocity field to understand how particles will advect in the field. Lagrangian and quasi-Lagrangian particles are not necessarily advected in a manner which is intuitive from a visual inspection of the velocity field alone.

In this paper we have combined surface radar observations and recent results from dynamical systems theory to identify a hidden dynamic structure of Monterey Bay. This structure, a highly convoluted repelling material line remains hidden both in instantaneous and averaged surface velocity plots. Yet the repelling LCS has a decisive influence on stirring in the bay: it repels nearby fluid parcels and hence induces qualitatively different behaviors for parcels released from its opposite sides. For pollutants, one of these behaviors, a quick escape to the open ocean, is highly desirable because it reduces the contamination of coastal areas.

As a particular use of our Lagrangian diagnostics, we have proposed a pollution release scheme that exploits the governing role of the repelling LCS in fluid transport. We assumed that pollution is released through a pipeline in the Moss Landing area, and showed how high-frequency radar data can be used to predict the position of the stable LCS relative to the pipeline outlet for a few days ahead of time. From this prediction, we have been able to determine environmentally friendly time windows of pollution release. These time windows last for about 100 hours, over which most of the pollution released from the pipeline will head towards the open ocean. When verified from actual “future” radar data, these pre-

dictions have proved very accurate: the error in predicting environmentally friendly time intervals of release remained consistently below 15%.

A general physical lesson from our analysis is that focused Lagrangian predictions for a geophysical flow can be feasible even if global Eulerian (i.e., velocity based) predictions are unrealistic. However, in the case that Eulerian velocity predictions are possible, through perhaps the prolongation of open-boundary modal coefficients (22) or data-assimilated hybrid models (35), the approach outlined in this article remains applicable, except that it is no longer necessary to make a prediction for the position of the LCS since the structures will have been computed from a “predicted” velocity field. The advantage of the approach outlined in this article is that the prediction of the LCS is one dimensional, whereas predicting the velocity field directly is a two or three dimensional problem. The accuracy and advantages of each approach need to be further investigated.

We need to stress, however, that the method presented in this article is based on near-surface HF velocity data. As a result, the pollution release scheme we described here only applies to contaminants that remain close to the ocean surface. A more general three-dimensional analysis could, in principle, be performed if velocity data at greater depths became available. Another assumption in this work is that the turbulent diffusive time scale for the contaminant is longer than the time of one recirculation in the bay. This assumption is to be verified via dye release studies and Lagrangian stochastic models for actual pollutants before a real-life implementation of our methods. Such an implementation would also require robustness with respect to measurement uncertainties and numerical errors. Recent results already show that Lagrangian coherent structures are remarkably robust, even under substantial errors, provided that the errors are deterministic and remain

localized in time (35).

In addition to finding optimal times to release pollutants so that the impact to the Elkhorn Slough or Monterey Bay is minimized, using the outlined dynamical systems approach has other ecological benefits. For example, the seawater sucked from the Elkhorn Slough and Moss Landing Harbor into the cooling intake system of the power plant is heated to approximately 20°C higher than natural temperatures and is expelled through a pipe that extends into Monterey Bay. The water contains billions of fish eggs, invertebrates and larvae. Some die and some live through the heating and cooling process. Where will the currents carry them once they are discharged into the bay? The methods presented here could help to answer what impact this daily relocation of dead and live species has on the local ecosystem of Monterey Bay, Moss Landing and the Elkhorn Slough.

An important conclusion of this paper is that it is possible to use nonlinear dynamical systems theory together with recent advances in current measurement techniques, such as HF radar or ADCPs, to analyze, understand, and predict where chemical contaminants, thermal pollution, and biological populations will be carried by the currents. This allows us to determine and mitigate the impact of various technologies on marine life in coastal zones.

Acknowledgments

The CODAR data used in this paper has been prepared by Michael Cook and Daniel Atwater. We are also grateful to Tricia Wilson (Elkhorn Slough National Estuarine Research Reserve), Thomas C. O'Reilly and Thomas Hoover (Monterey Bay Research Institute) for providing us with valuable information about the

sources and the transport of contaminants in the slough and the Moss Landing area.

