Time-Dependent Dynamical Systems and Geophysical Flows

Thesis by
Francois Lekien

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

California Institute of Technology
Pasadena, California

2003

(Defended March 12, 2003)

It's my thesis –

But you can read it.

Preface

This thesis was submitted at the *California Institute of Technology* on March 1st, 2003, as a partial fulfillment of the requirements for the degree of Doctor of Philosophy in Control and Dynamical Systems. The thesis was defended on March 12th, 2003, in Pasadena, CA, and approved by the following research committee:

- Pr. Jerry Marsden, *Control & Dynamical Systems*,
  California Institute of Technology,
- Pr. Richard Murray, *Control & Dynamical Systems* and *Mechanical Engineering*,
  California Institute of Technology,
- Pr. Anthony Leonard, *Graduate Aeronautical Laboratories*,
  California Institute of Technology
- Pr. George Haller, *Department Mechanical Engineering*,
  Massachusetts Institute of Technology,
- Pr. Igor Mezić, *Mechanical and Environmental Engineering*,
  University of California - Santa Barbara.

The thesis is intended to be a complete research report. However, an effort was made to keep the chapters independent from each other to facilitate their publication. Some chapters have or will be submitted for publication in their present forms:

- Lekien, F., C. Coulliette, and J. Marsden, Lagrangian Structures in High-Frequency Radar Data and Optimal Pollution Timing, in Proceedings of the 7th Experimental Chaos Conference, *American Institute of Physics*, Accepted for publication on August 26, 2002.

- Lekien, F., C. Coulliette, G. Haller, J. Marsden, and J. Paduan, Optimal Pollution Release in Monterey Bay Based on Nonlinear Analysis of Coastal Radar Data, *Journal of Physical Oceanography*, submitted.

- Lekien, F., C. Coulliette, A.J. Mariano, E. Ryan, L.K. Shay, and J. Marsden, Lagrangian Structures in Very High-Frequency Radar Data and Optimal Pollution Timing, *Physica D*, submitted.

- Coulliette C., Lekien, F. Haller G., Paduan J. and Marsden J, When Should Pollution be Released? *Environmental Engineering*, submitted.

- Lekien, F., C. Coulliette, R. Bank, and J. Marsden, Open Boundary Modal Analysis (OMA): A Complete Functional Basis to Interpolate, Extrapolate and Filter Experimental Eulerian Data, *Journal of Geophysical Research*, in preparation.

- Lekien, F., and J. Marsden, Why Warm Core Rings Are Always Anticyclonic?, *Physica D*, in preparation.

- Lekien, F., and J. Marsden, Separatrices in High-Dimensional Phase Spaces: Application to Van Der Waals Dissociation, *Journal of Chemical Physics*, in preparation.

Animations related to this thesis, related papers and supporting documents and slides used during various presentations of part of this work can be found at the following address: `http://www.cds.caltech.edu/~lekien/thesis/`.

Acknowledgements

I wish to extend warm thanks to my officemate and friend, Chad Coulliette, for hours of exciting work and enriching discussions. Most of our work has been joined over the past 4 years and this thesis belongs to him as much as it does to me. I am very grateful to Jerry Marsden and Richard Murray at the *California Institute of Technology* for their help and support, as well as to George Haller at the *Massachusetts Institute of Technology*, Randy Bank at the *University of California, San Diego* and Igor Mezić at the *University of California, Santa Barbara*.

Most of the numerical results presented in this report have been obtained using a 32 node *Beowulf Linux alpha* cluster funded by *Compaq* and the *Office of Naval Research*. Maintaining such a system at the highest performance level has been an exciting and challenging task that involved many people, including Michel Tanguay, Richard Murray, Tim Colonius, Chad Coulliette, Steve Wiggins, Steve Koonin, Anthony Leonard and Manuel Fiadero.

This research was supported by two *ONR* grants: N00014-01-1-0208 and the AOSN project N00014-02-1-0826. I want to thank project managers Manuel Fiadero, Resa Malek-Madani and Wen Master for their enthusiasm for our projects over the past 5 years. I am grateful for the tuition and stipend generously provided by a *Poincaré* fellowship.

I also want to thank the staff at *Caltech* who maintains the facilities to a maximum level of efficiency and particularly the staff of the *Department of Control and Dynamical Systems*, including Charmaine Boyd, Betty-Sue Herrala, Siarra Walker and Harpy Hindoyan. I am also thankful to RuthAnne Bevier, who watched our hardware network, fixed many problems and corrected numerous security problems.

This list would not be complete without a special thanks to my parents and my sister, Brigitte, who supported me and encouraged me during many years of studies. I also want to show my gratitude to my friends Katharine Ratnoff, Caroline Englebienne, Philippe Chatelain, Sofie De Pauw, Catherine De Pauw, Nicolas Husson, Jayaseeli Ferir, Raphael Robiette and David Speeckaert for their help and support in many situations.

Time-Dependent Dynamical Systems
and Geophysical Flows

by

Francois Lekien

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

Abstract

This thesis presents a dynamical systems approach to transport and mixing in geophysical flows. First, new algorithms are developed that allow one to study a dynamical system that is described in a variety of ways such as by means of observational data or numerical simulations of differential equations.

Next, methods available to study non-autonomous systems, such as hyperbolic trajectories and Lagrangian coherent structures, are developed. These concepts are applied to examples of interests: Monterey Bay, the coast of Florida and the circulation in the North Atlantic. Combining accurate current measurements and recent developments in dynamical systems theory provides new and original answers to many problems, such as the minimization of the impact of released contaminants in a coastal area or the optimization of the coverage by a group of drifters.

The appendices give details about MANGEN, a software package developed to produce the numerical results of this thesis. Some projects that make use of its algorithms, such as the dissociation rate of a molecule and efficient space mission design, are also described.

Contents

Chapter 1

Introduction

Over the past 10 years there has been much work in applying the approach and methods of dynamical systems theory to the study of transport in fluids from the Lagrangian point of view. Suppose one is interested in the motion of a *passive tracer* (e.g., dye, temperature, or any material that has a negligible effect on the flow), then, *neglecting molecular diffusion*, the passive tracer follows fluid particle trajectories which are solutions of

$$\frac{d}{dt}\mathbf{x} = \mathbf{u}(\mathbf{x}, t), \tag{1.1}$$

where $\mathbf{u}(\mathbf{x}, t)$ is the velocity field of the fluid flow, $\mathbf{x} \in \mathbb{R}^n, n = 2$ or 3. When viewed from the point of view of dynamical systems theory, the phase space of Eq. 1.1 is actually the physical space in which the fluid flow takes place. Evidently, "structures" in the phase space of Eq. 1.1 should have some influence on the transport and mixing properties of the fluid. Aref and El Naschie [1994]; Babiano et al. [1994] provide recent reviews of this approach.

To make the connection with the large body of literature on dynamical systems theory more concrete let us consider a less general fluid mechanical setting. Suppose that the fluid is two-dimensional and incompressible. We know that the velocity field can be obtained from the derivatives of a scalar valued function, $\psi(x_1, x_2, t)$, known as the *streamfunction*, as follows

$$\begin{aligned} \frac{d}{dt}x_1 &= \frac{\partial \psi}{\partial x_2}(x_1, x_2, t), \\ \frac{d}{dt}x_2 &= -\frac{\partial \psi}{\partial x_1}(x_1, x_2, t), \qquad\qquad (x_1, x_2) \in \mathbb{R}^2. \end{aligned} \tag{1.2}$$

In the context of dynamical systems theory, Eq. 1.2 is a time-dependent Hamiltonian vector field where the streamfunction plays the role of the Hamiltonian function. If the flow is time-periodic, then the study of Eq. 1.2 is typically reduced to the study of a two-dimensional area preserving *Poincaré map*. Practically speaking, the reduction to a Poincaré map means that rather than viewing a particle trajectory as a curve in continuous time, one views the trajectory only at discrete intervals of time, where the interval of time is the period of the velocity field. The value of making this analogy with Hamiltonian dynamical systems lies in the fact that a variety of techniques in this area have immediate applications to, and implications for, transport and mixing processes in fluid mechanics. For example, the persistence of invariant curves in the Poincaré map (KAM curves) gives rise to barriers in the flow, chaos and Smale horseshoes provide mechanisms for the "randomization" of fluid particle trajectories. An analytical technique, Melnikov's method, allows one to estimate fluxes as well as describe the parameter regimes where chaotic fluid particle motions occur. A relatively new technique, lobe dynamics, enables one to efficiently compute transport between qualitatively different flow regimes.

Dynamical systems techniques were first applied to Lagrangian transport in the context of two-dimensional, *time-periodic* flows. In recent years these techniques have been extended to include flows having arbitrary time dependence [Wiggins, 1992; Malhotra and Wiggins, 1998; Haller and Poje, 1998]. Recently, the dynamical systems approach has been extended to a number of geophysical fluid dynamics settings [Pierrehumbert, 1991a,b; Samelson, 1992; Duan and Wiggins, 1996]. These early works mainly involved kinematically defined velocity fields. Some of the first attempts to treat dynamically evolving velocity fields were done by del Castillo-Negrete and Morrison [1993] and Ngan and Shepherd [1997]. They considered special kinematic cases that could be argued to be dynamically consistent, and hence complication provided by dynamical consistency was not present [Coulliette and Wiggins, 2000].

The treatment of general dynamically evolving velocity fields became possible with the development of computational techniques to treat velocity fields which only had a numerical representation, e.g., the numerical solution of a partial differential equation. Early work along these lines is Shariff et al. [1992], Duan and Wiggins [1997] and Miller et al. [1997]. Recent work of this type in a geophysical fluid dynamics setting is that of Rogerson et al. [1999], which is concerned with fluid exchange across a barotropic meandering jet.

In order to begin applying dynamical systems theory to the study of transport we need the right-hand side of Eq. 1.2, i.e., the velocity field. Until recently, applications have been limited to the cases where the velocity field is expressed as analytical function of space and time [Dijkstra and Katsman, 1997]. Then we can compute velocity explicitly once the position and time is given.

This may not be the case when the velocity field is obtained through the solution of some fluid dynamical nonlinear partial differential equations of motion (e.g., the Navier-Stokes equations). In general, such nonlinear partial differential equations cannot be solved analytically, i.e., the right-hand side of Eq. 1.2 cannot be represented in the form of some elementary or special functions. However, they can often be solved numerically and the velocity field may be given as output of the model simulation at a discrete time sequence which may also be spatially discrete. Another way in which the right-hand side of Eq. 1.2 can be obtained is through observational methods. Modern remote sensing techniques (such as high-frequency (HF) radar arrays) have now been developed to the point where we can measure current fields at a fairly high resolution in space and time.

Whether we obtain the velocity field through numerical simulation of a nonlinear partial differential equation or through observations, the resulting velocity field (i.e., dynamical system) is given as a data file, with gaps in space and time. Moreover, it will only be known for a finite amount of time, which may be at odds with many notions from dynamical systems theory, which is often concerned with the asymptotic in time behavior. Consequently, the fact that the velocity field may only be known for a finite-time causes major difficulties with the application of dynamical systems techniques. The source used to describe a dynamical systems (model, numerical simulation or experimentally observed velocities) defined by Eq. 1.1 does not affect the dynamical constraints experienced by the system. As an effort to emphasize this concept, software and algorithms were written independently of the data source (see Fig. 1.1). This thesis is organized in seven chapters and the first chapter details how the ocean can be described as a dynamical system, that is, a fairly smooth differential equation. Increasingly accurate remote sensing techniques are available today and it is both appealing and unavoidable to be able to use the measured velocities directly to describe a dynamical system. Modal analysis has been used in the past 30 years (see Eremeev et al. [1992a,b, 1995b,a]; Lipphardt et al. [2000] for example) to extrapolate and filter noisy and incomplete data sets. Up until recently, the modes used did not correctly allow flow across an open-boundary of the domain. Open-boundaries are an important concept when the domain is not completely closed by a shoreline, which

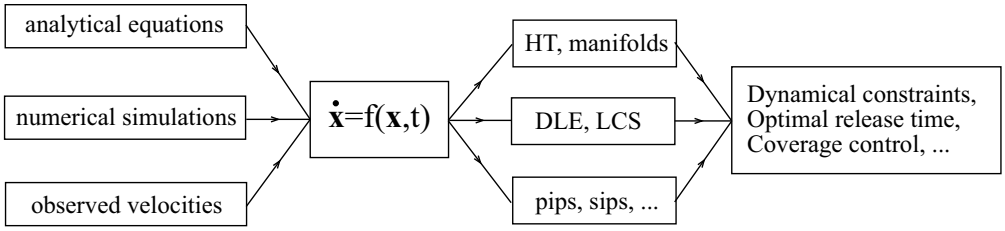

Figure 1.1: Different sources of input, such as analytical equations, numerical simulations and experimental data can be used to define a dynamical systems. Nevertheless, dynamical systems tools such as hyperbolic trajectories, direct Lyapunov exponents, invariant manifolds and lobes can be applied to any system.

is typical of coastal HF radar data[1]. Previous modal analysis methods project the data only onto closed-boundary modes, and then use an ad hoc procedure to add a zero-order mode to allow flow across the boundary. Open-boundary Modal Analysis (OMA) is the first method that allows flow across an open-boundary of the domain. We present the theory and a practical use of Open-boundary Modal Analysis (OMA), a complete set of eigenfunctions that can be used to interpolate, extrapolate and filter flows on an arbitrary domain with or without flow through a segment of the boundary.

Invariant manifolds, Lagrangian Coherent Structures (LCS) and Lobe dynamics provide a general theoretical framework for discussing, describing and quantifying organized structures in a fluid flow and determining their influence on transport. Earlier work [Malhotra and Wiggins, 1998] based on small perturbation of time-independent systems, developed the general mathematical framework of hyperbolic trajectories. To begin with, one can try to identify *Hyperbolic Trajectories* (HTs), i.e., "moving saddle points," whose stable and unstable manifolds divide the flow into different flow regimes. However, this approach does not extend to time-chaotic system with more than one dimension. A general method due to Haller [Haller and Poje, 1998; Haller, 2000; Haller and Yuan, 2000; Haller, 2001a] and based on an extension of Lyapunov exponents will be presented in Chapter 2. The theory of the Direct Lyapunov Exponents (DLE) method is discussed in Chapter 2. DLE and other criteria-based methods can be used to visualize quickly Lagrangian Coherent Structures (LCS) in the flow. We discuss the LCS for an analytical example and relate such structure to temperature fronts and other Lagrangian quantities of interest.

[1]When studying coastal system, the boundary of the domain of interest does not necessarily correspond to an actual shoreline. HF radar stations have a limited range and the boundary is made of a segment of shoreline and an open-boundary, an arbitrary curve in the ocean that closes the small region.

Combining accurate current measurement such as HF radar with recently developed methods in dynamical systems theory is a numerically challenging operation. However, the applications are multiple and the benefits are appealing. In coastal areas, pollutants released from the coast affect the marine environment in very subtle ways. By defining barriers and alleyways to transport, dynamical systems theory provides us with exceptional tools to determine interval of time where the impact of released contaminant is optimal. Numerical simulations in Monterey Bay are shown in Chapter 3 and we study the coast of Florida in Chapter 4. In both cases, a simple algorithm based on Lagrangian structure is able to reduce significantly the pollution and the peak of maximum concentration of pollutant generated by a numerically simulated source of pollution.

In Chapter 5, we detail and illustrate the *Manifold Deformation Framework* (MDF). The homoclinic or heteroclinic tangle framework is available for time-periodic systems and is typically used to describe chaotic stirring and lobe dynamics in those systems. MDF is a simpler version of the heteroclinic tangle, generalized to time-chaotic systems. The concept is applied to a simple model of the North Atlantic featuring characteristic similar to those of the Gulf Stream. In particular, the model is able to reproduce the apparition of rings along the jet. As observed experimentally, some of these rings contain fluid significantly warmer or colder than the surrounding water. It has been known for many years that these so called warm-core or cold-core rings are always anticyclonic[2]. Using the MDF and dynamical concepts, we show that cyclonic rings can indeed not have a cold or warm core.

The appendices mostly describe our software package MANGEN that was developed to produce the results shown in this thesis. Appendix B describes the latest developments of MANGEN and its graphical user interface. Major efforts were made to write a portable library of tools that can be used in many different problems. As an example, Appendix C uses MANGEN to study transport in a higher-dimensional systems. In this example inspired by Gillilan and Ezra [1991], an He atom interact with a Van der Waals potential created by a I_2 oscillator. Wang-Sang Koon, Shane Ross and Jerry Marsden recently pointed out that algorithms similar to those of MANGEN could be applied to the study of 3-body problems [Koon et al., 2000, 2001a,b]. Appendix D details a restricted 3-body problem and gives an overview of manifolds and transport in such systems.

[2]Cyclonic rings turn in the same direction as their containing gyre. Anticyclonic rings turn in the opposite direction.

Once you let the Ocean take you, there is no limit to the number of amazing and intriguing places where you may drift.

This has been my trajectory for the last 4 years.

Chapter 2

Open-Boundary Modal Analysis: A Complete Functional Basis to Interpolate, Extrapolate and Filter Experimental Eulerian Data

In collaboration with C. Coulliette, R. Bank and J. Marsden.

2.1 Introduction

In the last few years, technological advances in techniques such as high-frequency radar data provided us with fairly accurate measurements of the velocity vectors in certain coastal areas (see Coulliette and Wiggins [2001]; Lekien et al. [2003], for example). However, this data cannot be used directly as a dynamical systems, unlike a smooth ordinary differential equation (ODE) that can be integrated. For instance, there are missing gaps due to the radar range and wind conditions. As a result, oceanographers are in need of a way to extrapolate, interpolate and filter such data. The resulting smooth velocity field is often called *nowcast* in reference to the much more common forecasting operation. Nowcasting does not involve the extrapolation of the velocities in the future. It uses available data to determine the velocity everywhere in space at the same time data was collected.

Eigenvalue problems are known to provide a basis of functions over a domain. This fact has been used in fields such as electromagnetism and quantum mechanics. Recent works [Eremeev et al., 1992a,b, 1995a,b; Lipphardt et al., 2000] adapt these techniques to geophysical flows. In the case of a coastal oceanographic process, the boundary of the region studied may not correspond exactly to the shoreline (such as in Lipphardt et al. [2000]). In most cases, the region of interest is bounded by a segment of shoreline on one

side and an artificial boundary in the sea resulting from the limited range of action of the radar-antenna system. In such case, one needs to take into account the fact that portions of the boundary are permeable to water.

Many patches [Lipphardt et al., 2000] are available today to standard modal analysis techniques but they require an exact knowledge of the flux on the open-boundary. We present in this chapter a complete functional basis that can be used to project the flow on modes allowing flux through some segments of the boundary.

The method presented in this chapter is very similar to the method described in Lipphardt et al. [2000]. The fundamental difference is the addition of an infinite (possibly truncated) sequence of boundary modes allowing adequate degrees of freedom to project the experimental data on open-boundaries. Unlike other methods, the approach we highlight renders eigenfunctions on an unstructured (triangular) mesh with a general boundary.

2.2 Streamfunction and Relative Vorticity

2.2.1 State Equation Inside the Domain

The velocity nowcast will be referred to by a function $\bar{u}$ of a compact set $\Omega \in R^2$. If Ω is simply connected and its boundary is piecewise smooth, the Hodge decomposition [Eiseman and Stone, 1973] states that $\bar{u}$ can be written as the sum

$$\bar{u} = \bar{u}_\psi + \bar{u}_\phi, \tag{2.1}$$

where

$$\bar{u}_\psi = \nabla \times \psi \bar{k}, \tag{2.2}$$

is divergence-free and

$$\bar{u}_\phi = \nabla \phi, \tag{2.3}$$

is irrotational. Combining Eqs. 2.2 and 2.3 with Eq. 2.1 gives

$$\bar{u} = \nabla \times \left(\psi \bar{k}\right) + \nabla \phi, \tag{2.4}$$

where $\bar{k}$ is the unit vector orthogonal to the domain, pointing upwards. Applying $\nabla\cdot$ and $\nabla\times$ to Eq. 2.4 gives

$$\Delta \phi = \nabla \cdot \bar{u}, \tag{2.5}$$

and

$$\begin{aligned} \nabla \times \bar{u} &= \nabla \times \nabla \times \left(\psi \bar{k}\right) \\ &= -\Delta \left(\psi \bar{k}\right) \\ &= -\bar{k}\, \Delta \psi, \end{aligned}$$

so

$$\Delta \psi = -\bar{k} \cdot (\nabla \times \bar{u}) \,. \tag{2.6}$$

2.2.2 Boundary Conditions

We will let $\partial\Omega$ denote the boundary of the compact domain Ω. The unit vector normal to the boundary $\partial\Omega$ and pointing outward the domain is denoted by $\bar{n}$ (Fig. 2.1 (a)). The unit tangent vector is $\bar{t} = \bar{n} \times \bar{k}$. We assume that the boundary $\partial\Omega$ is continuous and piecewise C^1 (i.e., for all but a finite number of points x on the boundary, there exists an open set U_x containing x such that $U_x \cap \partial\Omega$ is then graph of a C^1 function). Multiplying Eq. 2.4 by either $\bar{t}$ or $\bar{n}$ gives

$$\begin{aligned} \bar{u} \cdot \bar{t} &= \bar{t} \cdot \left(\nabla \times \psi \bar{k}\right) + \bar{t}.\nabla\phi \\ &= \bar{t} \cdot \nabla\phi + \bar{n} \cdot \nabla\psi, \end{aligned} \tag{2.7}$$

and

$$\begin{aligned} \bar{u} \cdot \bar{n} &= \bar{n} \cdot \left(\nabla \times \psi \bar{k}\right) + \bar{n}.\nabla\phi \\ &= \bar{n} \cdot \nabla\phi - \bar{t} \cdot \nabla\psi, \end{aligned} \tag{2.8}$$

which can be used to establish the boundary conditions on the scalar functions ϕ and ψ. We will now assume that the boundary $\partial\Omega$ is made of the union of two different boundary pieces

$$\begin{cases} \partial\Omega = \partial\Omega_0 \cup \partial\Omega_1 \\ \partial\Omega_0 \cap \partial\Omega_1 = \emptyset, \end{cases} \tag{2.9}$$

where $\partial\Omega_0$ is a portion of real boundary (the shoreline) and $\partial\Omega_1$ is an artificial boundary through which a flux might exist (see Fig. 2.1 (b)). $\partial\Omega_1$ is included in our derivation because the basin of fluid may be too large to include the whole area in the computation. The user may have some measurements in a coastal area and cannot define Ω as the whole ocean. Instead he will join a segment of the real coastline $\partial\Omega_0$ to an artificial boundary $\partial\Omega_1$ (called the open-boundary) in the ocean to close the domain of interest. If $\partial\Omega = \partial\Omega_0$ our method reduces to the first two steps of the three-part algorithm described in [Lipphardt et al., 2000].

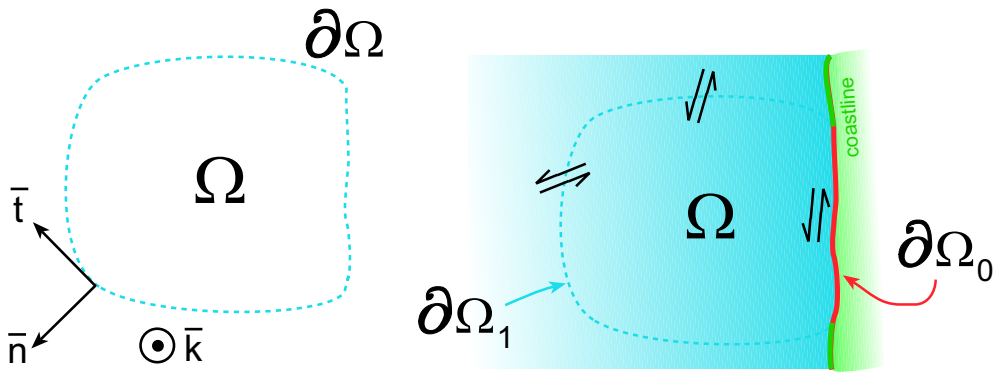

(a) (b)

Figure 2.1: (a) Domain Ω and its boundary $\partial\Omega$. $\bar{k}$ is the unit vector orthogonal to the ocean surface, pointing up to the reader. $\bar{n}$ is the unit vector normal to the boundary $\partial\Omega$ and the tangent vector is defined as $\bar{t} = \bar{n} \times \bar{k}$. (b) The portion $\partial\Omega_0$ is made of solid ground and has no normal flow while $\partial\Omega_1$ is subject to an arbitrary flux.

Eq. 2.7 and 2.8 are not sufficient to establish boundary conditions for the two functions ϕ and ψ. However, the decomposition given by Eq. 2.1 is not unique. Given a particular solution, one can add a divergence-free function to ϕ and a zero vorticity term to ψ. This degree of freedom allows us to make the following choice on ϕ and ψ: *The incompressible $\bar{u}_\psi$ part does not produce any flux through $\partial\Omega_0$, i.e., $\bar{n} \cdot \bar{u}_\psi = -\bar{t} \cdot \nabla\psi = 0$.*

This is a very natural choice in the sense that if the boundary was completely of type $\partial\Omega_0$, we would require $\psi = 0$ on the boundary (see [Eremeev et al., 1995b,a] or [Lipphardt et al., 2000] for example) which implies $\bar{t} \cdot \nabla\psi = 0$.
As a result, the second term of Eq. 2.8 disappears on $\partial\Omega_0$ and gives

$$\begin{aligned} \bar{n} \cdot \nabla\phi &= \bar{u} \cdot \bar{n} \text{ on } \partial\Omega_0 \\ &= 0 \text{ on } \partial\Omega_0, \end{aligned} \tag{2.10}$$

where we used the fact that $\bar{u}.\bar{n} = 0$ on $\partial\Omega_0$. We assume that $\partial\Omega_0$ is connected in $\partial\Omega$, i.e., there is only one connected segment of shoreline (see Fig. 2.1 (b), for example). Since we have $\bar{t}.\nabla\psi = 0$ on the segment $\partial\Omega_0$ and ψ can be added any arbitrary constant function (under the assumption that $\partial\Omega_0$ is a connected set) we can assume $\psi = 0$ on $\partial\Omega_0$. The

global equation for ϕ and ψ becomes

$$\begin{cases} \Delta\phi = \nabla.\bar{u} \text{ in } \Omega, \\ \bar{n}\cdot\nabla\phi = 0 \text{ on } \partial\Omega_0, \\ \bar{n}\cdot\nabla\phi = g_\phi(s) \text{ on } \partial\Omega_1, \end{cases} \tag{2.11}$$

and

$$\begin{cases} \Delta\psi = -\bar{k}.(\nabla\times\bar{u}) \text{ in } \Omega, \\ \psi = 0 \text{ on } \partial\Omega_0, \\ \psi = g_\psi(s) \text{ on } \partial\Omega_1, \end{cases} \tag{2.12}$$

where $g_\phi(s)$ and $g_\psi(s)$ are two unknown $L^2(\partial\Omega_1)$ functions of the arclength s on the boundary that depend on the velocity field $\bar{u}$. We can also write the problem in the following way

$$\begin{cases} \Delta\phi = \nabla\cdot\bar{u} \text{ in } \Omega, \\ \bar{n}\cdot\nabla\phi = g_\phi(s) \text{ on } \partial\Omega, \end{cases} \tag{2.13}$$

and

$$\begin{cases} \Delta\psi = -\bar{k}\,(\nabla\times\psi) \text{ in } \Omega, \\ \psi = g_\psi(s) \text{ on } \partial\Omega, \end{cases} \tag{2.14}$$

where we implicitly mean that g_ϕ and g_ψ vanish on $\partial\Omega_0$.

2.3 Homogeneous Solution with Inhomogeneous Boundary Conditions

The boundary modes or homogeneous modes are obtained by solving the homogeneous problem with inhomogeneous boundary conditions

$$\begin{cases} \Delta\phi^b - \int_{\partial\Omega} g_\phi(s)\,ds = 0 \text{ in } \Omega, \\ \bar{n}\cdot\nabla\phi^b = g_\phi(s) \text{ on } \partial\Omega, \end{cases} \tag{2.15}$$

2.3.1 Homogeneous Solution Basis

We assume that the set $\mathcal{F}$ of scalar functions on the boundary with support included in $\partial\Omega_1$ has a basis $\{g_i\}$. In this case any function g_ϕ can be written as a linear combination

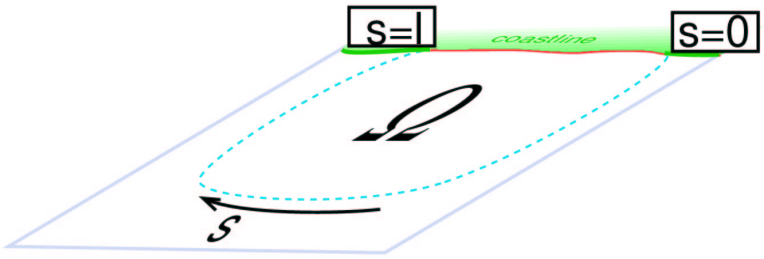

Figure 2.2: The parameter s is the arclength on the boundary and is defined in such a way that the open-boundary runs from $s = 0$ to $s = l$. The remainder of the boundary ($s > l$) is the shoreline.

of the basis function g_i. Let us define ϕ_i^b to be the solution of the problem

$$\begin{cases} \Delta\phi_i^b - \int_{\partial\Omega} g_i(s)\, ds = 0 \text{ in } \Omega, \\ \bar{n} \cdot \nabla\phi_i^b = g_i(s) \text{ on } \partial\Omega. \end{cases} \tag{2.16}$$

Providing that we add a condition such as

$$\int\int_{\Omega} \phi_i^b ds = 0, \tag{2.17}$$

the solutions ψ_i^b of Eq. 2.16 are unique. Indeed, this is seen by a standard argument. Suppose that ϕ_i^b and $\phi_i^{'b}$ are two solutions for Eq. 2.16 with the same g_i on the boundary. Their difference $\delta = \phi_i^{'b} - \phi_i^b$ must satisfy

$$\begin{cases} \Delta\delta = 0 \text{ in } \Omega, \\ \bar{n} \cdot \delta = 0 \text{ on } \partial\Omega. \end{cases} \tag{2.18}$$

Eq. 2.17 and 2.18 gives $\delta = 0$. Therefore, ϕ_i^b is unique.

2.3.2 Boundary Function Basis

A natural choice is a discrete Fourier basis of the function defined on the open-boundary

$$\{g_i(s)\} = \left\{1, \ldots, \sin\left(\frac{i\pi}{l}s\right), \cos\left(\frac{i\pi}{l}s\right), \ldots\right\}, \tag{2.19}$$

where l is the length of the open-boundary and we assumed that $s > l$ correspond to the shoreline (see Fig. 2.2). Note that this can be used to prove that there always exists a

discrete basis for the boundary functions and therefore that there always exists a way to obtain a discrete set of functions $\left\{\phi_i^b\right\}$, which we will assume in the following sections. This can be achieved by either using the Fourier basis or using numerical bump functions.

2.4 Inhomogeneous Modes with Homogeneous Boundary Conditions

The inhomogeneous modes satisfy the following equations

$$\begin{cases} \Delta\phi_i = \lambda_i^\phi \phi_i \text{ in } \Omega, \\ \bar{n}.\nabla\phi_i = 0 \text{ on } \partial\Omega, \end{cases} \tag{2.20}$$

$$\begin{cases} \Delta\psi_i = \lambda_i^\psi \psi_i \text{ in } \Omega, \\ \psi_i = 0 \text{ on } \partial\Omega, \end{cases} \tag{2.21}$$

where λ_i^ϕ and λ_i^ψ are the (unknown) eigenvalues and ϕ_i and ψ_i are the eigenfunctions.

According to Appendix A.4, the modes $\{\phi_i\}$ and $\{\psi_i\}$ form a basis of respectively $W^1(\Omega)$ and $W_0^1(\Omega)$. As a result, the associated vector fields $\{\nabla\phi_i\}$ and $\{\nabla\times\psi_i\}$ span, respectively, the spaces

$$\left\{\bar{u} \in L^2(\Omega)\times L^2(\Omega) \,\middle|\, \nabla\times\bar{u} = 0 \text{ and } \bar{n}.\nabla\bar{u}_{|\partial\Omega} = 0\right\}, \tag{2.22}$$

and

$$\left\{\bar{u} \in L^2(\Omega)\times L^2(\Omega) \,\middle|\, \nabla.\bar{u} = 0 \text{ and } \nabla\bar{n}.\bar{u}_{|\partial\Omega} = 0\right\} \tag{2.23}$$

Using the Hodge decomposition and the associated remark in Section 1, we conclude that the inhomogeneous modes $\{\nabla\phi_i, \nabla\times\psi_i\}$ span the set

$$\left\{\bar{u} \in L^2(\Omega)\times L^2(\Omega) \,\middle|\, \nabla\bar{n}.\bar{u}_{|\partial\Omega} = 0\right\} \tag{2.24}$$

2.5 Complete Basis

The nowcast velocity is given by a function $\bar{u}$ on Ω such that each component belongs to $L^2(\Omega)$ and $\bar{u}$ satisfies the boundary condition $\bar{n}.\bar{u} = 0$ on $\partial\Omega_0$. We assume that we have solved the boundary modes using Eq. 2.16 and we define the normal flow on the boundary

by

$$g(s) = \bar{n}.\bar{u}\,|_{\partial\Omega} \tag{2.25}$$

Since $\{g_i\}$ is a basis of $L^2[0,l]$, we can expand g as

$$g = \sum_{i=1}^{\infty} \alpha_i^b g_i, \tag{2.26}$$

We define

$$\bar{u}^H = \bar{u} - \sum_{i=1}^{\infty} \alpha_i^b \nabla \phi_i^b \tag{2.27}$$

and we remark that $\bar{u}^H \in L^2(\Omega)$ satisfies the homogeneous boundary conditions. In other words, $\bar{u}^H$ does not generate normal flow on either $\partial\Omega_0$ and $\partial\Omega_1$. As a result, $\bar{u}^H$ can be written as a unique sequence of homogeneous modes

$$\bar{u}^H = \sum_{i=1}^{\infty} \alpha_i^{\psi} \nabla \times \psi_i^H + \sum_{i=1}^{\infty} \alpha_i^{\phi} \nabla \phi_i^H, \tag{2.28}$$

and the velocity nowcast can be written as a unique sequence

$$\bar{u} = \sum_{i=1}^{\infty} \alpha_i^{\psi} \nabla \times \psi_i + \sum_{i=1}^{\infty} \alpha_i^{\phi} \nabla \phi_i + \sum_{i=1}^{\infty} \alpha_i^b \nabla \phi_i^b. \tag{2.29}$$

When the fluid is incompressible, the nowcast can be obtained with the incompressible modes only, because the projection on the other modes must vanish. Using only incompressible modes guarantees that the resulting nowcast will be incompressible. Assuming that the (finite or infinite) sequence of boundary modes that are incompressible is $\left\{\phi_{\beta_i}^b\right\}$, the nowcast $\bar{u}^H$ can be written using the truncated sequence

$$\bar{u} = \sum_{i=1}^{\infty} \alpha_i^{\psi} \nabla \times \psi_i + \sum_{i=1}^{\infty} \alpha_{\beta_i}^b \nabla \phi_{\beta_i}^b. \tag{2.30}$$

2.6 Example: The Unit Square

We consider the following domain

$$\Omega = [0,1] \times [0,1]\,, \tag{2.31}$$

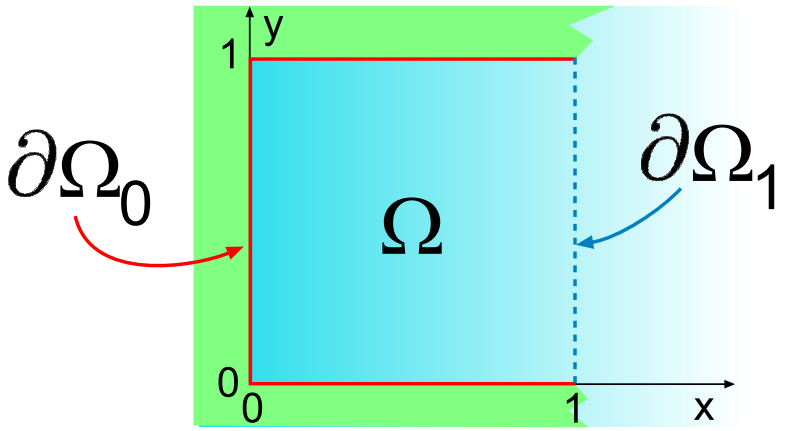

Figure 2.3: Sample domain for OMA modes. The domain Ω is the unit square. The eastern side of the square ($\partial\Omega_1$) is open and allow arbitrary flux. The three other sides of the square constitute the shoreline $\partial\Omega_0$.

where

$$\begin{aligned} \partial\Omega_0 &= \{0 \le x < 1 \text{ and } y = 0\} \\ &\cup \{0 \le x < 1 \text{ and } y = 1\} \\ &\cup \{0 \le y \le 1 \text{ and } x = 0\}, \end{aligned} \tag{2.32}$$

is the solid boundary (land) and

$$\partial\Omega_1 = \{0 \le y \le 1 \text{ and } x = 1\}, \tag{2.33}$$

allows transport (see Fig. 2.3). The inhomogeneous modes are given by:

$$\begin{aligned} \psi_{k,l} &= \sin(k\pi x)\sin(l\pi y), \\ \phi_{k,l} &= \cos(k\pi x)\cos(l\pi y), \end{aligned} \tag{2.34}$$

and the associated velocity modes

$$\begin{aligned} \bar{u}^{\psi}_{k,l} = \nabla \times \psi_{k,l} &= \pi \begin{pmatrix} k\sin(k\pi x)\cos(l\pi y) \\ -l\cos(k\pi x)\sin(l\pi y) \end{pmatrix}, \\ \bar{u}^{\phi}_{k,l} = \nabla \phi_{k,l} &= -\pi \begin{pmatrix} k\sin(k\pi x)\cos(l\pi y) \\ l\cos(k\pi x)\sin(l\pi y) \end{pmatrix}, \end{aligned} \tag{2.35}$$

Using a Fourier basis of the functions defined on $\partial\Omega_1$, we can compute the boundary

modes. We illustrate the process by computing the modes corresponding to

$$g_k(s) = g_k(y) = \pi cos(k\pi y). \tag{2.36}$$

One can verify that the unique solution to Eq. 2.16 is

$$\phi_k^b = \frac{e^{k\pi x} + e^{-k\pi x}}{e^{k\pi} - e^{-k\pi}} \cos(k\pi y), \tag{2.37}$$

and the corresponding boundary modes are

$$\bar{u}_k^b = \nabla.\phi_k^b = \frac{1}{k} \begin{pmatrix} \frac{e^{k\pi x} - e^{-k\pi x}}{e^{k\pi} - e^{-k\pi}} \cos(k\pi y) \\ -\frac{e^{k\pi x} + e^{-k\pi x}}{e^{k\pi} - e^{-k\pi}} \sin(k\pi y) \end{pmatrix} \tag{2.38}$$

For $g_k(y) = \sin(k\pi y)$, we could not find an analytical solution, only a numerical one. A detailed algorithm to compute OMA modes can be found in Section 2.8.

2.7 Extrapolation, Interpolation and Filtering

Eq. 2.29 can be rewritten in the more compact form

$$\bar{u} = \sum_{n=1}^{\infty} \alpha_n \bar{u}_n, \tag{2.39}$$

where $\bar{u}_n$ represent any of the linearly independent modes (boundary, incompressible or irrotational). Since the sequence of mode is infinite, the optimal coefficients α_n cannot be determined with a finite number of measurements. However, the smaller details are obtained for high eigenvalues and using a finite sequence as an approximation to Eq. 2.39 allows to filter the data and keep significant details only. The nowcast at a particular time is given by

$$\bar{u}' = \sum_{n=1}^{N} \alpha_n \bar{u}_n, \tag{2.40}$$

where N is the total number of modes (rotational, incompressible and boundary) and $\bar{u}_n$ represent any of these linearly independent modes.

