

# Using flow geometry for drifter deployment in Lagrangian data assimilation

By H. SALMAN<sup>1\*</sup>, K. IDE<sup>2</sup> and C. K. R. T. JONES<sup>1</sup>, <sup>1</sup>*Department of Mathematics, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA;* <sup>2</sup>*Department of Atmospheric and Oceanic Sciences & Institute of Geophysics and Planetary Physics, University of California, Los Angeles, CA, USA*

Manuscript received 21 November 2006; in final form 5 November 2007

## ABSTRACT

Methods of Lagrangian data assimilation (LaDA) require carefully chosen sites for optimal drifter deployments. In this work, we investigate a directed drifter deployment strategy with a recently developed LaDA method employing an augmented state vector formulation for an Ensemble Kalman filter. We test our directed drifter deployment strategy by targeting Lagrangian coherent flow structures of an unsteady double gyre flow to analyse how different release sites influence the performance of the method. We consider four different launch methods; a uniform launch, a saddle launch in which hyperbolic trajectories are targeted, a vortex centre launch, and a mixed launch targeting both saddles and centres. We show that global errors in the flow field require good dispersion of the drifters which can be realized with the saddle launch. Local errors on the other hand are effectively reduced by targeting specific flow features. In general, we conclude that it is best to target the strongest hyperbolic trajectories for shorter forecasts although vortex centres can produce good drifter dispersion upon bifurcating on longer time-scales.

## 1. Introduction

Increasing interest in Lagrangian data has led to recent developments in using such data for improving predictions of the ocean (Carter, 1989; Kamachi and O'Brien, 1995; Molcard et al., 2003; Özgökmen et al., 2003). These methods rely on reconstructing Eulerian velocity information from consecutive measurements of drifter positions which are then assimilated into Eulerian flow models. An alternative method, developed by Kuznetsov et al. (2003) and Ide et al. (2002), which employs an augmented state vector formulation together with an Extended Kalman filter, has been shown to maximize the information content in comparison to the above approaches for Lagrangian data assimilation (LaDA). In stark contrast to the aforementioned methods, the augmented state vector formulation allows one to evolve the error correlations between the Eulerian flow variables and drifter coordinates in a way that is entirely consistent with the evolution of error correlations in Eulerian data assimilation. The method of Kuznetsov et al. (2003) and Ide et al. (2002), therefore, bypasses the need to introduce approximations in the reconstruction of

Eulerian velocity information. The augmented state vector formulation has recently been applied, together with an Ensemble Kalman filter (EnKF), to a shallow water model of the ocean by Salman et al. (2006). The work of Salman et al. clearly illustrated the success of the method by establishing that the filter is stable provided that data is assimilated more frequently than the Lagrangian autocorrelation time-scale ( $T_L$ ). Methods that reconstruct Eulerian velocity information, on the other hand, report a maximum allowable time-interval of  $0.2T_L$ . The augmented state vector formulation is, therefore, indispensable in scenarios where infrequent measurements are available.

Salman et al. (2006) carefully analysed the performance of the method with respect to variations in different parameters, including the number of ensembles, the localization of the covariance matrix, and the drifter release sites. Of these parameters, the most difficult to understand, and yet perhaps most important, is the initial launch sites of the drifters. The root of the difficulty is two-fold; firstly it has been known for sometime now that simple unsteady two-dimensional velocity fields can give rise to chaotic advection in the motion of Lagrangian material particles (Aref, 1984; Aref and El Naschie, 1994). Such chaotic behaviour has come to be known as Lagrangian chaos and reflects a degree of unpredictability in the motion of material particles. Given that a good first order approximation that describes the motion of floats and drifters is to assume that they behave as material particles (Bower, 1991), understanding their motion in unsteady flows is extremely non-trivial. In light of these results, understanding

---

\*Corresponding author.

Present address: Department of Applied Mathematics and Theoretical Physics, Centre for Mathematical Sciences, University of Cambridge, CB3 0WA, UK.

e-mail: h.salman@damtp.cam.ac.uk

DOI: 10.1111/j.1600-0870.2007.00292.x

how drifter release locations affect the performance of our LaDA requires a more in depth understanding of what controls the trajectories of these drifters. The above difficulty is compounded by a second non-linear effect in that the velocity field that governs the motion of the drifters is part of the state vector that is being updated by our assimilation algorithm. This sets in a two-way coupling between the positions of the drifters and the their impact on the method. With such complications, the preliminary results presented in Salman et al. (2006) regarding the influence of different release sites on the performance of the method are not fully conclusive.

In this study, we aim to provide a more in depth and direct understanding of how different initial drifter placements can affect the performance of the LaDA method presented in Salman et al. (2006). Our approach will be to employ recent ideas from dynamical systems theory (Jones and Winkler, 2002) to extract Lagrangian coherent structures (LCS) from our numerically computed flows. These Lagrangian structures have been studied extensively for geophysical flows in recent years and are understood to orchestrate the evolution and motion of material particles (see Wiggins, 2005, for a recent review). They provide a geometric approach to analyse the motion of particles in unsteady 2-D velocity fields. By extracting such structures from our flow, the complications that arise from the non-linear dependence of the drifter positions on the underlying advecting flow field reduce to understanding the evolution of a set of material lines that delineate such flow structures.

While this approach has proven to be successful in a number of studies conducted to date (Poje and Haller, 1999; Coulliette and Wiggins, 2000; Jones and Winkler, 2002; Kuznetsov et al., 2002; Mancho et al., 2004), we are posed with an additional difficulty here due to the two-way coupling between the drifter positions and the velocity field that arises from assimilation. To simplify the problem in this study, we will focus on a twin-experiment configuration in which the flow field for the synthetic ‘truth’ is known. Under such idealized conditions, we can extract the LCS for the true state and use this as our template for the directed launch of drifters. The approach we will adopt will, therefore, build on the ideas used by Poje et al. (2002), and Toner and Poje (2004) to study how initial drifter placements influences the performance of our LaDA method. By targeting the Lagrangian flow structures of the ‘true’ system, we are able to simplify the problem drastically. At the same time, it enables us to focus on the structures of most relevance since it is these structures that control the motion of the drifter positions being assimilated into the model. This simplification also serves as a relevant benchmark case for future studies geared towards the design of an observing system for the optimal deployment of drifters for use with our LaDA formulation. In such a case, one is interested in where to release drifters to improve ocean forecasting without prior knowledge of the ‘true’ state of the system. The design of such a system introduces major challenges and it, therefore, helps to understand how the flow structures affect our data assimilation

method under the more idealized scenario of a twin-experiment setup. This in turn would allow us to determine how much improvement we can expect to gain from directed drifter placement in relation to random drifter placement before embarking on the involved design of such an observing system. A similar approach has been adopted in the study of Molcard et al. (2006) who analysed how directed drifter deployment affected the convergence of the data assimilation method developed by Molcard et al. (2003) and Özgökmen et al. (2003). As stated earlier, however, the formulation of our LaDA scheme is fundamentally different from these methods. The work presented here, therefore, provides an essential first step towards the development of a Lagrangian observing system for the methods developed by Ide et al. (2002), Kuznetsov et al. (2003) and Salman et al. (2006).

The remainder of the paper is organized into four main sections. We begin in Section 2 by stating the main elements of our LaDA method that is based on the use of the EnKF. The mathematical model used in this study together with the problem setup is briefly described in Section 3. We then describe how we compute the LCS for our double gyre flow and present results using two different methods. The first method computes relative dispersion, produced by integrating material particles in forward and backward time, to locate regions of maximal stretching within the flow. The second method employs the idea of straddling one of the stable or unstable directions of a distinguished hyperbolic trajectory (DHT) to compute the Lagrangian structures of interest. These Lagrangian templates are used to identify different release sites for sets of drifters that consequently have very different dispersion characteristics. Results are subsequently presented in Section 4 to quantify how targeting specific regions of the flow can affect our data assimilation method. We end in Section 5 with conclusions of the main results.

## 2. Ensemble Kalman filter for Lagrangian data

In this section, we will summarize the key elements of the formulation developed by Salman et al. (2006), Kuznetsov et al. (2003) and Ide et al. (2002) for assimilating Lagrangian data. We begin by writing our system of equations governing the  $(N_F \times 1)$  flow state vector,  $\mathbf{x}_F$ , and the  $(2N_D \times 1)$  drifter state vector,  $\mathbf{x}_D$ , of our  $N_D$  drifters in terms of the augmented state vector,  $\mathbf{x} = (\mathbf{x}_F^T, \mathbf{x}_D^T)^T$ . Using an ensemble forecast with  $N_E$  members, we have

$$\frac{d\mathbf{x}_j^f}{dt} = \mathbf{m}_j(\mathbf{x}_j, t), \quad j = 1, \dots, N_E. \quad (1)$$

Here, the subscript  $j$  denotes the ensemble member, the superscript  $f$  denotes the system forecast, and  $\mathbf{m}_j$  is the evolution operator. Using this augmented system of equations for each member of the ensemble, we can define the corresponding covariance matrix by

$$\mathbf{P}^f = \frac{1}{N_E - 1} \sum_{j=1}^{N_E} (\mathbf{x}_j^f - \overline{\mathbf{x}}^f)(\mathbf{x}_j^f - \overline{\mathbf{x}}^f)^T. \quad (2)$$

We note that the covariance matrix can be expressed in block matrix form as

$$\mathbf{P}^f = \begin{pmatrix} \mathbf{P}_{FF}^f & \mathbf{P}_{FD}^f \\ \mathbf{P}_{DF}^f & \mathbf{P}_{DD}^f \end{pmatrix}. \quad (3)$$

$\mathbf{P}_{FF}$ ,  $\mathbf{P}_{FD}$ ,  $\mathbf{P}_{DF}$ , and  $\mathbf{P}_{DD}$  are  $(N_F \times N_F)$ ,  $(N_F \times 2N_D)$ ,  $(2N_D \times N_F)$ , and  $(2N_D \times 2N_D)$  matrices, respectively. They correspond to correlations between the flow and drifter parts of the state vector. The mean of the state vector taken over the ensemble is denoted by an overline in eq. (2) and is defined as

$$\overline{\mathbf{x}^f} = \frac{1}{N_E} \sum_{j=1}^{N_E} \mathbf{x}_j^f. \quad (4)$$

When Lagrangian measurements become available, the analysis state of the system, denoted by  $\mathbf{x}^a$ , can be obtained using the perturbed observation EnKF (Evensen, 1994, 2003). The filter updates each ensemble member by

$$\mathbf{x}_j^a = \mathbf{x}_j^f + \mathbf{K} \mathbf{d}_j, \quad (5)$$

where the Kalman gain matrix is defined as

$$\mathbf{K} = \begin{pmatrix} \rho_{FD} \circ \mathbf{P}_{FD} \\ \rho_{DD} \circ \mathbf{P}_{DD} \end{pmatrix} (\rho_{DD} \circ \mathbf{P}_{DD} + \mathbf{R})^{-1}, \quad (6)$$

and the innovation vector  $\mathbf{d}_j$  is given by

$$\mathbf{d}_j = \mathbf{y}^o - \mathbf{K} \mathbf{x}_{D,j}^f + \tilde{\epsilon}_j^f. \quad (7)$$

The observation vector  $\mathbf{y}^o$  represents noisy spatial coordinates of the drifters in the zonal and meridional directions. Both the noise in the drifters' positions and the perturbations to the observations given by  $\tilde{\epsilon}_j^f$  are drawn from a gaussian distribution with a covariance matrix equal to  $\mathbf{R}$ . The latter is also made to satisfy the condition  $1/N_E \sum_{j=1}^{N_E} \tilde{\epsilon}_j^f = 0$ .

