

Enhanced Ocean Predictability through Optimal Observing Strategies

PI: A. D. Kirwan, Jr.
University of Delaware
Robinson Hall
Newark, DE 19716
phone: (302) 831-2977 fax: (302) 831-6521 email: adk@udel.edu

CO-PI: Bruce L. Lipphardt, Jr.
University of Delaware
Robinson Hall
Newark, DE 19716
phone: (302) 831-6836 fax: (302) 831-6521 email: brucel@udel.edu

CO-PI: Helga S. Huntley
University of Delaware
Robinson Hall
Newark, DE 19716
phone: (302) 831-1175 fax: (302) 831-6521 email: helgah@udel.edu

Award Number: N00014-07-1-0730

LONG-TERM GOALS

The effort funded under this project is part of our study of the dynamics determining transport characteristics of the ocean and their predictability. We are interested in identifying critical structures in the flow field, assessing their variability, and improving their predictability. This is a necessary element for making progress towards the overarching goal of forecasting the origin and fate of fluid particles in the ocean with an estimate of the forecast's uncertainty. Applications of this work vary widely, from biology to climate dynamics to specific Navy operational interests such as extended AUV missions and targeted sensor deployment.

OBJECTIVES

The main objective pursued in FY09 was the application of Lagrangian analysis techniques to an implementation of the Harvard Ocean Prediction System (HOPS) model for Ano Nuevo to identify strategies for effective sensor deployment, optimizing a variety of potential mission goals.

APPROACH

Often deployments of drifting or remotely steered sensors are designed to meet a specific observing target such as a particular geographic location or a length of time to stay within an area. By accounting for the ocean currents, such targets can be met more effectively and more efficiently. Models typically supply Eulerian ocean forecasts; yet the goals to be met are generally of a

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2009		2. REPORT TYPE		3. DATES COVERED 00-00-2009 to 00-00-2009	
4. TITLE AND SUBTITLE Enhanced Ocean Predictability through Optimal Observing Strategies				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Delaware, Robinson Hall, Newark, DE, 19716				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 13	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

fundamentally Lagrangian nature, i.e., they are related to the origin and fate of fluid patches. Consequently, it is appropriate to use a Lagrangian toolbox to analyze model output and generate products that are directly related to the Lagrangian goals. Over the years, we have developed a series of such analysis tools; in the effort described here, we focus primarily on their application to a particular region of the world of interest to the Navy. We also investigate the relative effectiveness of Eulerian- versus Lagrangian-based strategies to justify the additional computational cost of the latter.

Which type of Lagrangian analysis is appropriate depends on the particular question asked. However, they all rely on a set of calculated gridded trajectories. These can then be used to derive such products as:

Trajectory predictions – The simplest tool is a straight-forward answer to the question where an ocean particle – or a drifter with zero relative drag – is going to go. If the launch locations are pre-determined, this information can be useful for planning pick-up locations or length of desired observing window, for example.

Residence time maps – For optimizing time on station, it is important to know how long particles are expected to stay within a given area. Additionally, these maps provide insight into the location of sharp boundaries between patches of vastly differing behavior, which are evidence of dynamical structures directing the transport in the ocean.

Escape fate and origin maps – By dividing the boundary of the area of interest into individual color-coded segments, it is possible to partition the interior into regions of particles sharing a common fate or origin location along the boundary. This can be used for determining possible boundary launch locations for exploring a particular patch in the interior or planning recovery of launched sensors. These maps also provide limited information about the coverage of the corresponding trajectories.

Region-of-Interest launch (RoI) maps – If a specific region is targeted for observation within a larger area of interest, or if the possible launch positions do not necessarily include the geographical location to be observed, the RoI maps can be used to determine where a launch is advisable for a sensor that will pass through the desired target region. RoI maps as well display some of the intricate structure forming the backbone of the transport dynamics.

Relative dispersion maps – Relative dispersion has been defined in various ways with small variations in the computational scheme. However, in all cases, the result is an estimate of the stretching/compression behavior of particle trajectories. Ridges in these maps are an indication of the location of attracting or repelling manifolds in the dynamic system. They also divide the region into areas of high and low predictability, making them a valuable companion product for each of the other tools discussed here.

For each of these products, it is important to keep in mind that they are predictions based on the ocean model. Uncertainties in the Eulerian forecasts translate into uncertainties in the Lagrangian forecasts as well. Here, we do not discuss the assessment of the uncertainties specifically; see last year's Annual Report. However, we do explore spatial and temporal sensitivities of the results.