If we were using all the modes, the error on $\bar{u}$ will be zero but that would require an infinite number of measurements to determine the infinite number of coefficients α_n. Instead, we will assume that only N modes are computed. This means that we arbitrarily set the other coefficients to zero, i.e., that we project the real velocity in a subspace of the

space of all available L^2 velocities that are tangent to $\partial\Omega_0$

In order to compute those coefficients we will need at least N measurements. The actual coefficients α_n will be the optimum of a minimization problem where the function to minimize $\zeta(\alpha_n)$ is a function of the error $\bar{u} - \bar{u}'$.

Assuming that we have k measurements $\bar{u}_q^{mes}$ at the position $\bar{x}_q = (x_q, y_q)$ the function to minimize is

$$\epsilon = \sqrt{\sum_q \left\| \sum_n \alpha_n \bar{u}_n(\bar{x}_q) - \bar{u}_q^{mes} \right\|^2}, \tag{2.41}$$

where the norm can be $\|\|_2$ to get a least square minimization problem or any other norm depending on the user choice. A set of projection coefficients α_i minimizes the error, when

$$\frac{\partial \epsilon}{\partial \alpha_j} = 0, \forall j, \tag{2.42}$$

or

$$\sum_n \alpha_n \sum_q (\bar{u}_n(\bar{x}_q).\bar{u}_j(\bar{x}_q)) = n \sum_q \bar{u}_q^{mes}.\bar{u}_j(\bar{x_q}) \ , \forall j. \tag{2.43}$$

Eq. 2.43 is a linear system of k equations with N unknowns α_i. Usually, the number of modes N is much smaller than the number of measurements k and the linear system is over-determined. Many algorithms are available to find the solution of over-determined systems that minimizes the residue and most of them have a portable implementation. Here, we use the *GNU Scientific Library*[1].

2.8 Application to Monterey Bay

This application uses high-frequency (HF) radar technology [Paduan and Rosenfeld, 1996; Paduan and Cook, 1997; Prandle and Ryder, 1985; Goldstein and Zebker, 1987; Georges et al., 1996], which is now able to resolve time-dependent Eulerian flow features in surface currents along coastlines. Such a HF radar installation has been operating in Monterey Bay since 1994 [Paduan and Rosenfeld, 1996; Paduan and Cook, 1997]. In our study, we use data from this installation, acquired by the three HF radar antennas near Monterey Bay, CA. The observational data was collected in August 2000, binned every hour on a 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. 2.4. Also shown on Fig. 2.4 are the level sets of the percentage of available data in the bay for the month of August 2000. There

[1] http://www.gnu.org/gsl

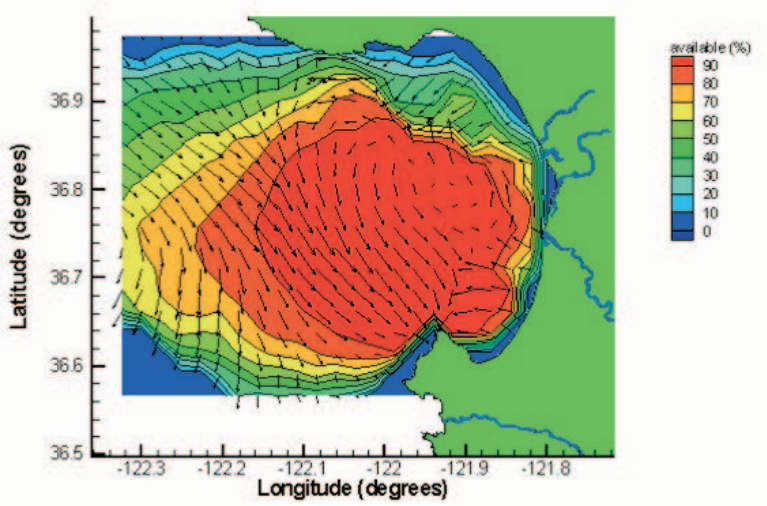

Figure 2.4: The black vectors show a footprint of the HF radar data in Monterey Bay on August 15 12:00 GMT. Also shown on this picture are the level sets of the percentage of available data during the month of August 2000. Regions in red have data almost all the time. Data is almost unavailable most of the time in the blue regions.

are many gaps in the data. Moreover, the missing data points are always located in the same region. The reason for this distribution is the range of the radar. The performance decreases with the distance. Also, some regions do not have ripples with enough amplitude to scatter the waves and the radar cannot determine the velocity. Our objective is to build an extrapolation method that filters and extrapolate such incomplete data sets.

A high-precision version of the shoreline for our domain was extracted on a topological map and is visible on Fig. 2.4. We used a numerical software package called $PLTMG$[2] to solve the mode equations (Eq. 2.20, 2.21 and 2.16). Fig. 2.5 shows a particular instance of the adaptive mesh used to compute one of the modes. The use of an unstructured mesh is necessary for applications such as the integration of particle or the computation of s Lagrangian structure near a complicated shoreline. Inadequate representation of the shoreline or a lack of precision in the velocity field near the coast often results in particles erroneously crossing the shoreline. Fig. 2.6 shows the unstructured mesh in a magnified region centered on Point Pinos, the southernmost part of the bay featured on Fig. 2.5, where separation of the flow between the bay and the ocean occurs. Fig. 2.7 shows the computed streamlines and velocities in the Point Pinos area. Such precise streamlines

[2]http://www.ucsd.edu/ rbank/pltmg

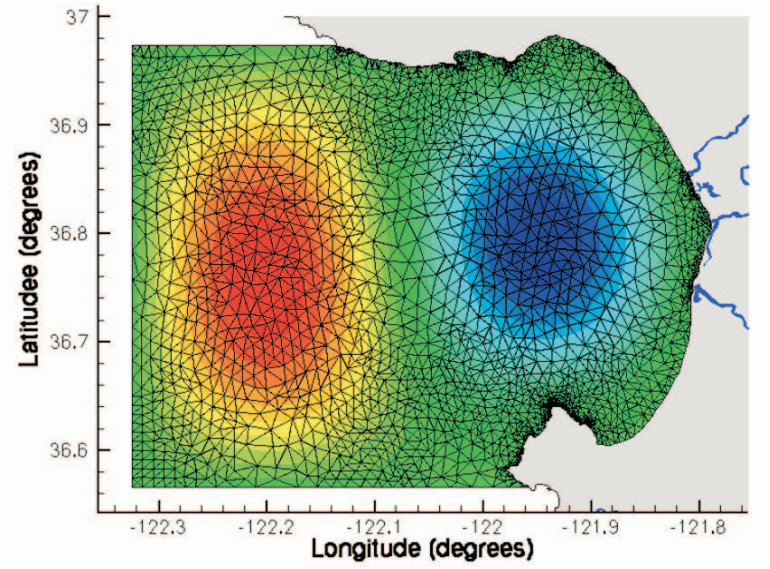

Figure 2.5: Adaptive mesh used for the computation of the 2nd incompressible mode ψ_2.

cannot be obtained with a structured mesh.

Fig. 2.8 shows two nowcasts realized with OMA. The upper nowcast uses only four incompressible modes and four irrotational modes. It cannot produce any flow normal to the boundary. It is similar to the nowcasts obtained after the first two steps of the algorithm described in [Lipphardt et al., 2000]. The lower nowcast uses eight boundary modes and is able to generate normal flow on the open-boundary (interaction with the Pacific ocean). As a result, the error between the HF radar data (red arrows on Fig. 2.9) and the nowcast, that does not use boundary modes (black arrows on Fig. 2.9), can be extremely large near the open-boundary. Fig. 2.9 also reveals that including 8 boundary modes in the nowcast decreases the error significantly near the open-boundary.

One might wonder if the effect of the boundary modes is only a minor correction of the nowcast near the open-boundary. In such cases, computations inside the bay (far away from the open-boundary) may not be affected by the boundary modes. Fig. 2.10 shows the nowcast on July 17, 1999 at 12:00 GMT. As previously, the upper nowcast does not use boundary modes. These are included in the lower nowcast of Fig. 2.10. One can notice that the two nowcast are dynamically different, even inside the bay, far away from the open-boundary. In this case, there is no separation point on the coastline. The flow spirals around the Bay. However, the open-boundary modes captured the separation point near Point Pinos.

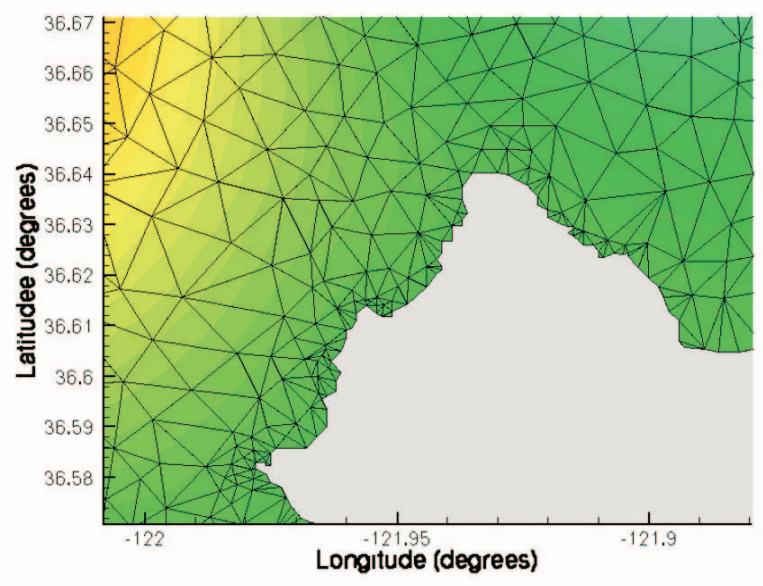

Figure 2.6: Unstructured mesh near Point Pinos, the southernmost point of the bay where the flow separates between the bay and the Pacific ocean.

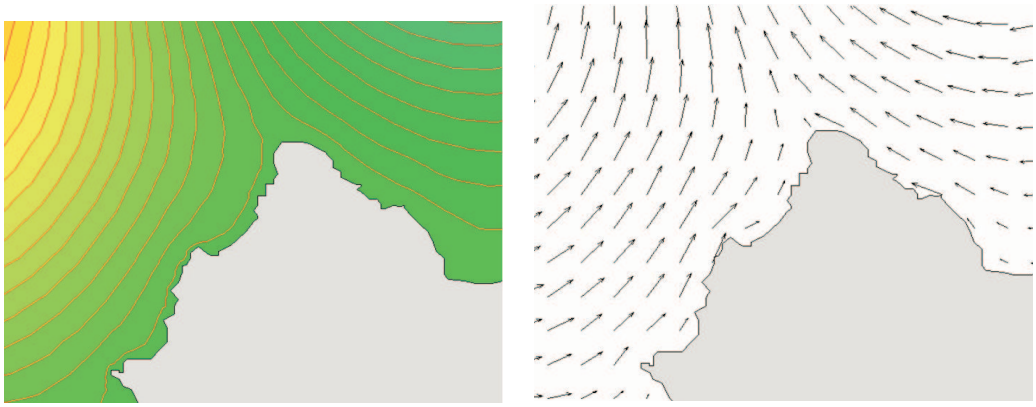

(a) (b)

Figure 2.7: (a) Accurate streamlines in the Point Pinos area. (b) Velocity in the Point Pinos area.

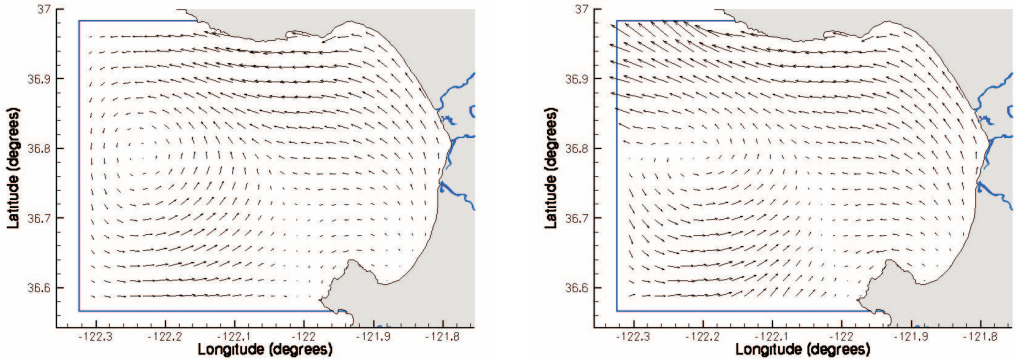

(a) (b)

Figure 2.8: Comparison between NMA and OMA nowcast on July 7, 2000 at 09:00 GMT. Both nowcasts use 4 incompressible modes and 4 irrotational modes. (a) does not use any boundary and cannot produce any flow through the open-boundary. (b) uses 8 boundary modes.

2.9 Conclusion

We presented a practical method to interpolate, extrapolate and filter experimental Eulerian data. This is the first modal analysis that includes a sequence of boundary modes. As a consequence, the modeler does not need to speculate on the open-boundary flux. Previous approaches [Lipphardt et al., 2000] require the use of a larger model or some assumptions to determine the flux through the open-boundary. In contrast, OMA adapts the flow near the boundary with the available data through Eq. 2.43. If the normal flow is known near the boundary, OMA uses this information and provides nowcasts similar to the adapted three-step algorithm in Lipphardt et al. [2000]. If, at some time, data is available only in the middle of the domain, far away from the boundary, OMA naturally projects the data on the boundary modes and finds the boundary flow that best fits the data.

There are many applications of such methods. Our group is primarily concerned with the interpolation, filtering and extrapolation of Eulerian velocities. OMA is also particularly well-suited for simple forecasting applications. Assuming that the data is collected

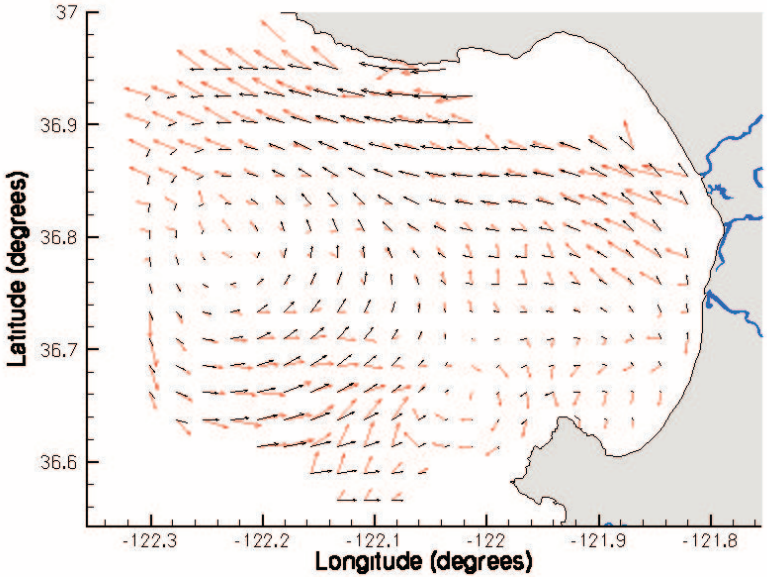

(a)

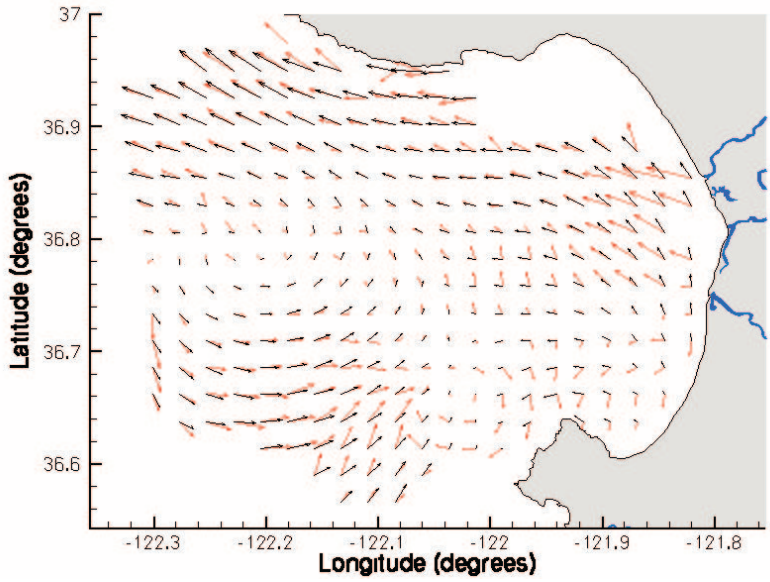

(b)

Figure 2.9: Comparison between NMA and OMA nowcast on July 7, 2000 at 09:00 GMT. On both images, the red vectors represent the available HF radar data for that time. The black arrows are the velocity at the same points evaluated by the OMA nowcast. (a) uses only 4 incompressible modes and 4 irrotational modes. Since the upper nowcast does not use any boundary mode, the error between the HF radar data and the nowcast near the open-boundary is large. (b) uses 8 additional boundary modes and is able to reduce the error between experimental data and extrapolated velocity.

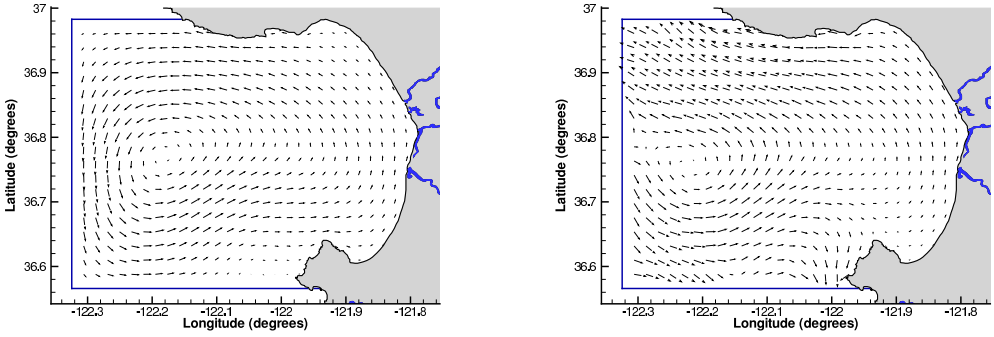

(a) (b)

Figure 2.10: Comparison between NMA and OMA nowcast on July 17, 1999 12:00 GMT. (a) uses 4 incompressible modes and 4 irrotational modes. (b) uses 4 additional boundary modes. Only the OMA nowcast (a) is able to reproduce the separation point on the shoreline (near Point Pinos) that was visible in the HF radar data. Certain important dynamical features can be wiped by removing the boundary modes.

at a constant rate (ΔT), OMA provides the nowcast as a sequence of coefficients

$$\begin{cases} \alpha_1(t_0)\,, \alpha_1(t_0+\Delta T)\,, \ldots\,, \alpha_1(t_0+i\Delta T)\,, \ldots \\ \alpha_2(t_0)\,, \alpha_2(t_0+\Delta T)\,, \ldots\,, \alpha_2(t_0+i\Delta T)\,, \ldots \\ \ldots \\ \alpha_k(t_0)\,, \alpha_k(t_0+\Delta T)\,, \ldots\,, \alpha_k(t_0+i\Delta T)\,, \ldots \\ \ldots \\ \alpha_N(t_0)\,, \alpha_N(t_0+\Delta T)\,, \ldots\,, \alpha_N(t_0+i\Delta T)\,, \ldots \end{cases} \tag{2.44}$$

Knowing $\alpha_k^p(t_0+i\Delta T)$ for $i=1\rightarrow N$, one can predict the evolution of the coefficient for $i>N$. The forecast velocity is given by Eq. 2.29.

Chapter 3

Lagrangian Coherent Structures

In collaboration with George Haller.

3.1 Introduction

A "moving" invariant manifold or material line can be defined as a time-dependent curve

$$f(x, y, t) = 0 \quad \forall t, \tag{3.1}$$

such that

$$\frac{\partial f}{\partial x}\dot{x} + \frac{\partial f}{\partial y}\dot{y} = -\frac{\partial f}{\partial t}. \tag{3.2}$$

It is not easy to define the equivalent of hyperbolic fixed points and stable or unstable invariant manifolds for time-dependent systems. The definition of "moving" invariant manifold can be used to define attracting or repelling lines in the flow, but most material lines are only hyperbolic for finite time. As a result the intersection, of those lines are hyperbolic trajectories only during short interval of times. More dramatically, some attracting lines become gradually less important and new attracting lines take over the general behavior of the flow. In some sense, it is hopeless to try to determine "the most influential" hyperbolic trajectories, because the influence of each trajectory changes over time.

To solve this problem, we turn to less elaborate but more robust methods. The definition of hyperbolic stagnation point or stable and unstable manifolds comes originally from the Lyapunov exponent theory. We give the definition of Lyapunov exponents for time-dependent systems in the next section and explain how they can be used to detect

Lagrangian structures in a flow. We then illustrate the method with a linear examples and discuss the sensitivity of Lyapunov exponents to anomalies in the flow.

3.2 Lyapunov Exponent for Time-Dependent Systems

A trajectory starting at time t_0 at the position $\mathbf{x}_0$ is located at the position $\mathbf{x}(t; t_0, \mathbf{x}_0)$ after a time $(t - t_0)$. The Lyapunov exponents of this trajectory are related to the norm of a perturbation of the trajectory. Let us assume that we perturb the initial condition to $\mathbf{x}_0 + \delta_x$. For infinitesimal perturbation $\delta_x(0)$, the position at time t of the perturbed trajectory will be given by

$$\mathbf{x}(t; t_0, \mathbf{x}_0) + \delta_x(t), \tag{3.3}$$

where

$$\delta_x(t) = \left(\frac{\partial \mathbf{x}(t; t_0, \mathbf{x}_0)}{\partial \mathbf{x}_0} \right) \delta_x(0), \tag{3.4}$$

and the norm of the perturbation becomes

$$\|\delta_x(t)\|^2 = \delta_x^T(0) \left(\frac{\partial \mathbf{x}(t; t_0, \mathbf{x}_0)}{\partial \mathbf{x}_0} \right)^T \left(\frac{\partial \mathbf{x}(t; t_0, \mathbf{x}_0)}{\partial \mathbf{x}_0} \right) \delta_x(0). \tag{3.5}$$

Let us define the Cauchy-Green strain tensor by

$$\mathbf{S}(t; t_0, \mathbf{x}_0) = \left(\frac{\partial \mathbf{x}(t; t_0, \mathbf{x}_0)}{\partial \mathbf{x}_0} \right)^T \left(\frac{\partial \mathbf{x}(t; t_0, \mathbf{x}_0)}{\partial \mathbf{x}_0} \right), \tag{3.6}$$

which allows us to rewrite Eq. 3.5 as

$$\|\delta_x(t)\|^2 = \delta_x^T(0)\, \mathbf{S}\, \delta_x(0). \tag{3.7}$$

The Lyapunov exponents are typically defined as the limit for $t \to +\infty$ of $(1/t - t_0)$ times the logarithm of the eigenvalues of

$$\left(\frac{\partial \|\delta_x(t)\|}{\partial \|\delta_x(0)\|} \right) = \frac{1}{2} \sqrt{\mathbf{S}}, \tag{3.8}$$

when $t \to \infty$ and where $\sqrt{\mathbf{S}}$ is well-defined since $\mathbf{S}$ is by construction positive definite. The factor $(1/t - t_0)$ forces the limit to converge for most initial conditions. Using the

definition of the logarithm of a positive definite operator, we can define

$$\mathbf{\Sigma}(t; t_0, \mathbf{x}_0) = \frac{1}{t - t_0} \ln \sqrt{\mathbf{S}(t; t_0, \mathbf{x}_0)}, \tag{3.9}$$

or equivalently

$$\mathbf{\Sigma}(t; t_0, \mathbf{x}_0) = \frac{1}{2} \frac{1}{t - t_0} \ln \mathbf{S}(t; t_0, \mathbf{x}_0), \tag{3.10}$$

and the Lyapunov exponents σ_1 and σ_2 are given the eigenvalues of

$$\lim_{t \to \infty} \mathbf{\Sigma}(t; t_0, \mathbf{x}_0). \tag{3.11}$$

3.3 Lagrangian Coherent Structures

To understand the evolution of fluid parcels, we use a geometric description of mixing from nonlinear dynamical systems theory. Two-dimensional time-periodic fluid flows have long been known to produce chaotic advection [Ottino, 1988], i.e., irregular stirring of fluid parcels. Instrumental in this stirring are stable and unstable manifolds of distinguished periodic fluid trajectories [Ottino, 1988]. Stable (respectively unstable) manifolds are material curves formed by fluid trajectories that converge to (respectively diverge from) an underlying periodic 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 blobs on their two sides to different spatial regions. For the same reason, unstable manifolds act as attracting material lines, targets along which fluid blobs spread out and form striations. We refer to attracting and repelling material lines jointly as hyperbolic material lines. Recent progress in nonlinear dynamical systems has extended the above geometric picture to velocity fields with general time dependence. As it turns out, families of hyperbolic material lines continue to organize finite-time mixing even in turbulent flows [Haller and Yuan, 2000]. Several numerical algorithms and theoretical criteria have been proposed to identify hyperbolic material lines in numerical and experimental data sets [Haller, 2000; Haller and Yuan, 2000; Haller, 2001a; Miller et al., 1997; Poje and Haller, 1999; Velasco Fuentes, 2001; Coulliette and Wiggins, 2001; Lapeyre et al., 2001; Joseph and Legras, 2001; Haller, 2001b]. Let us describe the Direct Lyapunov Exponent (DLE) algorithm [Haller, 2001a], based on the results of the previous section. It starts with the computation of the flow map, the map that takes an initial fluid particle position x_0 at time t_0 to their later position $\mathbf{x}(\mathbf{t}, \mathbf{x_0})$ at time t.

We take the largest singular value $\sigma_t(x_0, t_0)$ of the derivative of the flow map with respect to x_0. More specifically, we calculate the scalar field $\sigma_t(x_0, t_0)$ as the largest eigenvalue of the Cauchy-Green strain tensor. As argued in Haller [2001a], repelling material lines are local maximizing curves of $\sigma_t(x_0, t_0)$, which will allow us to capture these material lines at time t_0 as ridges of the scalar field $\sigma_t(x_0, t_0)$. The same procedure performed in backward time (i.e., for $t < t_0$) would render attracting material lines as ridges of $\sigma_t(x_0, t_0)$.

This algorithm takes into account Lagrangian hyperbolicity and is not based on instantaneous approximation of the Cauchy-Green tensor. The convergence time $(t - t_0)$ in Eq. 3.10 is an important parameter. It makes the Direct Lyapunov Exponent algorithm different from Lyapunov exponents (obtained for $t \to +\infty$). We know that a small convergence time will not produce satisfying results as it will not be Lagrangian enough (an Eulerian approximation can be computed for $t = t_0$. On the other hand, looking at infinite-time may not be the more efficient way to study non-autonomous systems. These can indeed switch between different dynamical modes or oscillate between different behaviors. An optimal convergence time $t - t_0$ needs to be determined for each application and the type of information seek (short or long term analysis and prediction). The next chapter includes a discussion about the minimum time $t - t_0$ that can be used for a specific application.

3.4 Order of Magnitude of the Exponents

We would like to show that factor $(t - t_0)^{-1}$ in Eq. 3.10 is important for numerical stability and also allow us to compute an approximation of σ. The analogy with Lyapunov exponents insures that the limit

$$\lim_{t \to +\infty} \sigma(t; t_0, \mathbf{x}_0), \tag{3.12}$$

converges and is equal to the Lyapunov exponent of the trajectory starting at $\mathbf{x}_0$ at time t_0. Omitting $(t - t_0)^{-1}$ creates an overflow problem, since

$$Eig\left\{\ln \sigma(t; t_0, \mathbf{x}_0)\right\} \sim (t - t_0), \tag{3.13}$$

when $t \to \infty$.

Consider the Jacobian of the velocity field

$$J = \begin{pmatrix} \frac{\partial u}{\partial x} & \frac{\partial u}{\partial y} \\ \frac{\partial v}{\partial x} & \frac{\partial v}{\partial y} \end{pmatrix}. \tag{3.14}$$

Under an infinitesimally small time δt, a trajectory starting at $\mathbf{x}_0$ at time t_0 moves to

$$\mathbf{x}(t) = \mathbf{u}\big|_{\mathbf{x}_0} \delta t. \tag{3.15}$$

A perturbed trajectory starting at t_0 in $\mathbf{x}_1 = \mathbf{x}_0 + \delta_x$ moves to

$$\mathbf{x}_0 + \delta_x + \mathbf{u}\big|_{\mathbf{x}_0} \delta t + \mathbf{J}\big|_{\mathbf{x}_0} \delta_x \delta t, \tag{3.16}$$

and the instantaneous perturbation is given by

$$\delta_x(\delta t) = \mathbf{J}\big|_{\mathbf{x}_0} \delta_x(0)\ \delta t. \tag{3.17}$$

Its growth is given by

$$\|\delta \mathbf{x}(\delta t)\| = \delta_x(0)\ \mathbf{J}\mathbf{J}^T\big|_{\mathbf{x}_0} \delta_x(0)\ \delta t, \tag{3.18}$$

and an approximation of the Lyapunov exponents are given by the eigenvalues of

$$\mathbf{S}(t; t, \mathbf{x}_0) = \mathbf{J}^T(\mathbf{x}_0, t)\mathbf{J}(\mathbf{x}_0, t). \tag{3.19}$$

3.5 A Simple Linear Example

We will apply the previous results to the "rotating saddle" defined by

$$\dot{\mathbf{x}} = \begin{pmatrix} \cos 2\omega t & \sin 2\omega t \\ \sin 2\omega t & -\cos 2\omega t \end{pmatrix} \mathbf{x}. \tag{3.20}$$

This system is linear and admits a stagnation point $\mathbf{x} = 0$. The eigenvalues are constant and equal to $+1$ and -1 and the system admits 2 eigenvectors associated with each eigenvalue

$$\mathbf{e}_{+1} = \begin{pmatrix} \cos \omega t \\ \sin \omega t \end{pmatrix}, \tag{3.21}$$

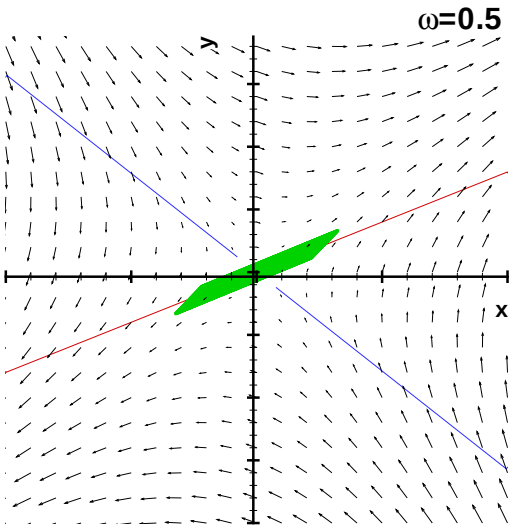

Figure 3.1: The rotating saddle for $\omega = \frac{1}{2}$. There exists a stable (blue) and an unstable (red) invariant manifold attached to the fixed point. A parcel (green) integrated will slowly spread to the unstable manifold.

and

$$\mathbf{e}_{-1} = \begin{pmatrix} -\sin\omega t \\ \cos\omega t \end{pmatrix}. \tag{3.22}$$

The eigenvectors are rotating with a constant angular speed ω. For $\omega = 0$ and ω small enough, we expect to find two invariant manifolds, one of each stability type. This situation is depicted in Fig. 3.1 for $\omega = 1/2$. We remark that the invariant manifolds are *not aligned with the instantaneous eigenvectors*. The unstable manifold is shifted by an angle ϕ behind the $\mathbf{e}_{+1}$ and the stable manifold is in advance with respect to $\mathbf{e}_{-1}$ with the same angle ϕ. The angle ϕ can be found as the solution of

$$\sin 2\phi = \omega. \tag{3.23}$$

Not surprisingly, Eq. 3.23 does not admit any solution when the rotation speed ω becomes too large. In that case, the rotating saddle behaves more like a center. Fig. 3.2 illustrates this phenomenon for $\omega = 3/2$.

Fig. 3.3 shows that when the rotation speed of the eigenvectors is not too large, the maximum ridge of the field σ captures the stable manifold of the system. Plotting $\sigma(-t; t_0, \mathbf{x}_0)$ would have outlined the unstable manifold.

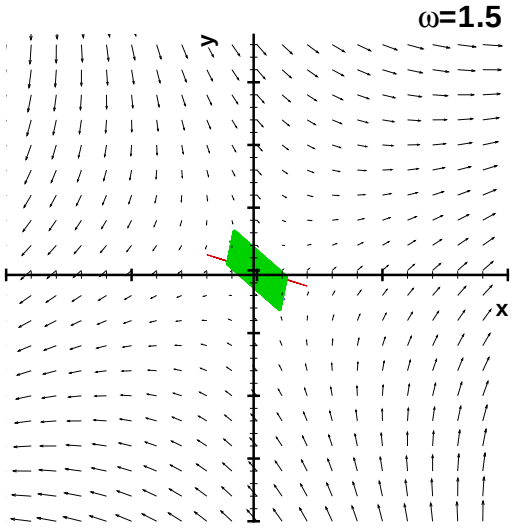

Figure 3.2: The rotating saddle for $\omega = \frac{3}{2}$. Even though the rotating saddle has a stable eigenvector with constant eigenvalue -1 and an unstable eigenvector with constant eigenvalue $+1$, the system does not have stable or unstable invariant manifold attached to the saddle because its eigenvectors are turning too fast. The trajectory at the origin is *not* hyperbolic and a parcel starting near the origin does not reveal any hyperbolic behavior.

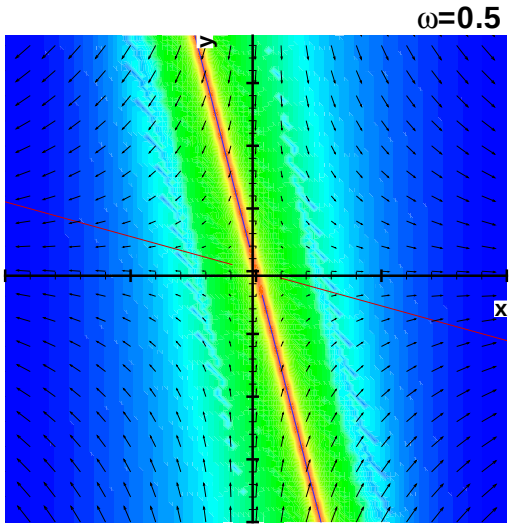

Figure 3.3: Direct Lyapunov Exponent Algorithm for the Rotating Saddle with $\omega = \frac{1}{2}$.

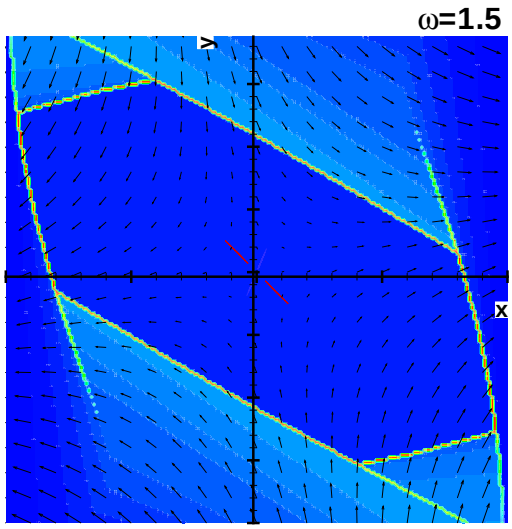

Figure 3.4: Direct Lyapunov Exponent Algorithm for the Rotating Saddle with $\omega = \frac{3}{2}$. Theoretically, the Lyapunov exponent is constant in space. The weak contrast and the angles are artifacts of neglecting trajectories outside the box.

Fig. 3.4 reports the same computation for $\omega = 3/2$. We can see that the DLE plot reveals no Lagrangian structure, as there are no stable or unstable manifolds.

It is important to remark that we would have expected a constant DLE plot for both values of ω. The system is linear and the Cauchy-Green tensor has constant eigenvalues in space and time. Instead we are able to pick up the unstable manifold for $\omega = 1/2$ and there are some shadows in the square for $\omega = 3/2$. This is a result of truncating the domain to a finite square. We decided to "ignore" trajectories that exit our domain because this is consistent with what we have to do for real data sets with open boundaries. The effect of this choice is to modify the velocity field outside of the domain (where the velocity becomes zero) and transform the open boundary to a solid boundary that acts as a source of stress. This seems to be an advantage in the region where there is a Lagrangian structure ($\omega = 1/2$), but creates parasite structures near the boundary even if there is not any structure ($\omega = 3/2$). Methods will be studied in the next chapter to avoid the parasite shadow effect of the boundary.

3.6 Relationship between LCS and Temperature Fronts

Thermal fronts in the atmosphere have been extensively studied [Lynch et al., 2001; Serreze et al., 2001]. In the ocean, many applications ranging from fishing to bioluminescence sampling would benefit from an efficient temperature front detection and prediction algorithm. Past successful efforts have been using the Thermal Front Parameter (TFP) defined as

$$\tau = -\nabla(|\nabla T|)\frac{\nabla T}{|\nabla T|}. \tag{3.24}$$

The TFP magnitude is large when there is a rapid change in the thermal gradient with a large component parallel to the direction of the (unitized) thermal gradient. If we restrict the system to a one-dimensional problem, a line parallel to the gradient of the temperature, this means that the second derivative of T is zero, i.e., the temperature "stops to increase dramatically." The minus sign is a convention and place the frontal boundary on the warm side of the concentrated level sets (ridge line in the field of TFP). Similarly, using $-\tau$ would detect cold fronts.

For a complete analysis, definition of TFP and temperature fronts as well as a review of all the possible cases, see Szoke et al. [1996]; Swenson et al. [1992]; Roden [1980]; Britz and Antonia [1987]. Notice that if a temperature front is directly created by advection of fluid coming from region with different initial temperature, there must also exists a LCS in the DLE field. To illustrate this phenomena, we used the ICON model of Igor Shulman (see Shulman et al. [2002] for a recent review). Fig. 3.5 shows Monterey Bay and the DLE field computed with the ICON model on July 1st 01:00 GMT. The black line superimposed on Fig. 3.5 indicates the Lagrangian structure in the bay. Fluid from both sides of this line is coming from different area and we expect that this line also corresponds to a temperature front. Recent experiments [Williams et al., 1997; Rothstein et al., 1999] show evidence of this for diffusive scalars. Fig. 3.6 shows the temperature field on the same day with the LCS superimposed. One can notice that the LCS corresponds almost exactly to the temperature front defined by τ.

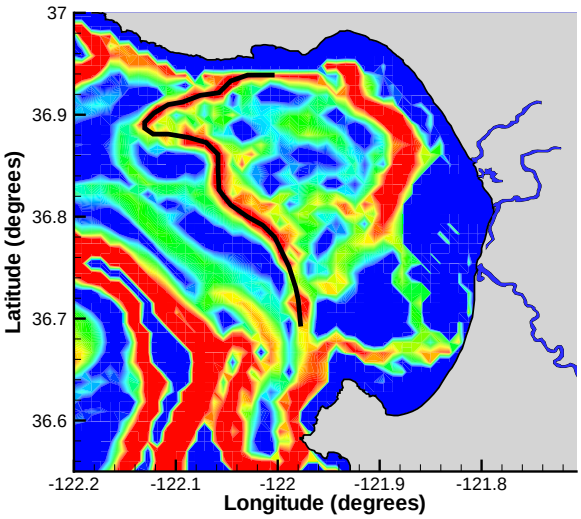

Figure 3.5: DLE field for the ICON model of Monterey Bay. A black line indicates the position of a DLE front.

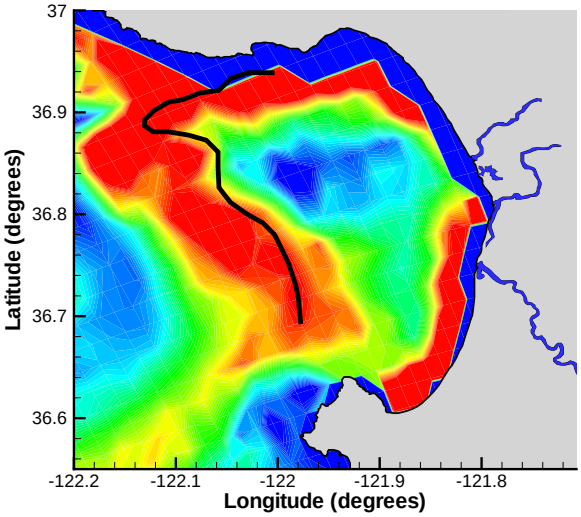

Figure 3.6: Temperature level sets in Monterey Bay on July 1st 01:00 GMT. The black line indicates the front in the DLE field (see Fig. 3.5) and also corresponds to a temperature front.