In eq. (6), the operator  $\circ$  denotes the Schur product between two matrices. The elements of the matrix  $\rho$  correspond to a distance-dependent correlation function. We have used the correlation function given by Gaspari and Cohn (1999) which is smooth and has compact support.

### 3. Flow model and experimental setup

The model we employ in this study is very similar to the one used in the work of Salman et al. (2006). The model is an idealized ocean model with a square domain configuration whose size in the zonal and meridional directions are denoted by  $L_x$  and  $L_y$ , respectively. As described in Cushman (1994) and Pedlosky (1986), the flow within this domain can be modeled by the reduced gravity shallow water system of equations which are given by

$$\frac{\partial h}{\partial t} = -\frac{\partial(hu)}{\partial x} - \frac{\partial(hv)}{\partial y}, \quad (8)$$

$$\begin{aligned} \frac{\partial u}{\partial t} &= -\mathbf{u} \cdot \nabla u + fv - g' \frac{\partial h}{\partial x} + F^u + \frac{1}{h} \left( \frac{\partial \tau_{xx}}{\partial x} + \frac{\partial \tau_{xy}}{\partial y} \right), \\ \frac{\partial v}{\partial t} &= -\mathbf{u} \cdot \nabla v - fu - g' \frac{\partial h}{\partial y} + \frac{1}{h} \left( \frac{\partial \tau_{yx}}{\partial x} + \frac{\partial \tau_{yy}}{\partial y} \right). \end{aligned} \quad (9)$$

$h$  is the surface height,  $(u, v)$  is the fluid-velocity vector,  $g'$  is the reduced gravity,  $F^u$  is a horizontal wind-forcing acting in the zonal direction, and  $f$  is the Coriolis parameter. We invoke the  $\beta$ -plane approximation allowing the Coriolis term to be expressed as

$$f = f_o + \beta y, \quad (10)$$

where  $f_o$  and  $\beta$  are constants. A zonal wind-forcing of the form

$$F^u = \frac{-\tau_o}{\rho H_o(t)} \cos(2\pi y/L_y), \quad (11)$$

$$H_o(t) = \frac{1}{L_x L_y} \int_0^{L_x} \int_0^{L_y} h(x, y, t) dx dy, \quad (12)$$

is employed in this work where  $x$  and  $y$  are the coordinates in the zonal and meridional directions measured from the western and southern boundaries of our flow domain, respectively,  $\tau_o$  is the wind stress,  $\rho$  is the density of water, and  $H_o(t)$  is the average water depth.

In contrast to the study presented in Salman et al. (2006), we have replaced the dissipation term by the form proposed by Schär and Smith (1993). This is given by

$$\tau_{ij} = \mu h \left( \frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} - \delta_{ij} \frac{\partial u_k}{\partial x_k} \right), \quad (13)$$

for all possible permutations of the indices  $i$  and  $j$ , which can take the two cartesian components  $x$  or  $y$ , and  $\mu$  denotes a (constant) eddy viscosity of the flow. This form of the dissipation term has been employed as it leads to a self-consistent formulation of the shallow-water system of equations as further discussed by Shchepetkin and O'Brien (1996) and Gent (1993). In particular, the resulting equations satisfy energy and momentum conservation principles that are otherwise violated with the form of the dissipation employed in Salman et al. (2006).

The above equations are supplemented with the boundary and initial conditions given by

$$u(x, y, t)|_{\partial\Omega} = 0, \quad v(x, y, t)|_{\partial\Omega} = 0, \quad (14)$$

$$u(x, y, 0) = 0, \quad v(x, y, 0) = 0, \quad h(x, y, 0) = H_o, \quad (15)$$

where  $\partial\Omega$  represents the boundaries of our flow domain. To setup a double gyre circulation, we solved the above equations with the model parameters given in Table 1.

Equations (8)–(9) were discretized on a uniform mesh with  $(n_x \times n_y) = (100 \times 100)$  grid points which corresponds to a grid spacing of  $(\Delta x \times \Delta y) = (20 \text{ km} \times 20 \text{ km})$ . This produces a set of  $N_F = 2(n_x - 1)(n_y - 1) + n_x n_y$  equations for the flow which correspond to the flow state vector  $\mathbf{x}_F$ . In addition, a set of

Table 1. Parameters prescribed for the reduced gravity shallow water system of equations

Property	Value
$L_x$	2000 km
$L_y$	2000 km
$f_o$	$6 \times 10^{-5} \text{ s}^{-1}$
$\beta$	$2 \times 10^{-11} \text{ m}^{-1} \text{ s}^{-1}$
$H_o$	500 m
$g'$	$0.02 \text{ ms}^{-2}$
$\rho$	$1000 \text{ kg m}^{-3}$
$\tau_o$	$0.05 \text{ N m}^{-2}$
$\Delta x$	20 km
$\Delta y$	20 km
$\Delta t$	12 min
$\mu$	$400 \text{ m}^2 \text{ s}^{-1}$

$2N_D$  equations corresponding to the zonal and meridional coordinates of the  $N_D$  drifters was integrated together with the flow. The drifters were assumed to be advected as material particles by the flow field. A twin-experiment setup was used for our assimilation experiments. The synthetic truth (control run) was obtained by integrating the model for 12 yr from an initial flow at rest using an average water depth of  $H_o = 500$  m. The flow for the forecast was made up of an ensemble consisting of 80 members. The ensemble was generated by perturbing the initial height field of each member so that the mean water depth taken over the ensemble is given by  $\overline{H_o} = 550$  m. A gaussian distribution with a variance of  $\sigma_h = 50$  m was used to generate  $H_o$  for the different members of the ensemble. Each member was then integrated forward in time for a 12 yr period to produce a set of flows corresponding to the different initial height fields. At the beginning of the 13th year, drifters were released into the true flow. Each drifter was initialized at a specified location within the flow domain and then integrated with our true flow over a period of 2 yr to generate a set of true trajectories. The observation errors corresponding to these drifter trajectories were taken to be distributed as independent gaussians with the same statistics; that is

$$E[\epsilon(t_k)\epsilon^T(t_l)] = \delta_{kl}\sigma^2I, \\ E[\epsilon(t_k)] = E[\epsilon(t_l)] = 0, \quad (16)$$

where  $\sigma$  was taken to be 200 m in this work. Each ensemble member of our forecast model was then integrated over the same 2 yr time interval using a corresponding set of perturbed drifters with the same error statistics as those given in eq. (16).

When performing our assimilation experiments, we fixed the assimilation time interval to be equal to 1 d. This is well within the Lagrangian autocorrelation time-scale of 10 d for this flow (Salman et al., 2006). Consequently, we avoided the potential problems of filter divergence that can arise with the use of larger time intervals.

## 4. Results

### 4.1. Computation of Lagrangian coherent structures

In this work, we will employ our knowledge of the ‘true’ velocity field to extract the desired LCS that control the motion of the observed drifter positions. The computation of LCS from a numerically computed velocity field with general time dependence is a topic that has received much attention in recent years (Wiggins, 2005). In general, the first step with any such method is to locate distinguished hyperbolic trajectories (DHTs) associated with the flow given over some finite time interval  $[t_0, t_1]$ . For time-periodic flows, such points can easily be identified as the fixed points of an appropriately chosen Poincaré map of the flow. However, for flows with general time-dependence, the DHTs are aperiodic in time. Identifying them, therefore, requires the use of alternative methods that have found varying degrees of success in different applications. The first approach is based on the idea that for flows having a Lagrangian time-scale much shorter than the Eulerian time-scale, the DHTs can be expected to lie close to an instantaneous saddle point (ISP). The precise conditions under which this property holds have been derived by Haller and Poje (1998) and applied to a double gyre flow of the ocean. In general, the separation between the Lagrangian and Eulerian flow time-scales is a property of many geophysical flows including the double gyre flow which we will consider in this work. This method will provide one way that we will employ to identify regions that can potentially have hyperbolic trajectories.