The GPS fix for the Duke power plant outfall and some of the pictures in the manuscript were taken on board of Moss Landing Marine Lab’s Research Vessel *John H. Martin*. We thank Capt. Lee Bradford (Moss Landing Marine Laboratories), Bruce Howe (University of Washington), and Fumin Zhang (Princeton University) for their assistance. The aerial pictures were taken by unmanned camera on board of the UV-18A *Twin Otter* of the Naval Postgraduate School of Monterey and the Center for Interdisciplinary Remotely-Piloted Aircraft Studies. We thank Steven Ramp and Todd Anderson (NPS) for providing us with this data.

We are also grateful to Naomi Leonard, Eddie Fiorelli and Fumin Zhang (Princeton University) for enlightening discussions throughout this work. Integrated particles and Lagrangian coherent structures for this paper were computed using ManGen².

This research was funded by the Office of Naval Research (grant no.N00014-01-1-0208 and ASAP MURI no.N00014-02-1-0826). We thank Thomas Curtin, Manuel Fiadero and Reza Malek-Madani for their support and helpful comments.

References

- (1) F. G. Prahl, E. Crecllus, and R. Carpenter. Polycyclic aromatic hydrocarbons in Washington coastal sediments: an evaluation of atmospheric and riverine routes of introduction. *Environ Sci Technol* **1984**, 18, 687–693.
- (2) D. W. Rice, C. P. Seltenrich, R. B. Spies, and M. L. Keller. Seasonal and annual distribution of organic contaminants in marine sediments from the

²<http://www.lekien.com/~francois/software/mangen>

Elkhorn Slough, Moss Landing Harbor and near-shore Monterey Bay, California. *Environmental Pollution* **1993**, 82, 79–91.

- (3) R. Smith. Optimal use of holding tanks to avoid pollution in narrow estuaries. *IMA J Appl Math* **1993**, 51, 187–199.
- (4) T. M. Georges, J. A. Harlan, and R. A. Lematta. Large-scale mapping of ocean surface currents with dual over-the-horizon radars. *Nature* **1996**, 379, 434–436.
- (5) R. M. Goldstein and H. A. Zebker. Interferometric radar measurement of ocean surface currents. *Nature* **1987**, 328, 707–709.
- (6) J. D. Paduan and L. K. Rosenfeld. Remotely sensed surface currents in Monterey Bay from shore-based hf radar (codar). *J. Geophys. Res.* **1996**, 101, 20669–20686.
- (7) J. H. Bikangaga and V. Nussehi. Application of computer modeling techniques to the determination of optimal effluent discharge policies in tidal water systems. *Water Res* **1995**, 29, 2367–2375.
- (8) R. T. Giles. Optimal strategies for discharging pollutants into narrow estuaries. *Water Res.* **1995**, 29, 563–569.
- (9) R. Smith. Using small holding tanks to reduce pollution in narrow estuaries. *J. Hydraulic Eng.* **1998**, 124, 117–122.
- (10) T. Webb and R. B. Tomlinson. Design procedures for effluent discharge into estuaries during ebb tide. *J. Env. Eng.* **1992**, 118, 338–362.

- (11) C. Coulliette and S. Wiggins. Intergyre transport in a wind-driven, quasi-geostrophic double gyre: An application of lobe dynamics. *Nonlinear Proc. Geophys* **2001**, 8, 69–94.
- (12) G. Haller. Finding finite-time invariant manifolds in two-dimensional velocity fields. *Chaos* **2000**, 10, 99–108.
- (13) S.C. Shadden, F. Lekien, and J.E. Marsden. Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows. *Physica D* **2005**, 212(3-4), 271–304.
- (14) G. Haller. Distinguished material surfaces and coherent structures in 3d fluid flows. *Physica D* **2001**, 149, 248–277.
- (15) A. C. Poje and G. Haller. Geometry of cross-stream mixing in a double- gyre ocean model. *J. Phys. Oceanogr* **1999**, 29, 1649–1665.
- (16) O. M. Velasco Fuentes. Chaotic advection by two interacting finite-area vortices. *Phys. Fluids A* **2001**, 13, 901–912.
- (17) S. P. Beerens, H. Ridderinkhof, and J. T. F. Zimmerman. Longitudinal dispersion in natural streams. *Chaos, Solitons and Fractals* **1998**, 4, 1011–1029.
- (18) J. R. Stirling. Transport and bifurcation in a non-area-preserving two-dimensional map with applications to the discharge of pollution in an estuarine flow. *Physica D* **2000**, 144, 169–193.
- (19) A. J. Watson, M. J. Messias, E. Fogelqvist, and K. A. V. et.al. Scoy. Mixing and convection in the Greenland Sea from a tracer-release experiment. *Nature* **1999**, 401, 902–904.