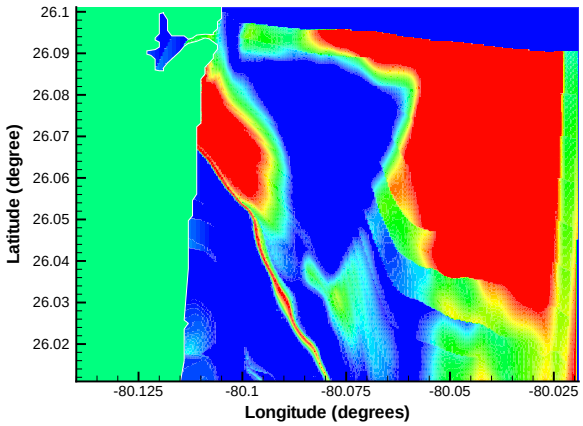

Figure 3.7: Average velocity in the x direction or x-patchiness plot for the RSMAS domain depicted in Chapter 5.

3.7 Other Methods

Extracting material lines in time-dependent flows is a complex problem and has been studied for many years. In this thesis we use mainly the method of the Direct Lyapunov Exponents. However, many other criteria have been used over the past. Early works [Malhotra et al., 1998; Poje et al., 1999] used the average x and y components of the velocity. We computed these fields (see Fig. 3.7 and 3.8), also called patchiness plots, for the small domain near the coast of Florida depicted in Chapter 5. The LCS extracted from the DLE plot (Fig. 3.9) corresponds exactly to the dividing line in the patchiness plots.

Other methods include relative and absolute dispersion [Weiss et al., 1998; Provenzale, 1999], finite strain [Haller, 2000; Haller and Yuan, 2000] and a new hyperbolicity criteria due to Iacono and Haller [2003].

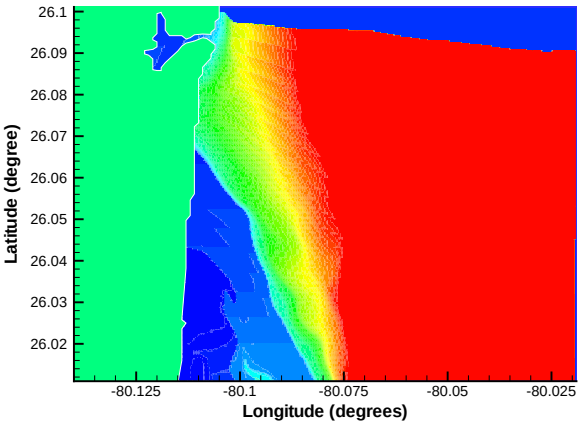

Figure 3.8: Average velocity in the y direction or y-patchiness plot for the RSMAS domain depicted in Chapter 5.

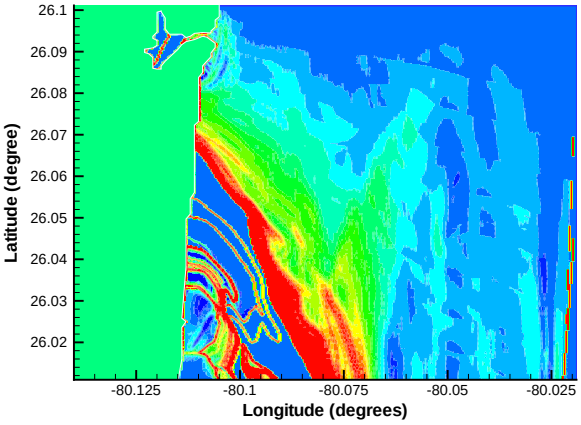

Figure 3.9: Direct Lyapunov Exponent field for the RSMAS domain depicted in Chapter 5.

Chapter 4

Optimal Pollution Release in Monterey Bay Based on Nonlinear Analysis of Coastal Radar Data

In collaboration with Chad Coulliette, George Haller, Jeffrey Paduan and Jerry Marsden.

4.1 Introduction

The release of pollution in coastal areas [Prahl et al., 1984; Rice et al., 1993; Verschueren, 1983] can lead to dramatic consequences for local ecosystems if the pollutants recirculate close to the coast rather than being transported out to the open ocean, where they are dispersed and then absorbed. This article shows how a combination of accurate current measurements and recent developments in nonlinear dynamical systems theory uncovers previously unknown flow structures that govern mesoscale ocean mixing. Knowledge of these Lagrangian (i.e., material) structures[1] can lead to predictions on a number of phenomena, ranging from the motion of plankton populations to the evolution of oil spills. The present article shows how Lagrangian flow structures can be exploited to reduce the damaging effects on coastal pollution. The focus of our study is the Elkhorn Slough, located near Moss Landing harbor of Monterey Bay. The Elkhorn Slough is a regular source of organic contaminants 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

[1] These Lagrangian structures are often called Lagrangian Coherent Structures or LCS.

hydrocarbons (PAHs) from the combustion of natural fossil fuels [Prahl et al., 1984; Rice et al., 1993; Verschueren, 1983].

Note that unlike other articles on pollution in this region, we are not implying that there is a pollution problem in Monterey Bay, but rather that predicting the optimal time window can significantly reduce the damage done in the coastal region by any amount of pollution, whether it is a small trickle from a stream or a huge oil spill from a tanker. For Monterey Bay, we will specifically show that surface current observations from coastal radar antennas located near Monterey Bay can be used to reduce the time which the aforementioned contaminants spend in the bay, and thus reduce the damage caused to the environment.

We examine high-frequency (HF) radar measurements [Paduan and Rosenfeld, 1996; Paduan and Cook, 1997; Prandle and Ryder, 1985; Goldstein and Zebker, 1987; Georges et al., 1996] of near-surface currents in Monterey Bay, and identify an LCS that governs the chaotic mixing of any Lagrangian particles, in this case we are specifically interested in Lagrangian contaminants over finite intervals of data. Specifically, we find a highly convoluted LCS— composed of a line of fluid particles— that repels nearby fluid parcels and hence acts as a barrier between two different types of motion: recirculation and escape from the bay. Release of pollution 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 highly desirable, while the former is to be avoided. 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 in any situation where sufficiently accurate wind or current velocity data is available, and the release of pollution can be contained until an appropriate release time. HF radar has also been demonstrated to work equally well in fresh water areas, but typically higher frequencies are necessary for the smaller regions, thus it is called Very High-Frequency (VHF) radar.

In contrast to earlier approaches to timed pollution release from holding tanks [Gould and Munro, 1981; Kay, 1990; Webb and Tomlinson, 1992; Smith, 1993; Bikangaga and Nussehi, 1995; Giles, 1995; Smith, 1998], we avoid the use of simplified models and target measured ocean data directly. This strategy accommodates constantly changing real-life 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 [Stirling, 2000; Haller, 2000; Haller and Yuan, 2000] for the analysis of HF radar data. The recent interest in the development and application of such methods stems from the realization that mixing in mesoscale geophysical flows is governed by coherent structures of finite lifespan [Haller, 2001a; Miller et al., 1997; Poje and Haller, 1999; Velasco Fuentes, 2001]. The presence of coherent features in measured geophysical flow data prevents the application of homogeneous and isotropic turbulence theory [Coulliette and Wiggins, 2000; Fischer et al., 1979] while the temporal irregularity and spatial complexity of such data renders the classic techniques of chaotic advection inapplicable [Watson et al., 1999; Zimmerman, 1986; Ridderinkhof et al., 1990; Beerens et al., 1998].

4.2 High-frequency Radar Measurements

Our analysis makes use of high-frequency (HF) radar technology [Paduan and Rosenfeld, 1996; Paduan and Cook, 1997; Prandle and Ryder, 1985; Goldstein and Zebker, 1987; Georges et al., 1996], which is now able to resolve time-dependent Eulerian flow features in surface currents along coastlines. Such an HF radar installation has been operating in Monterey Bay since 1994 [Paduan and Rosenfeld, 1996; Paduan and Cook, 1997]. In our study, we use data from this installation, acquired by the three HF radar antennas shown in Fig. 4.1. 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. 4.2.

We describe fluid particle motion in Monterey Bay as a dynamical system obeying

$$\dot{\mathbf{x}} = \mathbf{v}(\mathbf{x}, t). \tag{4.1}$$

To determine the velocity, the left-hand side of Eq. 4.1, we examine high-frequency (HF) radar measurements of near-surface currents in Monterey Bay. Ignoring measurement errors, the HF data is a footprint of the actual velocity in the bay as described by Eq. 4.1. 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. 4.3.

Using available HF velocity data, we advected the fluid particles using a fourth-order

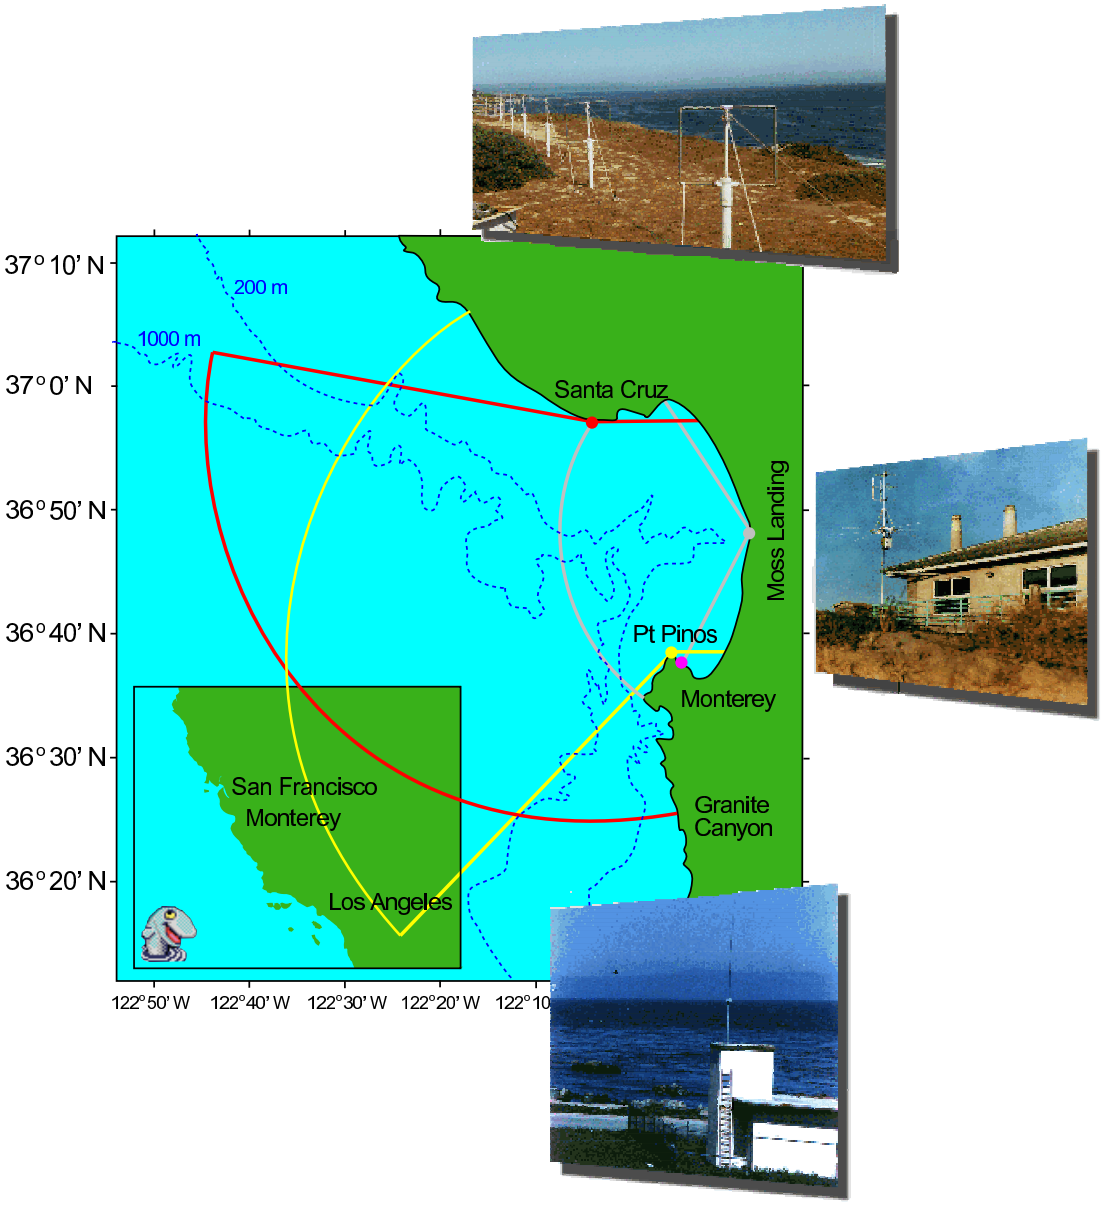

Figure 4.1: Monterey Bay HF radar network. Also shown are nominal coverage areas for the antennas at Santa Cruz, Point Pinos, and Moss Landing [Paduan and Rosenfeld, 1996; Paduan and Cook, 1997; Prandle and Ryder, 1985; Goldstein and Zebker, 1987; Georges et al., 1996].

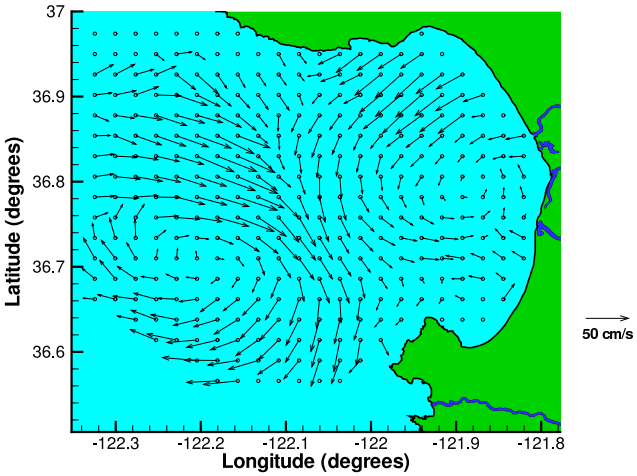

Figure 4.2: Instantaneous near-surface velocities (white arrows) at 08:00 GMT, August 8, 2000, obtained from the three HF antennas in Monterey Bay. Blue circles indicate all locations where continuous measurements were binned.

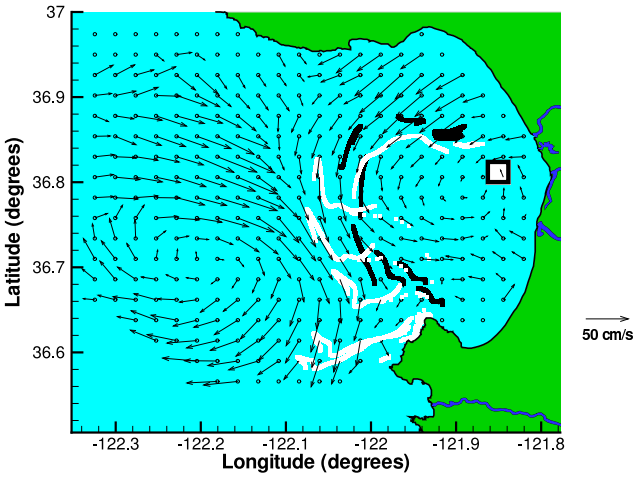

Figure 4.3: Evolution of two parcels of contaminants released from the same position near Moss Landing at 22:00 GMT, August 6, 2000 (black) and 09:00 GMT, August 7, 2000 (white). The motion of the two parcels is shown through daily snapshots over eight days. Note that the black parcel remains in the bay, while the white parcel departs from the bay.

Runge-Kutta algorithm combined with third-order polynomial interpolation in time and bi-cubic interpolation in space. We used these particle trajectories to approximate the flow map, which associates current positions to flow positions. We modelled the coastline as a free-slip boundary, and disregarded particles that crossed the linear fluid boundaries of the domain on the northern, southern and western edges. All these numerical algorithms have been compiled into a software package, MANGEN[2], described in Appendix B.

Note that the black contaminant parcel remains in the bay, whereas the white parcel exits the bay and moves towards 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 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. Such patterns are known to be delineated by attracting material lines, or finite-time unstable manifolds [Miller et al., 1997; Poje and Haller, 1999; Coulliette and Wiggins, 2001; Ridderinkhof and Zimmerman, 1992; Lapeyre et al., 2001; Haller, 2001b]. Here we shall use a recently developed nonlinear technique, Direct Lyapunov Exponent [Haller, 2001a] (DLE) analysis, which identifies repelling material in flow data as local maximizing curves of material stretching. We briefly recall this technique in the next section.

4.3 Lagrangian Coherent Structures

To understand the evolution of fluid parcels, we use a geometric description of mixing from nonlinear dynamical systems theory. Even time-periodic fluid flows have long been known to produce chaotic advection [Aref, 1984], i.e., irregular stirring of fluid parcels. Instrumental in this stirring are stable and unstable manifolds of distinguished periodic fluid trajectories [Ottino, 1988]. Stable (resp. unstable) manifolds are material curves formed by fluid trajectories that converge to (resp. diverge from) the distinguished 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 blobs on their two sides to different spatial regions. For the same reason, unstable manifolds act as attracting material lines, targets along which fluid blobs spread out and form striations. We refer to attracting and repelling material lines jointly as hyperbolic material lines.

[2]http://www.transport.caltech.edu

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. As it turns out, families of hyperbolic material lines continue to organize finite-time mixing in such flows, even when the flow becomes turbulent [Haller, 2000]. Several numerical algorithms and theoretical criteria have been proposed to identify hyperbolic material lines in general velocity data sets [Stirling, 2000; Haller, 2000; Haller and Yuan, 2000; Haller, 2001a; Miller et al., 1997; Poje and Haller, 1999; Velasco Fuentes, 2001; Lapeyre et al., 2001; Joseph and Legras, 2001; Haller, 2001b]. Here we use the Direct Lyapunov Exponent (DLE) algorithm [Haller, 2001b], which 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. We then compute the largest singular value $\sigma_t(\mathbf{x}_0, t_0)$ of the spatial gradient of the flow map. More specifically, we compute the largest eigenvalue of the Cauchy-Green strain tensor

$$\Sigma_t(\mathbf{x}_0, t_0) = \left[\frac{\partial \mathbf{x}(t, \mathbf{x}_0)}{\partial \mathbf{x}_0}\right]^T \left[\frac{\partial \mathbf{x}(t, \mathbf{x}_0)}{\partial \mathbf{x}_0}\right], \tag{4.2}$$

with the superscript T referring to the transpose of a matrix. Note that the scalar field $\sigma_t(\mathbf{x}_0, t_0)$ is related to the usual maximal finite-time Lyapunov exponent $\Lambda_t(\mathbf{x}_0, t_0)$ by the formula

$$\Lambda_t(\mathbf{x}_0, t_0) = \frac{1}{2t} \log \sigma_t(\mathbf{x}_0, t_0). \tag{4.3}$$

Repelling material lines are local maximizing curves or ridges of the scalar field $\sigma_t(\mathbf{x}_0, t_0)$ [Haller, 2001a, 2002]. The same procedure performed backward in time (i.e., for $t < t_0$) would render attracting material lines at t_0 as ridges of $\sigma_t(\mathbf{x}_0, t_0)$.

Composed of fluid particles, these curves are hidden to naked-eye observations of unsteady current plots, yet they fully govern global mixing patterns in the fluid. Such Lagrangian structures in measured ocean data have previously been inaccessible due to lack of an efficient extraction method.

4.4 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

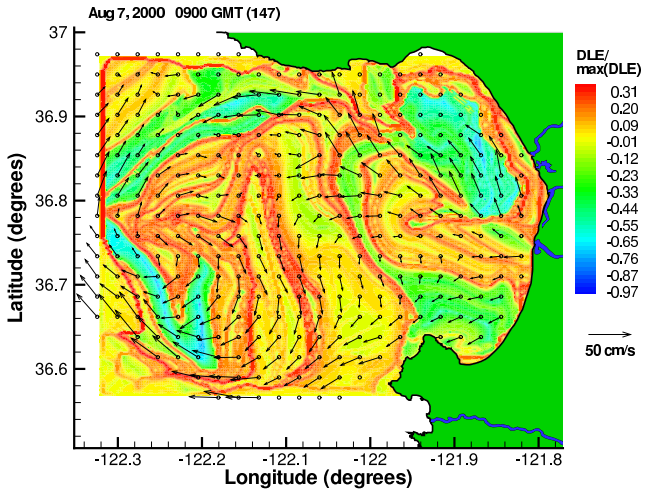

Figure 4.4: Lagrangian coherent structures in Monterey Bay at 06:00 GMT, August 8, 2000. Shown in the figure is the normalized distribution of $DLE_t(\mathbf{x}_0, t_0) = log[\sigma_t(\mathbf{x}_0, t_0)]$ for the initial time t_0 =06:00 GMT, August 8, 2000. The difference between the time t and the initial time t_0 is 200 hours. Local maximizing curves (ridges) of the scalar field indicate repelling material lines.

computation, we used a fourth-order Runge-Kutta algorithm combined with third-order polynomial interpolation in time and bi-cubic interpolation in space. We modelled the coastline as a free-slip boundary, and disregarded particles that crossed the open parts of the bay boundary. A sample result of such a computation is shown in Fig. 4.4, where the scalar distribution $DLE_t(\mathbf{x}_0, t0)$—the logarithm of $\sigma_t(\mathbf{x}_0, t0)$— is calculated over the initial grid of particles.

In agreement with the above general discussion, local maximizing curves, or ridges[3], 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. This structure can also be viewed as a stable manifold, a curve of fluid particles converging to an attachment point moving along the coast. The significance of this stable manifold is enormous: it divides the bay into two regions of different parcel behavior.

[3]By ridge, we mean a C^1 line in the flow similar to a water-dividing line in atmospheric science. In more technical terms, a ridge $R(\mathbf{x}, t)$ of a field $\sigma(\mathbf{x}, t)$ is a gradient curve of σ (i.e., $\forall t : \nabla_x L.\nabla_x \sigma = 0$) that has a maximum curvature in the orthogonal direction (i.e., $\left[\nabla\left(\nabla\sigma.\frac{\nabla L}{\|\nabla L\|}\right)\right].\frac{\nabla L}{\|\nabla L\|}$ is a local maximum in the $\frac{\nabla L}{\|\nabla L\|}$ direction).

Fluid parcels south of the stable manifold will recirculate in the bay after they pass by the coastal attachment point. At the same time, parcels north of the stable manifold clear to the open ocean after passing by the attachment point. This is in fact the reason underlying the different parcel behaviors in Fig. 4.3: the same release location fell on different sides of the stable manifold on August 6 and August 7. Fig. 4.5 illustrates this point by superimposing the instantaneous positions of the stable manifold on snapshots of parcel positions. Recall that the behavior of the white parcel is highly desirable for the evolution of pollutant blobs.

4.5 Prediction of Optimal Release Time Intervals

An important application 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 LCS can achieve a significant reduction of the impact of a contaminant in a coastal area, without reducing the total amount of contaminant released. This approach implicitly assumes that a sufficiently accurate prediction about the position of the LCS can be made, based on its previous position. Fortunately, prediction of Lagrangian quantities, such as the position of an LCS appears to be a much easier and reliable process than prediction of Eulerian data, such as the velocity.

Based on the analysis of the previous section, we now propose a pollution release scheme that minimizes the effect of contaminants on the coast of Monterey Bay. Assume that a pipeline carries contaminants from the Moss Landing area to an offshore release site shown in Fig. 4.6. (For consistency, this release site is the same as the location of release for the white and black parcels featured in Fig. 4.3.) Fig. 4.6 also shows the instantaneous intersection of the stable manifold—marked by a ridge of the DLE field—and the axis of the pipeline.

The motion of the intersection point along the axis of the pipeline is complicated, as evidenced by the time history of the intersection location in Fig. 4.7. Superimposed on this plot are the release times and release location of the white and black parcels of Fig. 4.3. Recall that the reason for their different future behaviors is the difference in their initial position relative to the curve of Fig. 4.7. In particular, the white parcel clears the bay quickly because it is released when the red curve is below the blue horizontal line, i.e., the point of release lies between the DLE peak and the coastline.

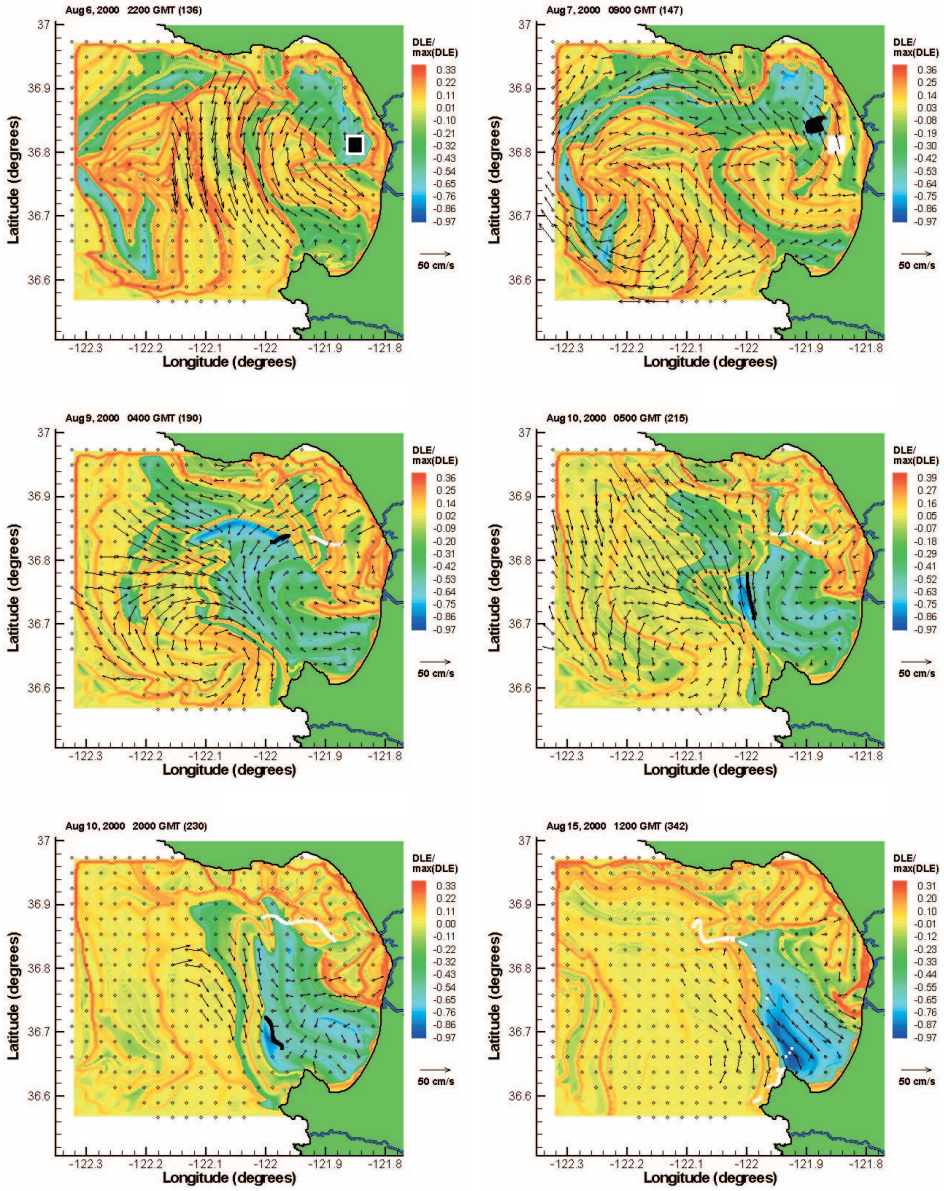

Figure 4.5: Two parcels of contaminants released from the same position near Moss Landing at 22:00 GMT, August 6, 2000 and at 09:00 GMT, August 7, 2000. The black arrows show instantaneous surface velocities captured by the HF radars. The ridges of the DLE field reveal the hidden Lagrangian structure of the bay at the same time instants.

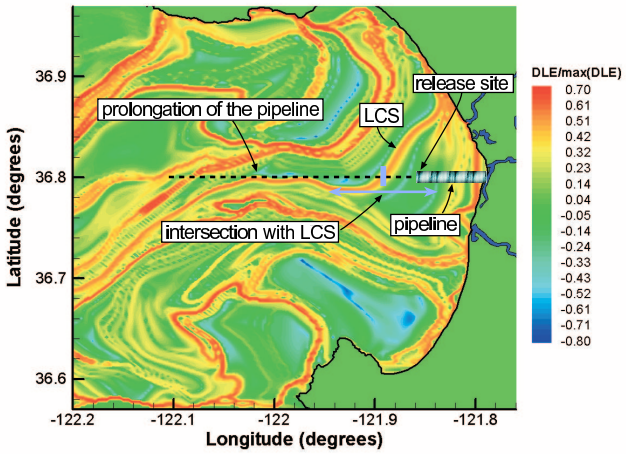

Figure 4.6: A pipeline carries contaminants to be released in the bay from the Moss Landing area. Also shown is the instantaneous intersection point between the Lagrangian coherent structure (LCS) and the axis of the pipeline.

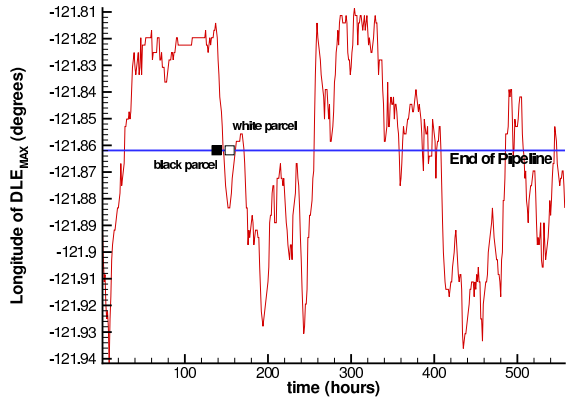

Figure 4.7: Oscillations of the DLE peak along the axis of the pipeline. The zero reference time corresponds to 07:00 GMT, August 1, 2000. The horizontal line marks the location of the outlet of the pipeline. The black and white squares represent the release time and release longitude of the parcels featured in Fig. 4.3.

It is tempting to think that the intersection curve in Fig. 4.7 predicts times of pollution release that will lead to quick clearance from the bay: Why not simply release pollution when the red curve is well below the blue line marking the outlet of the pipeline? As in the case of the white parcel, such a release would certainly guarantee that the contaminant blob is initially north of the stable manifold and hence leaves the bay quickly.

The above argument is flawed for practical applications, because any point of the red intersection curve in Fig. 4.7 is constructed from future velocity data over the next 200 hours. Such future data is clearly unavailable at any possible time of release. Trying to predict the velocity field in the bay-a necessity for advecting particles from the present into the future-is unrealistic because of the spatial and temporal complexity of the flow. Instead, we propose a focused Lagrangian prediction: we wish to predict the present and near-future location of the DLE peak, i.e., the maximum of $DLE_t(\mathbf{x}_0, t_0)$ field along the axis of the pipeline.

As a first step, we modify our calculation of the DLE field. We fix t = 22:00 GMT, Aug 6, 2000 as the present time when we would like to make our prediction. For any earlier time t0, we calculate the DLE peak from the field $DLE_t(\mathbf{x}_0, t_0)$; 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 modification results in a gradual–albeit surprisingly slow–growth of error between the actual DLE peak (computed with a constant 200 hour future window) and the real-time DLE peak (computed with a shrinking future window). The actual and the real-time DLE peak locations, as functions of time, are plotted in Fig. 4.8.

Remarkably, the real-time DLE peak curve approximates the actual curve within an error of 10% up until 15 hours before the present time. Close to the present time, however, the error of the approximation becomes substantial. We therefore have to stop our DLE calculation a few hours before the present time. To make our approach applicable to arbitrary data sets, we need a universal estimate for the time at which to stop.

To derive such a general estimate, we recall that the dimension of the Cauchy-Green strain tensor is $[velocity^2/length^2]$. From this we obtain that DLE_t is of the order of $\log((1/T)2)$, with T denoting a characteristic timescale over which the DLE field converges with sufficient accuracy. Denoting the average value of the DLE peak over the time interval $[t_0, t-T]$ by σ_{max}, we require $\sigma_{max} = \log((1/T)2)$, or, equivalently, $T = exp(-\sigma_{max}/2)$. Solving this general equation numerically for our particular choice of t_0 and t, we obtain T = 8 hours. Consequently, we need to stop our real-time DLE calculation about eight

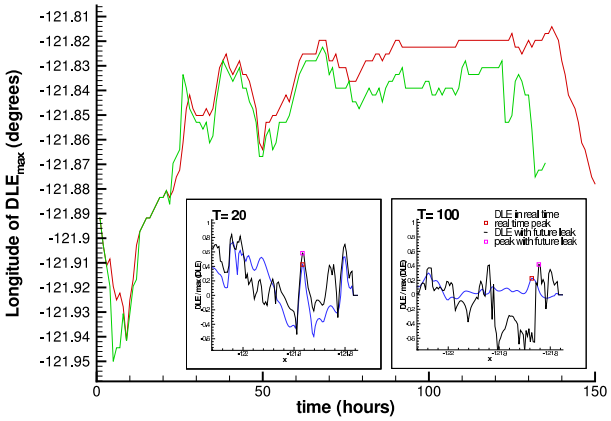

Figure 4.8: Oscillation of the DLE peak along the axis of the pipeline up to the present time $t =$ 22:00 GMT, Aug 6, 2000. The green curve is the real-time curve based on information up to the present time, with the DLE peak located from a numerical maximization along the pipe axis. The red curve is the actual DLE peek location computed with a constant 200-hour future time window. The inserts show a slice of the DLE contours along the axis of the pipeline. Time zero corresponds to 07:00 GMT, August 1, 2000.

hours before the present time to avoid substantial errors.

As a second step, we identify the main frequency components of the real-time DLE peak curve over the shortened time interval $[t_0, t - T]$. Shown in Fig. 4.9, the power spectrum density of the real-time DLE peak curve highlights seven dominant frequency components, with the importance of each frequency determined by the area under the corresponding spike in the spectrum. Surprisingly, the most influential component 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.

To complete our prediction procedure, we now use all the dominant Fourier modes of Fig. 4.9 to generate a prediction for the DLE peak location along the axis of the pipeline. 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 predicted and real-time DLE values. Fig. 4.10 shows the predicted DLE peak location together with the actual and the real-time locations. Note how faithfully the predicted curve reproduces the main features of the actual DLE peak oscillation.

In particular, Fig. 4.10 predicts that releasing contaminants from the pipeline between

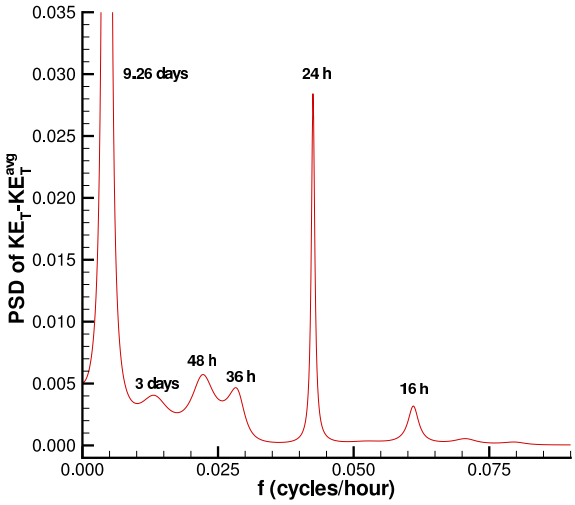

Figure 4.9: Power spectrum density of the real-time DLE peak oscillations shown in Fig. 4.7. The spikes at 48 hours and 4 days indicate harmonics associated with the 24-hour tidal oscillation. The importance of each frequency is proportional to the area below the corresponding spike.

5 hours and 110 hours from the present time, $t = 135$ hours, will cause most of the pollution to exit Monterey Bay without recirculation. Pollution released between 160 and 175 hours will not leave the bay immediately due to a short-lived excursion of the actual DLE peak curve into longitudes on the coastal side of the pipe outlet (see Fig. 4.10). However, assuming a constant rate of pollution release throughout the interval $[140h, 250h]$, one finds that recirculating contaminants constitute less than 15% of the total amount of contaminants released.

To illustrate the efficiency of the above pollution release scheme, we repeated the same prediction procedure for a different "present time," $t = 300$ hours. In this second case, the period of the dominant mode in the DLE peak oscillation was 9.25 days, which lead to the prediction shown in Fig. 4.11. As the prediction correctly reveals, the present time in this case in undesirable for pollution release. The predicted next "green" period is $[370h, 479h]$, while the actual green period turned out to be $[385h, 486h]$. This means that the error in predicting the earliest release time was approximately 15% of the actual green period, while the prediction error on the end time of the optimal pollution release was below 8%.

Beyond these predictions, the following general principles emerge from Fig. 4.10 and 4.11:

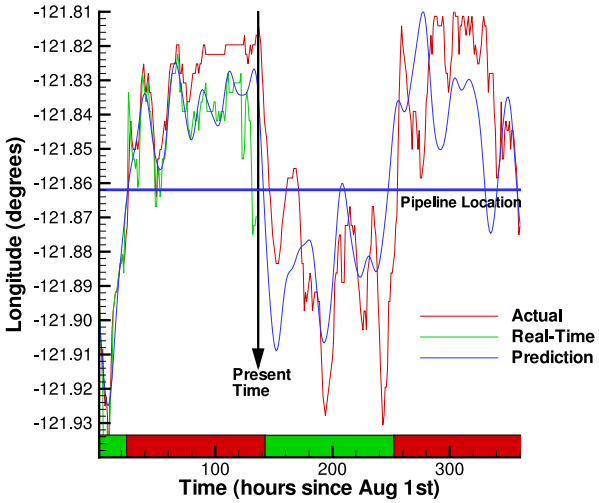

Figure 4.10: Actual, real-time, and predicted DLE peak location along the ax is of the pipeline. The horizontal line marks the location of the outlet of the pipe. The colorbar indicates the periods of desirable releases (green) and the periods to avoid (red).

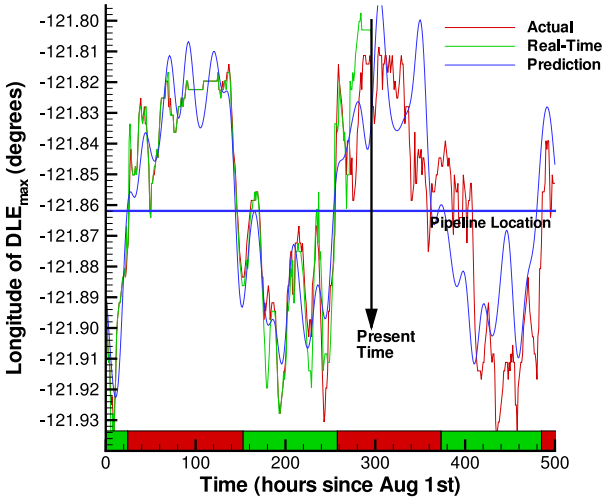

Figure 4.11: Actual, real-time, and predicted DLE peak location along the axis of the pipeline. The colorbar indicates the periods of desirable releases (green) and the periods to avoid (red).

- A longer pipeline is not necessarily better: selecting a shorter pipeline will raise the horizontal line (outlet position) in Fig. 4.10, which in turn leads to longer time windows for optimal release. At the same time, very short pipelines do not lead to quick clearance, because they will fall between secondary DLE peaks (visible in Fig. 4.6) and the coastline.

- For best performance, the holding tank must be able to hold contaminants produced over approximately 5 days. Using such a tank, one can wait, if needed, for the entire disadvantageous half-period of the main DLE peak mode to pass.