A second method that has proven useful in locating DHTs is to compute the relative dispersion (Jones and Winkler, 2002) or direct Lyapunov exponent (DLE) (Haller, 2001) in forward and backward time for a given flow field. Coherent structures are then delineated by ridges in the DLE (or relative dispersion) fields. These methods have been demonstrated to work well in regions of the flow not dominated by shear. In shear dominated regions, however, the computed DLE fields can give rise to spurious flow structures. This was illustrated with a particular example by Haller (2002) who showed that a shear flow with an inflection point can produce a ridge in the DLE field. The problem is attributed to the fact that both relative dispersion and the DLE identify regions of large stretching. Since a shear flow with an inflection point has an anonymously large stretching near the inflection point, spurious structures tend to emerge. Another disadvantage of these methods is that they are computationally expensive requiring the integration of a large number of particles. This is traded off, however, by the fact that they provide a direct method to identify DHTs for a general class of flows. Our experience shows that it is useful to use both the straddling and DLE (or relative dispersion) methods to gain information about the LCS of a given flow (Salman et al., 2007). For these reasons, we have opted to employ both methods in this work. We will begin by computing the DLE that provides a global portrait of the distribution of the coherent structures in our flow. We are

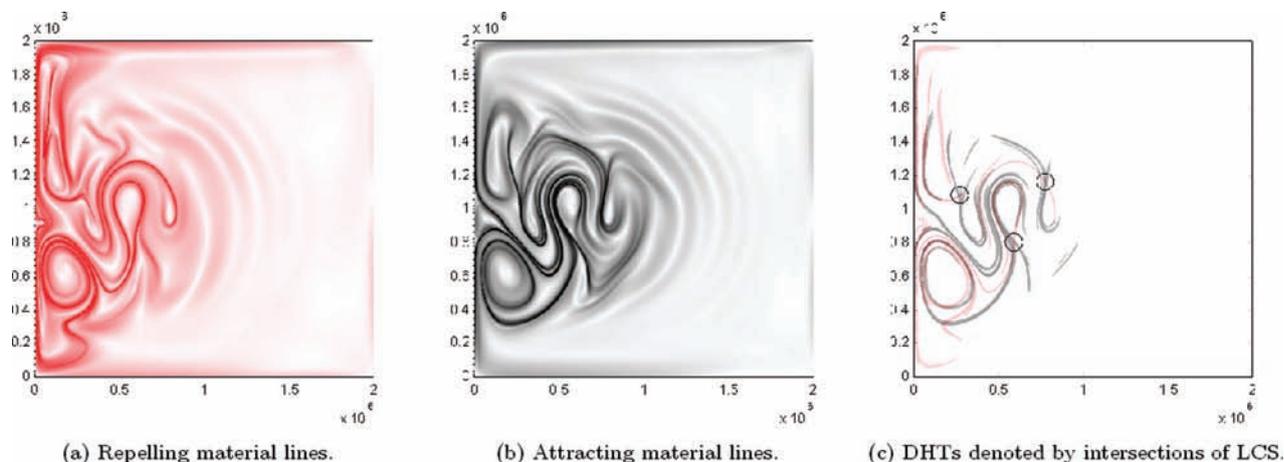


Fig. 1. Finite time DLEs computed in forward and backward time on the 60th day to extract repelling and attracting material lines, respectively. The intersections of the two sets of lines are also shown and indicate regions of distinguished hyperbolic trajectories (DHTs). Axes represent length in metres.

ultimately interested in targeting a few structures, however, that can produce the maximal dispersion characteristics for our drifters. We will, therefore, also employ the first method described above to identify the strongest ISPs (those with the largest—in absolute terms—eigenvalues). This can give a good indication of the strength of the DHT as we will show in the results that we will present below.

Figure 1 shows plots of the DLEs computed in forward and backward time at the 60th day of the 13th year. The time interval for integration was set to 50 d. Integration made in forward time is used to reveal repelling material lines whilst integration made in backward time is used to reveal attracting material lines. These attracting/repelling material lines are approximate finite-time generalizations of unstable/stable manifolds of time-periodic flows. They orchestrate how water parcels are advected by the underlying flow field and reveal regions that can lead to rapid dispersion of the drifters. From the figures, we observe that most of the underlying structure is located close to the western boundary within the mid-latitude jet. We also notice that significant shearing is present in the flow which gives rise to spurious structures (those with significantly lower values of the DLE field). To clarify the structures of most relevance, we have re-plotted the two fields in Fig. 1c. A lower threshold has been applied in this case to remove most of the spurious structures present in Fig. 1a and b. Here the pink values depict repelling material lines whilst grey depict attracting material lines. The two DLE fields have been overlaid on one another to identify regions where the extrema of the two fields intersect. Such intersections give a potential indication of where DHTs are located within the flow. We have employed this approach to identify three hyperbolic trajectories as indicated in Fig. 1c. We note that, in general, the repelling and attracting material lines will intersect at many locations. In selecting the intersections shown in Fig. 1c, we have made use of the fact that for our double gyre flow, we expect the DHTs to be located close to the ISPs. This is illustrated in

the corresponding plots of the instantaneous height (streamline field), shown in Fig. 2a, in which the ISPs are indicated by the cross-hairs. The cross-hairs have been scaled according to the magnitude of their respective eigenvalues and indicate that the three trajectories we have identified correspond to regions of strong hyperbolicity. We can, therefore, expect rapid dispersion of the drifters that straddle the respective repelling material lines of these DHTs.

In order to refine the computations of the attracting and repelling material lines presented in Fig. 2a, we have used the method of straddling used by Malhotra and Wiggins (1998) and also employed in the work of Salman et al. (2007) in combination with the DLE method. The idea here is to initialize a small material line segment close to the DHT and straddling one of its corresponding stable or unstable directions. For attracting material lines, we initialize a segment straddling the stable direction and integrate forward in time. The process is reversed for computing repelling material lines with the integration being performed in backward time. Haller and Poje (1998) showed that for geophysical flows, in which a clear separation exists between the Lagrangian and Eulerian flow time-scales, the stable and unstable directions of the DHTs are approximately aligned with the stable and unstable eigenvectors of the nearby ISP. We have, therefore, used this property to initialize our line segments close to the ISPs and straddling the respective eigenvectors as required. We point out that while this approach can potentially introduce an error in terms of the initial placement of the material line with respect to the structures of interest, the errors converge to zero at an exponential rate with respect to the time of integration. This property is a characteristic of DHTs and can be exploited in their computation.

Figure 2b shows plots of the attracting and repelling material lines corresponding to the DHTs identified in Fig. 2a. Also presented are the motion of drifters that have been initialized in six groups of three. Drifters in three of the groups are initialized at

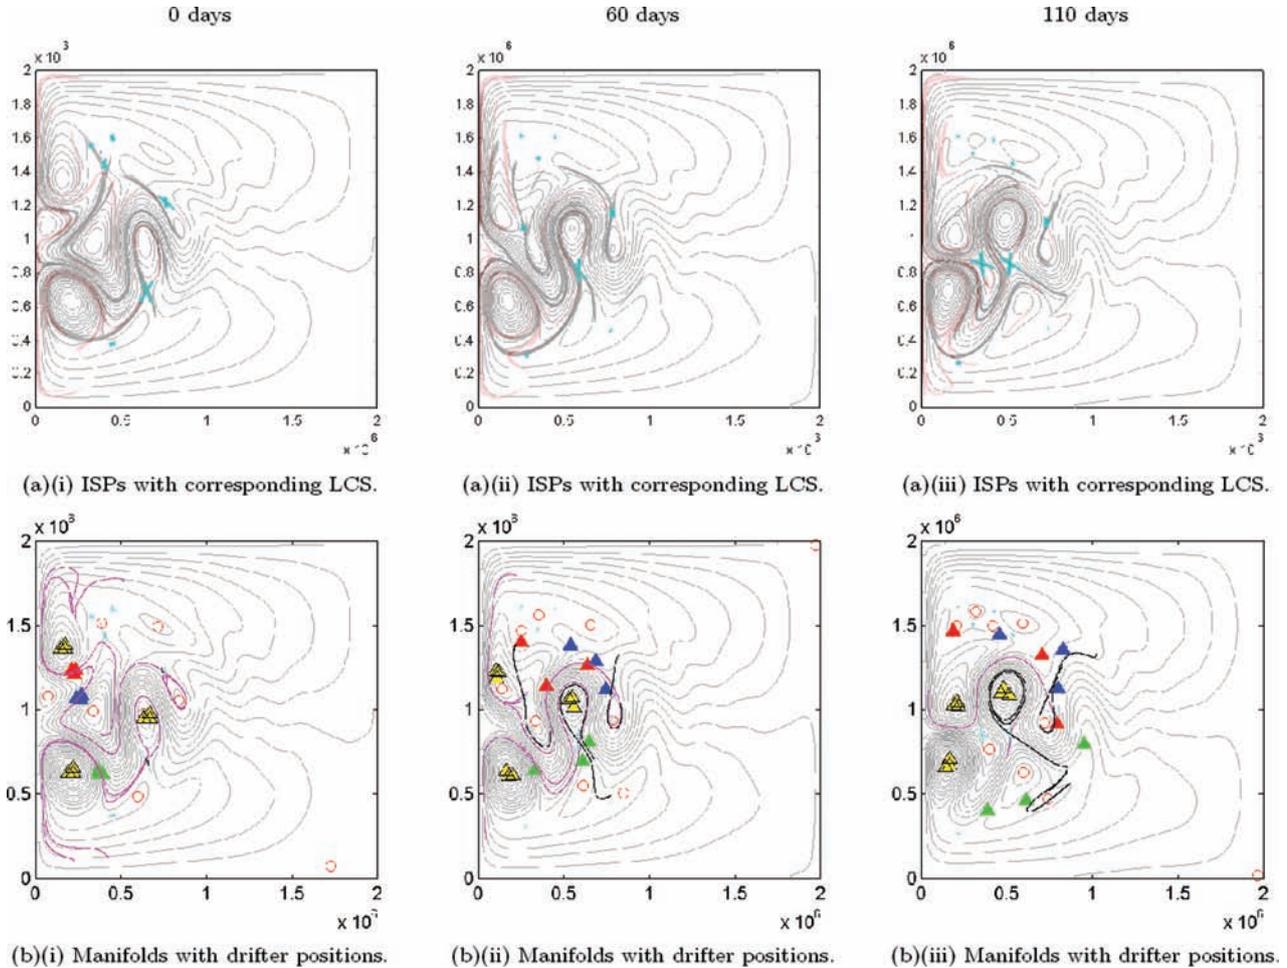


Fig. 2. Contour plots of height (streamline) field with corresponding LCS at different times. Instantaneous saddle points are indicated with cross-hairs, and elliptic points are indicated with a circle. The length of the cross-hairs are proportional to the magnitude of the eigenvalues of the saddle points. Triangles indicate the different sets of drifter positions at various times. (a)(i)–(iii), show LCS computed from the DLE method with repelling material lines plotted in red and attracting material lines plotted in grey. The LCS shown in (b)(i)–(iii) were computed using the straddling method with repelling material lines plotted in magenta and attracting material lines plotted in black. LCS shown in (a)(ii) correspond to those presented in Fig. 1(c). Axes represent length in metres.

positions straddling a repelling material line so that after some time they disperse rapidly as they encounter the respective DHT. The other three sets of drifters are initialized within the cores of three vortices. These drifters remain in these vortices throughout most of the first year. We will use the sets of drifters shown in Fig. 2b to determine how they affect the performance of our LaDA method presented in Section 2.

