

RESEARCH ARTICLE

Energy-optimal path planning in the coastal ocean

10.1002/2016JC012231

Deepak N. Subramani¹ , Patrick J. Haley Jr.¹, and Pierre F. J. Lermusiaux¹ 

¹Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA

Key Points:

- Fundamental energy-optimal path planning, rigorously utilizing realistic multiscale ocean forecasts
- Using stochastic optimization showcased energy-optimal paths in data-assimilative simulations
- Analyzed effects of tides, wind events, jets, shelfbreak front, and local circulations on the paths

Correspondence to:

P. F. J. Lermusiaux,
pierrel@mit.edu

Citation:

Subramani, D. N., P. J. Haley, Jr., and P. F. J. Lermusiaux (2017), Energy-optimal path planning in the coastal ocean, *J. Geophys. Res. Oceans*, 122, 3981–4003, doi:10.1002/2016JC012231.

Received 9 AUG 2016

Accepted 11 APR 2017

Accepted article online 21 APR 2017

Published online 15 MAY 2017

Abstract We integrate data-driven ocean modeling with the stochastic Dynamically Orthogonal (DO) level-set optimization methodology to compute and study energy-optimal paths, speeds, and headings for ocean vehicles in the Middle-Atlantic Bight (MAB) region. We hindcast the energy-optimal paths from among exact time-optimal paths for the period 28 August 2006 to 9 September 2006. To do so, we first obtain a data-assimilative multiscale reanalysis, combining ocean observations with implicit two-way nested multiresolution primitive-equation simulations of the tidal-to-mesoscale dynamics in the region. Second, we solve the reduced-order stochastic DO level-set partial differential equations (PDEs) to compute the joint probability of minimum arrival time, vehicle-speed time series, and total energy utilized. Third, for each arrival time, we select the vehicle-speed time series that minimize the total energy utilization from the marginal probability of vehicle-speed and total energy. The corresponding energy-optimal path and headings are obtained through the exact particle-backtracking equation. Theoretically, the present methodology is PDE-based and provides fundamental energy-optimal predictions without heuristics. Computationally, it is 3–4 orders of magnitude faster than direct Monte Carlo methods. For the missions considered, we analyze the effects of the regional tidal currents, strong wind events, coastal jets, shelfbreak front, and other local circulations on the energy-optimal paths. Results showcase the opportunities for vehicles that intelligently utilize the ocean environment to minimize energy usage, rigorously integrating ocean forecasting with optimal control of autonomous vehicles.

1. Introduction

In recent years, autonomous platforms that can operate in harsh dynamic environments with little or no human intervention are becoming ubiquitous for a variety of ocean exploration missions. For example, propelled Autonomous Underwater Vehicles (AUVs) and gliders have been employed for scientific research, oil and gas discovery, search and rescue, security and surveillance, and monitoring of sensitive coastal ecosystems [Bellingham and Rajan, 2007]. Expanding user bases and mission complexities raise the needs of longer endurance and optimization [Lermusiaux et al., 2016]. The task of providing rules for navigating such autonomous agents in a time-, energy-, data-, and/or safety-optimal fashion is called path planning. A traditional focus of path planning, however, has been on the motion of autonomous agents in static environments [e.g., Latombe, 1991]. In comparison, ocean vehicles are significantly affected by the dynamic ocean motions, including currents and waves. Often, the strong currents are comparable to the speed of propelled AUVs and common currents are 1–3 times that of glider speeds [Rudnick et al., 2004]. This situation provides an opportunity to optimize time and energy requirements by intelligently utilizing the ocean environment, especially for long-endurance autonomy and sustained sampling. For such planning, it is advantageous to leverage modern ocean modeling and forecasting systems. Our first objective here is to demonstrate a novel methodology that rigorously combines ocean field predictions with fundamental optimal planning Partial Differential Equations (PDEs) to compute energy-optimal paths between two locations in the coastal ocean. In Subramani and Lermusiaux [2016], we introduced this PDE-based stochastic Dynamically Orthogonal (DO) level-set optimization and illustrated it for planning paths in canonical flow scenarios including a double-gyre barotropic quasi-geostrophic circulation. In the present work, we focus on energy-optimal path planning for realistic missions in the complex multiscale coastal ocean. Specifically, we plan energy-optimal paths in the Middle-Atlantic Bight (MAB) and shelfbreak front region, using realistic multiscale ocean reanalyses obtained from the Multi-disciplinary Simulation, Estimation, and Assimilation System (MSEAS) [Haley and Lermusiaux, 2010; Haley et al., 2015a]. Our second objective is to illustrate and analyze how energy-optimal paths in the region are affected by wind-driven currents, shelfbreak front features, tidal flows, coastal jets, and other shelf circulations.

Utilizing ocean forecasts to plan optimal paths has been tried with techniques such as Rapidly Exploring Random Trees (RRTs), nonlinear optimization, evolutionary algorithms, potential field methods, and Lagrangian coherent structures based methods with varying degrees of success. *Rao and Williams* [2009] used RRTs with energy-based path cost and heuristics for generating feasible paths of underwater gliders navigating in the southern end of the East Australian Current, off the coast of New South Wales in the Pacific Ocean. They report energy improvements of 2–6% over grid paths when currents are not very strong, and no improvement when currents are very strong. *Alvarez et al.* [2004] used genetic algorithms to plan paths that minimize the energy required to overcome the total drag losses (modeled as a path integral of the cube of vehicle velocity) in a simulation of the Sicily channel in the Mediterranean Sea, using a forecast from the Harvard Ocean Prediction System (basis of MSEAS). *Kruger et al.* [2007] planned navigation paths in a simulation of the Hudson River around Manhattan using nonlinear optimization techniques and a cost function that includes obstacle avoidance, time of travel, target visitation, and the energy used by a time-dependent engine thrust. *Witt and Dunbabin* [2008] combine the use of potential fields for obstacle avoidance with an ad hoc heuristic-based optimization to minimize the energy needed by an AUV to overcome the drag force (quadratic nominal velocity) and acceleration force (rate of change of nominal velocity) in hindcast simulations for Brisbane’s Moreton Bay. Lagrangian Coherent Structures (LCSs) estimated from Monterey Bay HF-radar data were shown to be close to energy-optimal paths computed from heuristic nonlinear programming [*Inanc et al.*, 2005; *Zhang et al.*, 2008]. *Hsieh et al.* [2012] and *Michini et al.* [2014] describe collaborative tracking of LCSs using ocean data for the Santa Barbara Channel along the California coast. For a more general review of energy-optimal path planning, we refer to *Subramani and Lermusiaux* [2016] and *Subramani* [2014].