- Perhaps contrary to one's expectations, pollution release tied to high tide is typically not optimal: the dominant period in the DLE peak oscillation (approximately 9 days) lies far from the periods of the tidal oscillation and its harmonics.

4.6 Application to Pollution Control

Based on the above analysis, we propose a pollution release scheme that minimizes the effect of contaminants on the coast of Monterey Bay.

Although building a pipeline is not necessary for our method, 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 pipeline and release location are shown in Fig. 4.2.

For any given time, we consider a portion of the previously discussed LCS as it descends along the coastline of the bay from Santa Cruz, meandering southward past Moss Landing. The meandering of the LCS causes it to intersect the axis of our pipeline in several points. We refer to the first such intersection point (measured along the LCS starting from its coastal attachment point) as the barrier point.

According to the above discussion, if the barrier point is east of the pipeline outlet, between the outlet and the coast, then the pollution released from the outlet will re-circulate in the bay. However, if the barrier point falls to the west of the pipeline outlet, such that the pipeline outlet is between the aforementioned LCS and the coastline, the pollution released will exit the bay without re-circulating.

Clearly, the latter scenario is most desirable, because the contaminants can be more safely absorbed in this manner. The motion of the barrier point along the axis of the pipeline is complicated, as evidenced by Fig. 4.3.

Superimposed on this plot are the release time and release location of the white and black parcels of Fig. 4.1. The white parcel exits the bay quickly because it is released when the red curve is below the blue horizontal line, i.e., the pipeline outlet lies between the barrier point and the coastline. Using maximum entropy transforms, we have identified the main frequency components of the barrier point motion. Surprisingly, the most influential component of this motion is not the tidal oscillation (with a period of 24 hours), but rather a component with a period of 8.6 days.

To minimize the effect of coastal pollution, we propose using a holding tank that stores contaminants produced over a five day period. 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 is on west side of the pipeline outlet. Perhaps contrary to expectations, pollution release tied to high tide is typically not optimal: the dominant frequency in the oscillation of the barrier point lies far from frequencies associated with the tidal oscillation and its subharmonics.

Our pollution release algorithm relies on predicting the barrier point motion as follows: we chose August 6, 2000 22:00 GMT as the "present time" and computed the DLE plots (shown in Fig. 4.1) and the barrier point. No data after the "present time" is used in these computations. The result is a "real-time" barrier curve, as shown in Fig. 4.4.

We used all the significant frequencies of the spectrum of this curve to predict the barrier point location along the axis of the pipeline into the near future. Fig. 4.4 shows the predicted barrier point together with the actual and the real-time locations of the barrier point. Note how faithfully the predicted curve reproduces the main features of the actual DLE peak oscillations. From this prediction, we determined 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.

4.7 Conclusion

In this chapter, 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–or stable manifold–remains hidden both in instantaneous and averaged surface velocity plots. Yet the stable manifold 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. 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 our Lagrangian diagnostics, we have proposed a pollution release scheme that exploits the governing role of the stable manifold 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 manifold 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 predictions have proved very accurate: the error in predicting green 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. We need to stress, however, that our Lagrangian prediction is based on near-surface velocity data. As a result, the pollution release scheme we described here only applies to the part of the pollution that remains close to the ocean surface. A more detailed three-dimensional analysis could, in principle, be performed if velocity data at greater depths became available. Another assumption in this work is that the diffusive timescale for the contaminant is longer than the time of one recirculation in the bay, a property to be verified for actual pollutants before a real-life implementation of our scheme. Such an implementation would also require robustness with respect to measurement uncertainties and processing errors. Recent theoretical results do guarantee a high degree of robustness for Lagrangian coherent structures even under substantial errors, provided that the errors are deterministic and remain localized in time.We believe that laboratory experiments will be able to assess the effect of three-dimensionality, diffusion, and processing errors, which in turn could lead to applications of hidden flow structures in coastal pollution control, or in other problems involving mesoscale ocean mixing.

We have shown the existence of a convoluted set of attracting material lines in radar data obtained from Monterey Bay. We have also shown how these material lines can be used to minimize the effect of coastal pollution on the bay by predicting optimal release times. A general physical lesson from our analysis is that focused Lagrangian predictions for a geophysical flow, such as the prediction of material barrier locations, can be feasible even if global Eulerian (i.e., velocity based) predictions are unrealistic. We are not implying

Monterey Bay has a pollution problem that needs to be resolved, but simply highlighting an approach for reducing the impact of any pollution released in the Moss Landing area or any other coastal region. This approach can be used for making predictions about the trajectories of buoyant contaminants or the trajectories of nearly Lagrangian tracers. The data source can be HF 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. A real-time experimental realization of our pollution release will be most important, and is planned for future work.

Chapter 5

Lagrangian Structures in Very High-Frequency Radar Data along the Coast of Florida and Automated Optimal Pollution Timing

In collaboration with Chad Coulliette, George Haller, Arthur Mariano, Edward Ryan, Lynn Shay and Jerry Marsden.

5.1 Introduction

The release of pollution in coastal areas [Prahl et al., 1984; Rice et al., 1993; Verschueren, 1983] can lead to dramatic consequences for local ecosystems, especially if the pollution recirculates close to the coast rather than being transported out to the open ocean. Depending on their release position and release time, identical parcels of fluid can end up, or not, in environmentally sensitive areas with vastly different concentrations of contaminants. It is well known that Lagrangian motion and tracer dispersion in the ocean are sensitive to initial conditions at release time, thereby leading to vastly different patterns in the concentration of released contaminants in coastal regions. Using a combination of accurate surface current measurement [Shay et al., 2000] and recent developments in nonlinear dynamical systems theory [Stirling, 2000; Haller, 2001a], previously unknown flow structures[1] that govern the mesoscale transport of pollutants are documented. Knowledge

[1] By flow structure we mean Lagrangian structures, i.e., sets of distinguished fluid particles moving along with the flow.

of these Lagrangian structures should lead to reliable prediction for oceanic and coastal phenomena ranging from the motion of plankton populations to the evolution of oil spills. The focus of this chapter is on how Lagrangian structures can be exploited to reduce the concentration of coastal pollution.

Fluid particle motion is studied using particles that obey the dynamical system

$$\dot{\mathbf{x}} = \mathbf{v}(\mathbf{x}, t), \tag{5.1}$$

where $\mathbf{v}$ is the Eulerian velocity field of the flow. In contrast to earlier approaches to optimal pollution release in simple flow models [Webb and Tomlinson, 1992; Smith, 1993; Bikangaga and Nussehi, 1995; Giles, 1995; Smith, 1998], we rely on real-time velocity data obtained directly from coastal Doppler radar systems. The velocity on the left-hand side of Eq. 5.1 is given by Very High-Frequency (VHF) radar measurement of surface currents along the narrow shelf off southeast Florida [Shay et al., 2000]. Except for measurement errors, on the order of 5 cm/s, the VHF data shown in Fig. 5.1 is a discrete space-time representation of the actual velocity as described by Eq. 5.1.

5.2 VHF Radar Data along the Coast of Florida

Recent surface current observations from Ocean Surface Current Radar (OSCR) using the Very High-Frequency (VHF) mode reveal complex surface current patterns in this region [Shay et al., 2000; Peters et al., 2002]. Given the narrow shelf off Ft. Lauderdale and the strong Florida Current that intrudes onto the shelf on one-to-three day timescales [Peters et al., 2002], the site of what is called the "Four-Dimensional Current Experiment" 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 m s^{-1} just 8 km offshore. The average ambient relative vorticity is $4f$, where f is the local Coriolis parameter, and maximum relative vorticity exceeds $10f$. The dominant period in the velocity data increases from ten hours near-shore to 5 days offshore. There is also significant energy at 27 hours, the inertial period.

While the use of radio frequencies to measure ocean surface currents has received attention in recent coastal oceanographic experiments, using Very High-Frequency (VHF) radar [Prandle, 1987; Shay et al., 1995] provides velocity measurements at a horizontal resolution of 250 m and at a temporal resolution of 20 minutes, which is a very-high

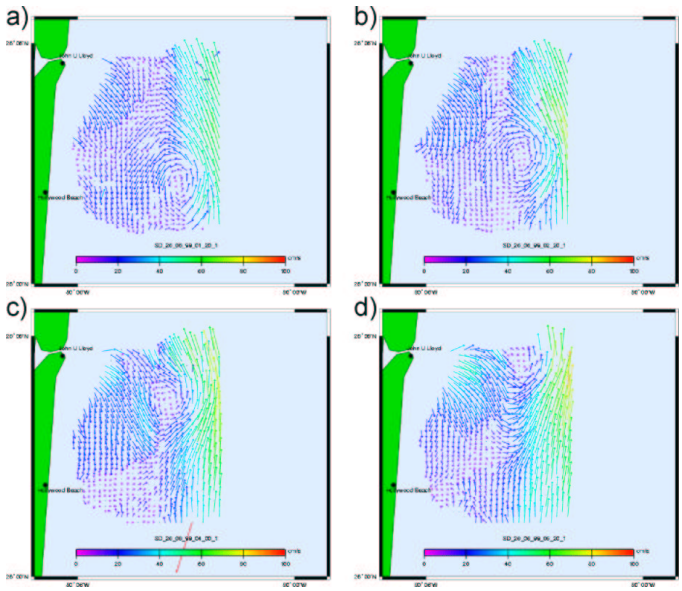

Figure 5.1: Surface velocity maps obtained by VHF radar along the southeast coast of Florida, near Fort Lauderdale, during the SFOMC 4-D Current Experiment on June 26, 1999: a) 01:20 GMT; b) 02:20 GMT; c) 04:00 GMT; d) 05:20 GMT. The sequence shows a northward propagating submesoscale vortex. The translation speed of the vortex is about 30 cm s^{-1} and its horizontal scale is 2-3 km.

resolution surface velocity data set. The OSCR VHF system was deployed for the Southern Florida Ocean Measurement Center (SFOMC) Four-Dimensional Current Experiment from June 25-August 25, 1999. Radio waves are backscattered from the moving ocean surface by surface waves of one-half of the incident radar wavelength. This Bragg scattering effect [Stewart and Joy, 1974] 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 [Shay et al., 2000].

The VHF system at SFOMC has two transmit/receive stations operating at 50 MHz that send electromagnetic signals which are then scattered from surface gravity waves with 3 m wavelengths. Coastal ocean currents were mapped over a 7 km $\times$ 8.5 km domain at 20 minute intervals with a horizontal resolution of 250 m at 700 grid points. The radars 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.

During this experiment, surface current observations (Fig. 5.2) 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 June 26, 1999 (Fig. 5.1) 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 hr period. See [Shay et al., 2000] for a more detailed analysis of this vortex.

5.3 Numerical Experiments

The temporal complexity of the currents becomes evident from tracking different realizations of a fluid parcel–a model for a blob of contaminant– released at the same time, but at a slightly different location. The results for two such numerical experiments are shown in Figs 5.3 and 5.4. 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 fourth-order Runge-Kutta-Fehlberg algorithm (RKF45) combined

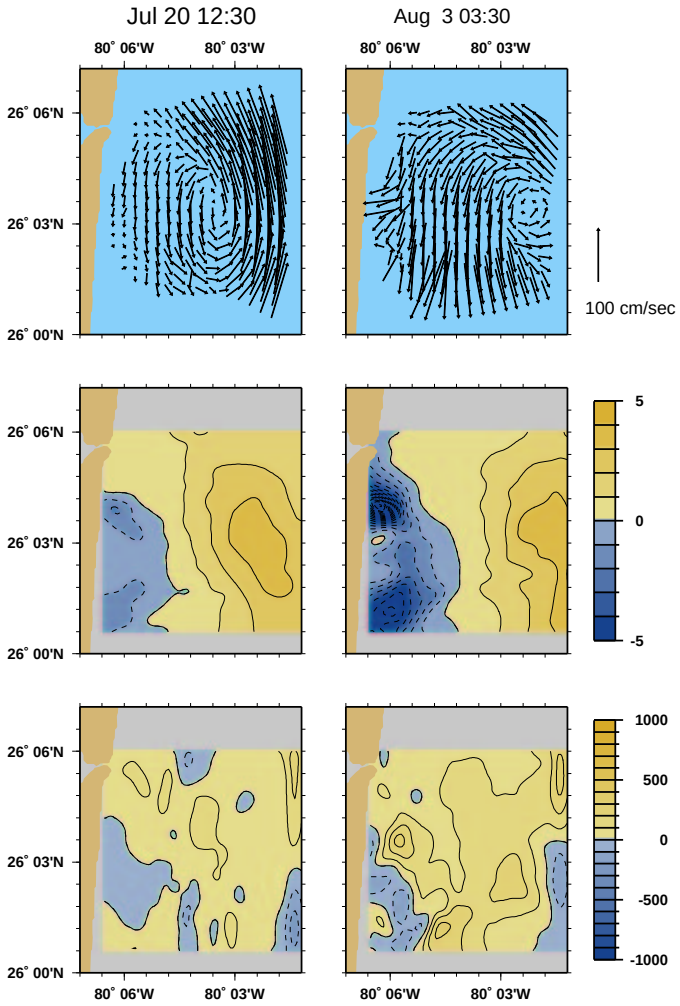

Figure 5.2: (Top panel) The velocity pattern obtained by HF radar along the coast of Florida, near Fort Lauderdale, for two different time periods, July 20 and August 3, 1999. A submesoscale vortex is evident in each velocity map. (Middle panel) The corresponding normalized, by the local Coriolis parameter, 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.

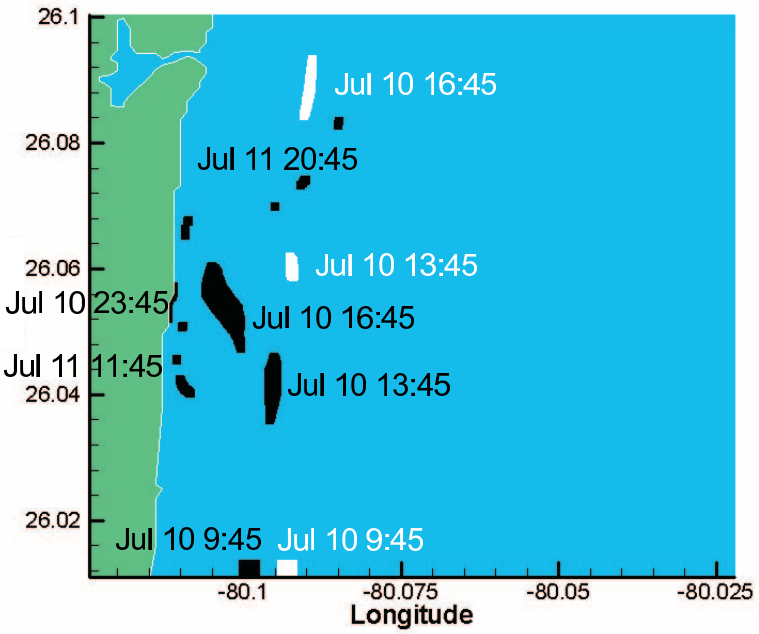

Figure 5.3: Two parcels of contaminant released at exactly the same time, but at slightly different initial locations on July 10, 1999, at 09:45 GMT. The white parcels leave the domain quickly as they are advected by the northward flowing Florida Current. The black parcels re-circulate near the coast for more than 36 hours.

with third-order tri-cubic interpolation in both space and time[2].

Note that large concentrations of the black parcels remain near the coast for relatively long times, 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. This observation inspires us to understand and predict different 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 [Ridderinkhof and Zimmerman, 1992; Miller et al., 1997; Malhotra and Wiggins, 1998; Poje and Haller, 1999; Mezic and Wiggins, 1999; Poje et al., 1999; Coulliette and Wiggins, 2001; Lapeyre et al., 2001; Haller, 2001b; Yuan et al., 2002]. A recently developed nonlinear technique called Direct Lyapunov Exponent (DLE) analysis [Haller,

[2]The resulting local interpolator provides a C^1 velocity field in extended phase space.

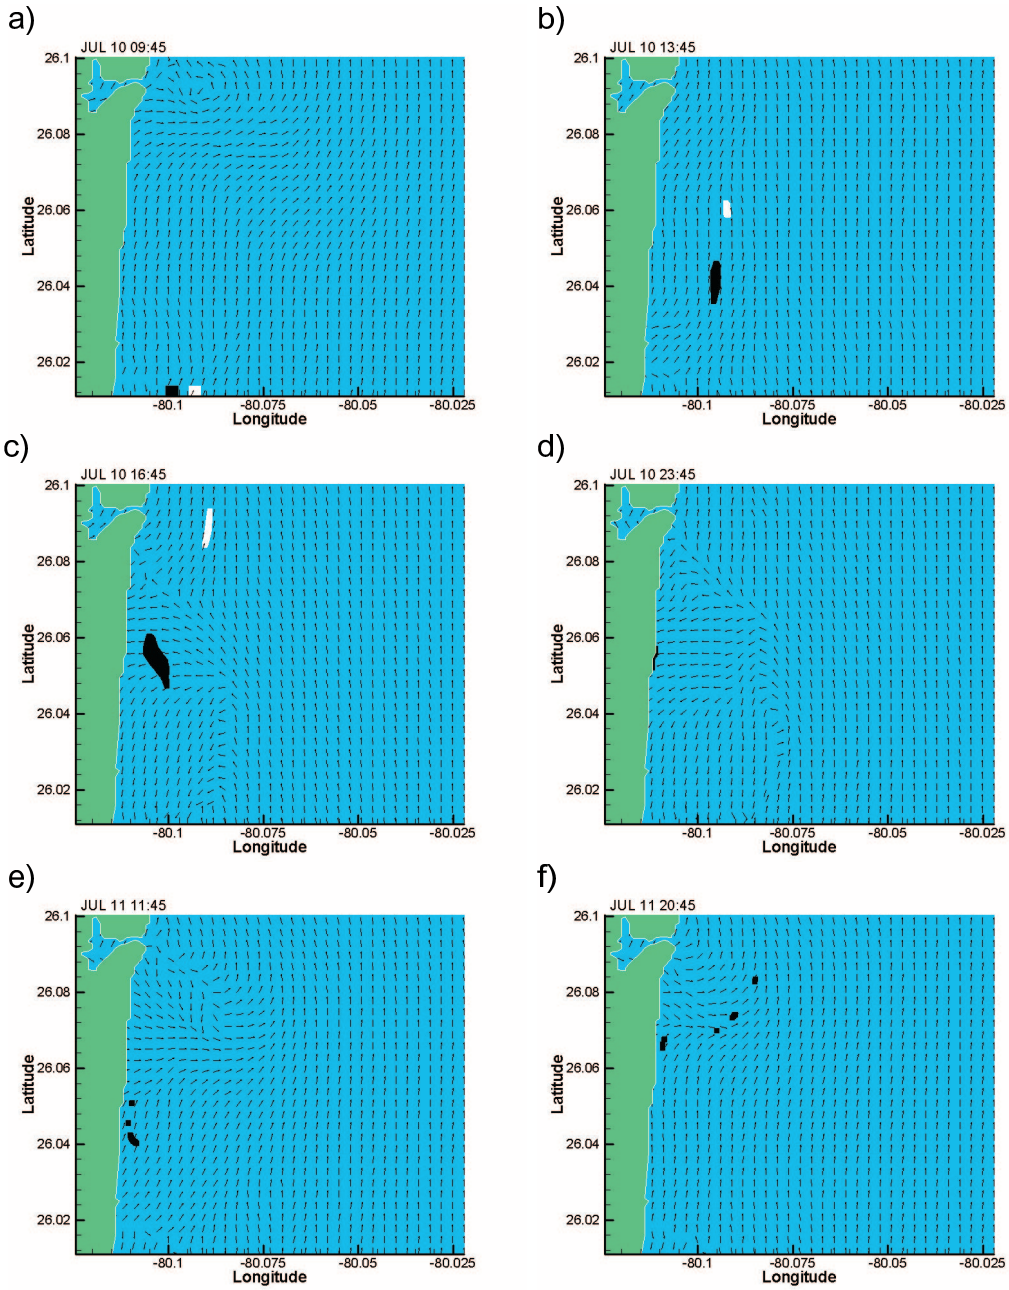

Figure 5.4: A time sequence of the motion of two parcels released almost at the same position on July 10, 1999, at 09:45 GMT. The interpolated velocity from the radar and the position of the parcels is shown for (a) July 10, 09:45 GMT, (b) July 10, 13:45 GMT, (c) July 10, 16:45 GMT, (d) July 10, 23:45 GMT, (e) July 11, 11:45 GMT and (d) July 11, 20:45.

2001a], identifies repelling or attracting material lines associated to Lagrangian Coherent Structures (LCS) in velocity data by means of maximizing curves of material stretching.

5.4 Lagrangian Coherent Structures

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 fourth-order Runge-Kutta-Fehlberg algorithm and a third-order interpolation for $t - t_0 = 25$ hours. These particle trajectories are used to approximate the flow map, which associates current positions to positions at time t_0. The coastline is modelled as a free-slip boundary. Particles that cross the open-boundaries of the domain on the northern, eastern and southern edges are disregarded. All of these numerical algorithms have been compiled into a software package, MANGEN, that is available from the authors upon request.

The largest singular value $\sigma_t(\mathbf{x}_0, t_0)$ of the spatial gradient of the flow map [Haller, 2001b] is computed. More specifically, the largest eigenvalue of the Cauchy-Green strain tensor

$$\Sigma_t(\mathbf{x}_0, t_0) = \left[\frac{\partial \mathbf{x}(t, \mathbf{x}_0)}{\partial \mathbf{x}_0}\right]^T \left[\frac{\partial \mathbf{x}(t, \mathbf{x}_0)}{\partial \mathbf{x}_0}\right], \tag{5.2}$$

with the superscript T referring to the transpose of a matrix, is calculated. Note that the scalar field $\sigma_t(\mathbf{x}_0, t_0)$ is related to the usual maximal finite-time Lyapunov exponent $\Lambda_t(\mathbf{x}_0, t_0)$ by the formula

$$\Lambda_t(\mathbf{x}_0, t_0) = \frac{1}{2t} \log \sigma_t(\mathbf{x}_0, t_0). \tag{5.3}$$

Repelling material lines are local maximizing curves or ridges of the scalar field $\sigma_t(\mathbf{x}_0, t_0)$[Haller, 2001a, 2002]. The same procedure performed backward in time (i.e., for $t < t_0$) would render attracting material lines at t_0 as ridges of $\sigma_t(\mathbf{x}_0, t_0)$.

These curves, composed of fluid particles, are not apparent to naked-eye observations of unsteady current plots, yet they fully govern global mixing patterns in the fluid. 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.

5.5 Analysis of the Data

Direct Lyapunov Exponents are used to analyze the Lagrangian trajectories in the SFOMC 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 maximal Lyapunov exponent are shown in Fig. 5.5. During the 2 months of 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 Lagrangian barrier between the coastal recirculating zone (southwest of the material line) and the Florida Current (northeast of the same material line). The material line is a barrier in the sense that particles cannot cross it. Superimposed on Fig. 5.5 are the two parcels used in Figs 5.3 and 5.4. A quick analysis reveals that any particle northeast of the barrier (white parcel) is flushed out of the domain in only a few hours. In contrast, parcels starting southwest of the barrier (black parcel) typically re-circulate several times near the Florida coast before they finally rejoin the current. It is important to realize that without the use of DLE or a similar method, the Lagrangian structure would still be there, but could not be seen or made use of in this way. We prefer to think of the currents as not influencing particle paths directly, but rather that the currents influence the Lagrangian structure, such as causing transport barriers and pathways, and the Lagrangian structure directly influences the particle paths.

5.6 Minimization of the Effect of Pollution

We remark that the location of the base of the structure (on the coast) can be used as a criteria to minimize the effect of coastal pollution. We will refer to the intersection of the coastline and the Lagrangian structure as the barrier point. For our analysis region and time period, factories and sewage discharge pipes along the coast should not release anything if the barrier point is located North of them. To illustrate how an efficient pollution release algorithm can be set up, we imagine a source of pollution with a fixed position along the coast. Using the DLE plots of Fig. 5.5, we identify zones of (green) favorable release[3] and (red) dangerous release[4].

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-

[3] when the manifold is below the position of the factory

[4] when the manifold is above the factory

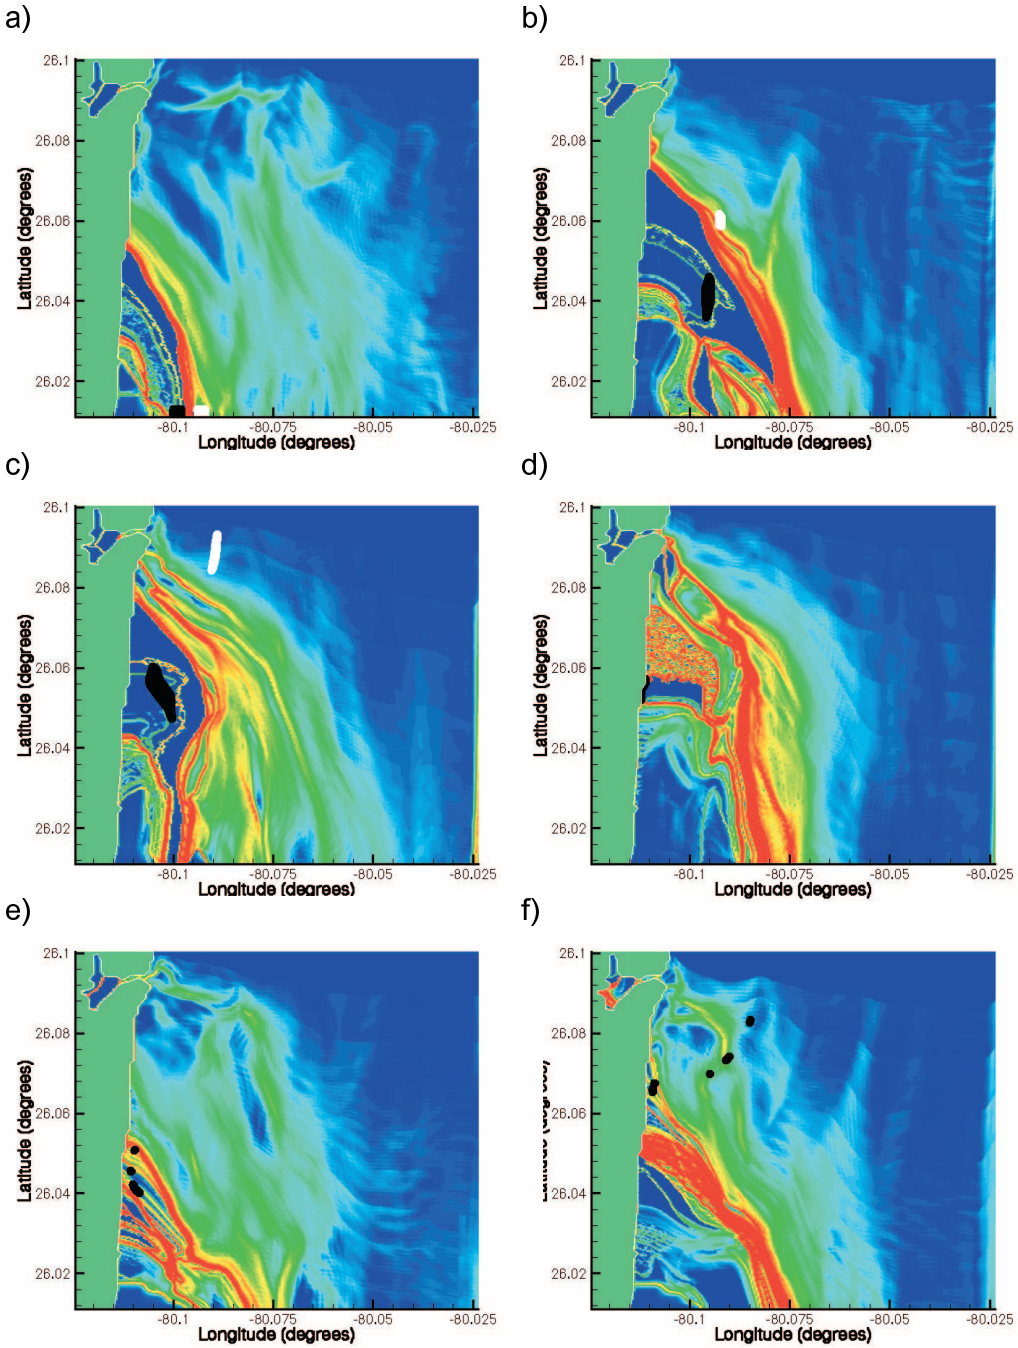

Figure 5.5: Level Sets of Direct Lyapunov Exponents σ along the coast of Florida on (a) July 10, 1999, 09:45 GMT, (b) July 10, 13:45 GMT, (c) July 10, 16:45 GMT, (d) July 10, 23:45 GMT, (e) July 11, 11:45 GMT and (d) July 11, 20:45 GMT. The simulation clearly shows repelling material lines (i.e., repelling material line) attached to the coast near Fort Lauderdale. Superimposed on each figure panel are the respective positions of the two parcels from Fig. 5.4 Every particle north of the manifold 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.

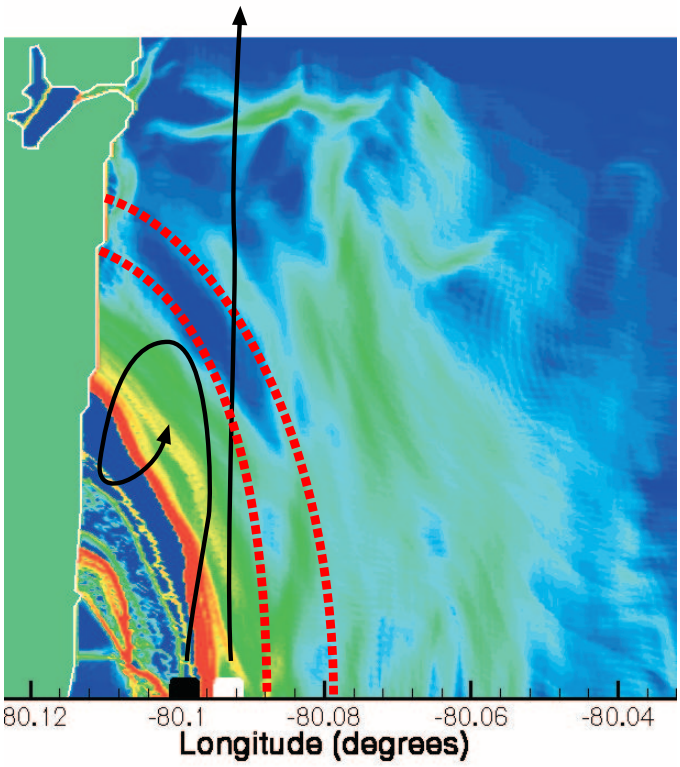

Figure 5.6: Level Sets of Direct Lyapunov Exponents along the coast of Florida on July 15, 1999, at 9:45 GMT. The dashed line represents future positions of the bright red line, the stable invariant manifold. The manifold is a line of particles and cannot be crossed. Every particle north of the manifold flows through the northern open-boundary (white parcel). It is dangerous to release contaminants below the branch of the manifold because it will persist between the coast and the manifold for a long time (black parcel).

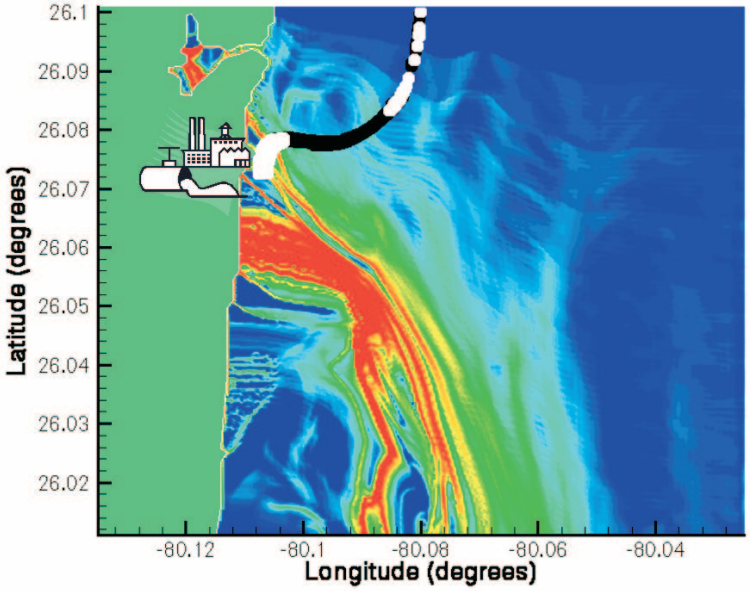

Figure 5.7: An imaginary factory rejects pollution along the Florida coast. The black spots are the resulting contaminants from a factory releasing at a constant rate. Superimposed on this figure are the white spots of a factory releasing only during environmental friendly time windows.

period of the barrier point oscillation, during which contaminants should not be released. The contents of the tank are released once the barrier point passes south of the source of pollution. Fig. 5.7 shows such a setting. The black spots 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 during environmental friendly time windows determined using our DLE algorithm and shown on Fig. 5.8. Also shown on Fig. 5.8 is the total mass of contaminant in the coastal area. Both pollution sources release the same mass of pollutant over the whole interval of time. However, Fig. 5.8 shows that, by getting information from the DLE plots, the white factory (dashed curve) is able to reduce, by a factor of three, the effect of the pollutant in the shallow coastal area.

However, in many cases the damage to the environment is a function of the maximum concentration of contaminant. From this point of view our algorithm (releasing nothing during "red" zones and as much as possible during "green" zone) does not seem to be efficient. The peak of maximum concentration for the white factory has only decreased

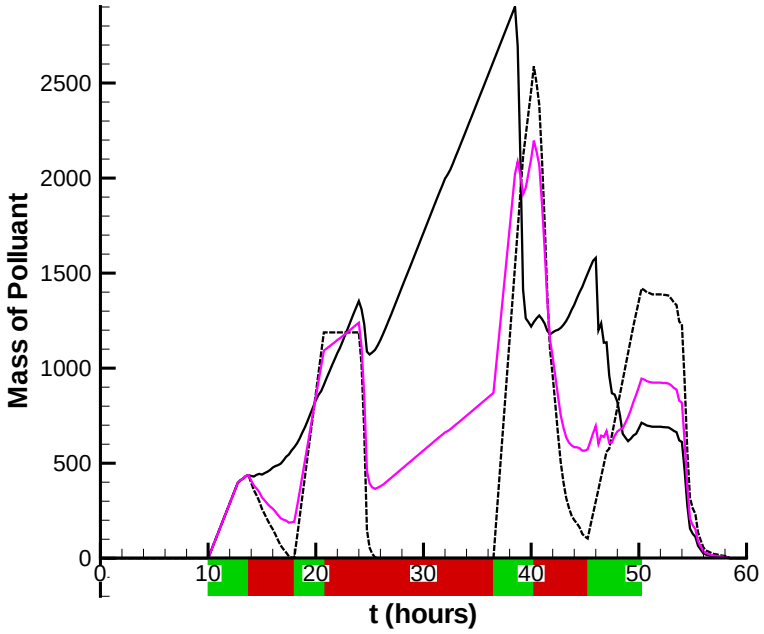

Figure 5.8: Three different release strategies for pollutants at a factory on the coast. The black line is for a uniform release in time. The dashed line is for releasing pollutants at times determined by the DLE analysis. The purple line is a strategy that is based on the DLE analysis and the criterion of minimizing peak values.

by a small amount. Fig. 5.8 reveals that a long "red" zone has the adverse effect of accumulating a large amount of pollutant in the tank. If such a zone if followed by a short "green" zone, a large amount of pollutant has to be released in a short amount of time, creating a high peak of concentration in the coastal area. To set up a more elaborate algorithm, we release a minimum flux of pollutant in the ocean independently of the type of zone (green or red). We define a new degree of freedom α, the percentage of incoming created contaminant that will be released during "red" zones. From this point of view, the solid curve of Fig. 5.8 corresponds to $\alpha = 30\%$, the dashed curve to $\alpha = 0\%$. The purple curve shows the result of a simulation using $\alpha = 33\%$ (i.e., 1/3 of the pollutant produced is always released directly in the ocean). Fig. 5.8 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.

5.7 An Automated Pollution Control Algorithm

The objective of this section is to show that it is possible to implement such a pollution release scheme based on the motion 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 hours 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 hours in the future.

5.7.1 Prediction Algorithm

Many predictive methods could be used. One might think of using a model of the area, use it to predict the velocity field and then compute the DLE field based on predicted velocities. However, the values of the DLE field are sequences in time and can also be predicted. We elected a simple prediction scheme. Since we are only interested in the longitude of the moving barrier, for each time, we computed the longitude of the barrier up to the latest time in the past where DLE can be computed (i.e., 8 hours before the

present time). This forms a time sequence that we would like to predict for at least 8 hours.

The spectrum of each sequence has been computed using the last 50 hours; a few samples are shown in Fig. 5.9. A complete animation of the spectrums computed at each time can be found on *http://transport.caltech.edu/florida*. 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. 5.10 shows the significant frequencies that were identified at each time step. Other simulations using longer time sequences (e.g., Fig. 5.11) reveal that the computed frequencies are not sensitive to the length of the time sequence used. Within an appropriate range (40-500 hours), using more data does not necessarily provide more accurate predictions.