The three co-PIs each contributed to all aspects of this project in a close collaboration.

WORK COMPLETED

During FY09, our effort was focused on applying the Lagrangian tools discussed above to a particular ocean model to demonstrate the feasibility and potential uses and to study what the analysis reveals about the dynamics of the chosen region of the ocean. More specifically, we

- processed HOPS model data for Ano Nuevo;
- computed an archive of gridded trajectories;
- analyzed synoptic Lagrangian maps, including residence time and escape fate maps;
- created a sample of Region-of-Interest maps;
- studied relative dispersion maps;
- explored the temporal and spatial sensitivities of deployment strategy recommendations based on the available Lagrangian products; and
- assessed the competitiveness of Eulerian products for similar recommendations.

RESULTS

The model underlying the analyses described here is the HOPS (Harvard Ocean Prediction System) model with 1.5 km horizontal resolution, implemented for Ano Nuevo north of Monterey Bay run on a rotated grid parallel to the California coast. There are multiple layers, although we have restricted our attention to the two-dimensional dynamics within a single layer. The model contains tidal components and is data-assimilating. For details, see for example *Robinson (1999)*. HOPS is accepted as a state-of-the-art ocean model (*Kantha and Clayson, 2000*). We thank Pierre Lermusiaux from MIT for providing the model data.

Traditionally, it has often been assumed that average velocity fields provide an adequate picture of the transport properties of the ocean. Figure 1 illustrates why such an Eulerian view is not enough. The vector field shows the model velocities averaged over ten days. The field is smooth and shows the northward coastal current as well as some off-shore eddy activity. With a stationary velocity field, it is straight-forward to mentally integrate to estimate particle paths, indicated in blue in the figure. Yet the actual trajectories from the time-varying field – shown in red – do not agree very well with these approximations. Neither trajectory length, nor exit location, nor even initial direction are consistently well predicted from the Eulerian picture alone. The accurate description of all of these important characteristics of the flow rely on the Lagrangian integration of the instantaneous velocities.




Figure 1: Comparison of trajectories from mean velocities (blue) and those from time-varying velocities (red) within the analyzed model domain at 30m depth. Launch positions are indicated as green diamonds. Trajectories were integrated from 07/28/2006 00:00 through 08/07/2006 00:00 or until they left the domain. The velocity field shown is the mean field over that period. None of the trajectory pairs agree very well. Note differences in initial direction, end location, path length, and smoothness.

For the purposes of designing optimal deployment strategies based on Lagrangian analyses, we have considered three different objectives:

1. Maximize time spent on station.
2. Targeted exit or entry location.
3. Targeted region of interest.

1. Maximizing time spent on station:

Here we have considered a “station” defined as a fairly large area, essentially coincident with the model domain. Of course, a similar analysis can be carried out for any subdomain as well.

An intuitive approach to the optimization problem at hand, given a model forecast for a velocity field, is to steer towards the instantaneous stagnation points (ISPs), i.e., to locations with instantaneous zero velocity. A study of residence time maps, however, shows that one can do much better. Figure 2 exhibits an example. Comparing the placement of “ideal” launch locations based on ISPs versus the strategy relying on maximizing the residence time of fluid particles, we see that they do not agree very well. This is generally true: ISPs are not a good predictor for long residence times due the compound effects of variations over time in the velocity field. In fact, in this example, the longest achievable residence time exceeds the length of the available model data record of 31 days, whereas particles launched at the stagnation points stay in the domain at most just over 3 days. 47% of the domain has residence times longer than that, so that a random choice of launch location is likely to perform better than an ISP: The mean residence time over ISP launches is less than a day, that over the entire domain more than four days. Other examples show slightly better results for ISP launches, but we have never found them to be good estimates of maximal residence times. Note that there is no significant difference between residence times associated with elliptic and hyperbolic stagnation points.




Figure 2: Residence time color map for 7/27/2006 00:00 at 30 m depth. Also shown are instantaneous stagnation points – yellow circles for hyperbolic, green circles for elliptic ISPs. The 23 points on the refined grid with trajectories that remain in the domain for the length of the available model velocity record are indicated by blue diamonds. Naturally, the blue diamonds fall within the dark green regions of high residence times; the ISPs, on the other hand, are all in regions of at best mediocre residence times.