In what follows, in section 2, we summarize the stochastic DO level-set optimization methodology for planning energy-optimal paths, rigorously utilizing ocean forecasts. In section 3, we describe the regional synoptic multiscale ocean circulation in the MAB region, as analyzed by the multiresolution MSEAS modeling system. We also study the atmospheric and tidal-forcing fields and their effects on the regional currents. In section 4, we provide the results of our path planning. We analyze the effects of wind-driven currents, tidal flows, and shelf circulations on the energy-optimal paths and quantitatively assess the energy savings. We also briefly discuss extensions for energy-optimal adaptive replanning, optimal sampling, and planning under uncertain forecasts. Finally, in section 5, we provide conclusions and future directions.

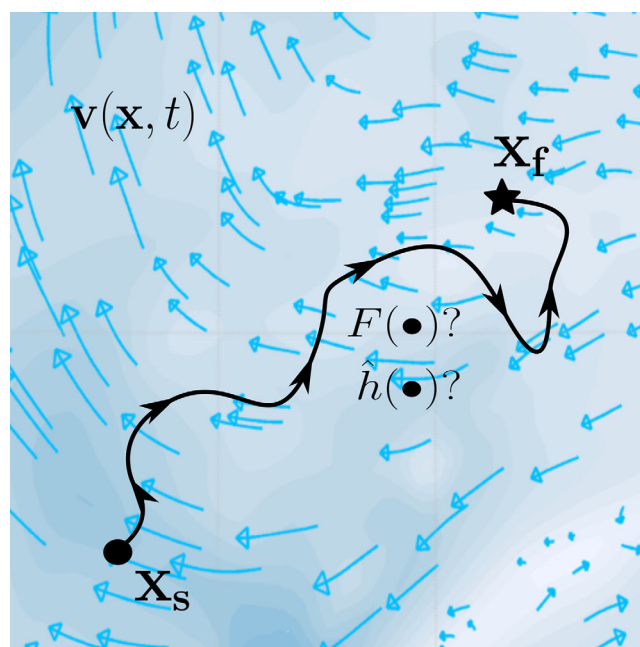


Figure 1. Consider planning the path of a vehicle between \mathbf{x}_s and \mathbf{x}_f in a flow field $\mathbf{v}(\mathbf{x}, t)$. For each arrival time, our goal is to compute a minimal energy path, among the group of time-optimal paths each corresponding to a different vehicle-speed time series. Adapted from *Subramani and Lermusiaux* [2016].

2. Theory and Methodology

We now outline the mathematical theory and solution method for energy-optimal path planning in complex ocean current fields. Details on the methodology including computational costs are provided in Appendix A. The goal is to compute truly energy-optimal paths, vehicle-speed time series, and headings to be followed by a vehicle navigating from a start point to an end point in the presence of a strong, dynamic, and deterministic environmental flow (see Figure 1). Following *Subramani and Lermusiaux* [2016], let us consider a vehicle moving with a to-be-optimized time-dependent speed function $F(\bullet)$ from a start point \mathbf{x}_s to an end point \mathbf{x}_f in the domain $\Omega \subseteq \mathcal{R}^2$ (Figure 1). The motion of the vehicle is affected by the dynamic currents $\mathbf{v}(\mathbf{x}, t) : \Omega \times [0, \infty)$. Let us assume that the vehicle’s power

requirement is a function of its speed, e.g., a power-law dependence on its speed $p(t) = F(t)^{n_p}$, $n_p \geq 0$ [Athans and Falb, 2007]. Thus, the total energy requirement for the vehicle is $E(\bullet) = \int_0^{T_a(\mathbf{x}_f; F(\bullet))} p(t) dt$, where $T_a(\mathbf{x}_f; F(\bullet))$ is the arrival time at the target. For a unique specified speed function $F(t)$, the corresponding energy usage E is minimized if the arrival time is the minimum arrival time at the target \mathbf{x}_f for that $F(t)$, i.e., if $T_a(\mathbf{x}_f; F(t)) = T(\mathbf{x}_f; F(\bullet))$. If $T_a(\mathbf{x}_f; F(t))$ was longer, more energy would be used than needed for that $F(t)$ and the path would not be energy-optimal. Hence, for a specified $F(t)$, the energy-optimal path is the time-optimal path. If one searches among all speed functions, one thus obtains the exact energy-optimal paths and headings, for a range of arrival times. Of course, for some speed functions, the end point cannot be reached, while for some arrival times, there can be multiple optimal paths. The present methodology naturally accounts for all such particulars [Lolla et al., 2014a; Subramani and Lermusiaux, 2016]. We also note that slowing the vehicle-speeds would decrease the instantaneous power requirement, but increase the optimal arrival time. A faster vehicle-speed, on the other hand, would have opposite effects. Changing the vehicle-speeds in accordance with the flow $\mathbf{v}(\mathbf{x}, t)$ can thus potentially decrease both the power requirement and arrival time resulting in a lower energy consumption.

Integrating the above remarks, to obtain truly energy-optimal paths in a dynamic flow $\mathbf{v}(\mathbf{x}, t)$, we solve an optimization problem that computes the energy-optimal vehicle-speed $F^*(t)$ from among all speed time series $F(\bullet)$ of vehicles that reach the target in minimum time. In other words, we seek energy-optimal paths from among all time-optimal paths for a vehicle navigating between two locations in the presence of a strong, dynamic, and deterministic flow. Mathematically, the problem is stated as follows:

$$\min_{F(\bullet)} \left\{ E(\bullet) \equiv \int_0^{T(\mathbf{x}_f; F(\bullet))} p(t) dt \right\}, \tag{1a}$$

$$\text{s. t. } \frac{\partial \phi(\mathbf{x}, t)}{\partial t} = -F(\bullet) |\nabla \phi(\mathbf{x}, t)| - \mathbf{v}(\mathbf{x}, t) \cdot \nabla \phi(\mathbf{x}, t), \tag{1b}$$

$$\text{in } (\mathbf{x}, t) \in \Omega \times [0, \infty),$$

$$T(\mathbf{x}_f; F(\bullet)) = \min_t \{ t : \phi(\mathbf{x}_f, t) \leq 0 \}, \tag{1c}$$

$$\phi(\mathbf{x}, 0) = |\mathbf{x} - \mathbf{x}_s|, \tag{1d}$$

$$p(t) = F(t)^{n_p}, n_p \geq 0, \tag{1e}$$

where equation (1a) is the optimization and the constraints equations (1b)–(1d) arise due to the fact that all the vehicles with different $F(\bullet)$ travel exactly in a time-optimal fashion.