Based on the relevant frequencies T_i, we used a simple predictive algorithm based on a Fourier decomposition. If we represent the motion of the Lagrangian barrier as a finite sequence of Fourier modes

$$y = A_0 + \sum_{i=1}^{N} A_i \cos\left(\frac{2\pi}{T_i}t\right) + B_i \sin\left(\frac{2\pi}{T_i}t\right), \tag{5.4}$$

the best model is obtained when one minimizes the error

$$\epsilon = \sqrt{\sum_{k} \left(y_k - A_0 + \sum_{i=1}^{N} A_i \cos(\frac{2\pi}{T_i}t_k) + B_i \sin(\frac{2\pi}{T_i}t_k) \right)^2}, \tag{5.5}$$

where y_k are the K positions of the Lagrangian barrier measured in the last 50 hours (at times t_k). The error is minimum when

$$\frac{\partial \epsilon}{\partial A_i} = 0, \tag{5.6}$$

and

$$\frac{\partial \epsilon}{\partial B_i} = 0, \tag{5.7}$$

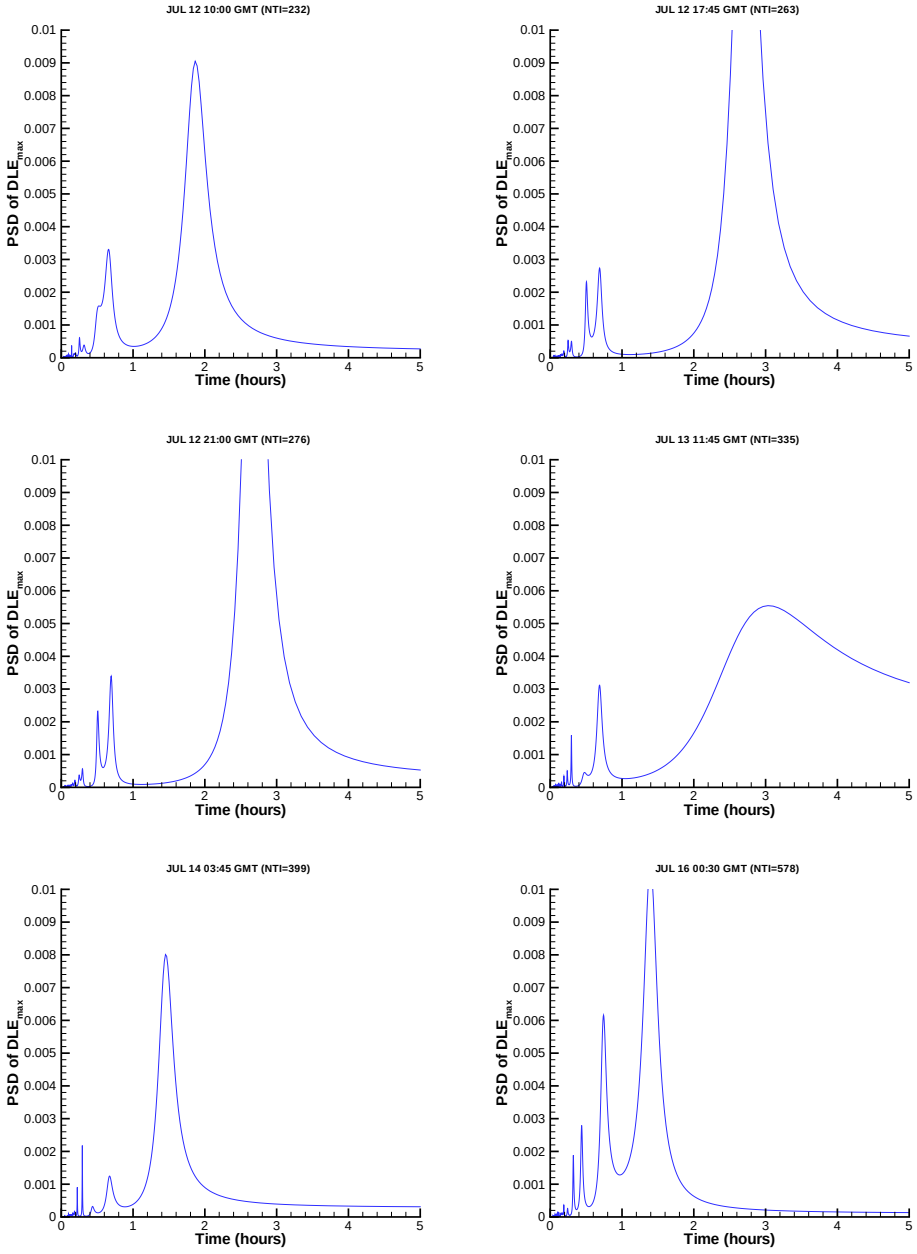

Figure 5.9: Entropy spectrum of the time sequence of the longitude of the Lagrangian barrier computed using 50 hours of data at selected time steps

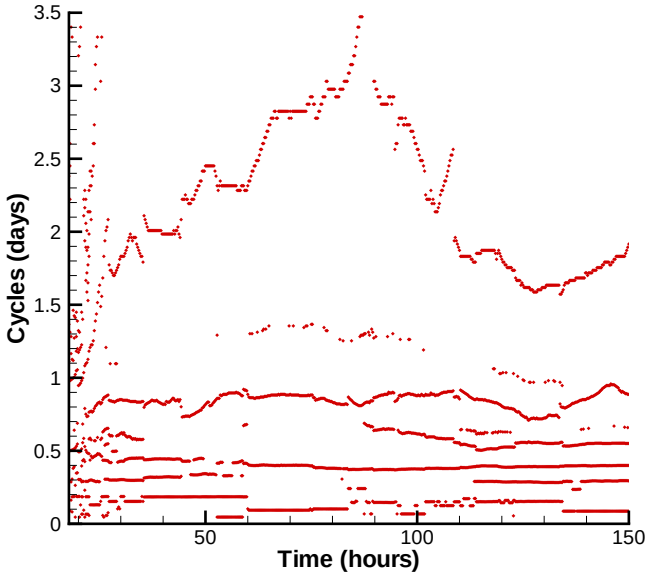

Figure 5.10: Spectral peaks in the time sequence of the longitude of the lagrangian barrier computed using 50 hours of data.

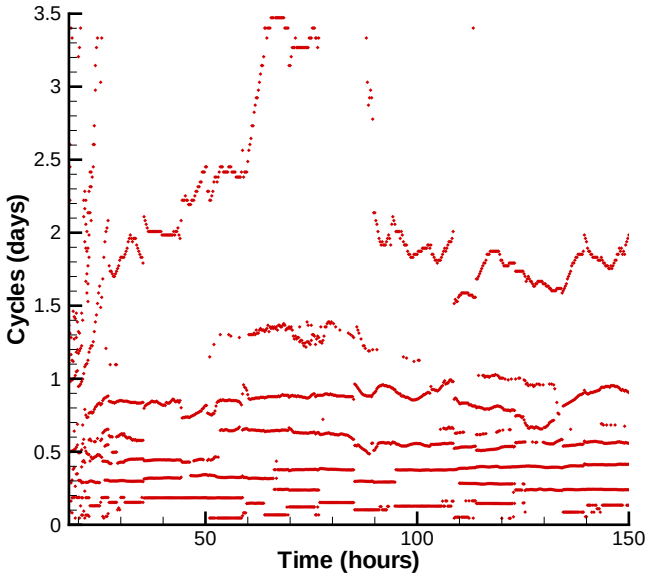

Figure 5.11: Spectral peaks of the time sequence of the longitude of the Lagrangian barrier computed using 75 hours of data.

for all i, which leads to the linear system in $(2K+1)$ unknowns[5]

$$\left(\; Q_{ij} \; \right) \begin{pmatrix} \vdots \\ A_j \\ \vdots \\ B_j \\ \vdots \end{pmatrix} = \begin{pmatrix} \vdots \\ \sum_k y_k \cos \frac{2\pi t_k}{T_j} \\ \vdots \\ \sum_k y_k \sin \frac{2\pi t_k}{T_j} \\ \vdots \end{pmatrix}, \tag{5.8}$$

where

$$Q_{ij} = \sum_k f_i(t_k) f_j(t_k), \tag{5.9}$$

$$f_i(t) = \begin{cases} \cos\left(\frac{2\pi t}{T_i}\right) & \text{if } 0 \leq i \leq K, \\ \sin\left(\frac{2\pi t}{T_i}\right) & \text{if } K < i \leq 2K+1, \end{cases} \tag{5.10}$$

The linear system given by Eq. 5.8 was solved using the *GNU Scientific Library (GSL)* for maximum 13 modes. The system is ill-conditioned when some frequencies are close to each other (e.g., if $t_k \approx t'_k$, the determinant of the linear system is close to zero, in which case the coefficients A_k, B_k, A'_k and B'_k can grow arbitrarily large). To avoid this problem, we ignore groups of close peaks (relative distance is smaller than 1%) and combine them automatically in a single resonant frequency. Moreover, we allowed the algorithm to select the best number of modes for each time step. The linear system solves the system give by Eq. 5.8 several times with only the jth most significant modes ($3 \leq j \leq 13$) and the best result is kept at each time slice.

At each time step, once the number K of modes is established and the coefficients A_k and B_k are computed, the sequence of Eq. 5.4 is used to extrapolate the position of the barrier during the past 8 hours were DLE is unavailable as well as in the future. Fig. 5.12 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 Jul 13 11:45 or have a strong divergence such as July 12 21:00 GMT. This phenomena has been observed several times

[5]The $2K+1$ variables are $A_0, A_1, A_2 \ldots A_K$ and $B_1, B_2, \ldots B_K$.

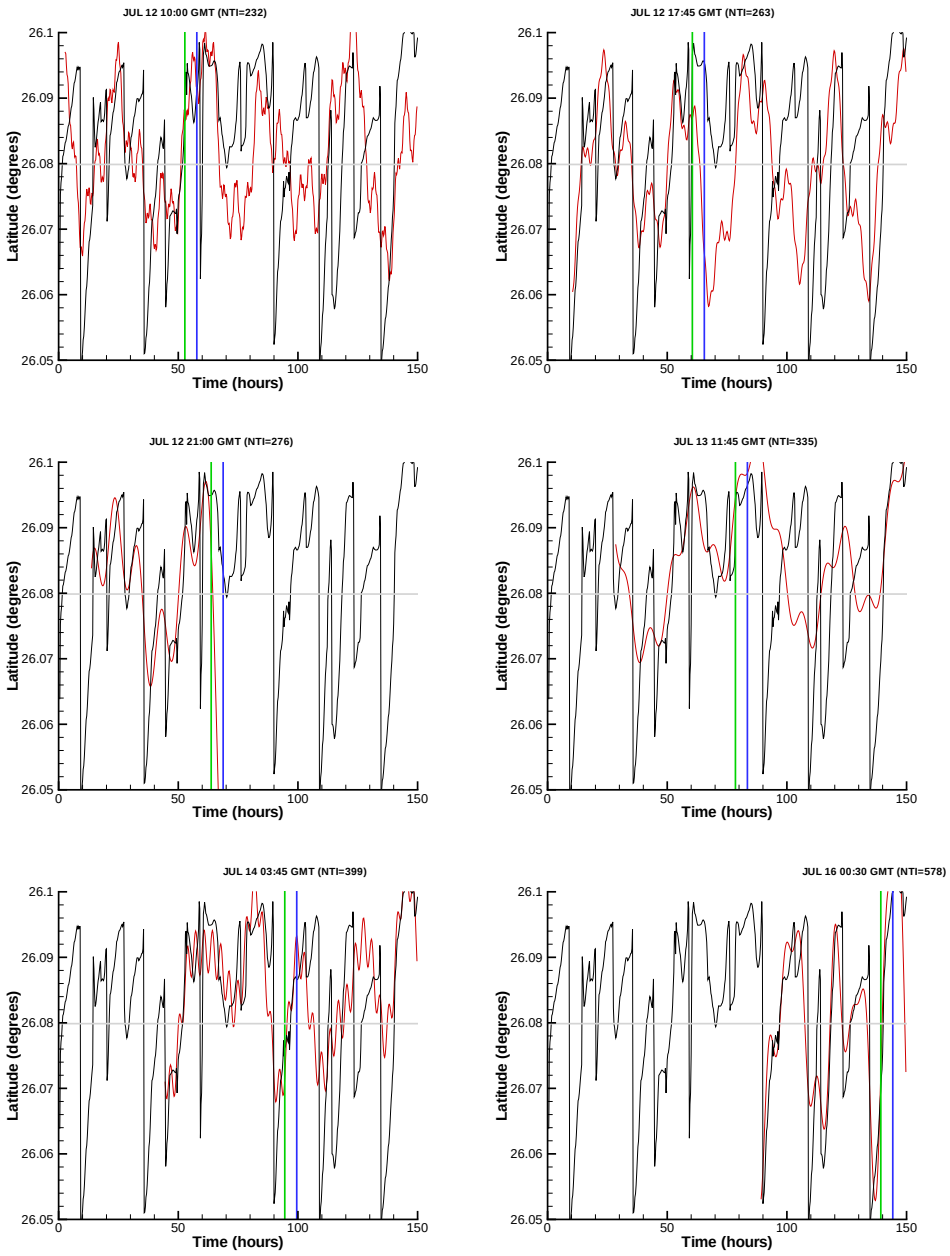

Figure 5.12: Predictive motion computed using 50 hours of data at selected time slices. The black curve represents the actual time sequence that can be computed at the end of the experiment (truth). The blue vertical line in each figure represents the actual time and the green vertical line represents 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 hours preceding the current time and in the future.

and is a consequence of the presence of close frequencies in the spectrum. Removing lower modes usually provides a sufficient solution.

We did not try to improve the prediction method. Our objective was to study the possible advantages of a pollution algorithm and we decided to let it use erroneous data as input, even though it was sometimes unable to predict the position of the barrier accurately. Our goal is to show that even in the presence of errors, a significant reduction of the concentration of pollution can be achieved.

To determine a favorable release-time intervals, we only consider the relative position of the predicted barrier with respect to the longitude of the release site. The release site is indicated by a the horizontal line on Fig. 5.12 and a favorable release interval occurs when the longitude of the barrier point is below the release site. To assess the quality of our prediction algorithm, we used the predicted curve 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[6]. 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. We ran the algorithm with different release sites and all of them had a success rate of about 70%-80%.

5.7.2 Decision Algorithm

To take into account the high error rate on the raw decision algorithm (favorable release zones), the factory or other release center is expected to run a more complex decision algorithm to decide how much pollution is to be release and how much is to be stored in the adjacent tank. We decided to use the data from the 4 last predictions (the last hour) to make a more accurate decision. 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

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

the tank. If the tank is full, the production is released in the ocean but the tank is not emptied.

- 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 not certain that the release will have a low impact on the environment, it is better not to take any risk and keep the pollutants produced. However, if the tank is more than half full, it is a much higher risk to store more pollutant because a very 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 seems 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 pollution has to be released in the ocean at all time to achieve effective reduction of the peak of maximum concentration.

The result of such an experiment is presented on Fig. 5.13. The factory was producing 1.2 tons of pollutant per hour and at least 0.4 tons/hours were to be dumped in the ocean ($\alpha = 33\%$). The tank had a maximum capacity of 30 tons and the maximum release rate from the tank was set to 3.6 tons / hour.

5.8 Conclusion

We have shown the existence of a set of repelling material lines in radar data obtained from the Fort Lauderdale area on the Florida coast. 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 making predictions about the trajectories of buoyant contaminants or the trajectories of nearly Lagrangian tracers. 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

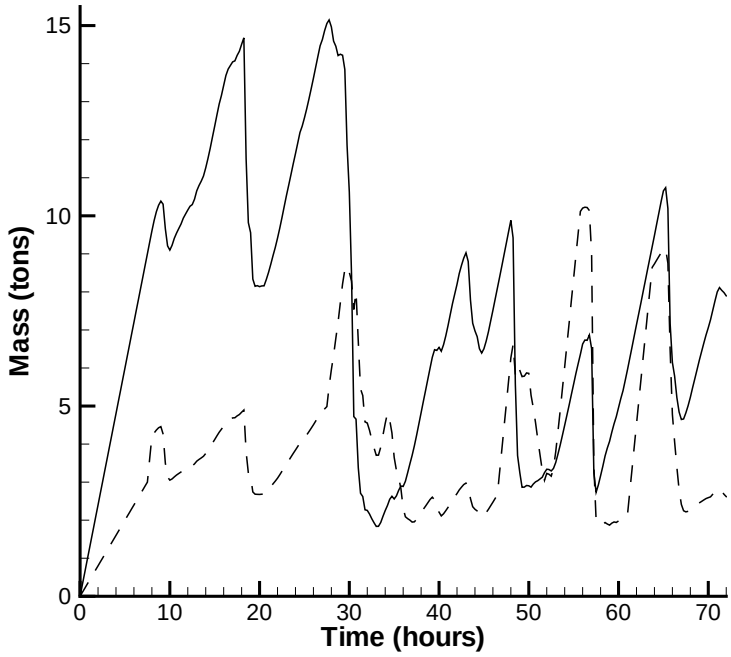

Figure 5.13: Mass of pollutant in the coastal area. The solid line represents the simulation of a source releasing pollutant at a constant rate. The dashed line is the simulation of a source of pollution releasing the same total amount of pollutant, but using the pollution control algorithm to determine environmental friendly release windows.

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.

Chapter 6

Why Are Warm Core Rings Always Anticyclonic?

In collaboration with Jerry Marsden

6.1 Introduction

In this chapter we describe the transport between the two gyres of a numerical circulation model of the North Atlantic. Our main result is a theorem explaining why cyclonic rings have always the same properties and temperature as the containing gyre. As a consequence only anti-cyclonic rings can have a core containing fluid with different properties than its environment.

To motivate our study, we present the sea temperature level sets in the North Atlantic during the first week of June 1984 on Fig. 6.1. The data used to construct this image was obtained with an AVHRR[1] sensor carried on a NOAA[2] satellite.

A radiometer on board of the satellite measures the amount of radiation emitted from the surface of the ocean which is directly related to the surface temperature. Each temperature is then assigned a different color and an image of SST is produced.

On Fig. 6.1, blue represents the coldest temperatures (between 0-15℃). Orange and red denote warmer temperatures (between 22-32℃). The Gulf Stream, located along the eastern coast of the United States is easily visible as the warmest water on the image. It is a fast, intense current known as a western-boundary current. These currents are located on the western side of every ocean basin. The Gulf Stream is a result of the wind

[1] Advanced Very High Resolution Radiometer
[2] Nation Oceanic and Atmospheric Administration

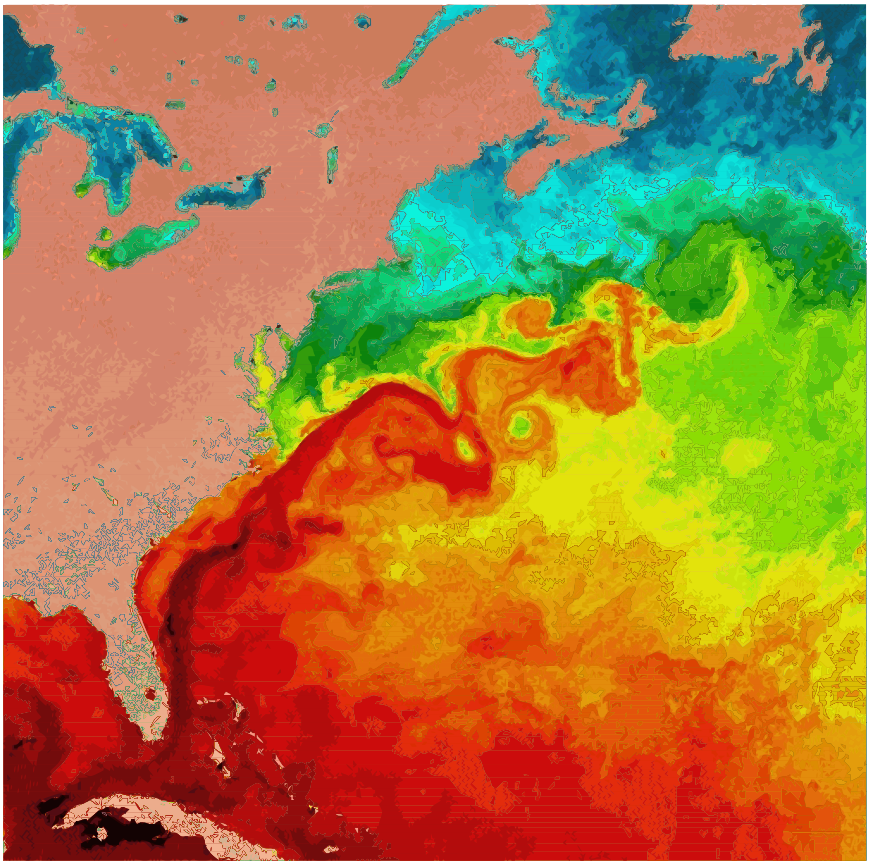

Figure 6.1: Shown as a bright red band, the Gulf Stream is about 27°C in this sea surface temperature image of the western North Atlantic during the first week of June 1984. This image is based on data from NOAA-7 Advanced Very High Resolution Radiometer (AVHRR) infrared observations, analyzed at the Rosenstiel School's Remote Sensing Laboratory. Warmer hues denote warmer temperatures and each pixel represents 1 km x 1 km of ocean. The Gulf Stream is an energetic, unstable flow. It meanders vigorously, with the meanders occasionally cutting off to form rings. Anticyclonic warm core rings are found north of the Gulf Stream and anticyclonic cold core rings are found south of the Gulf Stream. The rings migrate westward and occasionally remerge with the Gulf Stream. Credit: O. Brown, R. Evans and M. Carle, University of Miami Rosenstiel School of Marine and Atmospheric Science, Miami, Florida.

pattern acting on most of the North Atlantic Ocean. The combination of the Trade Winds (10°-25°N) blowing to the West and the Westerlies (35°- 55°N) blowing to the East cause

the North Atlantic to rotate clockwise. This basin wide, clockwise flow is referred to as the Subtropical Gyre. The Earth's rotation constraints the poleward flow in the western Atlantic to a narrow current on the western boundary of the ocean basin and generates the Gulf Stream. While western boundary currents exist in all the major ocean basins, the two in the northern Hemisphere (including the Kuroshio Current in the Pacific) are better developed and more intense than their southern Hemisphere counterparts.

Under the trade winds the broad, slow North Equatorial Current flows to the West. As the current approaches the Caribbean Sea, water is funnelled through the many channels of the Antilles and Caribbean Islands and through the Yucatan Channel - where it narrows and gains strength - and then loops into the Gulf of Mexico. This is where we can first observe an organized flow on satellite images. This flow enters the Gulf of Mexico at the Yucatan Channel and exits the Gulf at the Straits of Florida. The current is commonly referred to as the Loop Current while it is in the Gulf of Mexico and the Florida Current when it leaves the Gulf between Florida and Cuba.

To the North of the Subtropical Gyre is the Subpolar Gyre. The Subpolar Gyre which flows counterclockwise consists of much colder water than that found in the Subtropical Gyre. Fig. 6.1 also shows that the Gulf Stream does not follow a straight path. It has many twists and turns called meanders. If a meander becomes really sharp, it may pinch off and form what is called a ring. This is much like the formation of an oxbow lake by a river. Rings can be formed either to the North or to the South of the stream. For those rings formed to the North, the water in the center of the ring may be warmer than the surrounding water and these rings are called warm core rings. For those rings formed to the South of the stream, the center contains water that could be cooler than the surrounding water and they are called cold core rings. There are one warm core ring and two cold core rings on Fig. 6.1. The rings detach from the jet and move backwards to be re-injected in the jet.

Fig. 6.2 shows a magnification of two such rings on the satellite data gathered on May 29, 1997. The picture was extracted from the TOPEX/Poseidon mission, a joint venture between CNES and NASA to map ocean surface topography, launched on August 10, 1992 from the ESA launch facility at Kourou, French Guiana, aboard an Ariane 42P launch vehicle. Major observations were made using TOPEX/Poseidon data, including high-resolution (1-3 km) pictures of temperature fields.

The main observation to make on Fig. 6.1 and 6.2 is that the direction of rotation of a ring seems to determine wether or not it can have a warmer or colder core. Because of the

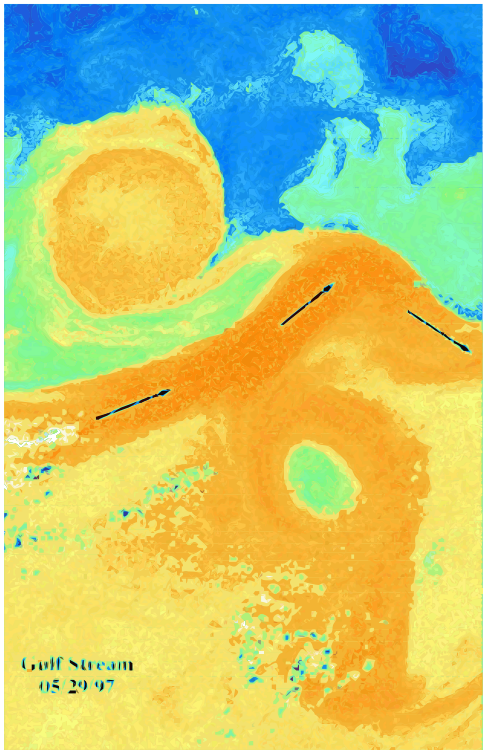

Figure 6.2: AVHRR sea surface temperature image of a portion of the Gulf Stream on May 29, 1997 extracted from the Topex/Poseidon mission. A warm core ring and a cold core ring are easily visible and were both formed between May 20-29. The warm ring sitting in the counterclockwise Subpolar Gyre is anticyclonic, as well as the cold counterclockwise ring in the clockwise Subtropical Gyre. The main result of this chapter is to explain why only anticyclonic rings can be involved in intergyre transport, i.e., only anticyclonic rings can have a warm or cold core. It is worth to remark that cyclonic rings are usually unnoticeable on SST images. They contain fluid from the containing gyre and there is no contrast for these rings in the SST image.

Gulf Stream, the North Atlantic is divided in two giant gyres turning in opposite direction: the cold and counterclockwise Subpolar Gyre in the North and the warmer and clockwise Subtropical Gyre in the South. Rings turning in the same direction as their containing gyre are said to be cyclonic. Rings turning in the opposite direction are anticyclonic. Satellite imaging of the North Atlantic shows clear evidence that warm core rings in the Subpolar Gyre are always oriented clockwise. On the other hand, cold core rings in the Subtropical Gyre have always a counterclockwise motion. As a result, rings containing

"foreign" fluid with a temperature similar to the other gyre are always anticyclonic. A clear evidence of this fact on SST pictures is that the radiometer is completely blind to cyclonic rings. There are no visible cyclonic rings on Fig. 6.1 and 6.2. As a matter of fact, cyclonic rings contain water with the same temperature, salinity and other properties than the containing gyre and are unnoticeable. However, assuming that the surface layer of the ocean is reasonably incompressible, there must be roughly as many rings of each type.

We show that under general assumption the direction of rotation of a cold or warm core ring is constrained by the presence of hyperbolic structures. These dynamical constraints force rings containing "foreign" fluid to turn in an anticyclonic direction. We use a three-layer quasi-geostrophic wind-driven model from Clark Rowley [Rowley, 1996] to illustrate this fact and give numerical evidence of the motion of the Lagrangian structures associated with the formation of an anticyclonic ring. This model is incompressible, but the main result, which states that cyclonic rings always have the same properties as their containing gyre, does not depend on this assumption.

6.2 The Model

We consider the model of Rowley [1996]. The data set used here is obtained by a numerical simulation of the mid-latitude ocean circulation in a rectangular basin geometry based on a three-layer, eddy-resolving quasi-geostrophic model subject to a wind stress curl. The basic circulation pattern in the upper layer is a double-gyre structure (Subtropical and Subpolar gyres) separated by an eastward jet (Gulf Stream) shooting off from the confluence point of the northward and southward western boundary currents (Fig. 6.3).

Depending on the ocean basin size and the wind stress amplitude τ, the circulation exhibits rich time-dependent dynamics [Dijkstra and Katsman, 1997] for a more detailed analysis) with realistic parameter values for the mid-latitude ocean [Lozier and Riser, 1990]. In a 1000 km $\times$ 2000 km rectangular domain with relatively mild amplitude $\tau \approx 0.16 - 0.17$ dyn/cm^2, unsteadiness of the circulation is spatially confined in the vicinity of the eastward jet near the western boundary. In this case, main Lagrangian transport occurs between the two gyres across the jet and it is governed by the invariant manifolds of the persistent and nearly stationary one-dimensional hyperbolic trajectory on the western and eastern boundaries. Such intergyre transport using the techniques of the Lagrangian lobe dynamics has been studied in Coulliette and Wiggins [2000].

In a larger 2000 km $\times$ 2000 km square basin with higher wind stress amplitude at

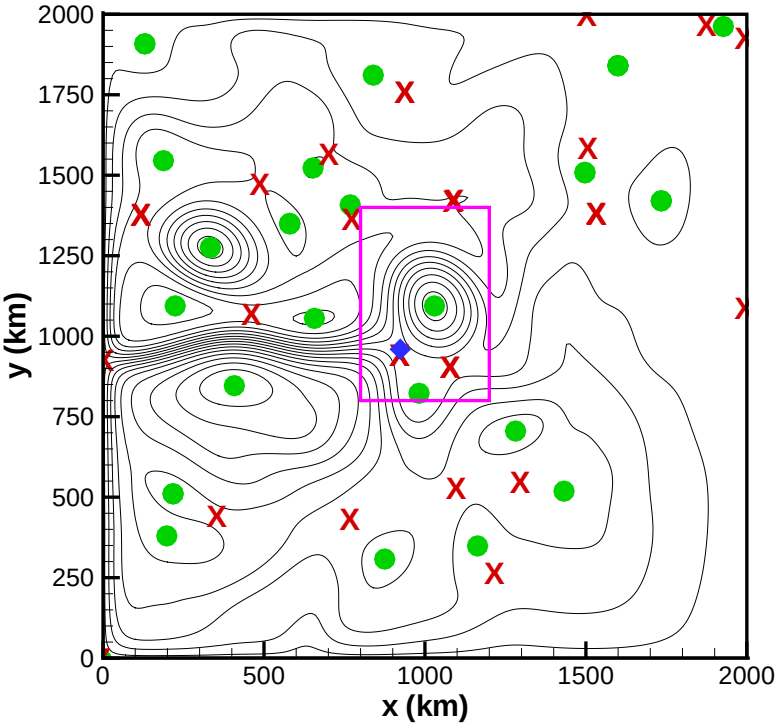

Figure 6.3: Instantaneous streamfunction field in the upper layer at the beginning of the 3 months period with $\tau = 0.5$ dyn/cm^2. Instantaneous hyperbolic and elliptic SPs are plotted by (red) X and (green) circle.

$\tau = 0.5$ dyn/cm^2, the circulation becomes more energetic and irregular time dependence spreads throughout the ocean basin. In a time sequence of the instantaneous streamline field, we observe that both cyclonic and anticyclonic eddies are formed by pinching off from the downstream jet (see purple rectangle on Fig. 6.3) and may migrate westward due to the interaction with each gyre's recirculation cell. Computations for the ring formation process in this paper have been made using $\tau = 0.16$ dyn/cm^2, a weaker wind stress amplitude (see Fig. 6.4). This choice has been made as an effort to produce simpler geometries to illustrate the theoretical framework underneath the fundamental processes. Similar conclusions arise with a higher (more realistic) wind stress.

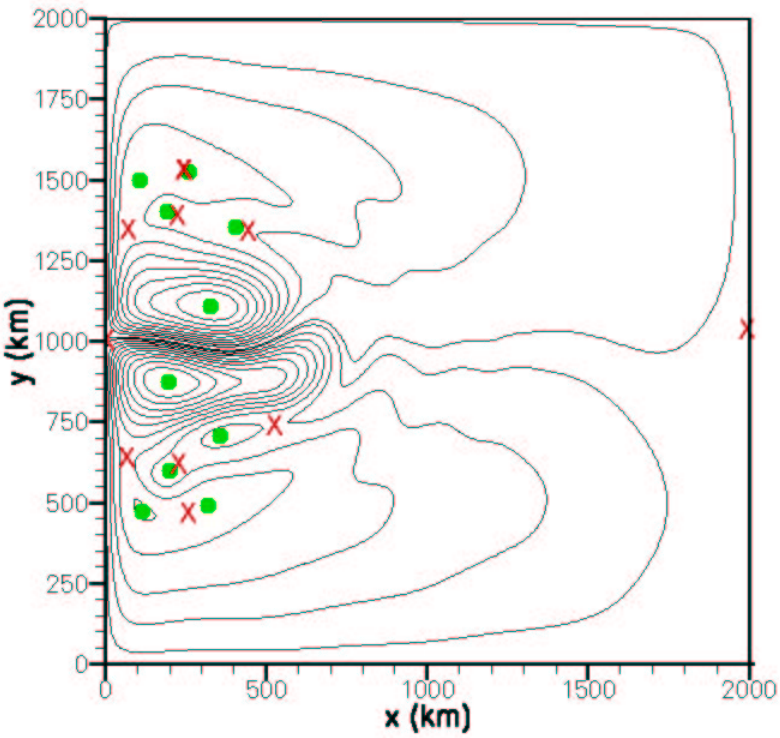

Figure 6.4: Instantaneous streamfunction field in the upper layer at the beginning of the 3 months period with $\tau = 0.16$ dyn/cm^2. Instantaneous hyperbolic and elliptic SPs are plotted by (red) X and (green) circle.

6.3 Lobe Dynamics and Inter-Gyre Transport

As a consequence of the double-gyre structure of the flow, there are two hyperbolic trajectories, one on the eastern boundary, whose unstable manifold propagate Eastwards in the basin and one on the eastern boundary whose stable manifold propagates Westwards [Coulliette and Wiggins, 2001]. These two manifolds form a subtle boundary between the southern and northern gyres (see Fig. 6.4). As depicted in Coulliette and Wiggins [2001], transport between these gyres is completely governed by the evolution and geometry of these manifolds. The hyperbolic trajectory on the western boundary is the (moving) point at which the flow along the boundary converges from the southern and northern gyres and subsequently moves into the interior of the flow. The existence of a material line attached to the western boundary and dividing the flow in two gyres has been verified by comput-

ing the ridge of maximum stretching described in Haller [2001a] (see Fig. 6.5). Similarly, on the eastern boundary there exists a (moving) point at which the flow separates as it collides with the eastern boundary, resulting in some fluid moving northwards along the eastern boundary and some moving South. This trajectory on the boundary has an associated stable manifold that extends into the interior of the flow. The unstable manifold

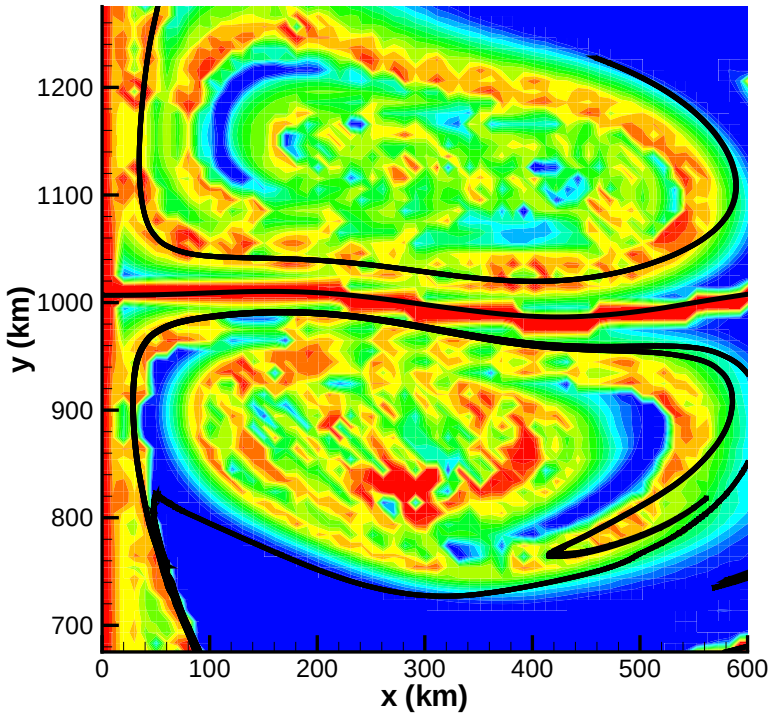

Figure 6.5: Comparison between the unstable invariant manifold and the ridge of maximum finite-time Lyapunov exponent. The hyperbolic trajectory on the western boundary correspond to both the trajectories of Ide et al. [2002] and the trajectory with maximum σ in Haller [2001a]

emanating from the western boundary may intersect the stable manifold emanating from the eastern boundary. An intersection point is primary (pip) if the segment of the unstable manifold from the western boundary to the intersection point does not intersect the segment of stable manifold from the intersection point to the stable hyperbolic trajectory. All other intersection points are called secondary intersection points (sip) [Romkedar, 1990;

Romkedar and Wiggins, 1990; Romkedar et al., 1990]. It is important to distinguish primary and secondary intersection points because it is the sequence of primary intersection points (pips) and the segment of unstable and stable manifolds between them that define lobes (Fig. 6.6). If the time dependence of the flow is periodic, then an infinite number of lobes is created. An extensive analysis of transport in two-dimensional maps can be found in Rom-Kedar and Wiggins [1991]; Romkedar [1994, 1995]. In temporally chaotic cases, we found numerical evidence both in numerical simulations and satellite images of an arbitrarily large number of lobes.

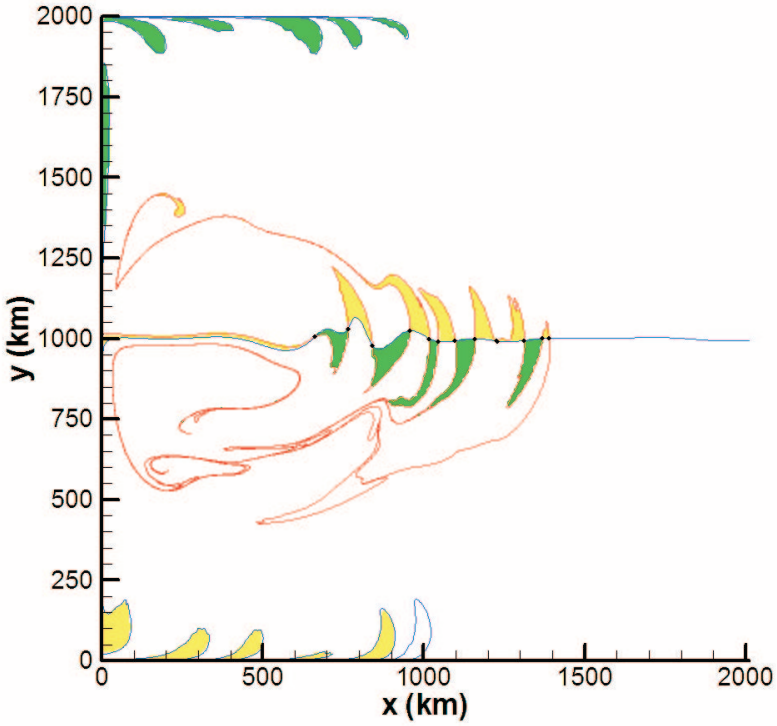

Figure 6.6: Unstable invariant manifold (in red) of the hyperbolic trajectory on the western boundary. Stable invariant manifold of the hyperbolic trajectory (in blue) on the eastern boundary. The two manifolds intersect transversally and between two primary intersection points (black dots), the manifolds define lobes (green and yellow surfaces). Green lobes appear to come from the northern gyre and are transported to the southern gyre. The opposite transport phenomena occurs for the yellow lobes.

At the initial time we define a boundary between the two gyres by taking the pieces of the stable and unstable manifold only up to some (arbitrary) pips. Under time evolution segments of the unstable manifold lengthen, segment of the stable manifold shorten, and intersection points must remain intersection points. The boundary defined above between the southern and northern gyres has deformed under time evolution. At the new time, we choose a new intergyre boundary by selecting a new pip (Fig. 6.7).

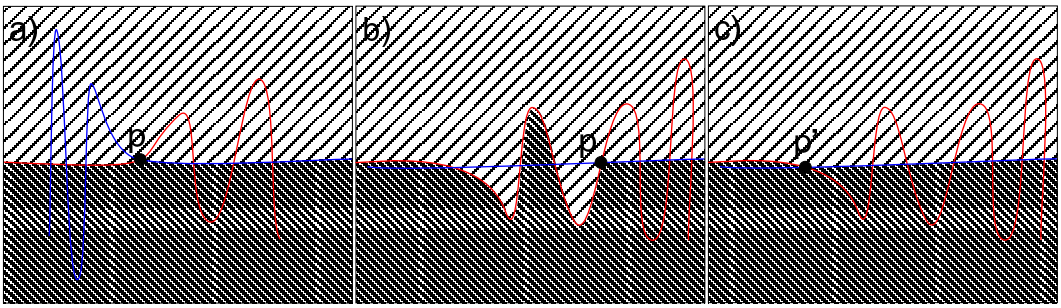

Figure 6.7: Transport in a double-gyre structure (schematics). (a) At some time t, the primary intersection point (pip) p is elected boundary intersection point (bip) The material boundary between the two gyre is the union of the segment of unstable manifold from the hyperbolic trajectory on the western boundary to p and the segment of stable manifold from p to the hyperbolic trajectory on the eastern boundary. (b) At time $t' > t$, point p has moved to the East and the boundary $U[ht_w, p] \cup S[p, ht_e]$ is not physically acceptable anymore. (c) another pip p' is elected as the bip for time t'. The new boundary is more realistic and between time t and t', the fluid transported from one gyre to the other is exactly the lobes between p and p'.

The amount of fluid that has moved from the southern gyre to the northern gyre is the area of the hatched lobes between the previous chosen pip and the new one. Once a boundary is chosen, consider a set of lobes that is defined by moving to the right along the stable manifold and stopping at some other pip. The area of these lobes is then the amount of fluid that has crossed the boundary between the southern and northern gyres at the present time since the earlier time.

6.4 Ring Formation

During an average year 10 to 15 rings are formed. The rings are usually between 150-300 km in diameter and the velocity at their edge can reach 1 m/s. The rings generally drift to

the South West at approximately 5 km/day. Warm core rings are always located between the Gulf Stream and the continental slope. Fig. 6.8 and 6.9 show an example of warm core ring that appeared during our simplified simulation. The life cycle of a warm core ring is generally a few months up to a year. These rings usually are re-absorbed by the Gulf Stream, although they also can break apart if they move onto the continental shelf. Warm core rings travel to the West, so many warm core rings are reabsorbed into the Gulf Stream near Cape Hatteras as they are squeezed between the Gulf Stream and the continental shelf. This process of rings being reabsorbed back into the Gulf Stream is known as coalescence. Warm core rings are easily observed in satellite images due to the large thermal contrast between the ring and the cold surrounding waters.