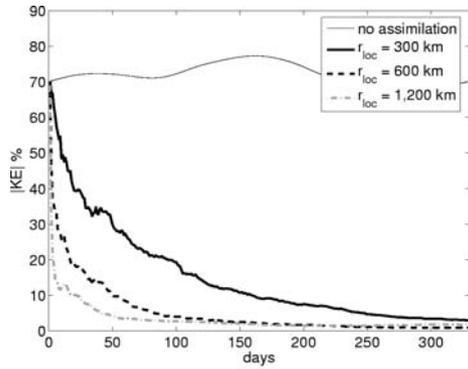
#### 4.2. Lagrangian data assimilation

Before presenting assimilation results for our drifter deployment strategy, we first need to prescribe several key parameters for our data assimilation simulations. Clearly, the method will strongly depend on the localization radius employed in our EnKF. The study of Salman et al. (2006) indicated that a localization radius of  $r_{\text{loc}} = 600$  km provided optimal results. To verify that

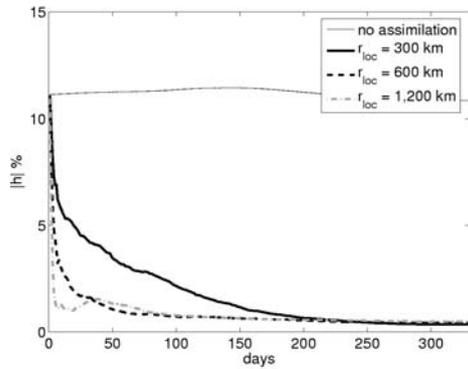
this figure continues to hold for our modified form of the governing flow equations (i.e. with different form of the dissipation term), we have repeated these assimilation experiments. As in Salman et al., we have tested three values of  $r_{\text{loc}}$  and used 36 drifters uniformly distributed throughout the flow domain in each case.

Results are presented in Fig. 3 in terms of the time history for the kinetic energy and height field norms. Here we define the error norms for the flow field and the drifters by

$$|\text{KE}|^t = \left[ \frac{\sum_{i=1}^{n_x-1} \sum_{j=1}^{n_y-1} \left( \overline{u_{i,j}^f} - u_{i,j}^t \right)^2 + \left( \overline{v_{i,j}^f} - v_{i,j}^t \right)^2}{\sum_{i=1}^{n_x-1} \sum_{j=1}^{n_y-1} \left( u_{i,j}^t \right)^2 + \left( v_{i,j}^t \right)^2} \right]^{1/2} \quad (17)$$



(a) Kinetic energy norm.



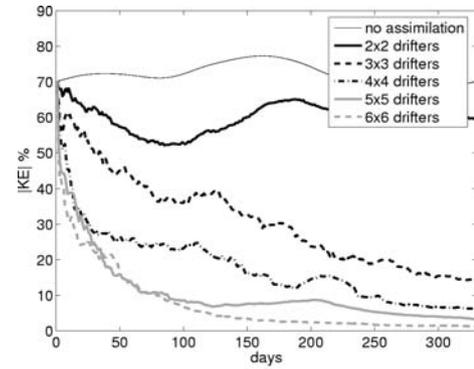
(b) Height field norm.

Fig. 3. Effect of localization radius on the convergence of the error norms.

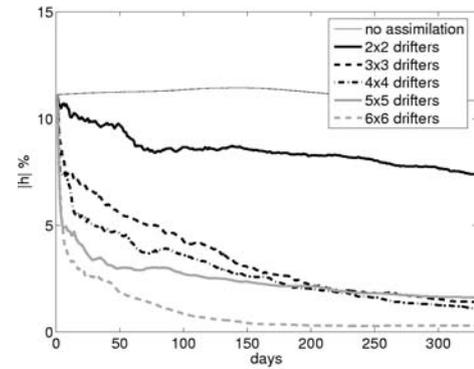
$$|h|^t = \left[ \frac{\sum_{i=1}^{n_x} \sum_{j=1}^{n_y} (\overline{h_{i,j}^f} - h_{i,j}^t)^2}{\sum_{i=1}^{n_x} \sum_{j=1}^{n_y} (h_{i,j}^t)^2} \right]^{1/2} \quad (18)$$

$$|\mathbf{x}_D|^t = \left[ \frac{\sum_{i=1}^{N_D} (\overline{x_{D,i}^f} - x_{D,i}^t)^2 + (\overline{y_{D,i}^f} - y_{D,i}^t)^2}{\sigma^2 N_D} \right]^{1/2}, \quad (19)$$

where overbars denote ensemble mean quantities. We note that, for the results to be presented in the remainder of this study, day 0 corresponds to the beginning of the 13th year. All three assimilation experiments show a reduction in the error norms with respect to the case without assimilation. The results also clearly indicate that as  $r_{loc}$  is increased, the initial rate of convergence in both norms also increases. However, while the convergence for both  $r_{loc} = 300$  km and  $r_{loc} = 600$  km is more or less monotonic over the length of the simulation, this is clearly not the case for  $r_{loc} = 1200$  km. This is best seen in the norm of the height field which reveals a gradual increase at around 25 d. Moreover, the norm in the kinetic energy field at 330 d with  $r_{loc} = 1200$  km is greater than the case with  $r_{loc} = 600$  km. As pointed out by



(a) Kinetic energy norm.



(b) Height field norm.

Fig. 4. Effect of number of drifters used on the convergence of the error norms.

Hamill et al. (2001), this degradation in the convergence of the method for large  $r_{loc}$  occurs as a result of spurious noisy correlations that exist in the covariance matrix. We will, therefore, use  $r_{loc} = 600$  km in the remainder of this work which we have clearly shown to be the optimum value for our modified system of equations.

A second key parameter that can have an important impact on the convergence of the method is the number of drifters released into the flow. The previous study conducted by Salman et al. (2006) indicated a value of 36 drifters uniformly released into the flow provided optimal results. Increasing the number of drifters beyond this value did not significantly alter the results. Given that we are interested in using our drifter deployment strategy to obtain the best results with the fewest possible number of drifters, we have explored how reducing the number of uniformly distributed drifters affected our results. Figure 4 shows the convergence of the kinetic energy and height flow fields for five different cases in which the total number of drifters uniformly released into the flow was varied from 4 to 36. The results clearly indicate that best results are obtained with 36 drifters. However, the gain made in the convergence of the method diminishes with an increasing number of drifters. In particular, there is a significant improvement in increasing the drifters from four to nine in comparison to improvements made beyond nine drifters. Given this observation, we have chosen to focus on the use of nine

Table 2. Initial drifter locations used in the uniform launch and directed launch simulations

Launch method	Drifter number	Coordinates (km)
Saddle (red)	1	(220.5,1211.0)
	2	(230.5,1231.0)
	3	(205.5,1231.0)
Saddle (blue)	4	(238.5,1066.8)
	5	(263.5,1062.0)
	6	(270.5,1087.0)
Saddle (green)	7	(365.0,617.5)
	8	(390.0,622.5)
	9	(365.0,642.5)
Centre (yellow)	10	(143.4,1361.5)
	11	(168.4,1361.5)
	12	(168.4,1386.5)
	13	(188.5,624.7)
	14	(218.5,624.7)
	15	(218.5,654.7)
	16	(629.5,950.0)
	17	(664.5,950.0)
	18	(664.5,980.0)
Uniform	19	(500.0,500.0)
	20	(1000.0,500.0)
	21	(1500.0,500.0)
	22	(500.0,1000.0)
	23	(1000.0,1000.0)
	24	(1500.0,1000.0)
	25	(500.0,1500.0)
	26	(1000.0,1500.0)
	27	(1500.0,1500.0)

drifters for our optimal deployment strategy. Our aim is to see how a directed drifter deployment strategy with only nine drifters will compare with the results obtained with a uniform drifter deployment strategy.

The sets of drifters used in our directed drifter deployment studies are those shown in Fig. 2b with their corresponding initial release locations given in Table 2. Also included in the table are the initial drifter deployment locations used in the uni-

form launch with nine drifters. As stated earlier, six groups of drifters were employed in the directed drifter deployment with each group containing three drifters. Three of these groups were deployed so as to straddle repelling material lines of associated DHTs whilst the remaining three groups were released within vortex centres. We have divided our directed drifter deployment data assimilation experiments using these drifters into three different cases. In the first case, we considered the set of drifters targeting DHTs. In the second case, we considered the set of drifters targeting only vortex centres. Finally, in the third case, we employed a drifter released within each of the vortex cores and two of the three groups of drifters targeting DHTs. This included drifters 1, 2, 3, 7, 8, 9, 10, 13 and 16. A total of nine drifters was used in each one of these assimilation experiments.

To illustrate the distinctive nature of the drifter motion in the four assimilation experiments (including the case with nine uniformly released drifters), we have plotted their trajectories over the time interval  $[0, 300]$  d. These are shown in Fig. 5 for each of the uniform, saddle, centre, and mixed drifter launch cases. Figure 5a, which corresponds to the uniform case, shows a good coverage of the flow by the drifter trajectories and can be directly attributed to the initially disparate drifter placements. The saddle case shown in Fig. 5b illustrates that DHTs can produce rapid drifter dispersion leading to good flow coverage even though the drifters were initially placed in clusters of three. The centre launch (Fig. 5c) shows the least spread in the drifter trajectories with drifters being trapped within their respective vortices. The mixed case is a combination of the centre and saddle cases and produces an intermediate level of dispersion. We will now begin to analyse how such different drifter trajectories influence the convergence of our LaDA method.

Figure 6 presents results over an assimilation time interval of 330 d for the three cases of directed drifter deployment described above. Also included for comparison is the case of uniformly released  $3 \times 3$  drifters and a case without data assimilation. Results for the kinetic energy norm shown in Fig. 6a indicate that in all three cases in which a directed drifter deployment was used, the convergence is initially much faster than in the case of uniformly released drifters. After a period of about 200 d, however, the vortex centre launch shows a gradual rise in the norm of the kinetic

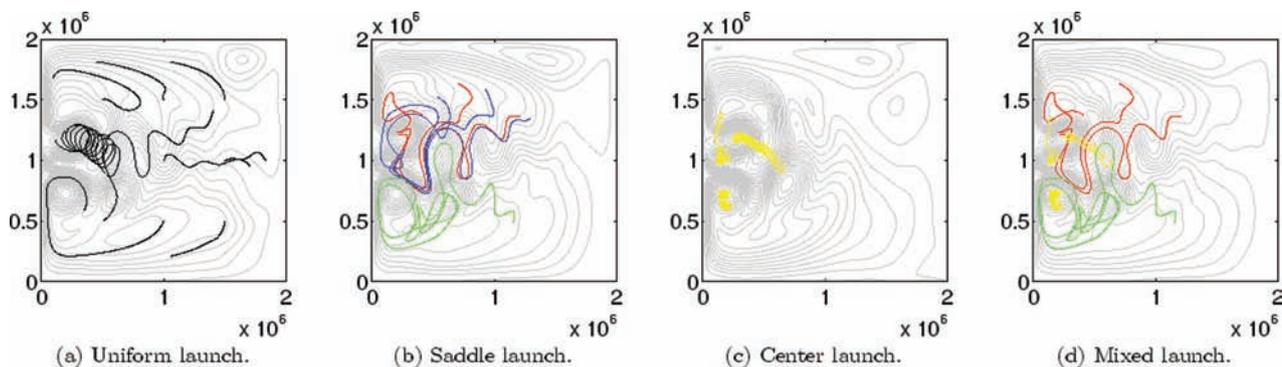


Fig. 5. Contour plots of the assimilated mean height (streamline) field at 300 d with corresponding drifter trajectories obtained by integrating over a time interval of 300 d. Axes represent length in metres.