The scalar field $\phi(\mathbf{x}, t)$ is a reachability front-tracking level-set function and the viscosity solution of the exact level-set Hamilton-Jacobi equation (1b) with terminal and initial conditions (equations (1c) and (1d)). Only the zero level-set is needed, but for numerical convenience, an open boundary condition can be used at the numerical domain boundaries $\delta\Omega$, e.g., $\frac{\partial^2 \phi(\mathbf{x}, t)}{\partial \mathbf{n}^2} |_{\delta\Omega} = 0$, where \mathbf{n} is the outward normal to $\delta\Omega$. The subsequent solution to the backtracking (equation (2)),

$$\frac{d\mathbf{x}^*}{dt} = -\mathbf{v}(\mathbf{x}^*, t) - F(\bullet) \frac{\nabla \phi(\mathbf{x}^*, t)}{|\nabla \phi(\mathbf{x}^*, t)|}, \tag{2}$$

$$0 \leq t \leq T(\mathbf{x}_f; F(\bullet)) \quad \text{and} \quad \mathbf{x}^*(T) = \mathbf{x}_f,$$

yields the continuous-time history of the time-optimal vehicle heading angles, $\theta^*(t)$ [Lolla et al., 2014a; Lolla and Lermusiaux, 2016; Subramani and Lermusiaux, 2016].

Energy-optimal paths can be computed by solving the optimization problem (equation (1)) with the current field $\mathbf{v}(\mathbf{x}, t)$ provided by an ocean modeling system. To solve equation (1), a stochastic optimization and dynamically orthogonal equations are employed.

2.1. Stochastic Dynamically Orthogonal Level-Set Optimization

Summarizing [Subramani and Lermusiaux, 2016], we treat $F(t)$ as a random variable, i.e., $F(t; \omega)$, where ω is a random event. This treatment converts the level-set PDE (equation (1b)) into a stochastic-PDE (S-PDE) written in the Langevin form as

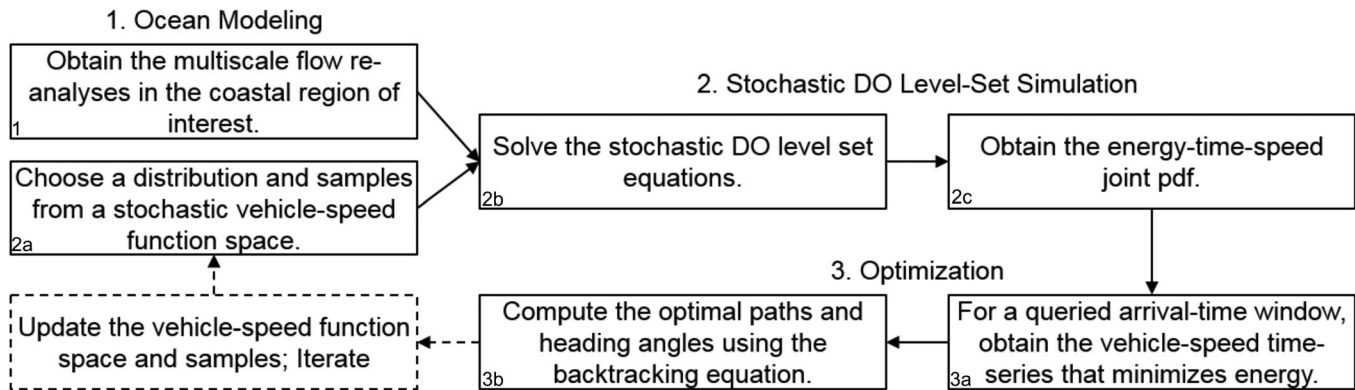


Figure 2. The flowchart outlining the three stages of the stochastic DO level-set optimization methodology for energy-optimal path planning, rigorously combining optimal planning PDEs with ocean flow forecasts (and/or reanalyses when replanning).

$$\frac{\partial \phi(\mathbf{x}, t; \omega)}{\partial t} = -F(t; \omega) |\nabla \phi(\mathbf{x}, t; \omega)| - \mathbf{v}(\mathbf{x}, t) \cdot \nabla \phi(\mathbf{x}, t; \omega). \tag{3}$$

For $F(t; \omega) \geq 0$, we solve the S-PDE equation (3) until the first time instant t such that $\phi(\mathbf{x}_f, t; \omega) \leq 0$, starting from the deterministic initial conditions $\phi(\mathbf{x}, 0; \omega) = |\mathbf{x} - \mathbf{x}_s|$ and with the boundary condition $\frac{\partial^2 \phi(\mathbf{x}, t; \omega)}{\partial \mathbf{n}^2} |_{\partial \Omega} = 0$. Such a stochastic simulation yields the distribution of the optimal arrival time $T(\mathbf{x}_f; F(\cdot; \omega))$ for a chosen distribution $F(\cdot; \omega)$.

To efficiently integrate the S-PDE equation (3), we employ its stochastic Dynamically Orthogonal (DO) level-set equations [Subramani and Lermusiaux, 2016], here summarized in Appendix A. These equations are simply a very efficient way to solve the S-PDE equation (3): we showed that they are 3–4 orders of magnitude faster than a direct Monte Carlo (MC) method. Moreover, the computational speedup comes at a modest cost of less than 2% error when compared to the MC solution and these MC-DO differences are commonly of the order of discretization errors [Subramani and Lermusiaux, 2016].

Once $T(\mathbf{x}_f; F(\cdot; \omega))$ is computed from $F(\cdot; \omega)$ and equations (A3–A5), the distribution of energy utilized is obtained as: $E(\omega) = \int_0^{T(\mathbf{x}_f; F(\cdot; \omega))} p(t; \omega) dt$. Finally, for a particular arrival time window, the speed function $F^*(t; \omega)$ that minimizes the total energy used $E(\omega)$, i.e., $F^*(t; \omega) = \arg \min_{F(\cdot; \omega)} E(\omega)$, can be obtained by searching among the marginal distribution of vehicle-speeds that reach the target in the queried arrival time window. If required, the vehicle-speed function space $F(\cdot; \omega)$ is updated and another iteration may be performed. The notation, DO equations, algorithm, and computational costs are outlined in Appendix A.