Unfortunately, in practice, designing an optimal deployment strategy for this objective is not as easy as the above analysis suggests. Aside from the problem of model errors, which we ignore here but recognize as an important limitation, residence time can be highly sensitive to perturbations in space or time. Figure 3a shows a residence time map with five trajectories overlaid. The launch locations are just 500 m between neighbors, a distance that is easily within the practical deployment margin of error and beyond the accuracy that can be expected from the model. The range of residence times over the total of 2 km is 3 to 11 days. That is the difference between an average and an excellent on-station time. The sensitivity to launch time is similarly great, as demonstrated in Figure 3b. Here particles were launched hourly over one day from the same location. Residence times for the resulting trajectories range from 0.5 to 16.5 days.




Figure 3: (a) Residence time color map for 7/28/2006 00:00 at 30 m depth. Superimposed are five trajectories launched on a line with 500 m spacing near a gradient in residence time, each colored a different shade of grey with black markers every 12 hours. While the tracks remain parallel for the first 12 hours, they quickly diverge thereafter, resulting in vastly different fates. (b) 24 trajectories, launched hourly starting at 7/28/2006 00:00 from the same location at 30 m depth. The endpoint of each trajectory is colored according to its residence time. The launch location is indicated by a magenta diamond. About half the particles leave the domain at the northern end, while the rest exit at the southern end. The latter group includes both those with shortest and those with longest residence times.

In light of these sensitivities, it is advisable to target deployment locations that are not necessarily coincident with the longest possible residence times but those that lie within persistent and sizable regions of high residence time.

One final challenge with using residence time maps is the need for velocity forecasts for at least the length of the longest trajectory in the area or the length of the desired on-station time. Velocity forecasts, however, are currently judged to be reliable only out to at most 72 hours. It is still possible to construct residence time maps with this limited data, but much of the detailed structure of the field is lost. Figure 4 shows the degradation as the forecast length shrinks from 15 to 5 and 3 days.




Figure 4: Residence time maps created by restricting the available velocity forecast to 15 days (cf. Figure 2), 5 days, and 3 days, respectively. Note the change in color scale. Increasingly more detail is lost. Instantaneous stagnation points are indicated as black circles to demonstrate that even with the shortest forecast window, the map still clearly shows that ISPs are not ideal launch locations for maximizing residence time

2. Targeted exit or entry location:

For this objective, the task is to find possible launch locations inside the domain, from where the drifters will traverse the domain and exit at a desired location – or to determine what portion of the domain can be explored given a particular launch location.

For the escape fate problem, the boundary was divided into 21 color-coded segments, each 10.2 km long. Then the interior of the domain is colored according to exit location of a particle launched at each particular location. Large regions of the domain are divided into coherent patches of uniform escape fate, with neighboring patches exiting through neighboring boundary segments. However, a portion of the field displays extremely high variability, where neighboring launch locations may send the particles out at opposite ends of the domain. See Figure 5.




Figure 5: An escape fate map for launches from 8/1/2006 12:00 at 30 m depth. The boundary is color-coded in 10.2 km long segments. Four representative trajectories are also plotted in black. High variability in escape fate can be seen in the center of the domain, while the northwestern part is especially cleanly organized.

The consequence for the deployment strategy hence is that large coherent patches should be targeted for launches, which will allow for greater certainty in escape fate. Note that it is often (though not always) possible to launch particles far away from the desired exit location, which guarantees that a good portion of the domain will be sampled along the way. A conclusion for ocean dynamics is that the fate of particles tends to be well organized generally, but that even fairly smooth velocity fields, like the ones we are dealing with here, can give rise to patches of highly variable behavior in the particle trajectories. More work is needed to identify the determining mechanism for this phenomenon and to study its predictability.

The origin problem addresses temporal rather than spatial variability. As the location is fixed, the timing of the launch will determine the drifter's pathway. Figure 6 displays an example of launches from a particular point on the southeastern boundary of the domain. Synthetic drifters were launched hourly over the course of one day, and a subsample of the resulting tracks is shown. It can be seen that the trajectories remain fairly similar for about half the day, at which point a sharp transition occurs to a single trajectory which uniquely samples a good portion of the northeastern boundary, followed by another group of similar trajectories, each now marked by a loop in the center of the domain. See also Figure 3b for another example.

This result suggests that the transitions in transport characteristics can be very quick, which is not at all apparent from the evolution of the underlying Eulerian velocity field. On the other hand, these transitions separate longer periods of times of fairly constant behavior. For a deployment strategy then, one would want to target such windows of better predictability.




Figure 6: Trajectories launched in 4-hour intervals from the same location on the boundary are plotted on top of the corresponding residence time maps. Note how the behavior transitions midway through this time series. The trajectory launched at 7/28/2006 12:00 also uniquely has a residence time of 21.4 days, versus about 7 – 7.5 days for all the other launches.