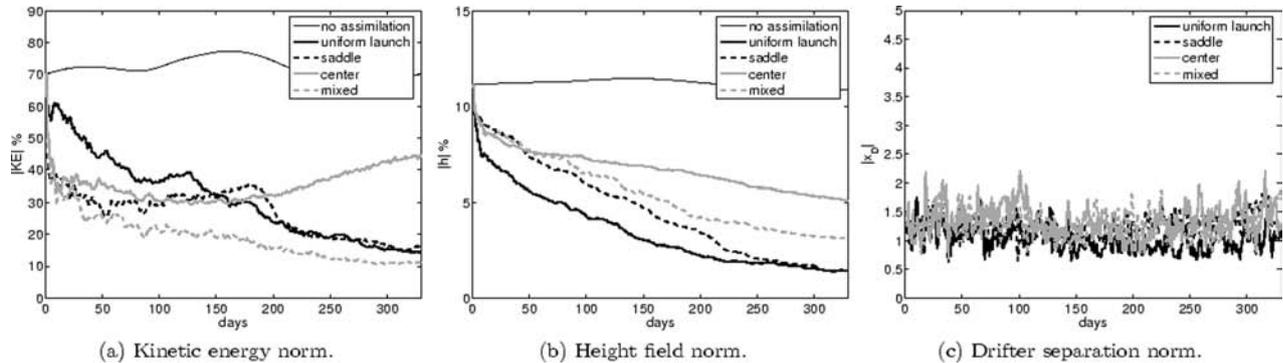


Fig. 6. Effect of initial drifter deployment on the convergence of the error norms.

energy and reveals the poorest performance of the four assimilation experiments considered. The cause of this gradual error increase is discussed later in this section. The case in which DHTs are targeted, reveals the same level of convergence as the uniform case at the end of 330 d. However, by far the best convergence is obtained with the mixed launch. Results for the convergence of the height norm presented in Fig. 6b reveal a different trend. In this case, the uniformly released drifters produce the fastest convergence with the saddle launch eventually resulting in similar norms after 330 d. The mixed launch is seen to produce a slightly poorer convergence but, as in the kinetic energy norms, the poorest results are obtained with the vortex centres. In all the cases considered above, the assimilated drifter trajectories are seen to track the true drifters as indicated by the non-dimensionalized drifter error norms presented in Fig. 6c.

To understand why the results for the kinetic energy and height field norms have different trends with respect to the various drifter launch sites, a more detailed inspection of the assimilated flow field is required. We begin by considering contour plots of the height field. Figure 7 presents contour plots of the height anomaly at three different times for all four assimilation experiments. Also shown in Fig. 7a are the corresponding plots for the control run (synthetic truth). Overlaid on the plots for the cases with assimilation are the ensemble mean positions of the nine drifters to illustrate how the flow field is updated by the motion of the drifters. Results for the case with uniformly released drifters are shown in Fig. 7b. The gradual convergence of the height field to that of the control run can be seen from the increasing coverage of the flow with negatively valued contours of the height anomaly. Results for the saddle launch are shown in Fig. 7c. At earlier times, we observe the corrections are confined closer to the western boundary where most of the drifters are initially released. As the drifters disperse, they provide more widespread coverage of the flow. Therefore, at later times, the height anomaly field is qualitatively similar to the uniform case and is consistent with the global error norms presented in Fig. 6b. Comparing these results with the case corresponding to the vortex centre launch presented in Fig. 7d reveals some stark differences. In particular, the drifters remain trapped within the eddies which themselves remain confined to the western bound-

ary region of the flow. This lack of dispersion in the drifters results in the height field being updated only locally. Consequently, the flow close to the eastern boundary remains with unrealistically high levels of the height anomaly field. This accounts for the poorer global trend in the norm of the height field observed in Fig. 6b. Finally, the results for the mixed launch reveal a trend that is more characteristic of the saddle case although the height anomalies are less well represented close to the upper boundary after 300 d.

To provide a clearer representation of how errors are reduced in the four different cases, we have computed contour plots of the root mean square errors of the height field. The results are presented in a similar arrangement to that used in Fig. 7. Considering the case of uniformly released drifters as shown in Fig. 8a, we observe the errors are initially reduced within the central part of the flow domain and far away from the boundaries. As time progresses and drifters disperse, the flow is gradually adjusted in remote regions including the boundaries. By 300 d, the error has been almost entirely eliminated. In comparison, the saddle launch presented in Fig. 8b reveals the errors are initially reduced closer to the western boundary. As the drifters rapidly disperse following their encounter with the respective DHTs, the flow is adjusted until the errors have been uniformly reduced throughout the flow. The poorer convergence with the case of vortex centres is very clearly illustrated in Fig. 8c. The results clearly show that a significant part of the flow has not been adjusted by the data assimilation method. Finally, the mixed case reveals results that are most similar to the saddle launch albeit that, after 300 d, the overall errors are slightly higher. The results presented here are fully consistent with those presented in Fig. 7 and clearly illustrate the impact that different drifter deployment strategies have on the assimilated flow field.

We will now consider how errors associated with the velocity field are influenced by assimilating different drifter data. In particular, we want to identify why the errors observed in Fig. 6a for the norms of the kinetic energy have a different trend from the norms of the height field presented in Fig. 6b. Contour plots of the root mean square errors of the kinetic energy field are presented in Fig. 9. A feature which immediately stands out from the plots presented here is that, even at the very early stages of the

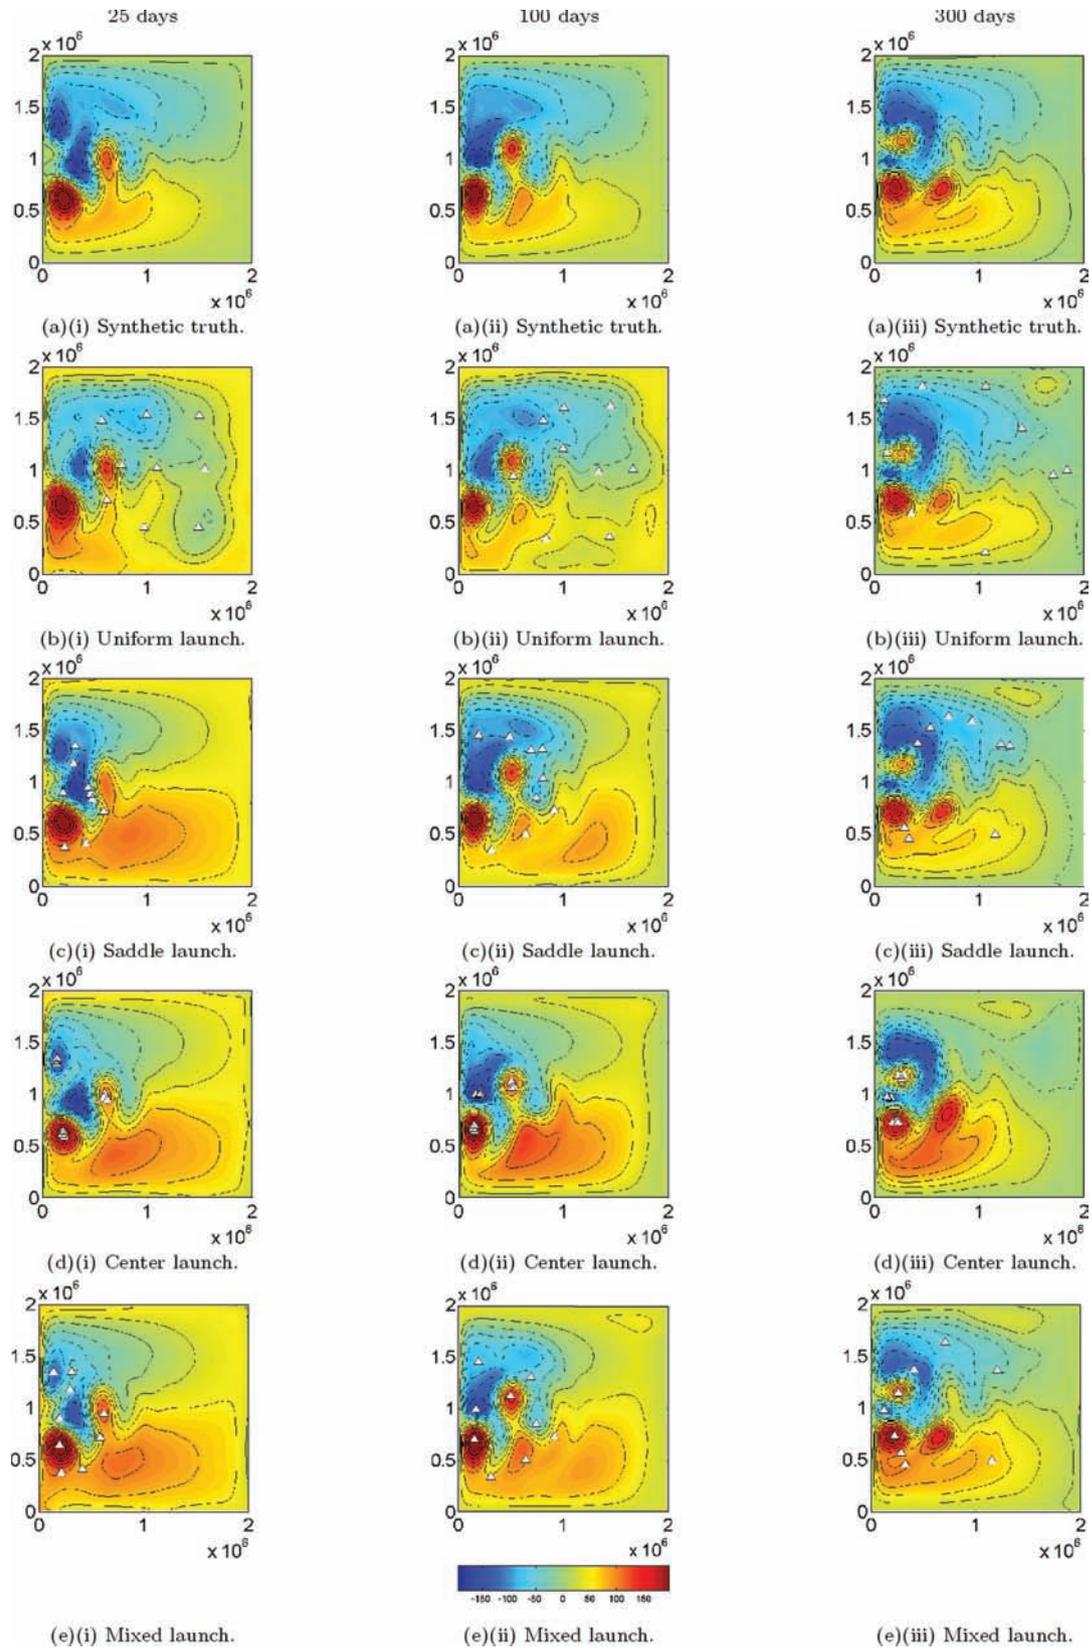


Fig. 7. Contour plots of true and assimilated height anomaly fields ( $h - H'_0$ ) for different deployment strategies with corresponding positions of the ensemble mean drifter locations. Axes represent length in metres.