In summary, the solution procedure simply consists of three distinct tasks as depicted in the flowchart in Figure 2. First is the ocean simulation to obtain the flow field for the planning horizon. Second is the stochastic simulation using the reduced-order DO level-set equations to obtain the distribution of optimal arrival time and energy for a stochastic class of vehicle-speeds. Third is the optimization to identify the energy-optimal vehicle-speed function for each arrival time window from the corresponding marginal distribution of energy and vehicle-speeds. Next, we apply this methodology to the Middle-Atlantic Bight-Shelf-break Front region.

3. Ocean Circulation in the Middle-Atlantic Bight and Shelfbreak Front Region

The first task in our path planning (Figure 2) is to obtain an estimate of the environmental flow field $\mathbf{v}(\mathbf{x}, t)$, to be used in equation (1b). Presently, $\mathbf{v}(\mathbf{x}, t)$ is obtained from a realistic ocean reanalysis of the MAB and shelfbreak front region for August–September 2006 during the real-time AWACS and SW06 exercises [WHOI, 2006; Lermusiaux et al., 2006; Newhall et al., 2007; Tang et al., 2007; Chapman and Lynch, 2010; Lin et al., 2010]. Next, we describe the ocean modeling system, the atmospheric and tidal forcing, and the multi-scale ocean flows encountered by gliders navigating in the MAB. Although we emphasize gliders, we note that our theory applies to other platforms such as propelled vehicles and surface crafts.

3.1. Multiresolution Data-Assimilative Ocean Modeling System

The ocean model employed is the MSEAS nonlinear free-surface hydrostatic primitive-equation (PE) model, configured with generalized-level vertical-coordinates and implicit two-way nested computational domains [Haley and Lermusiaux, 2010]. The domains are shown in Figure 3, overlaid on bathymetry. In the horizontal, they have a 3 km and 1 km grid resolution, respectively, and in the vertical, they employ 100 levels optimized to the thermocline and flow structures. Also shown in Figure 3 are the two pairs of start (circle) and end points (star) for which energy-optimal paths will be studied in detail. They are selected because they are representative of classic missions as well as of the results we obtained with many other start and end points (not shown). The main planning horizon is from 00 UTC 28 August 2006 to 00 UTC 9 September 2006.

The tidal-to-mesoscale ocean reanalysis is initialized with objectively analyzed temperature, salinity, and velocity fields for 14 August 2006. Two multiscale in-space analyses [Lermusiaux, 2002], inshore and offshore of the expected shelfbreak front are combined using a shelfbreak-front feature model [Lermusiaux, 1999; Gangopadhyay et al., 2003]. Each of these analyses fuses varied in situ synoptic AWACS-SW06 data, e.g., from gliders, conductivity-temperature-depth (CTD) profiles, etc., and historical data, e.g., from the National Marine Fisheries Service (NMFS), World Ocean Database, Gulf Stream Feature analyses, Buoy data, etc. The Gulf Stream is initialized using synoptic and historical CTD profiles as well as estimates of its position based on SST and NAVOCEANO feature analyses. Transport feature models are also used for the Gulf Stream and slope-recirculation gyre. Barotropic tides based on the high-resolution TPX07.2 surface-tide velocities and elevation [Egbert and Erofeeva, 2002] for 14 August 2006 are merged with the subtidal initial fields, following Haley et al. [2015a]. The reanalysis free-surface PE simulation is then integrated for 42 days. During integration, the ocean data collected during the AWACS and SW06 exercises as well as data of opportunity (NMFS, etc.) are assimilated. At the free surface, atmospheric-forcing fluxes are applied, optimally merging the Weather Research and Forecasting (WRF) fields into the larger Navy Operational Global Atmospheric Prediction System (NOGAPS) fields. At lateral ocean boundaries, the open boundary conditions (OBCs) employed consist of a new mixed sponge-radiation-tide formulation [Haley et al., 2015b]. The time-dependent barotropic tide velocities and elevation are used as forcing. The OBC relaxes the subtidal flow to an exponentially weighted time-running average, preventing the reflection of outward propagating internal tides by absorbing them in a sponge layer, but allowing the radiation of subtidal field variability such as fronts and eddies. Finally, the numerical and subgrid-scale parameters were tuned for the region by comparison of




Figure 3. The Middle-Atlantic Bight and shelfbreak front region where we compute energy-optimal paths. The two-way nested computational domains (1 km and 3 km resolution, respectively) are marked as white boxes, overlaid on bathymetry (color axis, in m). The AWACS/SW06 experiment occurred mostly in the smaller domain. Two pairs (lighter and darker gray) of start (circles) and end points (stars) for which energy-optimal paths are discussed in detail are also shown.

many PE simulations with independent in situ SW06 measurements. For the seasonal variability and mean circulation in the Middle-Atlantic Bight and Shelf-break front region, we refer to Kohut et al. [2004] and Lentz [2008].

Next, we describe features of the external forcing and internal flows that affect the energy consumption of gliders navigating in the region during 28 August 2006 to 9 September 2006.

3.2. External Forcing Fields

3.2.1. Atmospheric Forcing

The atmospheric forcing fluxes combining the WRF and NOGAPS reanalyses are illustrated in Figure 4 by the daily average of the hourly wind forcing (N/m^2). A major storm event is the tropical storm Ernesto which advects northward over Mid-