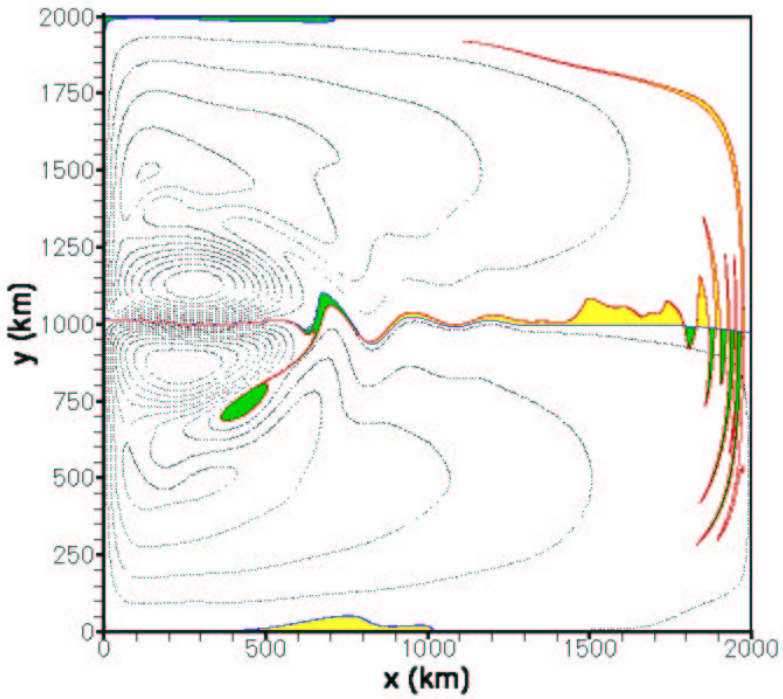

Figure 6.8: Creation of an anticyclonic warm core ring in the southern gyre. The green lobe deforms and becomes the warm core of the spinning ring.

From an Eulerian point of view, an anticyclonic ring detaching from the jet must create a hyperbolic stagnation point between the ring and the jet (see the red crosses on Fig. 6.3 and Fig. 6.4). We are expecting to find a hyperbolic trajectory and a Lagrangian coherent

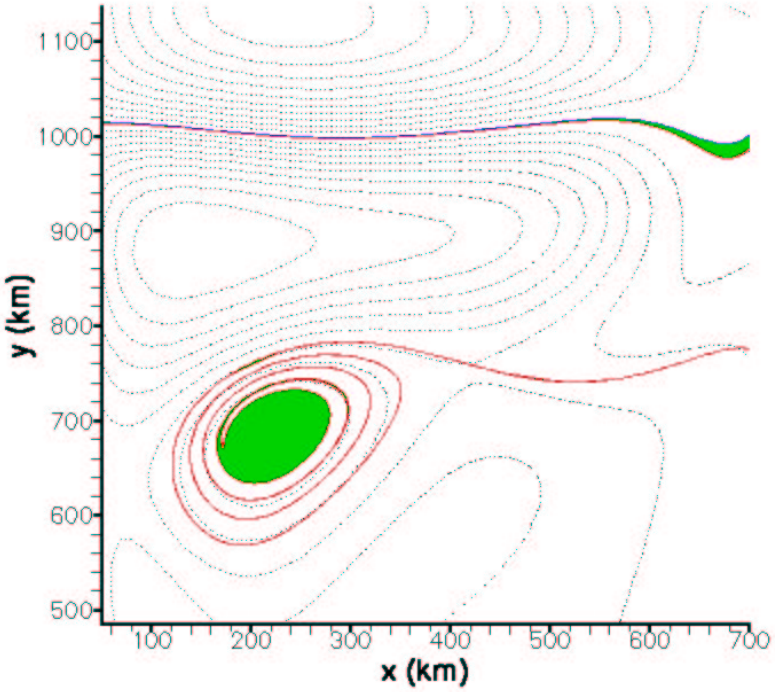

Figure 6.9: Enlargement of the ring shown in Fig. 6.8. The unstable manifold of the hyperbolic trajectory on the western boundary spins in a counterclockwise motion.

structure (see Haller [2001a] for more details on LCS) between the jet and the ring. This phenomena has been described and observed on a model of the Atlantic in Poje and Haller [1999].

6.5 Manifold Deformation Framework

The key to determining the direction of rotation of a ring is to understand the interaction between stable and unstable manifolds in a time-dependent flow. Intergyre transport is governed by the stable and unstable manifold of hyperbolic trajectories at the edges of the basin. The boundary of regions of foreign fluid (cold or warm core) are segments of these invariant manifolds [Poje and Haller, 1999]. The deformation of the invariant manifolds in the presence of other hyperbolic structures (such as the invariant manifolds of a hyperbolic trajectory associated with a ring) is studied in this section.

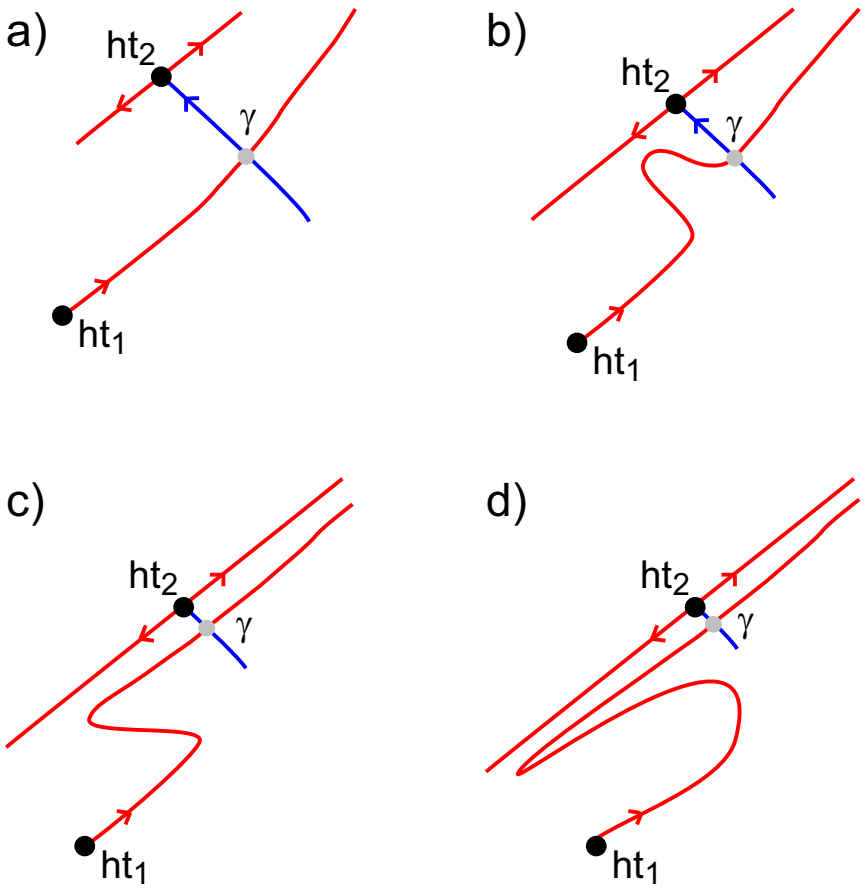

Figure 6.10: Manifold Deformation Process. The unstable manifold of ht_1 intersects the stable manifold of ht_2. This implies the deformation of the former manifold along the "deforming side" of the unstable manifold of ht_2

Figure 6.10 shows a typical situation in an intergyre transport problem. An unstable manifold is associated with a primary hyperbolic trajectory ht_1 and propagates inside a jet. An Eulerian ring can appear and create an Eulerian hyperbolic stagnation point, which is often associated with a finite time hyperbolic trajectory ht_2. We consider the simplest interaction possible where the stable manifold of the hyperbolic trajectory intersects the unstable manifold of ht_1. The intersection point γ must remain on both manifolds for the time that ht_1 and ht_2 remain hyperbolic. As a result, γ approaches the hyperbolic trajectory ht_2 at an exponential rate. On the other hand, trajectories on the unstable manifold of ht_1 that are sufficiently close to ht_1 separate exponentially and are responsible for a growth in length of the unstable manifold. Since γ cannot propagate further away

from ht_1 (it is "locked" by the stable manifold of ht_2), a loop must appear in the segment of manifold from ht_1 to γ. One can also explain this phenomenon by looking at the unstable manifold of ht_2. As the unstable manifold of the hyperbolic trajectory ht_1 approaches ht_2 on its stable manifold, the dynamics on the unstable manifold of ht_2 affects strongly the shape of the manifold[3]. Notice that if ht_2 travels at an exponential rate along the unstable manifold of ht_1, particles on the unstable manifold of ht_2 are still required to separate from ht_2 at an exponential rate. As a result, a sharp kink is also created and the same deformation framework takes place.

During this interaction, the unstable manifold of ht_1 is in contact with the two unstable manifolds of ht_2. One is oriented in the same direction and has no noticeable effect. The two manifolds continue together. The other segment is oriented in the opposite direction and pulls the manifold back to where it came from. However, this effect is local and only affects a small portion of the manifold before the intersection point γ. As a result, the unstable manifold experiences a sharp kink along the unstable manifold of ht_2 that does not have the same direction.

Theorem 6.5.1 (Manifold Deformation Framework) *If an unstable manifold is sufficiently close to another hyperbolic trajectory (called secondary ht) and intersects transversally one of its stable manifold, the time evolution of this unstable manifold will reveal at least one sharp kink along the unstable manifold of the secondary hyperbolic trajectory that is oriented in the opposite direction (called the deforming side). The kink lasts at least during the lifetime of the secondary hyperbolic trajectory (see Fig. 6.10).*

Planar autonomous systems do not allow transverse intersection of the manifolds and hence the theorem is irrelevant. In contrast, flows with periodic time-dependence are known to admit intersecting stable and unstable manifolds that create heteroclinic tangles [Romkedar, 1990; Romkedar and Wiggins, 1990]. Theorem 6.5.1 is a weak version of the heteroclinic tangle. In time-periodic systems, the intersection between the two manifolds, as described in theorem 6.5.1, implies an infinite number of intersections between the two manifolds and an infinite number of lobes. As time evolves, the lobes are stretched along both sides of the unstable manifold of the secondary hyperbolic trajectory, creating an infinite number of sharp kinks.

However, Theorem 6.5.1 provides an extension of the heteroclinic tangle framework to

[3]The influence of the unstable manifold of ht_1 on the manifolds of ht_2 is much less since ht_1 is far away.

quasi-periodic and time-chaotic dynamical systems. In particular, time-chaotic systems can develop very short time hyperbolic trajectories. In such cases, the sharp kink of theorem 6.5.1 may be the only visible deformation on the unstable manifold. Other cases with a finite number of intersections between the two manifolds lead to more kinks and more complex structures, all of them showing at least a sharp kink on the deforming side of the manifold. Sharp kinks have been noticed many times during our computations using the model described in the previous section as well as simulations based on HF radar data. This phenomena is a generic pattern of deformation of an unstable manifold under the presence of other hyperbolic trajectories and is not specific to a ring formation process. It also happens at different places and time in the flow.

A similar theorem can be established for stable manifolds intersecting a local unstable manifold. In this case, a sharp kink of the stable manifold is "predicted in the past."

6.6 Warm and Cold Core Rings

6.6.1 Anticyclonic Rings

The formation of anticyclonic rings is described in [Poje and Haller, 1999]. Figs. 6.11 and 6.12 illustrate the two possible situations when the anticyclonic ring is created near a cold-core lobe. In Fig. 6.11, the unstable manifold of the hyperbolic trajectory on the western boundary turns around the ring and becomes thinner and thinner. As a result, the anticyclonic ring of Fig. 6.11 does not have a cold core ring. On the other hand, Fig. 6.12 shows a more subtle behavior of the lobe. A stable manifold deforms the unstable manifold defining the lobe in such a way that a sharp kink spirals around the ring, creating a cold core in the center of the ring.

6.6.2 Cyclonic Rings

The process depicted in Fig. 6.11 is possible when the rotation is clockwise. As a result, a cyclonic ring can also be made of local fluid (from the same gyre) with a thin lobe spiralling in at the border of the ring. However, the case depicted in Fig. 6.12 is impossible for cyclonic rings due to geometric constraints generated by the MDF and it is the main object of this chapter.

Theorem 6.6.1 *Cyclonic rings cannot contain fluid from the other gyre.*

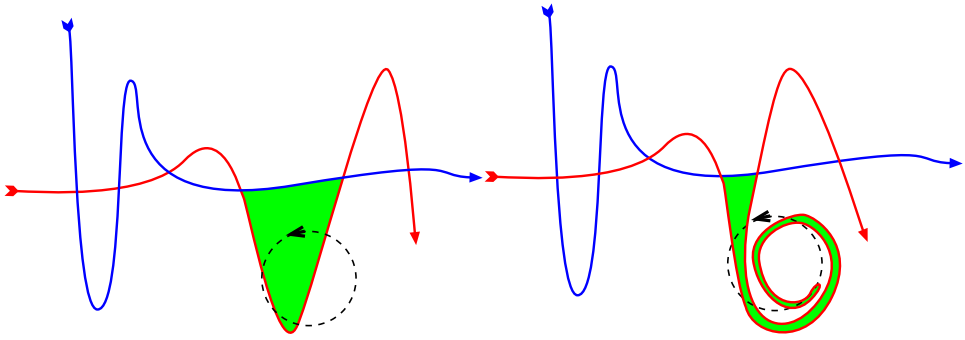

Figure 6.11: Anticyclonic ring formation in the southern gyre. In this simple process, the end of the lobe spirals inside the Eulerian ring. The interior of the ring is filled with fluid from the containing gyre.

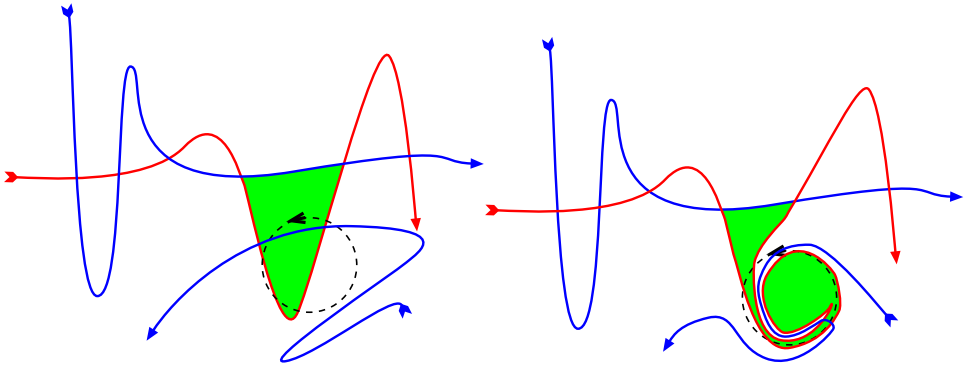

Figure 6.12: Anticyclonic ring formation in the southern gyre. The subtle mechanism described can only happen in anticyclonic rings. A stable manifold deforms the lobe and allows a patch of fluid from the other gyre to stay in the middle of the Eulerian ring. This ring contains fluid coming directly from the other gyre and is typical in measured data from the northern Atlantic.

Proof:

Let us consider the southern gyre. A similar proof can be derived for the other gyre. A region containing "foreign" fluid in the southern gyre is a green lobe of Fig. 6.8. The part of its boundary inside the southern gyre is a segment of unstable manifold turning in a counterclockwise direction. As seen on the dynamical screenshots of Fig. 6.12, the formation of a cold core ring involves a strong, hyperbolic deformation of this manifold, which is compatible with the presence of a hyperbolic trajectory near this manifold. If

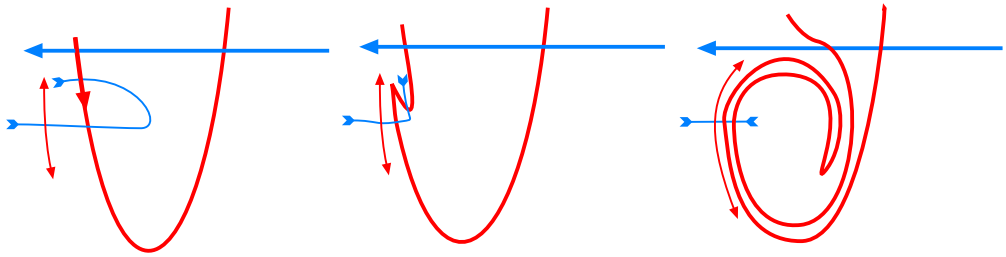

Figure 6.13: A hyperbolic trajectory outside the lobe creates a cyclonic ring that does not have a warm core.

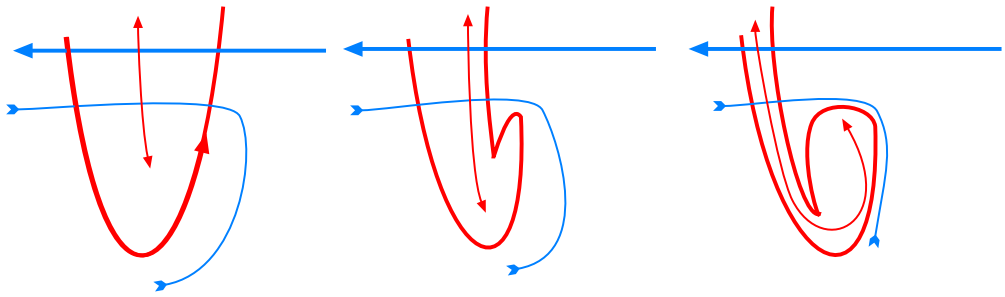

Figure 6.14: A hyperbolic trajectory inside the lobe can create an anticyclonic ring that has a warm core.

the secondary hyperbolic trajectory is outside the lobe, a kink develops and necessarily turns clockwise (see Fig. 6.13). However, in this case the kink itself is made of foreign fluid. The whole lobe deforms in a long clockwise spinning structure. On an SST image, the cyclonic ring does not appear to contain colder fluid. If the hyperbolic trajectory is inside the lobe (see Fig. 6.14 and Fig. 6.15), the ring can have a cold core. We will only consider the case where the two sides of the stable manifold of the ring intersect the unstable manifold that makes the border of the lobe [Poje and Haller, 1999]. There are two possible interactions allowed by Theorem 6.5.1. In Fig. 6.14, we consider the interaction of the manifold with the eastern stable manifold. The unstable manifold deforms and create a cold core that must spiral in a counterclockwise direction. This creates an anticyclonic cold core ring. One might notice that the interaction with the other side of the stable manifold is likely to generate a cyclonic cold core ring, so we must dismiss this case to prove Theorem 6.6.1. Fig. 6.15 reveals that this particular geometry implies new intersections of both the northern unstable manifold of the ring and the unstable manifold of the hyperbolic trajectory on the western boundary with the unstable manifolds of the ring. Moreover, each rotation of the ring implies new intersections between these manifolds. This case is

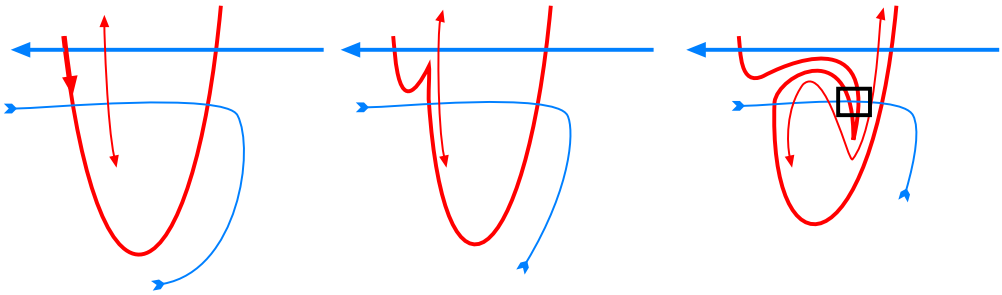

Figure 6.15: A hyperbolic trajectory inside the lobe cannot create a cyclonic ring that has a warm core without creating non-generic intersection between the manifolds.

highly non-generic. Intersection points between manifolds must be preserved and exist at all time. Such a complex structure has never been observed analysis.

Moreover, the unstable manifold that makes the boundary of the core ring intersects several times the two stable manifolds of the hyperbolic trajectory in the middle of the ring. As a result, the border of such a ring would shrink and become a long filament. An observer looking at an SST would not conclude that there is a cyclonic cold core ring. As a result, only the case in Fig. 6.14 is possible ad cold core rings are always anticyclonic.

6.7 Conclusion

We have shown that the transport between two gyres (or more generally between the two sides of a jet in a flow) does not involve cyclonic rings. We conclude that only anticyclonic rings can have a warmer or colder core and this result does not depend on incompressibility. This result explains the theoretical and numerical results found in Coulliette and Wiggins [2000]; Poje and Haller [1999]

We have extracted the Lagrangian structures involved in the proof for a simple double-gyre model and we have shown sea surface temperature plots of the northern Atlantic from satellite imagery to verify that cyclonic rings cannot contain foreign fluid.

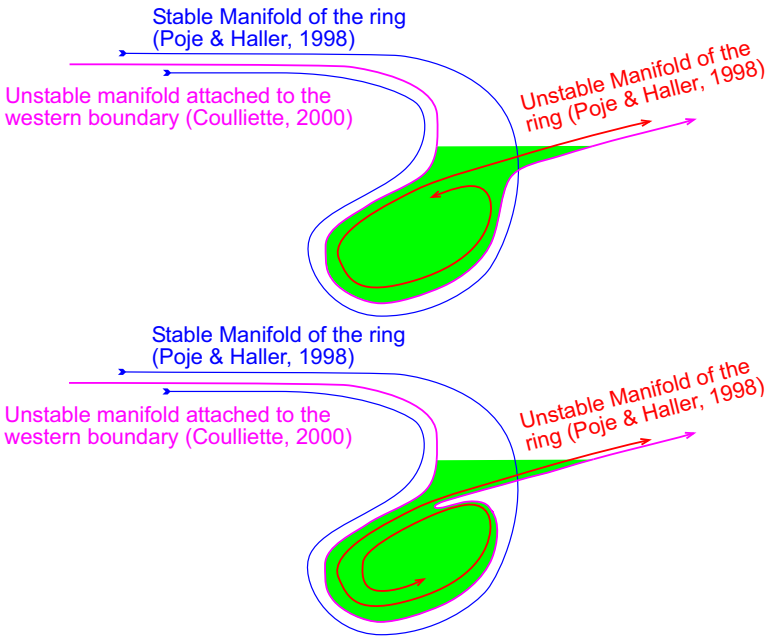

Figure 6.16: Anticyclonic ring formation in the southern gyre. The ring will contain a foreign (warmer) core because the stable manifold attached to the western boundary [Coulliette and Wiggins, 2000] intersects the loop of the stable manifold associated with the creation of the ring [Poje and Haller, 1999] before following the unstable manifold of the ring. The intersection of the two Lagrangian structures moves slowly with the ring but the length of unstable manifold between the western boundary and the intersection must increase exponentially. As a result, the unstable manifold from the western boundary spiral down with the unstable manifold of the ring, creating the boundary of the foreign core.

Chapter 7

Conclusion

The use of dynamical systems theory in geophysical flows is a relatively new and quickly growing area of research. In the past, many industrial and military applications benefited from a better knowledge of ocean dynamics. Up until recently, efforts were limited to Eulerian models and observations. However, many authors noticed that Lagrangian transport and coastal advection may be far more complicated than their corresponding Eulerian signatures. A simple analysis of Eulerian velocity fields is insufficient, while dynamical systems theory provides barriers, transport conduits, and a geometric view of the processes involved.

Major dynamical systems results for two-dimensional autonomous velocity fields and two-dimensional time-periodic velocity fields such as the Poincaré-Bendixon theorem, the heteroclinic tangle, lobe dynamics and Melnikov methods do not, unfortunately, apply directly to time-aperiodic velocity fields and these are the ones important in ocean dynamics.

Moreover, there are no realistic analytic solutions or perturbed solutions of the Navier-Stokes' equations on a complex domain such as a coastline. As a result, a dynamical system representing a geophysical flow is often described by the numerical output of an ocean simulation or experimental measurements. Thus, in this setting, the relevant dynamical systems are not defined by analytical equations, but rather by a finite data set containing discrete values.

This thesis contributes to these two problems by providing tools and some solutions. Our major accomplishments are:

- **Abstraction Layer for Dynamical Systems.** The generalization of normal mode analysis to domains with an open-boundary allows the representation of experimen-

tally observed velocities (such as HF radar data) as a finite sequence of modes. These modes can be interpolated everywhere in space with advanced techniques such as our C^1 tri-cubic interpolator. In our view, equations or finite-time data sets obey a similar dynamical definition, and dynamical systems tools can be applied to both indifferently.

- **Manifold Deformation Framework.** The MDF is a geometric description of non-autonomous transport in terms of invariant manifolds. In our view, MDF is a weak version of the heteroclinic tangle (or λ-lemma) that remains true for time-chaotic systems. Application of this framework explains, for example, why cold-core rings in the North Atlantic are always anticyclonic.

- **Examples of Geophysical Flows.** Based on HF radar collected in Monterey Bay and the coast of Florida, we carried out dynamical systems studies of flows defined by observational data. In both cases, hyperbolic trajectories, invariant manifolds, Lyapunov exponents, and Lagrangian Coherent Structures were computed.

- **Coastal Pollution Optimization Algorithm.** Based on our experience with the Monterey Bay and Florida data, we developed an algorithm to reduce the impact of released contaminants in a coastal area. This algorithm assumes the availability of HF radar stations or other Eulerian measurement devices, and uses a holding tank to store contaminants between timed releases. The algorithm applies directly to real-life pollution problems.

- ***MANGEN* Software Library.** *MANGEN* was designed as a re-usable package for transport analysis, containing several interpolation, advection and data manipulation routines.

Major challenges still remain in defining the principal players in time-dependent flow with arbitrary time dependence: hyperbolic trajectories, separation points, invariant manifolds and other Lagrangian structures. Moreover, the study of the interactions and dynamical constraints between these objects is expected to lead to analogs of the *Poincaré-Bendixon* theorem or Lobe dynamics for time-chaotic flows. Higher-dimensional systems (three-dimensional time-chaotic flows or classical mechanical systems with more than two degrees of freedom) are also the focus of increasing interest. Appendices C and D show some examples of such systems.

Appendix A

Eigenfunction Basis in Sobolev Spaces

A.1 Lebesgue Spaces and Properties

We assign the L^2 scalar product

$$(f, g) = \int_\Omega f(x)g(x)dx, \tag{A.1}$$

to the space $C_0^\infty(\Omega)$ of all the smooth functions with a compact support in the open set Ω. The associated L^2 norm is given by

$$|f| = \sqrt{(f, f)}. \tag{A.2}$$

Not all Cauchy sequences (f_k), $f_k \in C_0^\infty(\Omega)$ for the L^2 norm have a limit in $C_0^\infty(\Omega)$. However, for any g in $C_0^\infty(\Omega)$, (f_k, g) is a Cauchy sequence of scalars because

$$\|(f_k, g) - (f_l, g)\| = \|(f_k - f_l, g)\| \leq |f_k - f_l|\,|g|\,, \tag{A.3}$$

and has a unique limit. Let us define

$$< f, g >= \lim_{k\to\infty} (f_k, g). \tag{A.4}$$

Clearly, $f : g \to < f, g >$ is linear in g. In addition

$$g_k \xrightarrow{D} g \Longrightarrow g_k \xrightarrow{L^2} g, \tag{A.5}$$

so f is a distribution on $C_0^\infty(\Omega)$.

One can show that two Cauchy sequences of functions of $C_0^\infty(\Omega)$ that are equivalent (i.e., $\|f_k - f'_k\| \to 0$) determine the same distribution

$$f = \lim_{k\to\infty} f_k = \lim_{k\to\infty} f'_k. \tag{A.6}$$

As a result, we can define

Definition A.1.1 *(Lebesgue Space) $L^2(\Omega)$ is the space of distributions defined on $C_0^\infty(\Omega)$ that are the limit (for the norm D') of the Cauchy sequences (for the norm L^2) of $C_0^\infty(\Omega)$. We define the scalar product in L^2 by*

$$(f, g) = \lim_{k\to\infty} (f_k, g_k), \tag{A.7}$$

where (f_k) and (g_k) are two Cauchy sequence of $C_0^\infty(\Omega)$ respectively determining f and g.

$L^2(\Omega)$ has the following properties

Theorem A.1.1 *$L^2(\Omega)$ is a complete space (and therefore a Hilbert space).*

Theorem A.1.2 $C^0(\Omega) \in L^2(\Omega)$

Proof: Suppose that $f \in C^0(\Omega)$. The open set Ω is compact, so the L^2 norm of f is

$$\|f\| = \sqrt{\int_\Omega |f(x)|\, dx} < \infty, \tag{A.8}$$

One can then show that any $C^0(\Omega)$ function satisfying Eq. A.8 determines a distribution f belonging to $L^2(\Omega)$.

A.2 Sobolev Spaces and Properties

The space $W^1(\Omega)$ is the space of square integrable distributions on Ω, which derivatives are also square integrable on Ω.

Definition A.2.1

$$W^1(\Omega) = \left\{ f \in L^2 \,\middle|\, \frac{\partial f}{\partial x} \in L^2, \frac{\partial f}{\partial y}, \in L^2 \right\}, \tag{A.9}$$

using the scalar product W^1

$$(f, g)_1 = (f, g)_0 + (\nabla f, \nabla g)_0 \,, \tag{A.10}$$

where $(.,.)_0$ *is the standard scalar product in* L^2*.*

In addition we define the Sobolev space $W_0^1(\Omega)$ as the limit (using D' norm) of the Cauchy sequences (using W^1 norm) of functions $C_0^\infty(\Omega)$ using the scalar product W^1.

Definition A.2.2

$$W_0^1(\Omega) = \overline{C_0^\infty(\Omega)}, \tag{A.11}$$

where the adherence uses the norm W^1*.*

We have the following results

Theorem A.2.1 $W^1(\Omega)$ *is a Hilbert space.*

Theorem A.2.2 $W_0^1(\Omega)$ *is a closed subset of* $W^1(\Omega)$*.*

Theorem A.2.3 *(Meyers and Serrin)*

$$W^1(\Omega) = \overline{\{f \in C^1(\Omega), \|f\|_1 < \infty\}}. \tag{A.12}$$

A.3 Uniqueness of Boundary Problems

Let us define the following properties of a linear operator A in a Hilbert space:

Definition A.3.1 *A is symmetric if the domain of A is dense and*

$$\forall f, g \in D(A) : (f, Ag) = (Af, g). \tag{A.13}$$

Definition A.3.2 *A is coercive if there exists* $\rho > 0$ *such that*

$$\forall f \in D(A) : (Af, f) \geq \|f\|^2. \tag{A.14}$$

We have the following theorem

Theorem A.3.1 *If the linear operator A is symmetric and coercive on a Hilbert space H, the boundary problem* $Af = g$*,* $g \in H$ *has a unique solution in H.*

A.4 Orthogonal Basis of Eigenfunctions

Since Δ is symmetric and coercive in $W^1(\Omega)$ and $W_0^1(\Omega)$ we have the following theorems about the eigenfunctions of this operator

Theorem A.4.1 *The problem*

$$\begin{array}{rcll} -\Delta f & = & \lambda f & \textit{in } \Omega \\ f & = & 0 & \textit{on } \partial\Omega, \end{array} \tag{A.15}$$

has a complete orthogonal sequence of eigenfunctions in $W_0^1(\Omega)$.

Theorem A.4.2 *The problem*

$$\begin{array}{rcll} -\Delta f & = & \lambda f & \textit{in } \Omega \\ \bar{n}.\nabla f & = & 0 & \textit{on } \partial\Omega, \end{array} \tag{A.16}$$

has a complete orthogonal sequence of eigenfunctions in $W^1(\Omega)$.

Appendix B

Software Development: MANGEN

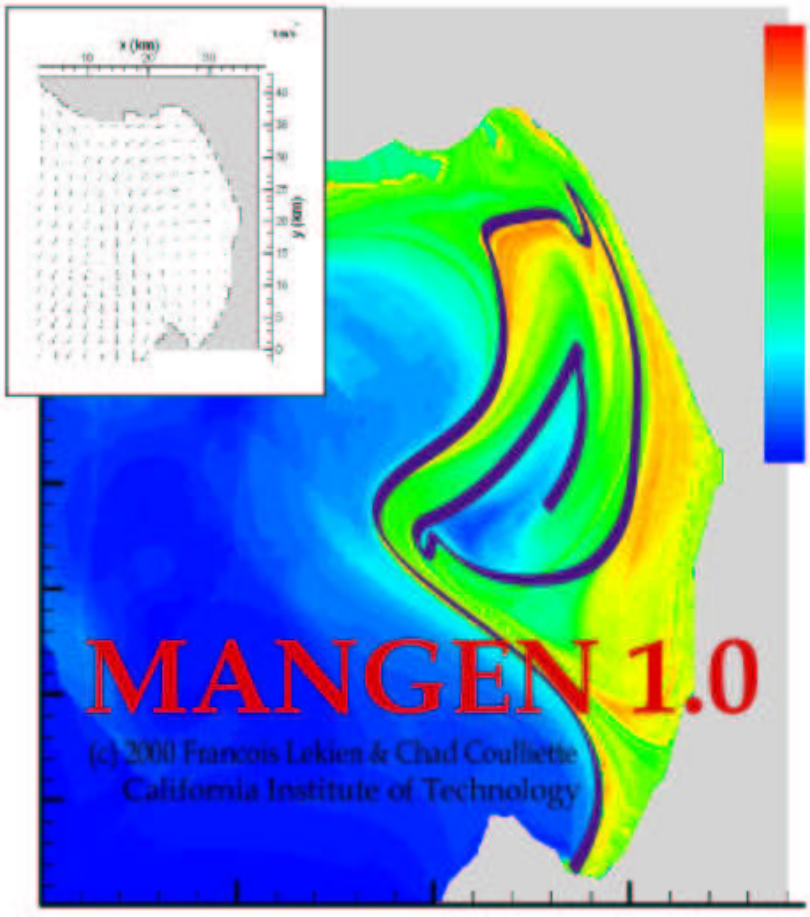

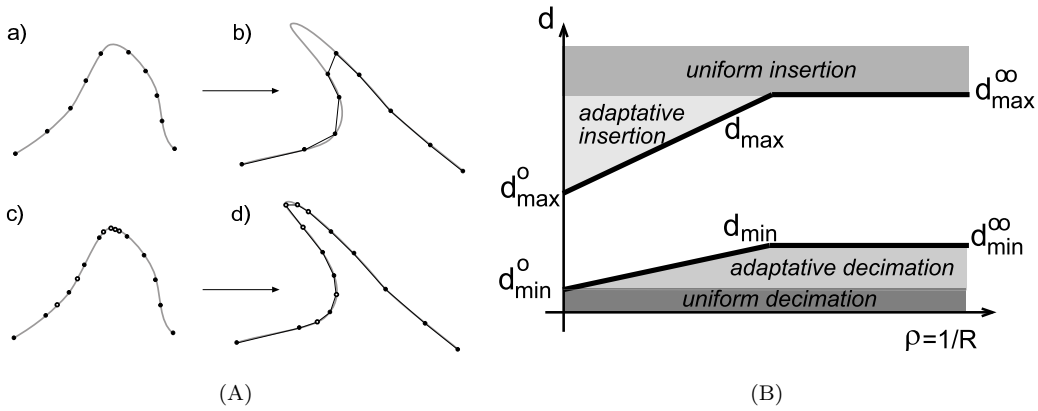

(A) (B)

Figure B.1: (A) More points are needed where the manifold has high curvature and less points are needed where the curvature is low. (B) *MANGEN* uses a robust and efficient algorithm to sample the curves. Points are inserted when the distance becomes larger than d_{max} and points are removed when the distance becomes smaller than d_{min}. Both d_{max} and d_{min} depend on the curvature of the curve.

B.1 Overview

While lobe dynamics has always been recognized as an *exact* transport theory and can theoretically give short- and long-term transport rates, computational issues limited its applications [Romkedar et al., 1990; Rom-Kedar and Wiggins, 1991]. The manifolds computed in such problems are typically windy. Furthermore, the length of these complicated curves grows quickly with the size of the time window of interest. The amount of points needed to describe long segments of manifolds can be prohibitively large if naive computational methods are used. This has forced many authors (for example, see Mackay et al. [1984, 1987] and Meiss and Ott [1986]) to incorporate a stochastic model (Markov chain). This approach ignores the fine structure of the lobes and manifolds.

MANGEN is able to compute very long segments of stable and unstable manifolds with high accuracy by *conditioning the manifolds adaptively*, for instance, by inserting more points along the manifold where the curvature is high, as depicted in Fig. B.1. As a result, the length and shape of the manifold is not an obstacle anymore and many more iterates of lobes than hitherto possible can be generated accurately.

It is important to notice that while the number of points used to represent the manifolds is large and the adaptive conditioning is a CPU expensive operation, it allows the computation of transport rates and escape times using the one-dimensional boundary of

sets of lobes. The same accuracy can only be reached in a brute force computation (using a two-dimensional grid of trajectories) if the number of points is so high that it covers all the details of the fine structures in phase space. It is our opinion that while the shape of the manifolds seemed to be an obstacle to the use of lobe dynamics during many years, carefully designed algorithms can overcome this problem and can give us an accurate picture of the manifolds. In such cases, lobe dynamics provides us with an exact transport rate with fewer points than a direct computation.

The algorithms detailed in the previous sections may be applied to analytical velocity fields fairly quickly. However, applying these techniques to real data sets requires more work. Large data files need to be converted into a binary format and special care has to be used not to overload the computer memory. *MANGEN* is precisely a library of such tools. In particular, data sets can be located on a distant computer and read through a TCP/IP connection.

In addition, *MANGEN* contains basic tools to interpolate between grid points, efficient and accurate integration methods and more specific algorithm to handle coastlines and open-boundaries.

B.2 Graphical User Interface

A major step in the distribution of MANGEN has been the development and release of a graphical user interface (GUI), that allows anyone to start working on his/her own project with minimal knowledge of the underlying algorithms (Fig. B.2, B.3 and B.4). The GUI allows many of MANGEN's functions to be controlled by the user with a mouse a few keystrokes. For example, to run MANGEN on the Monterey Bay data described in Chapter 4, the user opens the dialog of Fig. B.3 by clicking on the *Dynamical System* button of the *Simulation* → *Configure* menu. Filling this dialog box with the path to the data file, the type of data and the dimensions of the basin is enough for MANGEN to start the simulation. The *Simulation* → *Configure* menu contains other buttons to activate the different possible simulations. Several examples, including those presented in this report, have been included in the distribution and can be loaded using the *File* → *Open* menu. The user can load, modify, save and run any example or his/her own simulations.

Porting MANGEN to parallel supercomputers was an important step, that allowed us to work on problems with higher complexity and even more accuracy. However, the code for parallel applications is a lot more complex and it makes the code very difficult

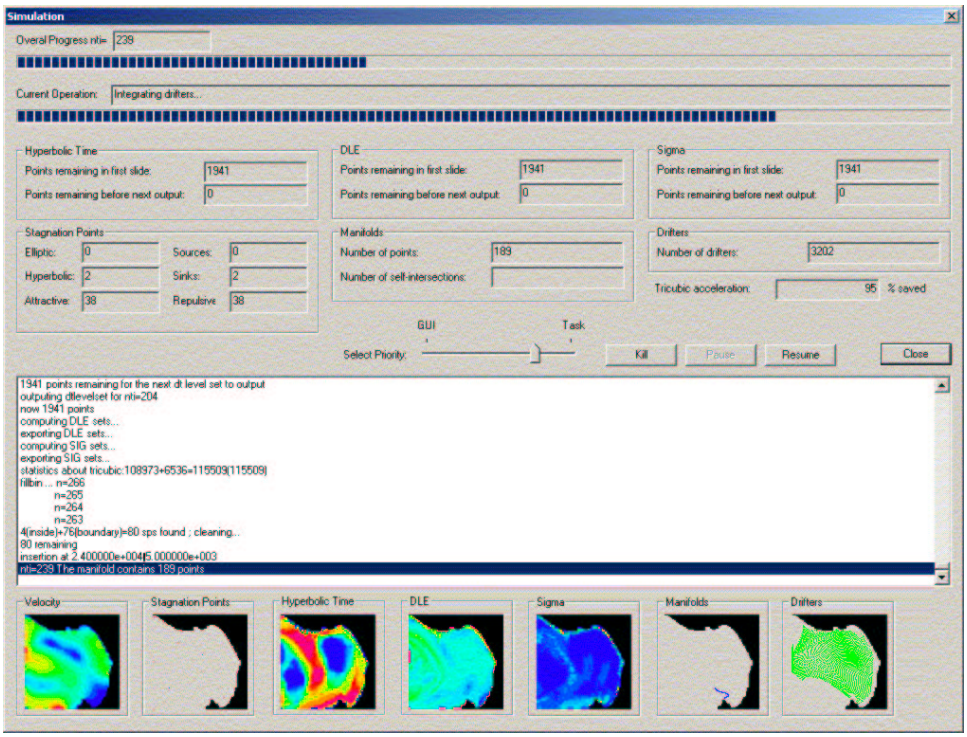

Figure B.2: Graphical User Interface (GUI) of MANGEN. Progress bars allow the user to follow the simulation and evaluate its execution time. Buttons are used to pause, stop or start the current simulation and the user has the ability to change the job priority interactively. The bottom of the GUI window contains a few panels that can be used to display intermediate results. In many cases, the use of these panels allowed us to stop a ill-conditioned simulation and saved us the time of waiting for the complete results.