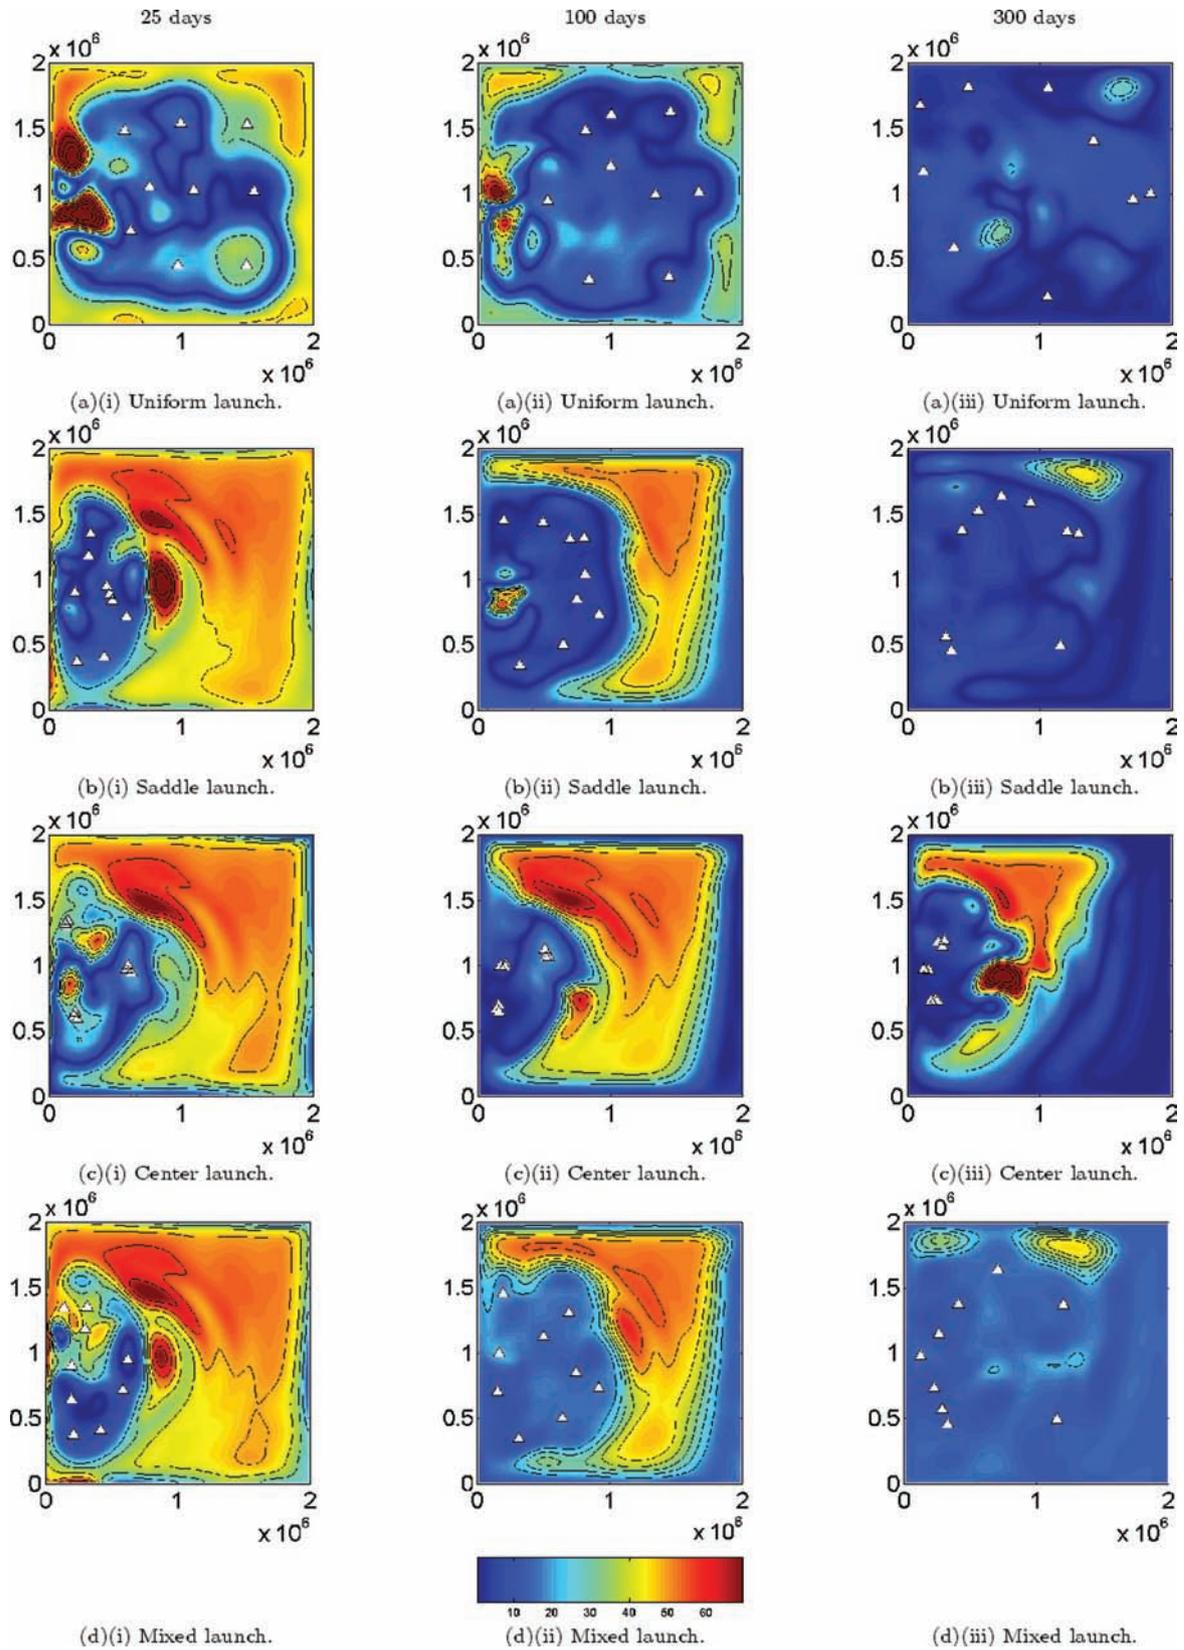


Fig. 8. Contour plots of the root mean square error of the height field for different deployment strategies with corresponding positions of the ensemble mean drifter locations. Axes represent length in metres.