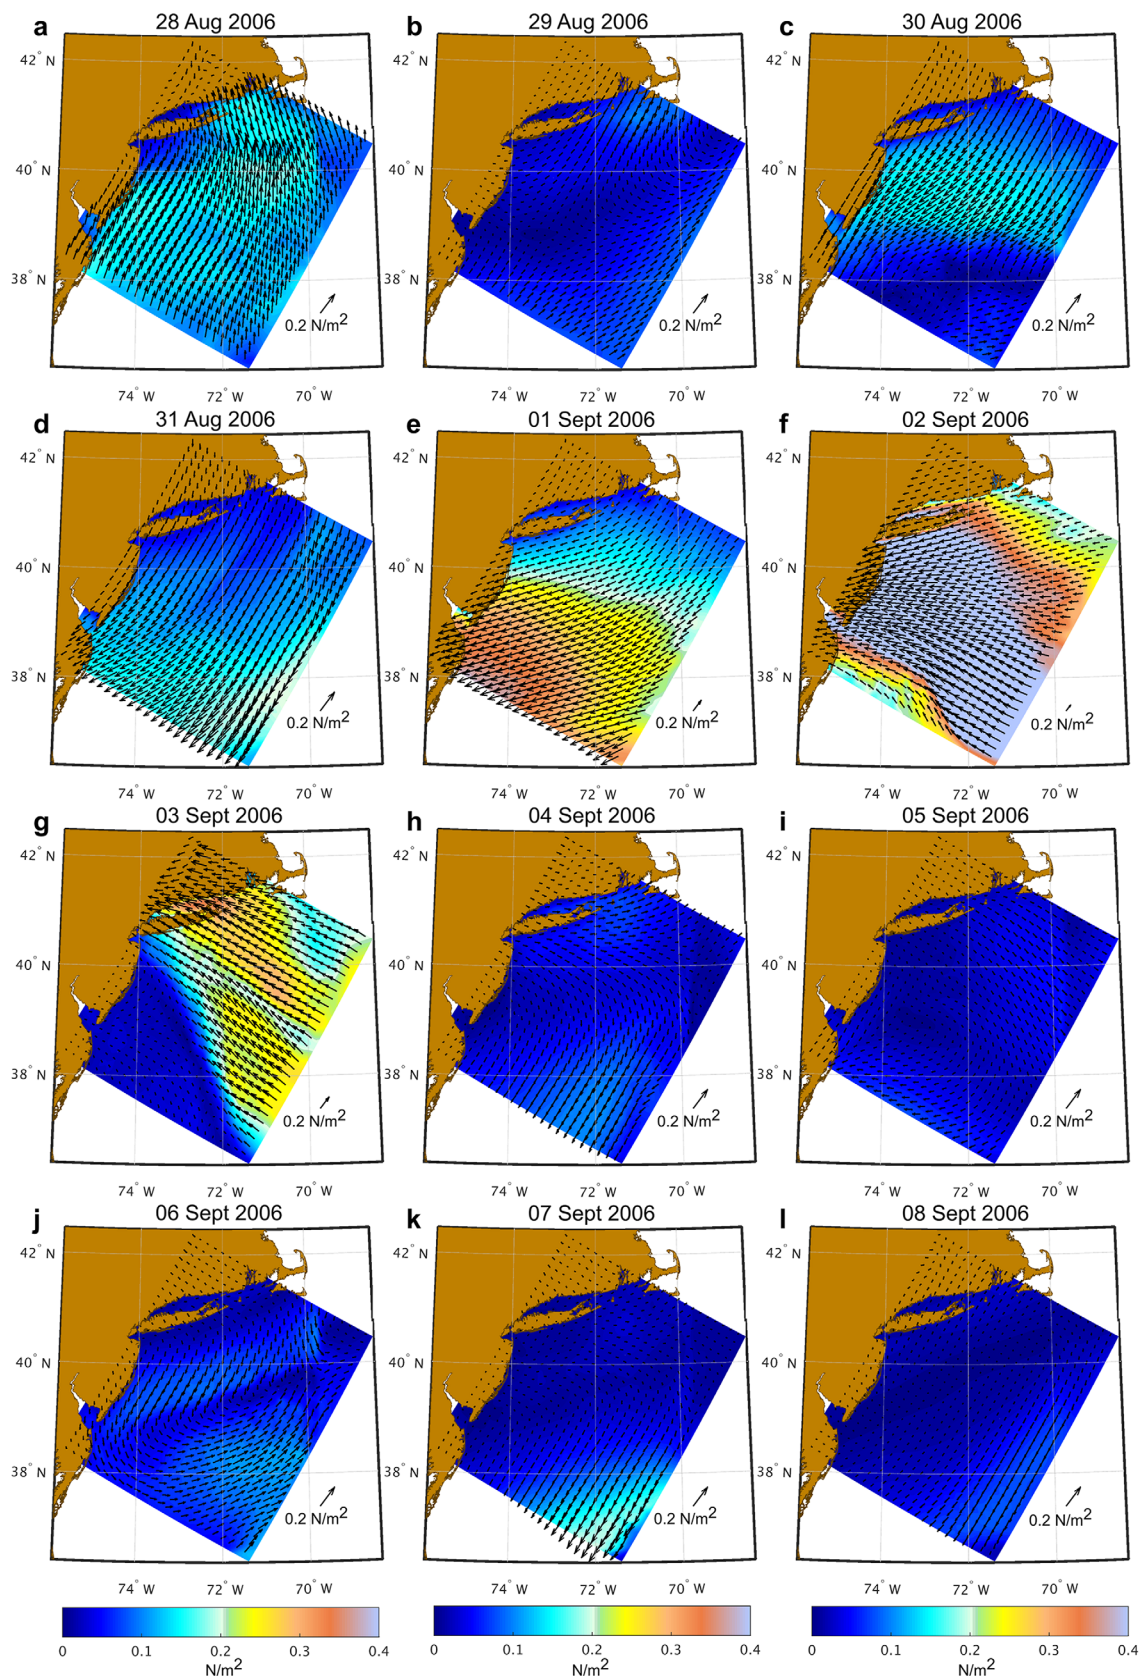


Figure 4. Daily averaged surface wind stress in N/m² (colored), overlaid with daily averaged wind vectors showing directions and amplitudes (scale arrow is 2 N/m²). The wind stress that actually forces the MSEAS ocean reanalyses is hourly and obtained from a blending of the WRF and NOGAPS reanalyses.




Figure 5. Color map of the dominant M2 tidal component (in cm/s), overlaid with subsampled M2 tidal ellipses.

ate initial tides and the lateral OBC tidal forcing. The model best fits the multicomponent TPX07.2 surface-tide velocities and elevations, correcting for the higher-resolution bathymetry and coastlines employed in the MSEAS simulations and for the specifics of the MSEAS PE bottom stress and numerics [e.g., *Haley and Lermusiaux, 2010, Appendix 2.3*]. For example, robust discrete tidal continuity in the MSEAS domain is ensured through an optimization procedure.