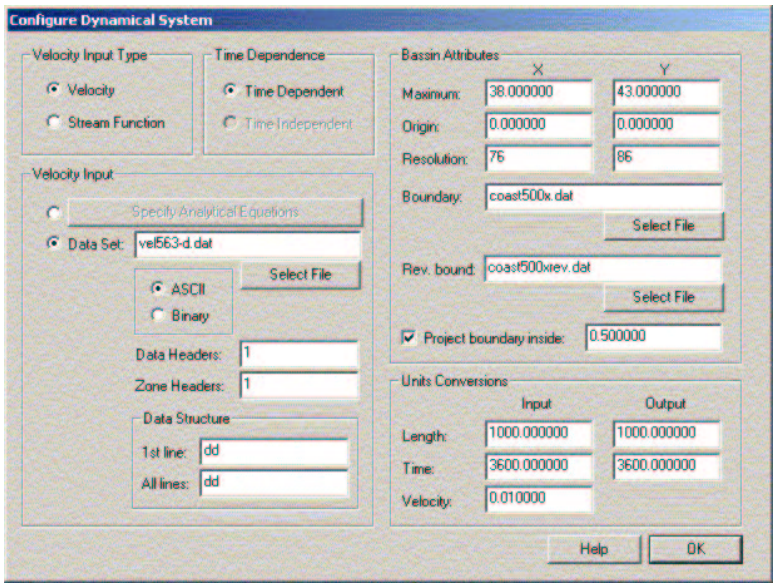

Figure B.3: The Data File Configuration Box is part of the GUI of MANGEN. It allows to work with a variety of ASCII and binary files as well as user defined analytical equations.

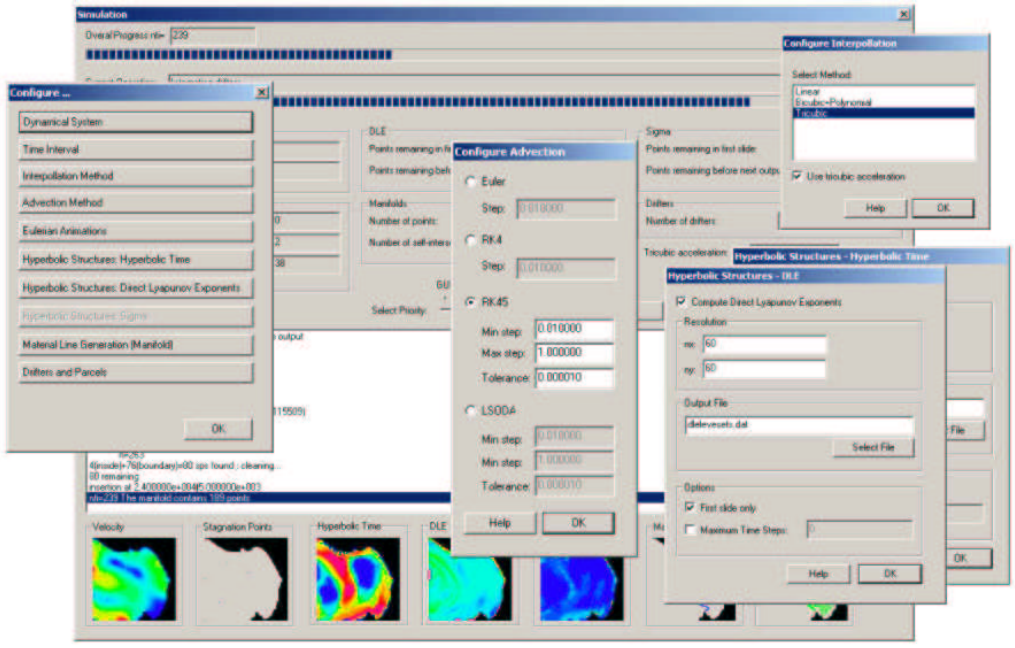

Figure B.4: A few dialog boxes of the MANGEN GUI that are used to configure a simulation.

to reproduce in a short amount of time. We are confident that providing toolboxes and a GUI will favor the understanding and the use of dynamical systems techniques, as well as promote research in this area.

For a novice in programming, MANGEN provides a friendly interface to all our methods and algorithms. For a programmer interested in developing new algorithms or criteria, MANGEN furnishes a toolbox that allows him to focus on his work only, letting MANGEN taking care of the data structure, the data files, interpolation, advection, boundaries and other complicated details of the computation. If a programmer decides to focus on single processor code, MANGEN will still work on parallel supercomputers. The new code will run in single processor mode, while the remaining of MANGEN will be executed at high speed, taking advantage of the multiple processors.

B.3 Availability and Installation

MANGEN version 1.2 has just been released for UNIX (tested on Linux alpha, Linux i386, Solaris and HPUX) and Windows workstations (tested on Windows NT, Windows 2000 and Windows XP). The code contains MPI support to run MANGEN on UNIX parallel supercomputers.

The new version of MANGEN fixes a few bugs from earlier versions and contains a fully automated installation method for each operating system. From the CD, MANGEN is installed on a Windows system using a standard Installation Wizard (Fig. B.5 and B.6) and on a UNIX machine by the commands:

```
cd /cdroot/source/UNIX/
configure
make UX
make
```

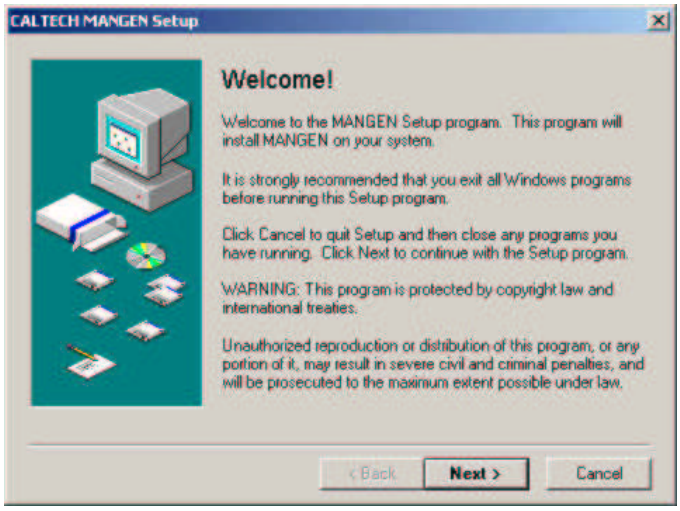

Figure B.5: Installation Wizard for the Windows versions of MANGEN.

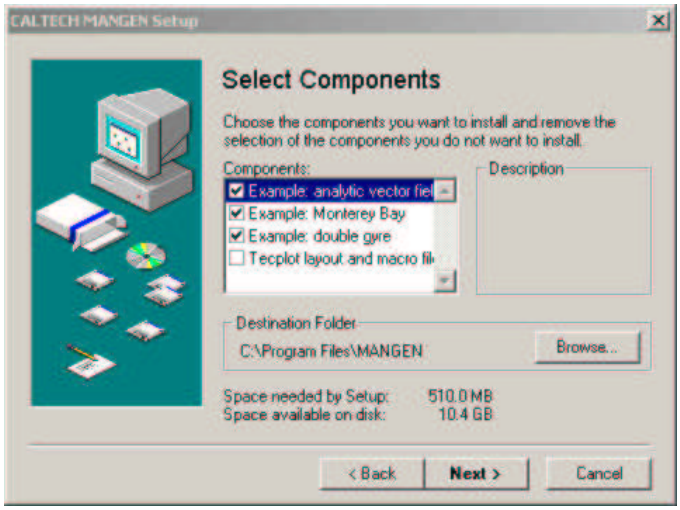

Figure B.6: Installation Wizard for the Windows versions of MANGEN (custom packages and choice of the examples).

Appendix C

Separatrices in High-Dimensional Phase Spaces: Application to Van der Waals Dissociation

This chapter is an attempt to reproduce transport and mixing studies in higher-dimensional phase spaces. The results explained in this chapter strongly rely on the fact that the map studied is autonomous and is a perturbation of a two-dimensional saddle combined with a two-dimensional oscillator. It should be regarded as a first step toward the analysis of arbitrary high-dimensional time-dependent systems.

C.1 The Model

We consider the model of Gillilan and Ezra [1991] corresponding to a free particle (He) plus free rotor (I_2), utilizing the zero impact parameter classical centrifugal sudden approximation, periodically kicked by the potential $V(r,\gamma)$

$$H = \frac{p_r^2}{2m} + \frac{p_\gamma^2}{2I} + TV(r,\gamma) \sum_{n=0,\pm1,\ldots} \delta(t - nT). \tag{C.1}$$

A four-dimensional symplectic map is obtained by integrating over one period T

$$\begin{aligned} p_r &\mapsto p_r' = p_r - T\tfrac{\partial V}{\partial r}, \\ p_\gamma &\mapsto p_\gamma' = p_\gamma - T\tfrac{\partial V}{\partial \gamma}, \\ r &\mapsto r + \tfrac{T}{m} p_r', \\ \gamma &\mapsto \gamma + \tfrac{T}{I} p_\gamma' \qquad \mod 2\pi. \end{aligned} \tag{C.2}$$

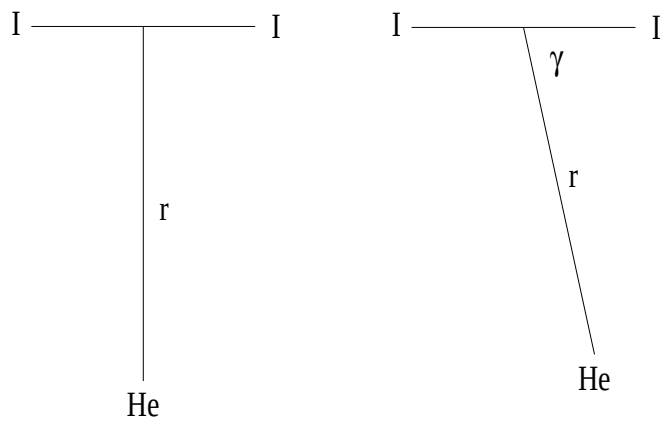

Figure C.1: Van Der Waals bond between He and I_2.

The map is periodic in the angle γ, as well as in the rotor angular momentum p_γ. The period Δp_γ is given by

$$\frac{T\Delta p_\gamma}{I} = \pi, \tag{C.3}$$

where, as we will see when we describe the potential in the next subsection, the absence of a factor 2 on the right-hand side of (C.3) is a result of the potential being periodic in γ with period π, rather than 2π.

The potential $V(r,\gamma)$ is the full three-dimensional potential evaluated with the I_2 bond length s fixed at its value at the kick ($V(r,\gamma) = V(r,\gamma,s=\bar{s})$, where $\bar{s}$ is the mean value of the oscillator s for example).

C.1.1 The Potential

The potential used by Gillilan and Ezra [1991], first proposed by Gaspard and Rice [1989], is given by

$$V(r,\gamma) = D\left((1+\alpha\cos 2\gamma)\, e^{-2\beta(r-r_e)} - 2e^{-\beta(r-r_e)}\right) \tag{C.4}$$

The potential parameters (D,α,β,r_e) are related to the well depth (W) of the potential at the T-shaped minimum, the value of r at the minimum (r_{min}), and the internal barrier height B as follows (B, $W > 0$)

$$\alpha = \frac{B/W}{2-(B/W)},$$

$$D = W\left(1-\alpha\right), \tag{C.5}$$

$$r_e = r_{min} - \left(\ln(1-\alpha)\right)/\beta.$$

A realistic set of parameter values is given in Table 1.

W (van der Waals well depth)	1.62×10^{-4}
B (height of barrier to internal rotation)	8.72×10^{-5}
r_{min} ($He-I$ distance at potential minimum)	7.0
β (Morse parameter)	0.6033
T (I_2 vibrational period for $v=0$)	10,986
T (I_2 vibrational period for $v=10$)	12,485
T (I_2 vibrational period for $v=20$)	14,717
T (I_2 vibrational period for $v=30$)	17,859
T (I_2 vibrational period for $v=40$)	31,502
s (equilibrium I_2 bond length)	5.717
m (reduced mass of $He-I_2$ system)	7,069.4
I (moment of inertia of I_2)	3,782,220.8

Table 1. Parameters for the $He-I_2$ Hamiltonian. All quantities are given in atomic units.

C.1.2 The Two-Dimensional Subsystem: T-Shaped $He-I_2$

If we fixe $\gamma = \frac{\pi}{2}$, $p_\gamma = 0$ in Eq. C.2, we find that these values of γ and p_γ do not change under iteration. This gives rise to a two-dimensional invariant subsystem of the four-dimensional map given by Eq. C.2

$$\begin{aligned} p_r &\mapsto p_r' = p_r - T\frac{\partial V}{\partial r}, \\ r &\mapsto r + \frac{T}{m}p_r', \end{aligned} \tag{C.6}$$

The resulting molecule is referred to as T-shaped $He-I_2$. The two-dimensional map (C.6) has a saddle type fixed point at $p_r = 0$, $r = \infty$ with separatrices.

C.1.3 Phase Space Structure of the Four-Dimensional Map

The potential (C.4) has the useful property that as $r \to \infty$ the asymptotic motion in the $r - p_r$ coordinates is independent of γ. This implies that $p_r = 0$, $r = \infty$ defines a two-dimensional invariant set parameterized by $\gamma - p_\gamma$ (which, since they are both angular variables, parameterize a two-torus). This saddle type two-dimensional invariant set has three-dimensional separatrices.

C.2 Fragmentation Dynamics for T-Shaped $He - I_2$

C.2.1 Lobe Dynamics and Decay Rates

To take advantage of the periodical kicks approximation, we worked with the Poincaré section

$$\mathcal{P}_T \ : \ x \to \mathcal{P}_T(x,t) = (x(t+T,t,x), t+T), \tag{C.7}$$

where T is the vibrational period. Its explicit form is

$$\mathcal{P}_T \ : \ \begin{pmatrix} r \\ p_r \end{pmatrix} \to \mathcal{P}_T \begin{pmatrix} r \\ p_r \end{pmatrix} = \begin{pmatrix} r + \frac{T}{m}(p_r - T(\frac{\partial V}{\partial r})) \\ p_r - T(\frac{\partial V}{\partial r}) \end{pmatrix}. \tag{C.8}$$

In the two-dimensional phase space (r, p_r), the Poincaré map has a hyperbolic fixed point at $(+\infty, 0)$. A stable invariant manifold and an unstable invariant manifold are attached to this saddle and are shown on Fig. C.2. In the range of frequencies that we studied, those two manifolds always have an infinite number of transverse heteroclinic intersections. As described in Gillilan and Ezra [1991], dissociation can also be studied with a sample set of points initially inside a separatrix. The fragmentation corresponds to the constant flux of particles crossing this separatrix. We chose a convenient primary intersection point p_0 and defined a separatrix as the union of the segments of the stable and unstable manifolds going from the fixed point to p_0 (Fig. C.2). With such a definition, the separatrix is a surface of non-return. Particles that escape from it never re-enter afterwards. This is a fundamental result that will not persist in higher dimensions.
We remark that a few KAM tori exist around the non-hyperbolic fixed point located inside the separatrix. Particles located in those tori are trapped forever and cannot dissociate. On the other hand, those regions can never be accessed from outside, so from a realistic point of view, we want to remove the KAM region from the separatrix area.

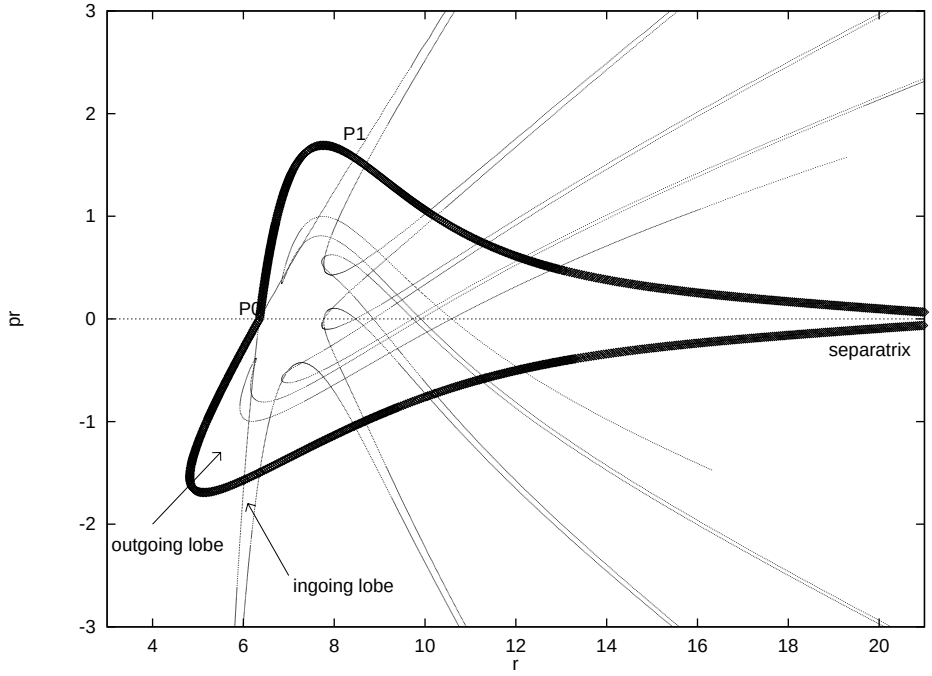

Figure C.2: Two-dimensional phase space structure for $\nu = 0$ ($T = 10986$) and associated separatrix.

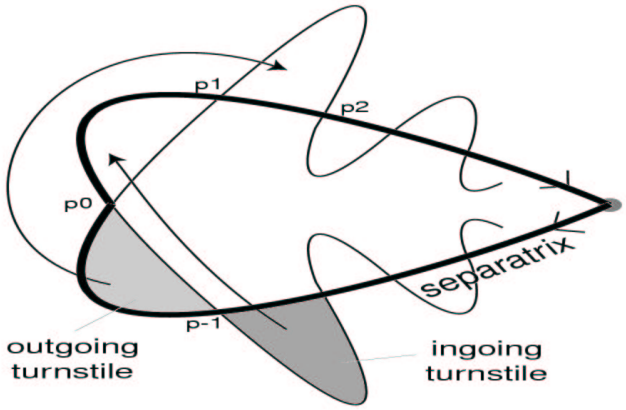

Figure C.3: Sketch of the structure of the phase space for the two-dimensional approximation.

C.2.2 Results

The fragmentation rate could be computed using a sample set of particles initially inside the separatrix. In those terms, flux across the separatrix is equivalent to the number of particle dissociating. However, transport across the separatrix only occurs through an *outgoing* and an *ingoing lobe* depicted on Fig. C.2. We can then compute the number of particles that escape the separatrix at each iteration using the area of the outgoing lobe. As time evolves, successive mappings of the empty ingoing lobe lead to more and more empty inner lobes inside the separatrix. Fig. C.5 pictures several steps of the dynamics and some of the empty lobes have one or more intersections with the outgoing lobe. Those regions contain no particles and have to be removed from the area of the outgoing lobe to have the exact decay. This process (Fig. C.3 and C.3) is extensively described in [Wiggins, 1992] and leads to a general expression of the flux during mapping i

$$F_i = A(L_{out}) - \sum_{j=1}^{i} A(L_{out} \cap L_{in}^j), \tag{C.9}$$

where $A(L_{out})$ is the area of the outgoing lobe and $A(L_{out} \cap L_{in}^j)$ is the area of the intersection between the outgoing lobe and the ith mapping of the ingoing lobe. This last formula gives a result identical to a direct simulation that would count particles leaving the separatrix. The major difference is that we compute the fragmentation using a one-dimensional segment of the manifold instead of two-dimensional sets of particles. The

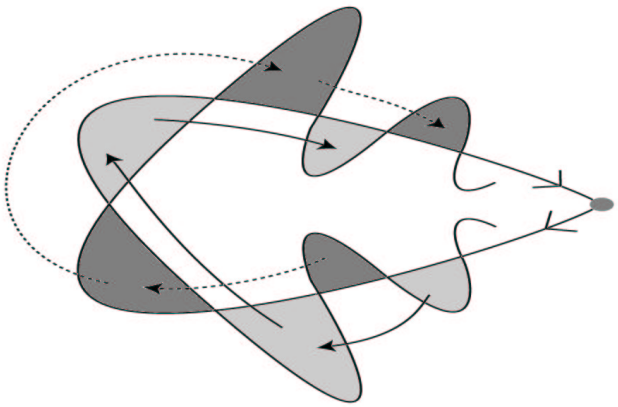

Figure C.4: Sketch of the lobe dynamics associated with the two-dimensional phase space.

computational effort is thus reduced by a few orders of magnitude.

The periods T used for each value of the quantum number ν of the oscillator are discussed in Gillilan and Ezra [1991]. The decay has been defined as the number of remaining particles divided by the number of initial particles inside the separatrix, that are subject to dissociation, i.e., those that do not belong to a KAM tori.

C.3 Fragmentation Dynamics for the Full Four-Dimensional Map

C.3.1 Theoretical Aspects of High-Dimensional Separatrices

No matter what the dimension of the space is, we are always looking for hypersurfaces of codimension 1. Those are the only sets that are likely to be dividing surfaces. Since the fixed points do not always have codimension 1 invariant manifolds attached to them, we may have to look at the invariant manifolds of structures of higher dimension, like the stable and unstable manifold of a fixed surface of codimension 2.

While the intersection between two invariant curves must be either empty, tangent or an infinite set of transverse intersections, the interaction between two hypersurfaces is usually more subtle. We will only discuss the generic cases and assume that the intersection between two hypersurfaces of codimension 1 is the union of sets of codimension 2. This excludes the possibility of a point of tangent intersection between the two hypersurfaces.

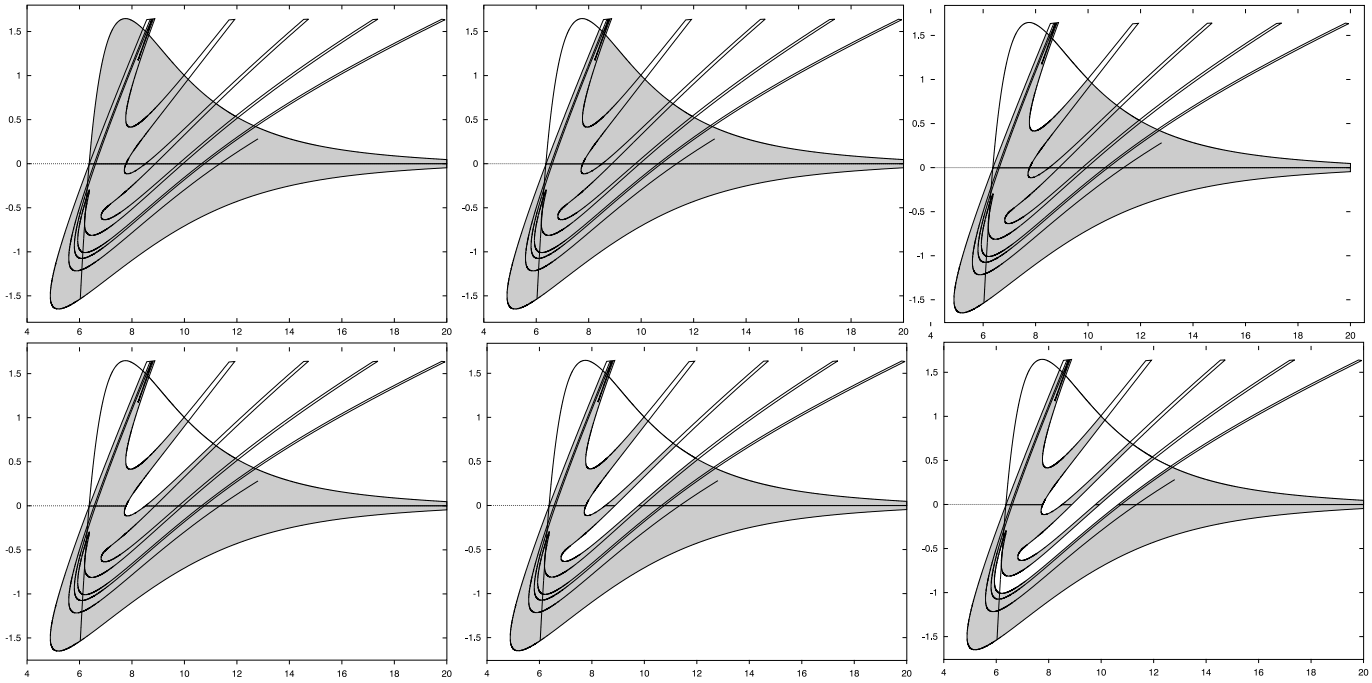

Figure C.5: Partially filled outgoing lobe

C.3.2 Lobe Dynamics and Decay Rates

Fig. C.7 and C.8 show two different three-dimensional sections of the manifolds in the four-dimensional space. Associated with each of these section is a sliced separatrix, that may be continuous or discontinuous.

C.3.3 Results

In order to give a clear representation of the method used to define the separatrix, we chose a point that should obviously be inside, so we can associate each intersection point with an angular position α (Fig. C.9).

The use of α is interesting because, as soon as we are only interested in intersection points, it reduces the two-dimensional (r, p_r)-subspace to a one-dimensional α-space. When two intersection points are mixing, their α value must be the same. The reverse proposition is however not true due to the presence of secondary intersection points at the edge of horseshoes. We usually ignore those points to obtain a simpler α-representation. For small p_γ, we can follow a single (dashed) line to turn around the torus. This curve is then sufficient to support a separatrix in the involved slices. Fig. C.10 shows slices of the

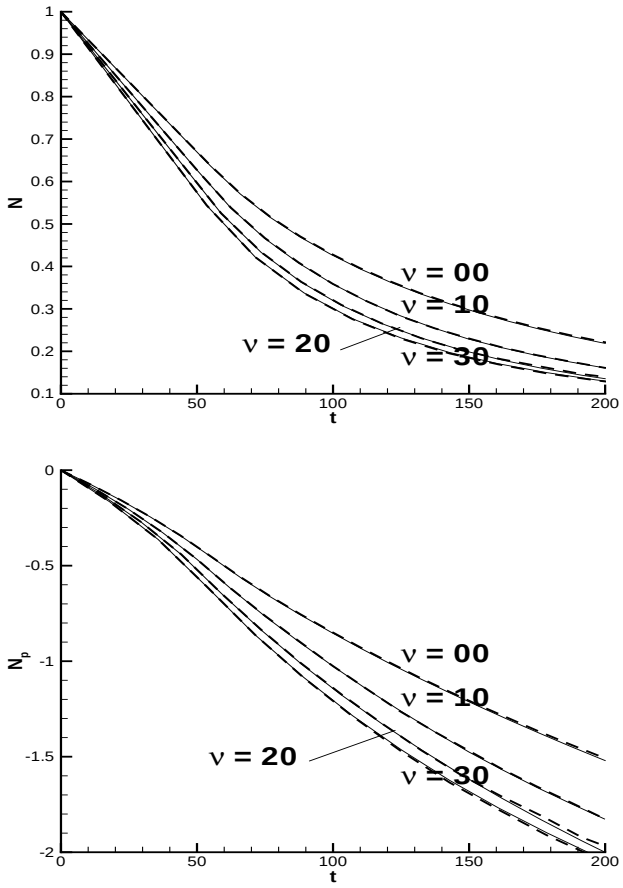

Figure C.6: Fragmentation rate for the two-dimensional approximation

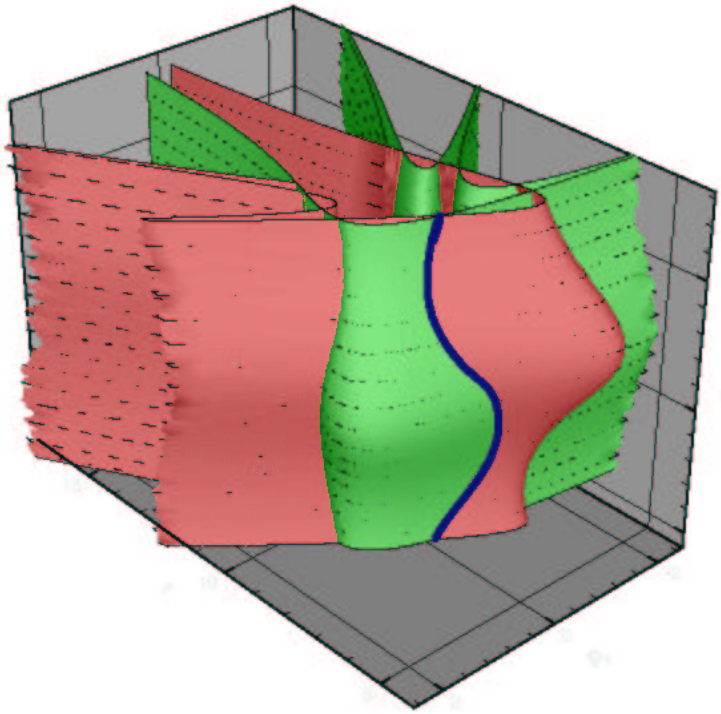

Figure C.7: Section ($\gamma = 0$) of the four-dimensional phase space structure for $\nu = 0$ and associated sliced separatrix.

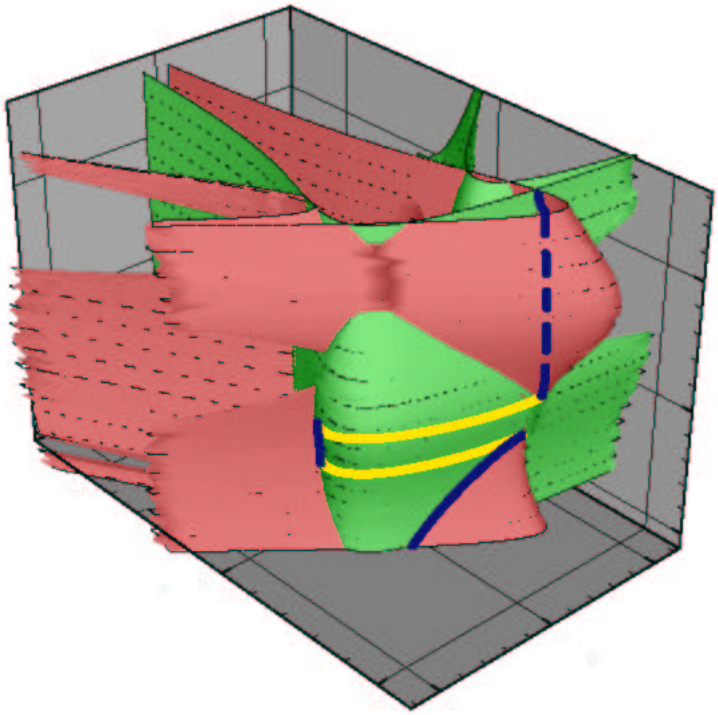

Figure C.8: Section ($\gamma = \frac{\pi}{2}$) of the four-dimensional phase space structure for $\nu = 0$ and associated sliced separatrix.

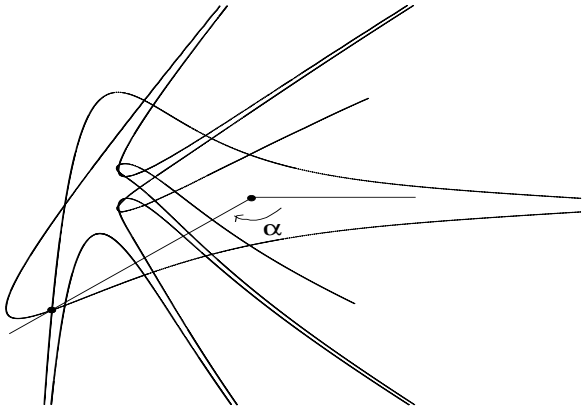

Figure C.9: Definition of α.

resulting separatrix. While increasing p_γ the curve will suffer a first bifurcation resulting in the mixing of the different branches of the manifold. The new branches spiral around the torus and never close. According to Beigie [1995], we introduced a discontinuity in the definition of the separatrix. The arrow indicates the point that we traced to find one of the branches of the manifold. From this point, we can go left and right until the curve makes a U-turn. A this point we switch to another branch and so on. It does not really matter whether or not we jump up or down, but the resulting curve must loop around the torus. Once the reactive curve has been successively defined for every value of p_γ, we have the intersection of a separatrix with every slice. This will allow us to compute the decay in the next section. Fig. C.11 shows a representation of the chosen part of the intersection manifold. It is actually a two-dimensional surface embedded in the four-dimensional phase space but the use of α allows us to show it in a three-d space. One can easily see the main surface and the holes inside due to its folding. The pieces on top and below complete the surface, that we want to use to define the three-dimensional separatrix. It is worths to remark that this surface has no physical meaning. It is only a subset of the intersection manifold that is suitable to define the separatrix.

Knowing the intersection of the separatrix with all constant γ and p_γ planes is sufficient to state whether or not a point is inside the separatrix. All one has to do is to find the nearest slice, project the point in this plane and check if the point is inside this 1-dimensional sliced separatrix. We are then able to compute the decay using a sample set of initial particles evolving in time.

Since the decay is relative to transport across the separatrix, the results are strongly dependent of the accuracy of the separatrix. Roughly speaking, the resulting decay is only interesting if the chosen separatrix is effectively a surface of non-return. For this purpose, we tracked the points to know how many of them could escape the separatrix and then re-enter it. Fig. C.12 shows the cumulative number of such points divided by the separatrix area, versus the r coordinate of the starting point. One can see that around $r = 9$, the number of points increases slowly. It corresponds to the critical zone where the stable and unstable manifolds join. The number of points then stabilizes when we leave this zone. There seems to be less and less *wrong* points. A huge step however appears around $r = 20$ due to the deformation of the tail of the separatrix around the fixed plane that should be at infinity. Since we are mostly interested in determining the accuracy of our separatrix with respect to the particular choice of pieces of intersection manifold, we also made a similar computation that ignored transport after $r = 20$. This number,

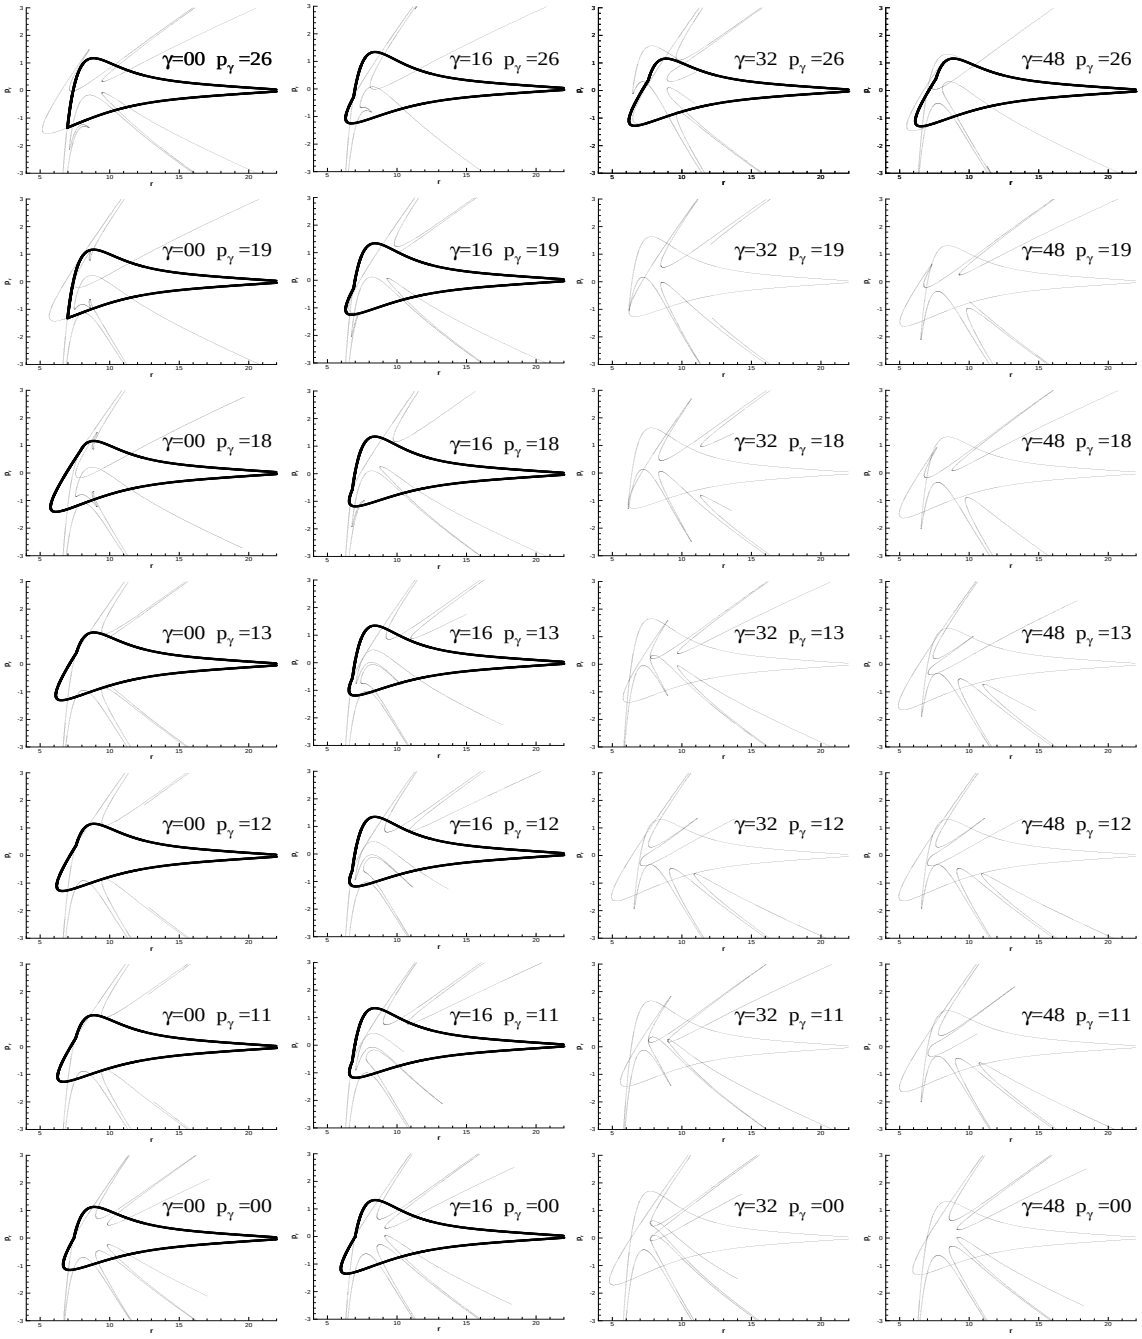

Figure C.10: Sliced manifolds.