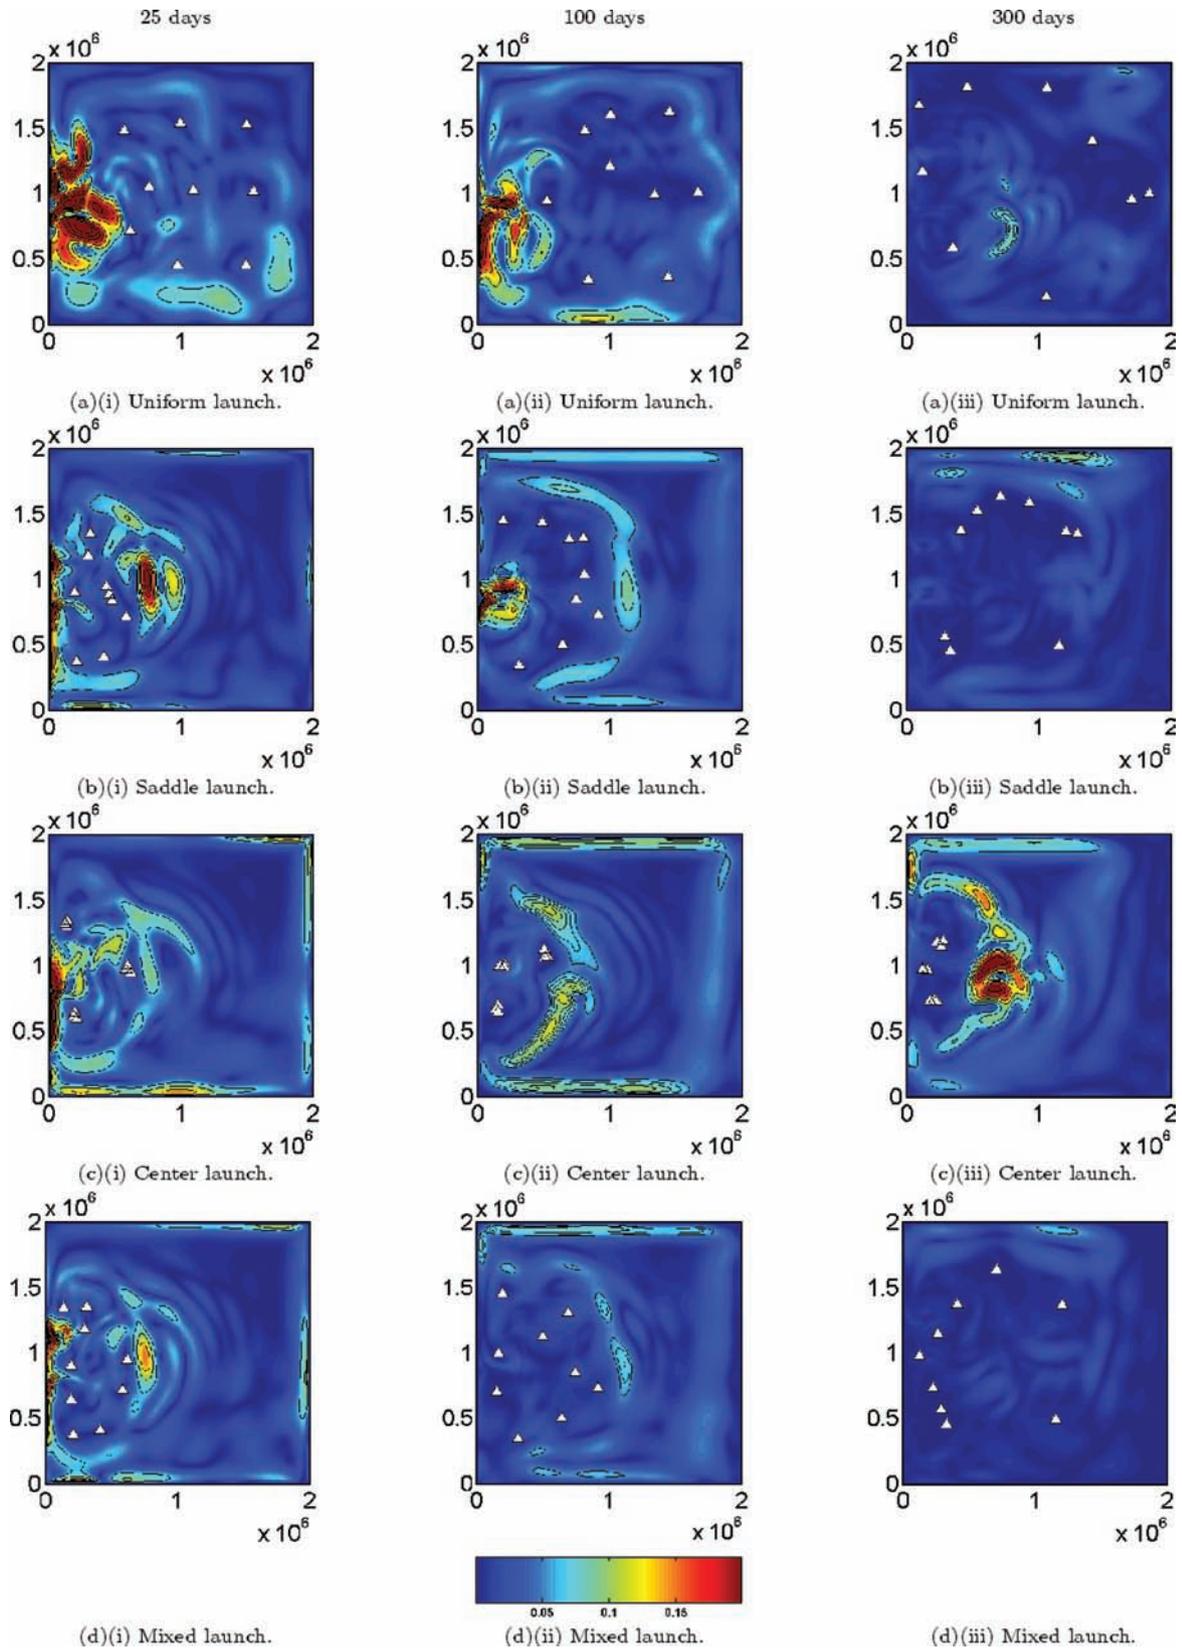


Fig. 9. Contour plots of the root mean square error of the kinetic energy field for different deployment strategies with corresponding positions of the ensemble mean drifter locations. Axes represent length in metres.

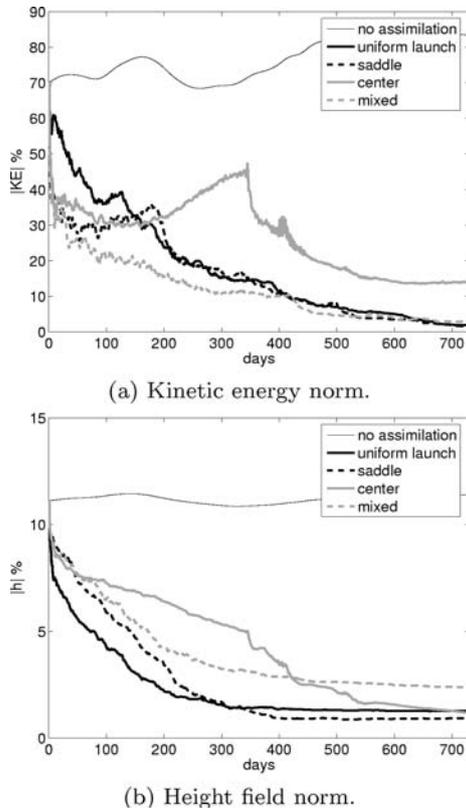


Fig. 10. Effect of initial drifter deployment on the convergence of the error norms for longer forecasts.

data assimilation, the errors are primarily localized close to the separation point of the western boundary currents. This region of the flow is highly active containing the formation of vortices from the unsteady meandering jet. Therefore, unlike the errors for the height field, the errors for the velocity are localized to energetic regions of the flow. Recognizing this key point helps to

explain the differences observed in the behaviour of the norms for these two quantities. Now focusing on the case of the uniform drifter launch shown in Fig. 9a, we observe that the drifters are initially placed in the open ocean far away from the energetic western boundary eddies. This results in a slight delay in the correction of this part of the flow, as shown in Fig. 9a(i) and (ii), and explains the slower initial convergence observed in Fig. 6a relative to the three directed drifter deployment cases. The saddle launch on the other hand, as with the centre and mixed launches, show that, after 25 d, a significant percentage of the error has been removed [see Fig. 9b(i), c(i) and d(i)]. In all three cases, the placement of the drifters closer to the western boundary results in improved convergence at earlier times. For longer times, the dispersion of the drifters in the saddle case helps to eliminate almost all of the errors. The centre launch on the other hand degrades after a sufficiently long time. The problem is attributed to an eddy that forms in the region where the maximum errors are seen in Fig. 9c(iii) [see also Fig. 7d(iii)]. Without sufficient drifter dispersion, the eddies containing the drifters are well represented while other energetic eddies remote from the drifters can not be correctly forecast. Best results for the kinetic energy errors are, however, predicted with the mixed launch as shown in Fig. 9d. In this case, drifters released within the eddies help to suppress errors within these energetic regions of the flow whilst errors in remote regions can be corrected by the dispersion of the remaining drifters. Therefore, while the saddle launch proves to be optimal in removing errors in the height field, the mixed launch is most effective at removing errors in the velocity field.

In all the results presented thus far, the centre launch has produced the poorest convergence history in the norms of the flow field. However, if we consider what happens over longer forecasts, we observe a sudden change in the trends of the error norms. This is vividly clear in the results presented in Fig. 10 in which a sudden and sharp fall in the norms is seen for both the

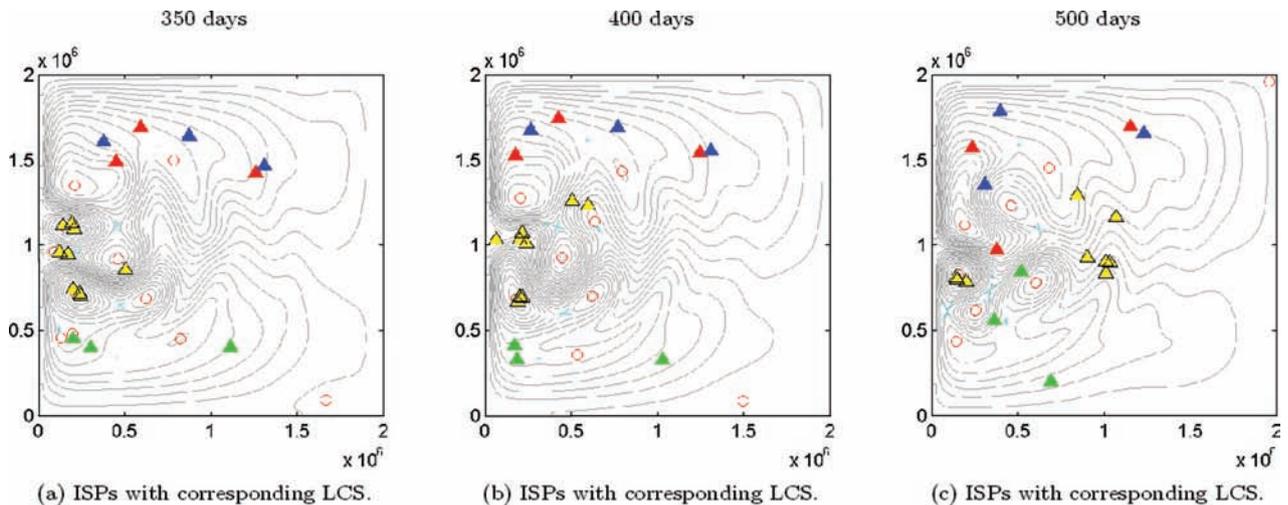


Fig. 11. Contour plots of the height (streamline) field of the synthetic truth with corresponding positions of the drifters to illustrate how bifurcations in the flow affect their evolution. Axes represent length in metres.

height and the velocity fields at around 350 d. After 730 d, despite the drastic reduction in the error of the velocity field, the errors are still substantially higher than the other three cases. The errors in the height field on the other hand converge to a point where they are of the same level as the uniform launch case. We have already observed that errors in the height field tend to be more global than errors in the velocity field. The trends observed in Fig. 10b for the height field, therefore, suggest a mechanism whereby the drifters within the vortex centres must have undergone rapid dispersion. This is indeed the case as evidenced from plots of the drifter locations at later times shown in Fig. 11. We observe that at 350 d, one of the drifters has escaped from the vortex within which it was initially released. The mechanism responsible for the escape of the drifter is triggered by a bifurcation of the vortex centre with a saddle point leading to the annihilation of the vortex. At 400 d, a similar fate befalls the drifters trapped within the second vortex. Therefore, by 500 d, the drifters have migrated eastwards assisted by the meandering jet. This expulsion of the drifters to the eastern part of the flow allows errors in this region of the flow to be corrected which could not be accomplished whilst the drifters were trapped within the vortices. These results show that forecasts made over times that are longer than the Eulerian flow time-scales are determined by transformations within the geometry of the flow due to finitely lived structures. Forecasts made from Lagrangian data over longer times will, therefore, be strikingly different from shorter forecasts.