Figure 5 shows the amplitude of the dominant M2 tidal component, overlaid with subsampled M2 tidal ellipses. The flow experienced by vehicles traveling over the shelf is clearly influenced by tides [e.g., *Keen and Glenn, 1995; He and Wilkin, 2006; Kelly and Lermusiaux, 2016*], especially when winds are weak. In the area where the present missions occur ($\sim 72.4^{\circ}\text{W}$ to 74°W , and 38.5°N to 40.5°N , see Figure 3), the maximum barotropic tidal flows vary between 6 and 15 cm/s for the duration under study. Later on, we will describe effects of this tidal forcing on both energy-optimal paths and constant-speed time-optimal paths (see section 4).

3.3. Ocean Flows and Regional Circulation Encountered by Autonomous Gliders

In the present study, all gliders follow the same yo-yo pattern in the vertical. These yo-yo patterns go from the near surface to either the local near-bottom or 400 m depth, whichever is shallower (for the missions considered, most of the paths occur on the shelf, within about 20–100 m). The effects of small vertical ocean velocities are assumed to be accounted for by the glider controller and the yo-yo patterns are assumed to be close to vertical when compared to environmental horizontal scales [*Lolla et al., 2014b; Subramani et al., 2015*]. The currents that a glider encounters during such yo-yo motions are then the horizontal currents integrated along its close-to-vertical path. These are the instantaneous currents used in our optimization. For our week-long or so missions, we focus next on their subtidal variability.

Figure 6 thus shows the daily averaged horizontal currents, depth-averaged from the surface to the local near bottom or 400 m (for depths shallower than 400 m, the flow shown is thus the barotropic flow up to the near bottom). We clearly see the Gulf Stream in the southeast corner of the larger domain where the ocean currents can reach 200 cm/s. However, for our missions, gliders do not navigate that far. Hence, we focus on the shelf and shelfbreak region. Its main persistent meandering jet is the shelfbreak front. It flows from the northeast to the southwest with a barotropic magnitude of ~ 15 cm/s. On 28 August 2006 (Figure 6a), the flow on the New Jersey shelf is a cyclonic gyre. On 29 August (Figure 6b), the shelf flow reverses due to southwestward winds, developing first along the coast of Long Island and intensifying to a midshelf

Atlantic states as an extra tropical cyclone. Our computational domain experiences the northeastern edge of the cyclone with winds blowing primarily onshore. The storm enters our domain on 31 August 2006 15 UTC from the south-southeast (SSE) and completely exits by 4 September 2006. The largest winds into the New Jersey coast have a maximum wind stress reaching 0.7 (N/m^2) during the peak of the storm (in the domain), on 2 September 2006 (Figure 4f). Before the storm, on 28 August, there are moderate winds throughout the domain (Figure 4a), parallel to the New Jersey coast and onshore by Long Island. A minor wind event occurs from 6 September 00 UTC to 6 September 17 UTC (Figure 4k). Barring this event, for glider missions, there is no other remarkable wind activity in the region from 4 September to 8 September 2006.

3.2.2. Barotropic Tidal Forcing

A shallow water tidal model, based on *Logutov and Lermusiaux [2008]* and *Logutov [2008]* but with nonlinear terms including nonlinear bottom stresses, is used to gener-