3. Targeted region of interest:

This last mission is closely related to the previous one, except that we are not interested in targeting a particular exit location but a smaller subregion which is to be explored. In the example shown in Figure 7, two RoIs are chosen. The window of exploration is set to a week. The possible launch locations from where a sensor would reach the target location in the target window are indicated in yellow and red, respectively, for the two regions. Orange indicates a launch location resulting in a sampling of both RoIs. Over time, the areas of possible launch evolve smoothly, sometimes overlapping, sometimes separating completely, sometimes intertwining.




Figure 6: Map of possible launch locations at 30 m depth on 7/31/2006 00:00 to reach a desired region of interest in the window 7/31/2006 00:00 through 8/7/2006 00:00. RoIs are indicated as black squares; the domain of all tested launches is marked with a white boundary. Launches reaching the RoI to the northwest are colored yellow, those for the RoI to the southwest are colored red; locations reaching both are colored orange. This example shows intertwining of the two colored regions but fairly little overlap.

Our investigation into the oceanographic aspects of this phenomenon is still somewhat preliminary. As before, we observe sharp spatial and temporal boundaries delineating coherent patches. However, it is not yet clear how these relate to more general transport properties. The RoI maps are highly

dependent, of course, on the chosen region(s) of interest and the target window. It remains to be seen whether the features of the maps are resilient to changes in these parameters, which could be a sign of a more fundamental structure organizing transport in the area.

IMPACT/APPLICATIONS

Here we have described our results in the context of sensor deployment mission, but this type of analysis has broader applicability. For example, residence time of pollutants is important in evaluating the environmental impact of ocean dumping or an oil spill. Trajectory predictions can assist in rescue operations. Escape fate, especially near a coastline, distinguishes between different flow regimes that may determine the water quality in the area.

On a more fundamental level, the synoptic Lagrangian maps of the type we are developing here allow the identification of flow structures that organize the ocean dynamics into distinct transport patterns. The sharp gradients in the Lagrangian maps demonstrate the need to get small scale dynamics right for reliable coarse scale descriptions of the transport. The high temporal sensitivity also confirms how challenging the forecasting problem is. Yet at the same time, the maps indicate where observations can be particularly useful for constraining the range of possible future states.

RELATED PROJECTS

The work reported here is closely related to three other ONR projects by the same principal investigators:

N00014-09-1-0703: How well do blended velocity fields improve the predictions of drifting sensor tracks? – This projects seeks to investigate two different methods of data blending and their effectiveness for improving trajectory predictions. Since trajectory predictions underlie all other Lagrangian analyses, enhancing their accuracy is immediately relevant for all applications of Lagrangian forecasts.

N00014-09-1-0559: Assessing the Lagrangian Predictive Ability of Navy Ocean Models – In order to make forecasts including an uncertainty window, an estimate of the model uncertainty is needed. This work attempts to develop such an estimate and study its characteristics for Lagrangian forecasts in particular by analyzing an ensemble of model runs from an operational Navy ocean model.

N00173-08-1-G009: Prediction of evolving acoustic sensor arrays – This effort is focused on demonstrating how Lagrangian analysis of Navy ocean model predictions can be performed at a Navy operational center and how Lagrangian products can be delivered to fleet operators on scene in near-real time to support tactical decision making.

REFERENCES

Kantha, L. H., and Carol Anne Clayson, *Numerical Models of Oceans and Oceanic Processes*. San Diego, CA: Academic Press, pp. 495-496, 2000.

Robinson, A. R., Forecasting and Simulating Coastal Ocean Processes and Variabilities with the Harvard Ocean Prediction System, *Coastal Ocean Prediction* (C. N. K. Mooers, ed.), AGU Coastal and Estuarine Studies Series, 77-100, American Geophysical Union, 1999.

PUBLICATIONS

Kirwan, A. D., Jr., Dynamics of “critical” trajectories, *Prog. in Oceanogr.*, 70: 448-465, 2006.

Lipphardt, B. L., Jr., D. Small, A.D. Kirwan, Jr., S. Wiggins, K. Ide, C. E. Grosch, and J. D. Paduan, Synoptic Lagrangian maps: Application to surface transport in Monterey Bay, *J. Mar. Res.*, 64: 221-247, 2006.

Lipphardt, B. L., Jr., A. Poje, A. D. Kirwan, Jr., L. Kantha, and M. Zweng, Death of three Loop Current rings, *J. Mar. Res.* 66: 25-60, 2008.