## 5. Conclusions

We have presented a thorough investigation to demonstrate how judiciously chosen drifter deployment sites can have a major impact on the convergence of a LaDA method that employs an augmented state vector together with an EnKF formulation. This is in part because the method we have developed corrects the flow field through correlations obtained from a covariance matrix of the augmented state vector. Since a correlation function is needed to suppress undesirable noisy correlations between drifters and remote regions of the flow, only the local flow field can be modified by a given drifter. Such an approach is becoming standard in almost all operational forms of the EnKF. Consequently, the dispersion characteristics of Lagrangian drifter trajectories, which are governed by the underlying flow geometry, will strongly influence the performance of the method.

We have attempted to quantify the dependence of the method on these drifter trajectories by extracting LCS of the underlying flow field. Such structures uncover the underlying flow geometry that directly control the motion of Lagrangian drifters. By using these structures together with a directed drifter deployment strategy that targets vortex centres or DHTs of the flow, we were able to realize drifter trajectories with contrasting dispersion characteristics. In this study, we compared four different launch sites; a uniform drifter deployment within the ocean basin, a saddle

launch strategy, a vortex launch strategy, and a mixed combination of saddle and vortex centre launches. In all cases, we have used a total of nine drifters. The different cases were tested on a twin experiment data assimilation configuration in which the model forecast was generated by perturbing the initial height field of the control run (synthetic truth) simulation. Our results showed that the convergence of the height and velocity fields produced by the different launch strategies had strikingly different time histories. In particular, the mixed launch produced the best convergence for the velocity field whereas the uniform and saddle launches were best at minimizing the errors in the height field. These contrasting behaviours are linked to the different nature of the errors in the two fields. In general, errors in the height field were associated with an elevated water depth on average over the entire ocean basin. This was a direct manifestation of the errors introduced in the initial conditions and lead to a globally distributed error in the height field. The errors in the velocity field, on the other hand, are directly correlated to the location of the western boundary currents, the meandering mid-latitude jet, and the Sverdrup gyres and are more local in nature. These results suggest that good dispersion characteristics for the drifters are needed to remove global errors while local errors are best removed by targeting specific energetic regions of the flow. We have also found that bifurcations of coherent flow structures in the underlying flow field can trigger unexpected rapid dispersion of drifters initially released within such coherent vortices. Forecasts made over longer time-scales can, therefore, produce results that are markedly different from those obtained with shorter forecasts.

The results presented in this study have been obtained in the context of a model problem based on a twin-experiment configuration in which the true flow is known a priori. This simplifying assumption was used to allow us to extract LCS associated with the true flow field which ultimately determine the dispersion characteristics of the drifters. For the results obtained in this study to be applicable in practice, this simplifying assumption needs to be relaxed. In particular, LCS of the forecast should be used to determine the optimum launch sites of the drifters. In addition, most operational models that are used in forecasting centres are three-dimensional and exhibit a greater degree of variability between the forecast and the ‘truth’. While further work is needed to address these issues, we point out that some recent results by Lermusiaux et al. (2005, 2006) indicate that robust LCS can persist even in the presence of significant uncertainties in the flow field. These results provide a positive sign that the ideas presented in this work may readily generalize to more complex flow scenarios in which the truth is not known a priori.

## 6. Acknowledgments

HS and CKRTJ were supported by the Office of Naval Research Grants N00014-93-1-0691 and N00014-03-1-0174; KI was

supported by the Office of Naval Research Grant N00014-04-1-0191. The authors would like to thank Dr. Pierre F.J. Lermusiaux and an anonymous referee for their valuable comments.

## References

- Aref, H. 1984. Stirring by chaotic advection. *J. Fluid Mech.* **143**, 1–21.
- Aref, H. and El Naschie, M. S. 1994. *Chaos Applied to Fluid Mixing*, Pergamon, New York.
- Bower, A. S. 1991. A simple kinematic mechanism for mixing fluid parcels across a meandering jet. *J. Phys. Oceanogr.* **21**, 173–180.
- Carter, E. F. 1989. Assimilation of Lagrangian data into numerical models. *Dyn. Atmos. Oceans* **13**, 335–348.
- Coulliette, C. and Wiggins, S. 2000. Intergyre transport in a wind-driven, quasi-geostrophic double gyre: an application of lobe dynamics. *Nonlinear Process. Geophys.* **7**, 59–85.
- Cushman-Roisin, B. 1994. *Introduction to Geophysical Fluid Dynamics*, Prentice Hall, Englewood Cliffs.
- Evensen, G. 1994. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.* **99**(C5), 10143–10162.
- Evensen, G. 2003. The Ensemble Kalman Filter: theoretical formulation and practical implementation. *Ocean Dyn.*, **53**, 343–367.
- Gaspari, G. and Cohn, S. E. 1999. Construction of correlation functions in two and three dimensions. *Q. J. R. Meteor. Soc.* **125**, 723–757.
- Gent, P. R. 1993. The energetically consistent shallow-water equations. *J. Atmos. Sci.* **50**, 1323–1325.
- Haller, G. and Poje, A. C. 1998. Finite time mixing in aperiodic flows. *Physica D* **119**, 352–380.
- Haller, G. 2001. Distinguished material surfaces and coherent structures in 3D fluid flows. *Physica D* **149**, 248–277.
- Haller, G. 2002. Lagrangian coherent structures from approximate velocity data. *Phys. Fluids A* **14**, 1851–1861.
- Hamill, T. M., Whitaker, J. S. and Snyder, C. 2001. Distance-dependent filtering of background error covariance estimates in an ensemble Kalman filter. *Mon. Wea. Rev.* **129**, 2776–2790.
- Ide, K., Kuznetsov, L. and Jones, C. K. R. T. 2002. Lagrangian data assimilation for point vortex systems. *J. Turbul.* **3**, 053.
- Jones, C. K. R. T. and Winkler, S. 2002. Invariant manifolds and Lagrangian dynamics in the ocean and atmosphere. In: *Handbook of Dynamical Systems*, Volume 2 (eds B. Fiedler). Elsevier, Amsterdam, 55–92.
- Kamachi, M. and O'Brien, J. J. 1995. Continuous assimilation of drifting buoy trajectory into an equatorial Pacific Ocean model. *J. Mar. Syst.* **6**, 159–178.
- Kuznetsov, L., Toner, M., Kirwan, A. D., Jones, C. K. R. T., Kantha, L. H., and co-authors. 2002. The Loop Current and adjacent rings delineated by Lagrangian analysis of near-surface flow. *J. Mar. Res.* **60**, 405–429.
- Kuznetsov, L., Ide, K. and Jones, C. K. R. T. 2003. A Method for Assimilation of Lagrangian Data. *Mon. Wea. Rev.* **131**, 2247–2260.
- Lermusiaux, P. F. J. and Lekien, F. 2005. Dynamics and Lagrangian Coherent Structures in the Ocean and their Uncertainty. in: *Extended Abstract in Dynamical System Methods in Fluid Dynamics, Oberwolfach Workshop* (eds E. Marsden Jerrold and Jurgen Scheurle), Mathematisches Forschungsinstitut Oberwolfach, July 31st–August 6th, 2005, Germany, 19–20.
- Lermusiaux, P. F. J., Chiu, C.-S., Gawarkiewicz, G. G., Abbot, P., Robinson, A. R., and co-authors. 2006. Quantifying uncertainties in ocean predictions. *Oceanography*. **19**, 92–105.
- Malhotra, N. and Wiggins, S. 1998. Geometric structures, lobe dynamics, and Lagrangian transport in flows with aperiodic time-dependence, with applications to Rossby wave flow. *J. Nonlinear Sci.* **8**, 401–456.
- Mancho, A. M., Small, D. and Wiggins, S. 2004. Computation of hyperbolic trajectories and their stable and unstable manifolds for oceanographic flows represented as data sets. *Nonlinear Process. Geophys.* **11**, 17–33.
- Molcard, A., Piterbarg, L.I., Griffa, A., Özgökmen, T. M. and Mariano, A. J. 2003. Assimilation of drifter observations for the reconstruction of the Eulerian circulation field. *J. Geophys. Res.* **108**(C3), 1 1–1 21.
- Molcard, A., Poje, A. C. and Özgökmen, T. M. 2006. Directed drifter launch strategies for Lagrangian data assimilation using hyperbolic trajectories. *Ocean Model.* **12**, 237–267.
- Özgökmen, T. M., Griffa, A. and Mariano, A. J. 2000. On the Predictability of Lagrangian Trajectories in the Ocean. *J. Atmos. Oceanic Technol.* **17**, 366–383.
- Özgökmen, T. M., Molcard, A., Chin, T. M., Piterbarg, L. I. and Griffa, A. 2003. Assimilation of drifter observations in primitive equation models of midlatitude ocean circulation. *J. Geophys. Res.* **108**(C7), 31 1–31 17.
- Pedlosky, J. 1986. *Geophysical Fluid Dynamics*, 2d Edition. Springer-Verlag, New York.
- Poje, A. and Haller, G. 1999. Geometry of cross-stream mixing in a double-gyre ocean model. *J. Phys. Oceanogr.* **29**, 1649–1665.
- Poje, A. C., Toner, M., Kirwan, A. D. and Jones, C. K. R. T. 2002. Drifter launch strategies based on Lagrangian templates. *J. Phys. Oceanogr.* **32**, 1855–1869.
- Salman, H., Kuznetov, L., Jones, C. K. R. T. and Ide, K. 2006. A method for assimilating Lagrangian data into a shallow-water equation ocean model. *Mon. Wea. Rev.* **134**, 1081–1101.
- Salman, H., Hesthaven, J. S., Warburton, T. and Haller, G. 2007. Predicting transport by Lagrangian coherent structures with a high order method. *Theor. Comp. Fluid Dyn.* **21**, 39–58.
- Schär, C. and Smith, R. B. 1993. Shallow-water flow past isolated topography. Part I: vorticity production and wake formation. *J. Atmos. Sci.* **50**, 1373–1400.
- Shchepetkin, A. F. and O'Brien, J. J. 1996. A method for assimilating Lagrangian data into a shallow-water equation ocean model. *Mon. Wea. Rev.* **124**, 1285–1300.
- Toner, M. and Poje, A. C. 2004. Lagrangian velocity statistics of directed launch strategies in a Gulf of Mexico model. *Nonlinear Process. Geophys.* **11**, 35–46.
- Wiggins, S. 2005. The dynamical systems approach to Lagrangian transport in oceanic flows. *Ann. Rev. Fluid Mech.* **37**, 295–328.