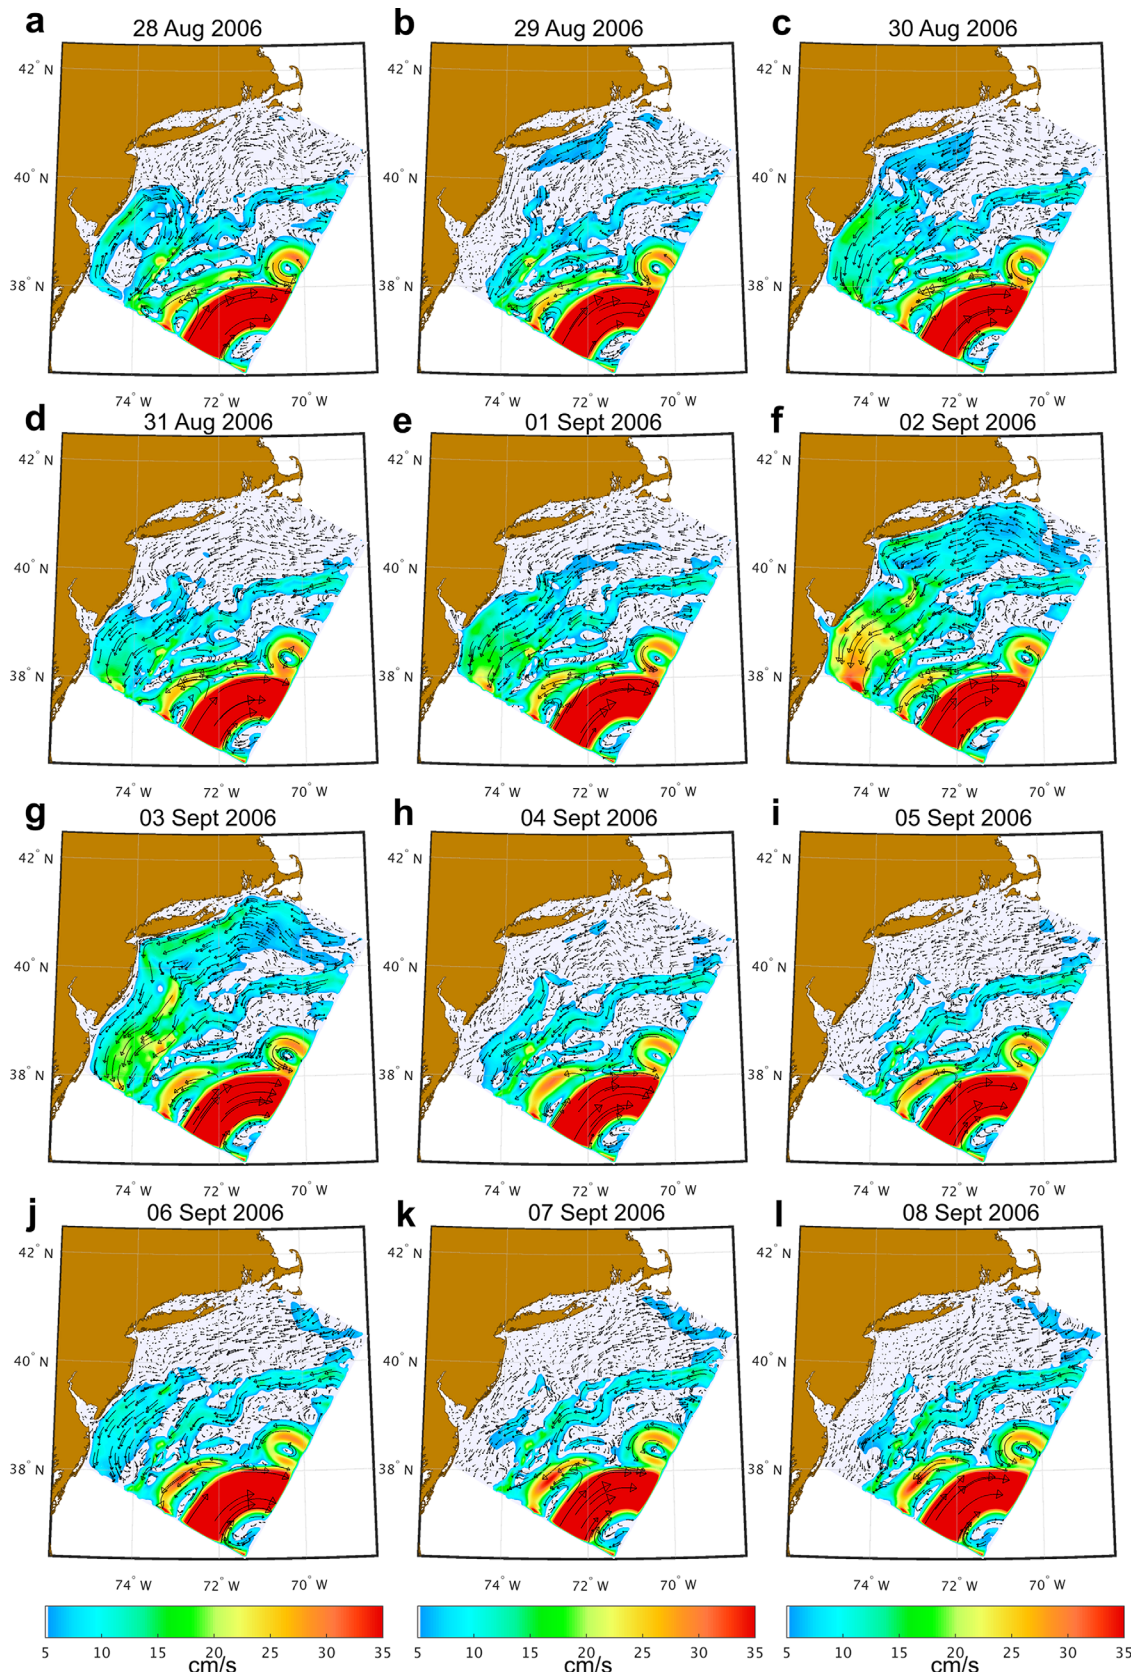


Figure 6. Daily vertically averaged horizontal ocean flow field in the Middle-Atlantic Bight and shelfbreak front region. The horizontal currents shown are those encountered by vehicles in a yo-yo pattern from the near surface to either the local near bottom or 400 m depth, whichever is shallower. The flow patterns are illustrated by their vectors, overlaid on a color plot of the flow magnitude (in cm/s).

south-southwestward flow by 30 August (Figure 6c). On 31 August and 1 September (Figures 6d and 6e), this flow weakens except south of the New Jersey shelf. The extra tropical storm Ernesto passes over the domain during 1 and 3 September. The domain sees mostly the northeastern edge of the storm as it moves north. The barotropic ocean response is also a southwestward and alongshore flow, but 2–3 times stronger than on 30 August (reaching 15–30 cm/s) and shelf wide with intensification near the coast and a correlation length of 400–500 km (Figures 6f and 6g). This flow can be intelligently utilized by gliders to reduce their energy requirement, as we will show later. From 4 September to 8 September (Figures 6h–6l), a midshelf jet flows south on the New Jersey shelf, parallel to the shelfbreak front. It starts northwest of the end point of mission 1 (39.76°N, 72.54°W) and south of the Hudson canyon, and reaches a barotropic magnitude of 5–10 cm/s. This midshelf jet is a remnant of the shelf response to Ernesto, partly supported by a local cross-shelf density gradient (not shown) that Ernesto established. This weak flow advects slower vehicles to the south away from the target as we will see later. On 6 September (Figure 6j), by the New Jersey coast, a barotropic shelf flow reaching 10–20 cm/s is forced by the minor wind event between 00 to 17 UTC on that day. Finally, density-driven eddies and currents occur on the shelf during the whole period but their barotropic flows are weak.

4. Energy-Optimal Paths and Their Analyses

To compute the energy-optimal paths of gliders navigating between two locations in the MAB region, we now complete stages 2 and 3 (Figure 2) of the methodology. For the illustrations and dynamical studies, we consider only two of the varied glider missions that were hindcast. The first missions have a start point off the coast of New Jersey and an end point north-northwest of the Hudson River canyon (lighter gray points in Figure 3). The second missions have a start point off the coast of New Jersey and an end point across the shelfbreak front within the AWACS/SW06 intensive region (see darker gray points in Figure 3). The main planning period is from 00 UTC 28 August 2006 to 00 UTC 9 September 2006. Next, we analyze the optimal paths describing the responses to tides, wind forcing, and local currents. We wrap up the section with a short discussion on adaptive replanning, adaptive sampling, and forecast uncertainty.

4.1. From the Coast of New Jersey to the Hudson River Canyon

The start point is (40.05°N, 73.8°W) and end point is (39.76°N, 72.54°W) as shown by the lighter gray circle (start) and star (end) in Figure 3. For illustrative purposes, we used two stochastic classes of vehicle-speeds (box 2a in Figure 2) in which the speed *switches* every 48 h (or 24 h) and maintains the new speed for 48 h (or 24 h). These time scales are in line with common time scales in the coastal ocean and represent what could commonly be done today. The *switch-sampling* algorithm used here is outlined in Appendix A and a few samples are shown in Figure A1. Based on common operational constraints and again for illustrative purposes, we assume that the minimum and maximum speeds of the glider are 10 and 25 cm/s, respectively. Other choices are also possible depending on operator preferences and mission requirements [e.g., Ramp et al., 2009; Leonard et al., 2010]. The numerical parameters are as follows: $dx=dy=3$ km; $dt=0.1$ h; $n_x=175$; $n_y=150$; $n_r=65,536$; and $n_{s,\phi}=120$ (See Appendix A for details on the notation). Results of the stochastic DO level-set optimization are shown in Figure 7.