- (20) J. T. F. Zimmerman. The tidal whirlpool— a review of horizontal dispersion by tidal and residual currents. *Netherlands J. Sea Res.* **1986**, 20, 133–154.
- (21) L. Shay, T. Cook, B. Haus, J. Martinez, H. Peters, A. Mariano, J. VanLeer, P. Edgar An, S. Smith, A. Soloviev, R. Weisberg, and M. Luther. VHF radar detects oceanic submesoscale vortex along Florida coast. *EOS, Transactions, American Geophysical Union* **2000**, 81(19), 209–213.
- (22) F. Lekien, C. Coulliette, R. Bank, and J. Marsden. Open-boundary Modal Analysis: Interpolation, extrapolation and filtering. *Journal of Geophysical Research – Oceans* **2004**, 109(C12004).
- (23) P.F.J. Lermusiaux, D.G.M. Anderson, and C.J. Lozano. On the mapping of multivariate geophysical fields: Error and variability subspace estimates. *Quarterly Journal of the Royal Meteorological Society* **2000**, 126(565), 1387–1429.
- (24) I. Shulman, C.R. Wu, J.K. Lewis, J.D. Paduan, L.K. Rosenfeld, J.C. Kindle, S.R. Ramp, and C.A. Collins. High resolution modeling and data assimilation in the Monterey Bay area. *Continental Shelf Research* **2002**, 22(8), 1129–1151.
- (25) H. Aref. Stirring by chaotic advection. *J. Fluid Mech.* **1984**, 143, 1–21.
- (26) J.M. Ottino. *The Kinematics of Mixing: Stretching, Chaos, and Transport*. Cambridge University Press, Cambridge, 1988.
- (27) G. Haller. Lagrangian structures and the rate of strain in a partition of two-dimensional turbulence. *Phys. Fluids A* **2001**, 13, 3368–3385.

- (28) B. Joseph and B. Legras. Relation between kinematic boundaries, stirring, and barriers for the Antarctic polar vortex. *J. Atmosph. Sci.* **2002**, 59(7).
- (29) G. Lapeyre, B. L. Hua, and B. Legras. Comments on ‘finding finite-time invariant manifolds in two-dimensional velocity fields’. *Chaos* **2001**, 11(2), 427–430.
- (30) P. D. Miller, C. K. R. T. Jones, A. M. Rogerson, and L. J. Pratt. Quantifying transport in numerically generated velocity fields. *Physica D* **1997**, 110, 1–18.
- (31) G. Haller. Lagrangian coherent structures from approximate velocity data. *Phys. Fluids A* **2002**, 14, 1851–1861.
- (32) F. Lekien, C. Coulliette, A. J. Mariano, E. H. Ryan, L. K. Shay, G. Haller, and J. Marsden. Pollution release tied to invariant manifolds: A case study for the coast of Florida. *Physica D* **2005**, 210 (1-2), 1–20.
- (33) F. Lekien and J. Marsden. Tricubic interpolation in three dimensions. *Int. Journal for Numerical Methods in Engineering* **2005**, 63(3), 455–471.
- (34) F. Lekien and G. Haller. Unsteady flow separation on slip boundaries. *J. of Fluid Mechanics* **2006**, Submitted.
- (35) P. F. J. Lermusiaux, C.-S. Chiu, G. C. Gawarkiev, P. Abbot, A. R. Robinson, P. J. Haley, P. J. Leslie, S. J. Majumdar, A. Pang, and F. Lekien. Quantifying uncertainties in ocean predictions. *Oceanography* **2006**, 19(1), 80–93.