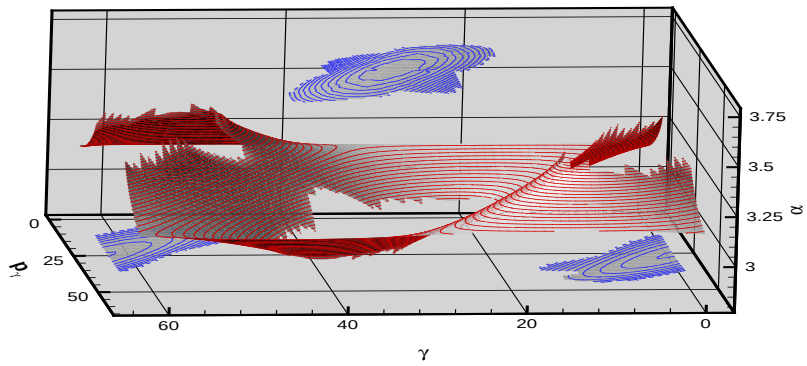

Figure C.11: Three-dimensional view of the piece of intersection manifold used to define the separatrix.

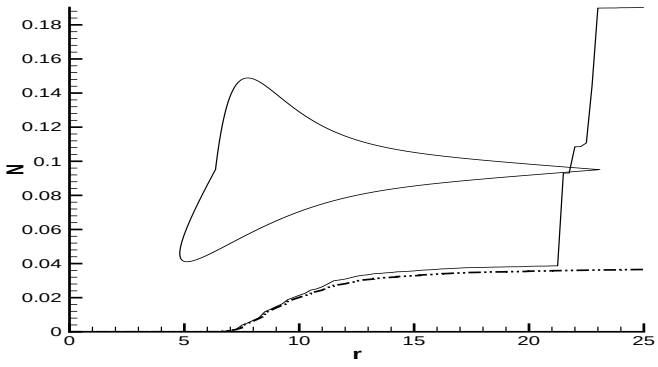

Figure C.12: Accuracy of the separatrix - The heaviest curve only accounts for the nose effect.

Table C.1: Separatrix accuracy - Percent of points returning inside the separatrix

	$four-dimensional$	$0th\ order$
nose only	0.038	2.225
total	0.190	2.700

(C.10)

referred as *nose accuracy*, best ranks the choice of a separatrix and the corrected curve is printed heavier on Fig. C.12. As given in *table C.1*, one can easily see that the biggest part of the error comes from the tail. This entails that we effectively have a surface of non-return. For comparison, we made the same simulations for a 0th order separatrix, built by the repetition of the two-dimensional separatrix for every slice of the double-torus. The number of illegal points is there more than 50 times higher.

C.3.4 Comparison of Decay Rates for the Two- and Four-Dimensional Map

The number of particles inside the separatrix is shown on Fig. C.13, as well as the same quantity for the two-dimensional approximation. We note that the discontinuity of the separatrix allows us to divide the flux Φ_{4D} in two contributions

$$\Phi_{4D} = \Phi_i + \Phi_r, \tag{C.11}$$

where Φ_i is the irreversible flux resulting from the lobe dynamics and Φ_r represents the exchange of particle through the walls that have been added to close the separatrix. In our case, the reversible flux was small with respect to the irreversible flux. Nevertheless, it creates a difference in the rate of dissociation. The flux for the two-dimensional approximation is completely irreversible, because the separatrix is a true barrier. At the beginning of the experiment, all the particles are seeded inside the separatrix and most of the particles will stay inside during the first iteration of the map. As a result, there are no or few particles outside the separatrix that can cross the wall toward the interior and the reversible flux will be mainly oriented toward the domain ($\Phi_r < 0$). This implies that $\Phi_{4D} < \Phi_{2D}$ during this phase.

After a few iterations, most of the molecules have dissociated and the population is higher outside the domain. As a result the reversible flux will be mainly in favor of

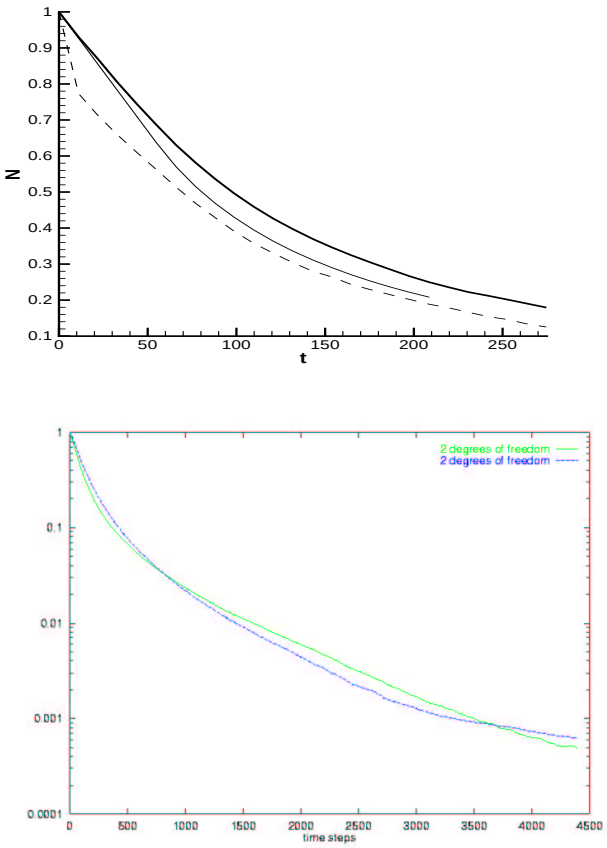

Figure C.13: Decay for the four-dimensional approximation (heaviest curve), the two-dimensional approximation (lighter) and the 0^{th}-order separatrix simulation (dashed curve)

dissociation and

$$\Phi_r > 0 \Longrightarrow \Phi_{4D} > \Phi_{2D}. \tag{C.12}$$

One can see on Fig. C.13 that this is indeed the case and that eventually the number of particles inside the four-dimensional separatrix becomes larger that the population in the two-dimensional separatrix because of the smaller flux. After a large number of iterations, the number of particles in both approximations becomes so small that the two curves seem to intersect and become identical.

Appendix D

Transport in the Restricted Three-Body Problem

In collaboration with Wang-Sang Koon, Shane Ross and Jerry Marsden.

D.1 Motivation

Our understanding of the solar system has changed dramatically in the past several decades with the realization that the orbits of the planets and some minor bodies are chaotic. In the case of planets, this chaos is of a sufficiently weak nature that their motion appears quite regular on relatively short timescales [Laskar, 1989]. In contrast, small bodies such as asteroids, comets, and Kuiper belt objects can exhibit strongly chaotic motion through their interactions with the planets and the sun, exhibiting Lyapunov times of only a few decades [Torbett and Smoluchowski, 1990; Tancredi, 1995].

The recent discovery of several extrasolar planetary systems has stimulated interest in the morphological and dynamical features that may be present in *generic* planetary systems. Interesting quantities to compute are: expected distributions of objects in the presence of dynamical sculpting due to planets and moons (e.g., generic circumstellar belts and circumsolar rings); rates of small body collision with a planet; and rates of capture and escape from one orbital resonance with a planet to another.

Since most important transport questions involve motion between different regions of the phase space, there have been attempts in the dynamical astronomy community to consider the entire global phase space picture of all mean motion resonances. Malhotra [1996] provides a good recent example. Motivated by numerical studies of the stability of

low-eccentricity and low-inclination orbits of small bodies in the trans-Neptunian Kuiper belt, Malhotra [1996] used the PCR3BP (Restricted 3-Body Problem) to describe the basic phase space structure in the neighborhood of Neptune's exterior mean motion resonances. This model provides a "baseline" and a reasonable explanation for much of the dynamical behaviors found in the large scale numerical studies [Levison and Duncan, 1993].

D.2 Introduction to the PC3RBP

In Malhotra [1996], Malhotra provides a direct visualization, in a two-dimensional surfaces-of-section, of a global *mixed phase space* structure of stable and chaotic zones, as seen in Fig. D.1(a).

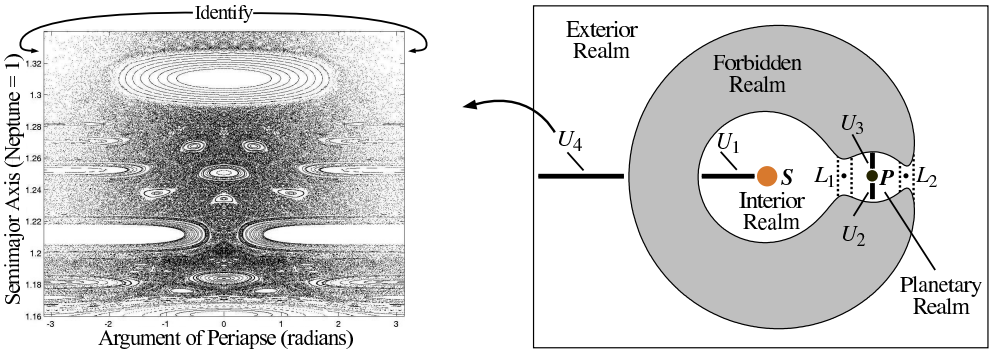

(a) (b)

Figure D.1: Poincaré section of the flow in the restricted three-body problem. (a) The mixed phase space structure of the planar circular restricted three-body problem (PCR3BP) is shown on this Poincaré section. KAM tori and the chaotic sea are visible. The x-axis is an angular coordinate and thus the right and left sides are to be identified. (b) The *interior, exterior* and *planetary realms* are connected by bottlenecks about L_1 and L_2 and particles can pass between realms only through these bottlenecks. The star and planet, denoted S and P, respectively, are fixed in this rotating frame. The location of four Poincaré sections that we use for our analysis of the PCR3BP, U_1, U_2, U_3, and U_4 are shown in the configuration space.

Recent work [Koon et al., 2000] applies dynamical systems techniques to the problem of heteroclinic connections and interior-exterior transitions in the PCR3BP, laying the foundation for *tube dynamics.* The point of view developed in Koon et al. [2000] is that the invariant manifold structures associated to periodic orbits about L_1 and L_2, the (Conley-

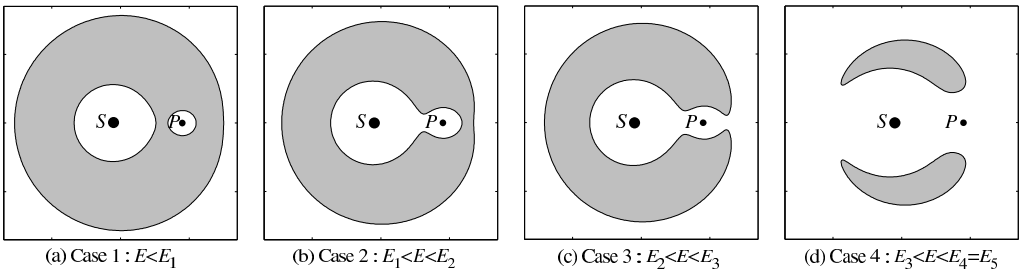

Figure D.2: There are five cases of allowable motion. The star and planet, denoted S and P, respectively, are fixed in this rotating frame. (a) In case 1, the particle is trapped either exterior or interior to the planet's orbit, or around the planet itself. It is energetically prohibited from crossing the *forbidden realm,* shown in gray. (b)-(d) As the energy E of the particle increases, the bottlenecks connecting the realms open. In case 5, not shown, the entire configuration space is energetically accessible.

McGehee) *phase space tubes* [Conley, 1968; Mcgehee and Meyer, 1974], are fundamental tools that can aid in understanding transport throughout the solar system as well as transport between the inside and outside of a planet's orbit and other phenomena. The *interior, exterior* and *planetary realms* are connected by bottlenecks about L_1 and L_2, as depicted in Fig. D.1(b). Particles can pass between realms only through these bottlenecks by being inside phase space tubes, regions bounded by pieces of the stable and unstable invariant manifolds of periodic orbits around L_1 and L_2. We can determine the flux between realms by monitoring the flux through these tubes. Under a Markov assumption, the rate of escape of particles temporarily captured by Mars can be computed [Jaffe et al., 2002]. Theory and direct numerical simulation are shown to agree to within 1%, showing the promise of a statistical theory for the computation of useful quantities in dynamical astronomy.

We suggest an alternative approach to computing transport quantities that will complement large scale numerical simulations by providing insight into the underlying dynamics.

D.3 Transport in the PCR3BP

The value of the energy is an indicator of the type of global dynamics possible for a particle in the PCR3BP, which can be broken down into five cases (see Fig. D.2). In case 1, shown in Fig. D.2(a), the particle is trapped either exterior or interior to the planet's orbit, or around the planet itself. For energy values greater than that of L_2 (case 3), there is a

bottleneck around L_1 and L_2, permitting particles to move between the three realms.

The study of transport within a particular realm of the PCR3BP can be reduced to the study of its associated Poincaré section. For purposes of our description, we will consider the Poincaré section in the exterior realm, located as shown in Fig. D.1(b) which, for example, would be a model for a scattered Kuiper belt object under the influence of the sun and Neptune.

Following Wiggins [1992], lobe dynamics theory states that the two-dimensional phase space M of the Poincaré map f can be divided into disjoint *regions* of interest

In our application, the Poincaré map of the PCR3BP within a particular realm possesses *resonance bands* consisting of alternating unstable and stable periodic points. For instance, suppose we wish to study the transport in the exterior realm, the Poincaré section of which is shown in Fig. D.1(a). In Fig. D.3, we show the invariant manifolds associated to a particular hyperbolic fixed point superimposed upon the top part of Fig. D.1(a). The two branches of the unstable manifolds are shown in red and the two branches of the stable manifolds are shown in green.

Suppose we consider a single resonance in the exterior realm, and we want to consider the phenomena of crossing a resonance and capture into the resonance using lobe dynamics. We can define the *resonance region*, and then obtain a distribution of times for particles to cross from one side of the resonance to the other and residence times within the resonance region. The manifolds shown in Fig. D.3, as well as the associated *primary intersection points* (pips), form the *lobes* that will determine the transport between three regions: R_2, the resonance region; R_1 and R_3, the regions below and above the resonance region, respectively. In Fig. D.4, the lobes for the upper and lower boundaries are shown. These boundaries each show a sequence of four lobes. To the best of our knowledge, this is the first example of a physical system that has a sequence of four lobes. Using the lobe dynamics framework, the transport of species between these regions.

The *MANGEN* software package has been used to run some preliminary simulations on the PCR3BP problem and we were able to easily detect dozens of lobes, as shown in Fig. D.4. We computed the area of each lobe in the area preserving Poincaré section to determine the quality of the simulation. Based on the preliminary results, we are confident that we can compute very long sequences of lobes with sufficient accuracy to compute transport rates and escape times across several resonances.

It is important to notice that while the number of points used to represent the manifolds is large and the adaptive conditioning is a CPU expensive operation, it allows the

computation of transport rates and escape times using the one-dimensional boundary of sets of lobes. The same accuracy can only be reached in a brute force computation (using a two-dimensional grid of trajectories) if the number of points is so high that it covers all the details of the fine structures in phase space. It is our opinion that while the shape of the manifolds seemed to be an obstacle to the use of lobe dynamics during many years, carefully designed algorithms can overcome this problem and can give us an accurate picture of the manifolds. In such cases, lobe dynamics provides us with an exact transport rate with fewer points than a direct computation.

D.4 Perspectives

The resulting global theory will allow us to tackle solar system transport problems. Dynamical systems theory has been used to give a qualitative description of the topological features of the PCR3BP phase space, such as resonance bands, chaotic layers, and KAM

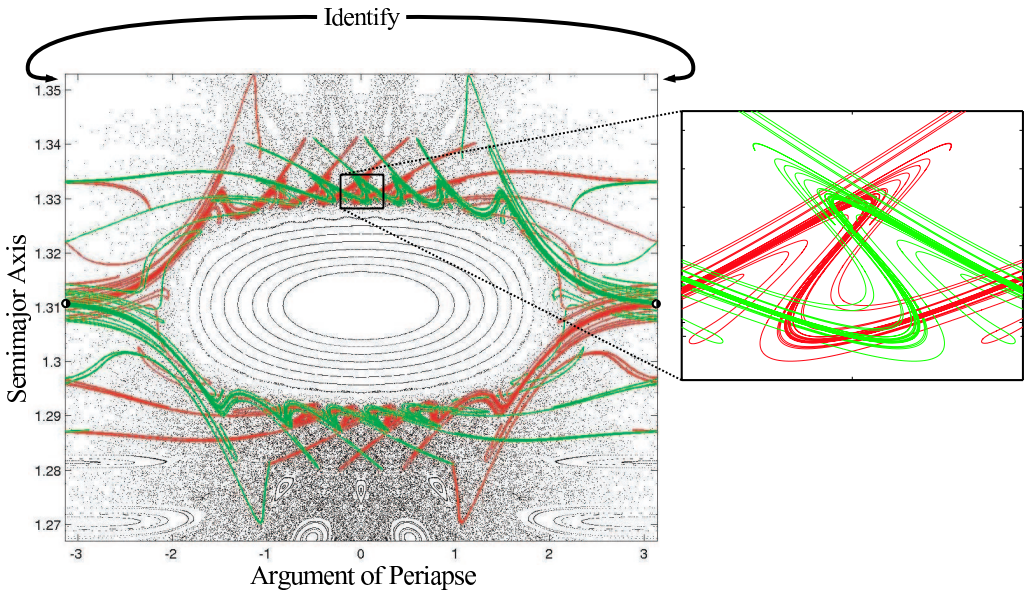

Figure D.3: Transport in a mixed phase space using lobe dynamics. We superimpose on the upper part of Fig. D.1(a) the dynamical structure that is "unobservable" in Malhotra [1996], but is crucial for applying lobe dynamics theory: hyperbolic fixed points (the half-filled circles, identified as the same point), their stable and unstable manifolds, pips, partial barriers and turnstiles, etc. An enlargement of the boxed region is shown at right.

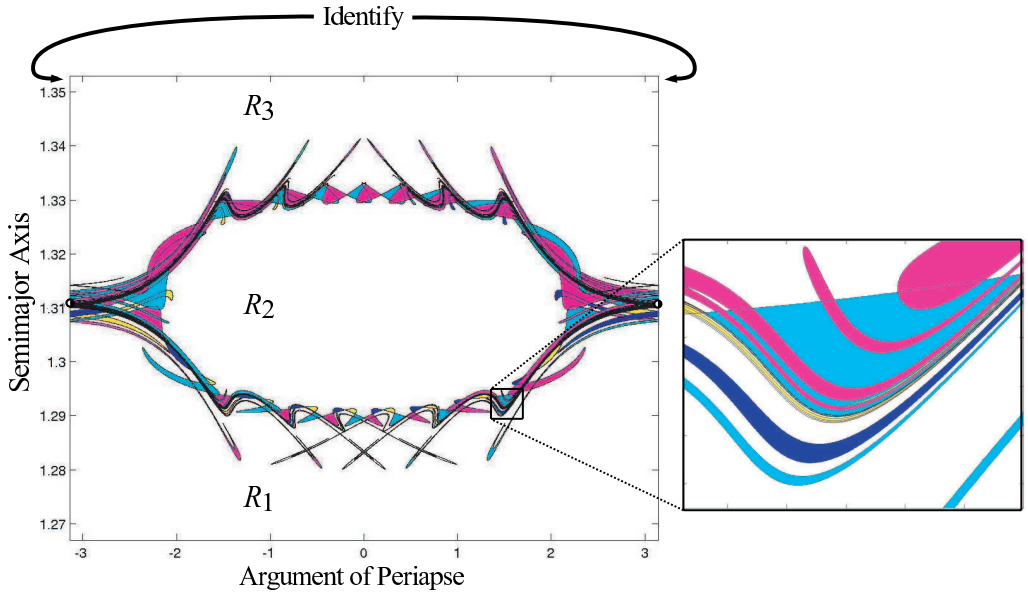

Figure D.4: Transport across a resonance. The lobes formed by both the upper and lower boundaries of a resonance region R_2 are shown. Using the lobe dynamics framework, the transport of species between these regions can be computed systematically. An enlargement of the boxed region is shown at right.

tori (Malhotra [1996] and references therein). But we plan to make this analysis more quantitative by computing statistical measures such as *residence times within regions of phase space* and *transport rates between various regions and across different realms*.

Since the CR3BP is such a fundamental model in studying small body orbital dynamics, our work will shed light on several interesting and important problems, such as those mentioned in the introduction. As mentioned previously, the techniques developed to study transport in the solar system may be applied to other areas of study, including fluid dynamics and physical chemistry.

Bibliography

Aref, H. [1984], Stirring by Chaotic Advection, *J. Fluid Mech.* **143**, 1–21.

Aref, H. and M. S. El Naschie [1994], Chaos Applied to Fluid Mixing, *Chaos, Solitons, and Fractals* **4(6)**, 1–380.

Babiano, A., A. Provenzale, and A. Vulpiani [1994], Chaotic Advection, Tracer Dynamics, and Turbulent Dispersion. Proceedings of the NATO Advanced Research Workshop and EGS Topical Workshop on Chaotic Advection, Conference Center Sereno di Gavo, Italy, 24-28 May 1993, *Physica D* **76**, 1–329.

Beerens, S. P., H. Ridderinkhof, and J. T. F. Zimmerman [1998], Longitudinal Dispersion in Natural Streams, *Chaos, Solitons and Fractals* **4**, 1011–1029.

Beigie, D. [1995], Multiple Separatrix Crossing in Multi-Degree-Of- Freedom Hamiltonian Flows, *Journal of Nonlinear Science* **5**, 57–103.

Bikangaga, J. H. and V. Nussehi [1995], Application of Computer Modeling Techniques to the Determination of Optimal Effluent Discharge Policies in Tidal Water Systems, *Water Res* **29**, 2367–2375.

Britz, D. and R. Antonia [1987], A Computer Algorithm for the Identification of Temperature Fronts in a Turbulent Shear-Flow, *Experiments in Fluids* **5**, 134–140.

Conley, C. [1968], Low Energy Transit Orbits in Restricted 3-Body Problem, *Siam Journal on Applied Mathematics* **16**, 732.

Coulliette, C. and S. Wiggins [2000], Intergyre Transport in a Wind-Driven, Quasi-Geostrophic Double Gyre: An Application of Lobe Dynamics, *Nonlinear Processes in Geophysics* **7**, 59–85.

Coulliette, C. and S. Wiggins [2001], Intergyre Transport in a Wind-Driven, Quasi-Geostrophic Double Gyre: An Application of Lobe Dynamics, *Nonlinear Proc. Geophys* **8**, 69–94.

del Castillo-Negrete, D. and P. J. Morrison [1993], Chaotic Transport of Rossby Waves in Shear Flow, *Phys. Fluids A* **5(4)**, 948–965.

Dijkstra, H. A. and C. A. Katsman [1997], Temporal Variability of the Wind-Driven Quasi-Geostrophic Double Gyre Ocean Circulation: Basic Bifurcation Diagram, *Geophys. Astrophys. Fluid Dynamics* **85**, 195–232.

Duan, J. and S. Wiggins [1996], Fluid Exchange Across a Meandering Jet with Quasi-Periodic Variability, *J. Phys. Oceanography* **26 (7)**, 1175–1188.

Duan, J. and S. Wiggins [1997], Lagrangian Transport and Chaos in the Near Wake of the Flow Around an Obstacle: A Numerical Implementation of Lobe Dynamics, *Nonlinear Processes in Geophysics* **4**, 125–136.

Eiseman, P. and A. Stone [1973], Hodge Decomposition Theorem, *Notices of the American Mathematical Society* **20**, A169–A169.

Eremeev, V., L. Ivanov, and A. Kirwan [1992a], Reconstruction of Oceanic Flow Characteristics from Quasi-Lagrangian Data .1. Approach and Mathematical- Methods, *Journal of Geophysical Research-Oceans* **97**, 9733–9742.

Eremeev, V., L. Ivanov, A. Kirwan, and T. Margolina [1995a], Amount of Cs-137 and Cs-134 Radionuclides in the Black-Sea Produced by the Chernobyl Accident, *Journal of Environmental Radioactivity* **27**, 49–63.

Eremeev, V., L. Ivanov, A. Kirwan, and T. Margolina [1995b], Analysis of Cesium Pollution in the Black-Sea by Regularization Methods, *Marine Pollution Bulletin* **30**, 460–462.

Eremeev, V., L. Ivanov, A. Kirwan, O. Melnichenko, S. Kochergin, and R. Stanichnaya [1992b], Reconstruction of Oceanic Flow Characteristics from Quasi-Lagrangian Data .2. Characteristics of the Large-Scale Circulation in the Black-Sea, *Journal of Geophysical Research-Oceans* **97**, 9743–9753.

Fischer, H. B., N. H. Brooks, J. Imberger, R. C. Y. Koh, and E. J. List [1979], *Mixing in Inland and Coastal Waters.* Academic Press Inc., New York.

Gaspard, P. and S. Rice [1989], Hamiltonian Mapping Models of Molecular Fragmentation, *Journal of Physical Chemistry* **93**, 6947–6957.

Georges, T. M., J. A. Harlan, and R. A. Lematta [1996], Large-scale Mapping of Ocean Surface Currents with Dual over-the-horizon Radars, *Nature* **379**, 434–436.

Giles, R. T. [1995], Optimal Strategies for Discharging Pollutants into Narrow Estuaries, *Water Res.* **29**, 563–569.

Gillilan, R. and G. Ezra [1991], Transport and Turnstiles in Multidimensional Hamiltonian Mappings for Unimolecular Fragmentation - Application to Van der Waals Predissociation, *Journal of Chemical Physics* **94**, 2648–2668.

Goldstein, R. M. and H. A. Zebker [1987], Interferometric Radar Measurement of Ocean Surface Currents., *Nature* **328**, 707–709.

Gould, D. J. and D. Munro [1981], Environmental Engineering, *Coastal Discharges*, 45–50.

Haller, G. [2000], Finding Finite-Time Invariant Manifolds in Two-Dimensional Velocity Fields, *Chaos* **10**, 99–108.

Haller, G. [2001a], Distinguished Material Surfaces and Coherent Structures in 3D Fluid Flows, *Physica D* **149**, 248–277.

Haller, G. [2001b], Lagrangian Structures and the Rate of Strain in a Partition of Two-Dimensional Turbulence, *Phys. Fluids A* **13**, 3368–3385.

Haller, G. [2002], Lagrangian Coherent Structures from Approximate Velocity Data, *Phys. Fluids A* **14**, 1851–1861.

Haller, G. and A. Poje [1998], Finite-Time Transport in Aperiodic Flows, *Physica D* **119 3-4**, 352–380.

Haller, G. and G. Yuan [2000], Lagrangian Coherent Structures and Mixing in Two-dimensional Turbulence, *Physica D* **147**, 352–370.

Iacono, A. and G. Haller [2003], Stretching and Alignment in Slowly Varying Velocity Fields, *Physics of Fluids* **in preparation**.

Ide, K., D. Small, and S. Wiggins [2002], Distinguished Hyperbolic Trajectories in Time-Dependent Fluid Flows: Analytical and Computational Approach for Velocity Fields Defined as Data Sets, *Nonlinear Processes in Geophysics* **9**, 237–263.

Jaffe, C., S. Ross, M. Lo, J. Marsden, D. Farrelly, and T. Uzer [2002], Statistical Theory of Asteroid Escape Rates, *Physical Review Letters* **89**, 011101.

Joseph, B. and B. Legras [2001], On the Relation Between Kinematic Boundaries, Stirring, and Barriers for the Antarctic Polar Vortex, *J. Atmosph. Sci.* **in press**.

Kay, A. [1990], Advection-Diffusion in Reversing and Oscillating Flows. 1: The Effect of a Single Reversal, *IMA J. Appl. Math* **45**, 115–137.

Koon, W., M. Lo, J. Marsden, and S. Ross [2000], Heteroclinic Connections Between Periodic Orbits and Resonance Transitions in Celestial Mechanics, *Chaos* **10**, 427–469.

Koon, W., M. Lo, J. Marsden, and S. Ross [2001a], Low Energy Transfer to the Moon, *Celestial Mechanics & Dynamical Astronomy* **81**, 63–73.

Koon, W., M. Lo, J. Marsden, and S. Ross [2001b], Resonance and Capture of Jupiter Comets, *Celestial Mechanics & Dynamical Astronomy* **81**, 27–38.

Lapeyre, G., B. L. Hua, and B. Legras [2001], Comments on 'Finding Finite-Time Invariant Manifolds in Two-Dimensional Velocity Fields', *Chaos* **11**, 427–469.

Laskar, J. [1989], A Numerical Experiment on the Chaotic Behavior of the Solar-System, *Nature* **338**, 237–238.

Lekien, F., C. Coulliette, and J. Marsden [2003], Lagrangian Structures in High-Frequency Radar Data and Optimal Pollution Timing. In *Proceedings of the 7th Experimental Chaos Conference*. American Institute of Physics.

Levison, H. and M. Duncan [1993], The Gravitational Sculpting of the Kuiper Belt, *Astrophysical Journal* **406**, L35–L38.

Lipphardt, Jr., B. L., A. D. Kirwan, Jr., C. E. Grosch, J. K. Lewis, and J. D. Paduan [2000], Blending HF Radar and Model Velocities in Monterey Bay through Normal Mode Analysis, *Journal of Geophysical Research* **105(C2), 3**, 425–450.

Lozier, M. and S. Riser [1990], Potential Vorticity Sources and Sinks in a Quasi-Geostrophic Ocean - Beyond Western Boundary Currents, *Journal of Physical Oceanography* **20**, 1608–1627.

Lynch, A., A. Slater, and M. Serreze [2001], The Alaskan Arctic Frontal Zone: Forcing by Orography, Coastal Contrast, and the Boreal Forest, *Journal of Climate* **14**, 4351–4362.

Mackay, R., J. Meiss, and I. Percival [1984], Transport in Hamiltonian Systems, *Physica D* **13**, 55–81.

Mackay, R., J. Meiss, and I. Percival [1987], Resonances in Area-Preserving Maps, *Physica D* **27**, 1–20.

Malhotra, N., I. Mezic, and S. Wiggins [1998], Patchiness: A New Diagnostic for Lagrangian Trajectory Analysis in Time-Dependent Fluid Flows, *International Journal of Bifurcation and Chaos* **8**, 1053–1093.

Malhotra, N. and S. Wiggins [1998], Geometric Structures, Lobe Dynamics, and Lagrangian Transport in Flows with Aperiodic Time Dependence, with Applications to Rossby Wave Flow, *Journal of Nonlinear Science* **8**, 401–456.

Malhotra, R. [1996], The Phase Space Structure near Neptune Resonances in The Kuiper Belt, *Astronomical Journal* **111**, 504–516.

Mcgehee, R. and K. Meyer [1974], Homoclinic Points of Area Preserving Diffeomorphisms, *American Journal of Mathematics* **96**, 409–421.

Meiss, J. and E. Ott [1986], Markov Tree Model of Transport in Area-Preserving Maps, *Physica D* **20**, 387–402.

Mezic, I. and S. Wiggins [1999], A Method for Visualization of Invariant Sets of Dynamical Systems Based on the Ergodic Partition, *Chaos* **9**, 213–218.

Miller, P. D., C. K. R. T. Jones, A. M. Rogerson, and L. J. Pratt [1997], Quantifying Transport in Numerically Generated Velocity Fields, *Physica D* **110**, 1–18.

Ngan, K. and T. G. Shepherd [1997], Chaotic Mixing and Transport in Rossby-wave Critical Layers, *Physica D* **110**, 105–122.

Ottino, J. [1988], *The Kinematics of Mixing: Stretching, Chaos, and Transport*. Cambridge University Press, Cambridge.

Paduan, J. D. and M. S. Cook [1997], Mapping Surface Currents in Monterey Bay with CODAR-type HR Radar, *Oceanography* **10**, 49–52.

Paduan, J. D. and L. K. Rosenfeld [1996], Remotely Sensed Surface Currents in Monterey Bay from Shore-Based HF Radar (CODAR), *J. Geophys. Res.* **101**, 20669–20686.

Peters, H., L. Shay, A. Mariano, and T. Cook [2002], Current Variability on a Narrow Shelf with Large Ambient Vorticity, *Journal of Geophysical Research-Oceans* **107**, 3087.

Pierrehumbert, R. T. [1991a], Chaotic Mixing of Tracer and Vorticity by Modulated Traveling Rossby Waves, *Geophys. Astrophys. Fluid Dyn.* **58**, 285–319.

Pierrehumbert, R. T. [1991b], Large-scale Horizontal Mixing in Planetary Atmospheres, *Phys. Fluids A* **3(5)**, 1250–1260.

Poje, A. C. and G. Haller [1999], Geometry of Cross-Stream Mixing in a Double-Gyre Ocean Model, *J. Phys. Oceanogr.* **29**, 1649–1665.

Poje, A., G. Haller, and I. Mezic [1999], The Geometry and Statistics of Mixing in Aperiodic Flows, *Physics of Fluids* **11**, 2963–2968.

Prahl, F. G., E. Crecellus, and R. Carpenter [1984], Polycyclic Aromatic Hydrocarbons in Washington Coastal Sediments: An Evaluation of Atmospheric and Riverine Routes of Introduction, *Environ Sci Technol* **18**, 687–693.

Prandle, D. [1987], The Fine-Structure of Nearshore Tidal and Residual Circulations Revealed by HF Radar Surface Current Measurements, *J. Phys. Oceanogr.* **17**, 231–245.

Prandle, D. and D. K. Ryder [1985], Measurement of Surface Currents in Liverpool Bay by High-Frequency Radar, *Nature* **315**, 128–131.

Provenzale, A. [1999], Transport by Coherent Barotropic Vortices, *Annual Review of Fluid Mechanics* **31**, 55–93.

Rice, D. W., C. P. Seltenrich, R. B. Spies, and M. L. Keller [1993], Seasonal and Annual Distribution of Organic Contaminants in Marine Sediments from the Elkhorn Slough, Moss Landing Harbor and Near-Shore Monterey Bay, California, *Environmental Pollution* **82**, 79–91.

Ridderinkhof, H. and J. T. F. Zimmerman [1992], Chaotic Stirring in a Tidal System, *Science* **258**, 1107–1111.

Ridderinkhof, H., J. T. F. Zimmerman, and M. E. Philippart [1990], Tidal Exchange Between the North-Sea and Dutch Walden Sea and Mixing Time Scales of the Tidal Basins, *Netherlands J. Sea Res.* **25**, 331–350.

Roden, G. [1980], On the Variability of Surface-Temperature Fronts in the Western Pacific, as Detected by Satellite, *Journal of Geophysical Research-Oceans and Atmospheres* **85**, 2704–2710.

Rogerson, A., P. D. Miller, L. J. Pratt, and C. K. R. T. J. Jones [1999], Lagrangian Motion and Fluid Exchange in a Barotropic Meandering Jet, *J. Phys. Oceanogr.* **29(10)**, 2635–2655.

Rom-Kedar, V. and S. Wiggins [1991], Transport in Two-Dimensional Maps: Concepts, Examples, and a Comparison of the Theory of Rom-Kedar and Wiggins with the Markov Model of MacKay, Meiss, Ott, and Percival, *Physica D* **51**, 248–266.

Romkedar, V. [1990], Transport Rates of a Class of 2-Dimensional Maps and Flows, *Physica D* **43**, 229–268.

Romkedar, V. [1994], Homoclinic Tangles — Classification and Applications, *Nonlinearity* **7**, 441–473.

Romkedar, V. [1995], Secondary Homoclinic Bifurcation Theorems, *Chaos* **5**, 385–401.

Romkedar, V., A. Leonard, and S. Wiggins [1990], An Analytical Study of Transport, Mixing and Chaos in an Unsteady Vortical Flow, *Journal of Fluid Mechanics* **214**, 347–394.

Romkedar, V. and S. Wiggins [1990], Transport in 2-Dimensional Maps, *Archive for Rational Mechanics and Analysis* **109**, 239–298.

Rothstein, D., E. Henry, and J. Gollub [1999], Persistent Patterns in Transient Chaotic Fluid Mixing, *Nature* **401**, 770–772.

Rowley, C. [1996], *A Modeling Study of the North Atlantic Current*, PhD thesis, University of Rhode Island, Kingston, RI.

Samelson, R. M. [1992], Fluid Exchange Across a Meandering Jet, *J. Phys. Oceanogr.* **22**, 431–440.

Serreze, M., A. Lynch, and M. Clark [2001], The Arctic Frontal Zone as Seen in the Ncep-Ncar Reanalysis, *Journal of Climate* **14**, 1550–1567.

Shariff, K., T. Pulliam, and J. Ottino [1992], *A Dynamical Systems Analysis of Kinematics in the Time-Periodic Wake of Circular Cylinder.* In Vortex Dynamics and Vortex Method, Proc AMS-SIAM Conf.

Shay, L. K.and Graber, H. C., D. B. Ross, and R. D. Chapman [1995], Mesoscale Ocean Surface Current Structure Detected by HF Radar, *J. Atmos. Ocean Tech.* **12**, 881–900.

Shay, L., G. Goni, and P. Black [2000], Effects of a Warm Oceanic Feature on Hurricane Opal, *Monthly Weather Review* **128**, 1366–1383.

Shulman, I., C. Wu, J. Lewis, J. Paduan, L. Rosenfeld, J. Kindle, S. Ramp, and C. Collins [2002], High-Resolution Modeling and Data Assimilation in the Monterey Bay Area, *Continental Shelf Research* **22**, 1129–1151.

Smith, R. [1993], Optimal Use of Holding Tanks to Avoid Pollution in Narrow Estuaries, *IMA J Appl Math* **51**, 187–199.

Smith, R. [1998], Using Small Holding Tanks to Reduce Pollution in Narrow Estuaries, *J. Hydraulic Eng.* **124**, 117–122.

Stewart, R. H. and J. W. Joy [1974], HF Radio Measurements of Surface Currents, *Deep-Sea Res* **21**, 1039–1049.

Stirling, J. R. [2000], Transport and Bifurcation in a Non-Area-Preserving Two-Dimensional Map with Applications to the Discharge of Pollution in an Estuarine Flow, *Physica D* **144**, 169–193.

Swenson, M., P. Niiler, K. Brink, and M. Abbott [1992], Drifter Observations of a Cold Filament off Point Arena, California, in July 1988, *Journal of Geophysical Research-Oceans* **97**, 3593–3610.

Szoke, S., V. Wertz, and E. Persoons [1996], Use of Temperature Front Parameters to Compute the Drying Front, *Journal of Agricultural Engineering Research* **65**, 313–323.

Tancredi, G. [1995], The Dynamical Memory of Jupiter Family Comets, *Astronomy and Astrophysics* **299**, 288–292.

Torbett, M. and R. Smoluchowski [1990], Chaotic Motion in a Primordial Comet Disk Beyond Neptune and Comet Influx to the Solar-System, *Nature* **345**, 49–51.

Velasco Fuentes, O. M. [2001], Chaotic Advection by Two Interacting Finite-Area Vortices, *Phys. Fluids A* **13**, 901–912.

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

Watson, A. J., M. J. Messias, E. Fogelqvist, and K. A. V. e. Scoy [1999], Mixing and Convection in the Greenland Sea from a Tracer-Release Experiment, *Nature* **401**, 902–904.

Webb, T. and R. B. Tomlinson [1992], Design Procedures for Effluent Discharge into Estuaries during Ebb Tide, *J. Env. Eng.* **118**, 338–362.

Weiss, J., A. Provenzale, and J. McWilliams [1998], Lagrangian Dynamics in High-Dimensional Point-Vortex Systems, *Phys. Fluids* **10**, 1929–1941.

Wiggins, S. [1992], *Chaotic Transport in Dynamical Systems*. Springer-Verlag, New York.

Williams, B., D. Marteau, and J. Gollub [1997], Mixing of a Passive Scalar in Magnetically Forced Two-Dimensional Turbulence, *Physics of Fluids* **9**, 2061–2080.

Yuan, G., L. Pratt, and C. Jones [2002], Barrier Destruction and Lagrangian Predictability at Depth in a Meandering jet, *Dynamics of Atmospheres and Oceans* **35**, 41–61.

Zimmerman, J. T. F. [1986], The Tidal Whirlpool— A Review of Horizontal Dispersion by Tidal and Residual Currents, *Netherlands J. Sea Res.* **20**, 133–154.