4.1.1. Stochastic DO Level-Set Simulation

Figure 7a shows the distribution of minimum arrival time, vehicle-speed time series, and total energy utilized, as computed by solving the stochastic DO level-set equations for these two stochastic classes. In Figure 7a, the y axis has been normalized with the maximum energy utilized and the x axis is the arrival time. Each gray dot corresponds to one of the 65,536 samples of the vehicle-speed distribution. The colored dots correspond to gliders with a constant relative speed throughout their mission: the coloring corresponds to that speed. They are explicitly represented in the figure as they will be used for our analysis in section 4.1.1. To further improve the computational efficiency, during the integration, we pruned those samples of vehicle-speed time series that use more energy than constant-speed gliders for the same reachable front [Subramani and Lermusiaux, 2016]. Our methodology computes all the points represented in Figure 7a by just two DO simulations: one for time-dependent vehicle-speeds and another for constant vehicle-speeds.

4.1.2. Optimization

We perform an optimization to obtain the energy-optimal vehicle-speed time series for all arrival times. As examples, we show five such paths in Figure 7b, corresponding to the arrival times marked in blue in Figure 7a (the numbers 1–5 in Figure 7b correspond to those of Figure 7a). The region shown in Figure 7b is magnified

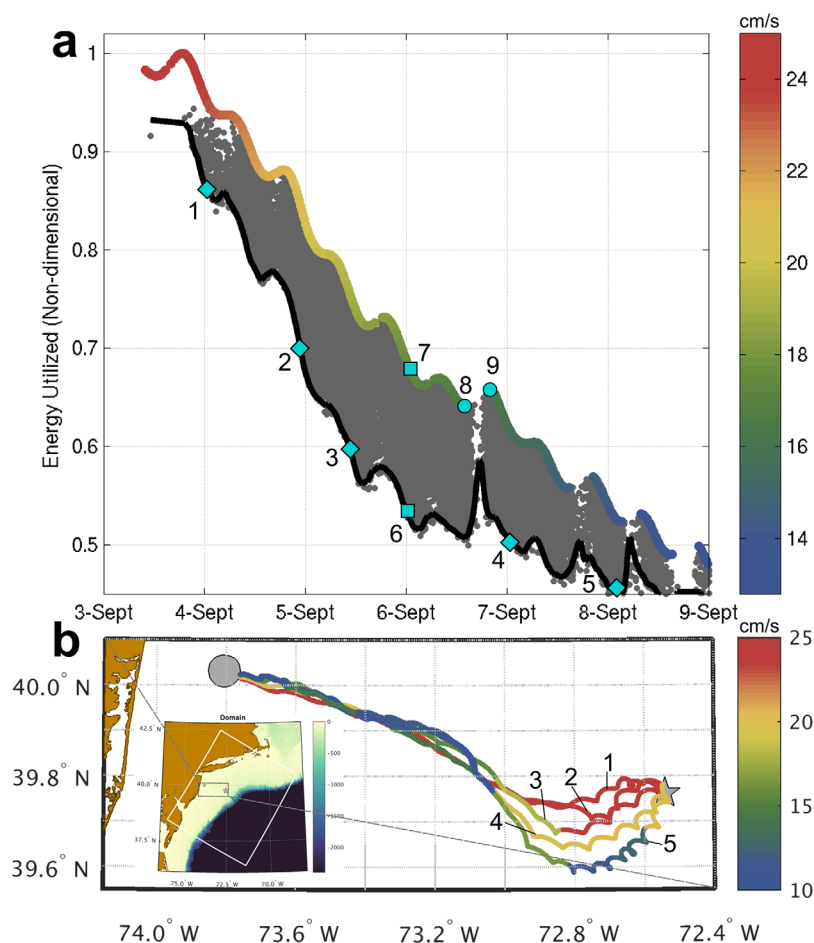


Figure 7. Energy-optimal path planning for gliders navigating from the coast of New Jersey to the Hudson River canyon: (a) Distribution of optimal arrival time (x axis) and energy utilized (y axis) for the samples (represented as dots in the cloud) of vehicle-speed time series, $F(t; r)$. The colored dots are constant-speed samples with the color representing their vehicle-speeds in cm/s (color axis to the right). The gray dots are realizations that consume less energy than constant-speed realizations but with same arrival times (only these “gray dot” realizations are shown since the others are pruned by our algorithm). The lowest energy envelope is shown by a thick black line. The paths corresponding to dots highlighted with numbers 1–8 are used as examples in our analysis. (b) Energy-optimal paths 1–5 shown on a magnified view of a region of the computational domain (in the inset). Paths are colored by their instantaneous speeds in cm/s (color axis to the right). These paths are computed using the vehicle-speed time series that minimizes the energy requirement for that arrival time from the above distribution. Paths 1 and 2 are closer to a straight line joining the start (circle) and end points (star) than paths 3–5, which are advected to the south by ocean currents. Points 6–9 are used for analysis later.

from the computational domain, as shown in the inset. Paths 1 and 2 are closer to a straight line joining the start (circle) and end points (star) than paths 3–5, which are advected to the south by ocean currents.

4.1.2.1. Energy-Optimal Path of a Glider Reaching the Target on 6 September 2006

We now describe and analyze the reachable sets, fronts, and paths between the start and end points for a glider that travels with the energy-optimal time series computed above, focusing on the single arrival time 02 UTC 6 September 2006. To evaluate the energy savings, we need a baseline for comparison. We choose this baseline as the energy requirement of constant-speed gliders traveling in a time-optimal fashion. The reason for this choice is as follows. If the gliders do not change their speeds along their paths, traveling in a time-optimal fashion will use the least energy (see section 2). Our method computes the vehicle-speed time series that would optimize the energy usage among all exact time-optimal paths. Hence, arriving at the same time as time-optimal constant-speed gliders is the benchmark to beat for an energy-optimal glider.

Before proceeding, we provide two remarks. First, it is not straightforward to identify the constant speed with which a glider must travel in order to reach the target at a specific arrival time window. To simulate such a large number of constant-speed gliders, the DO level-set method is very efficient (colored points in Figure 7a). Second, choosing as benchmark a constant-speed time-optimal glider is very demanding: the

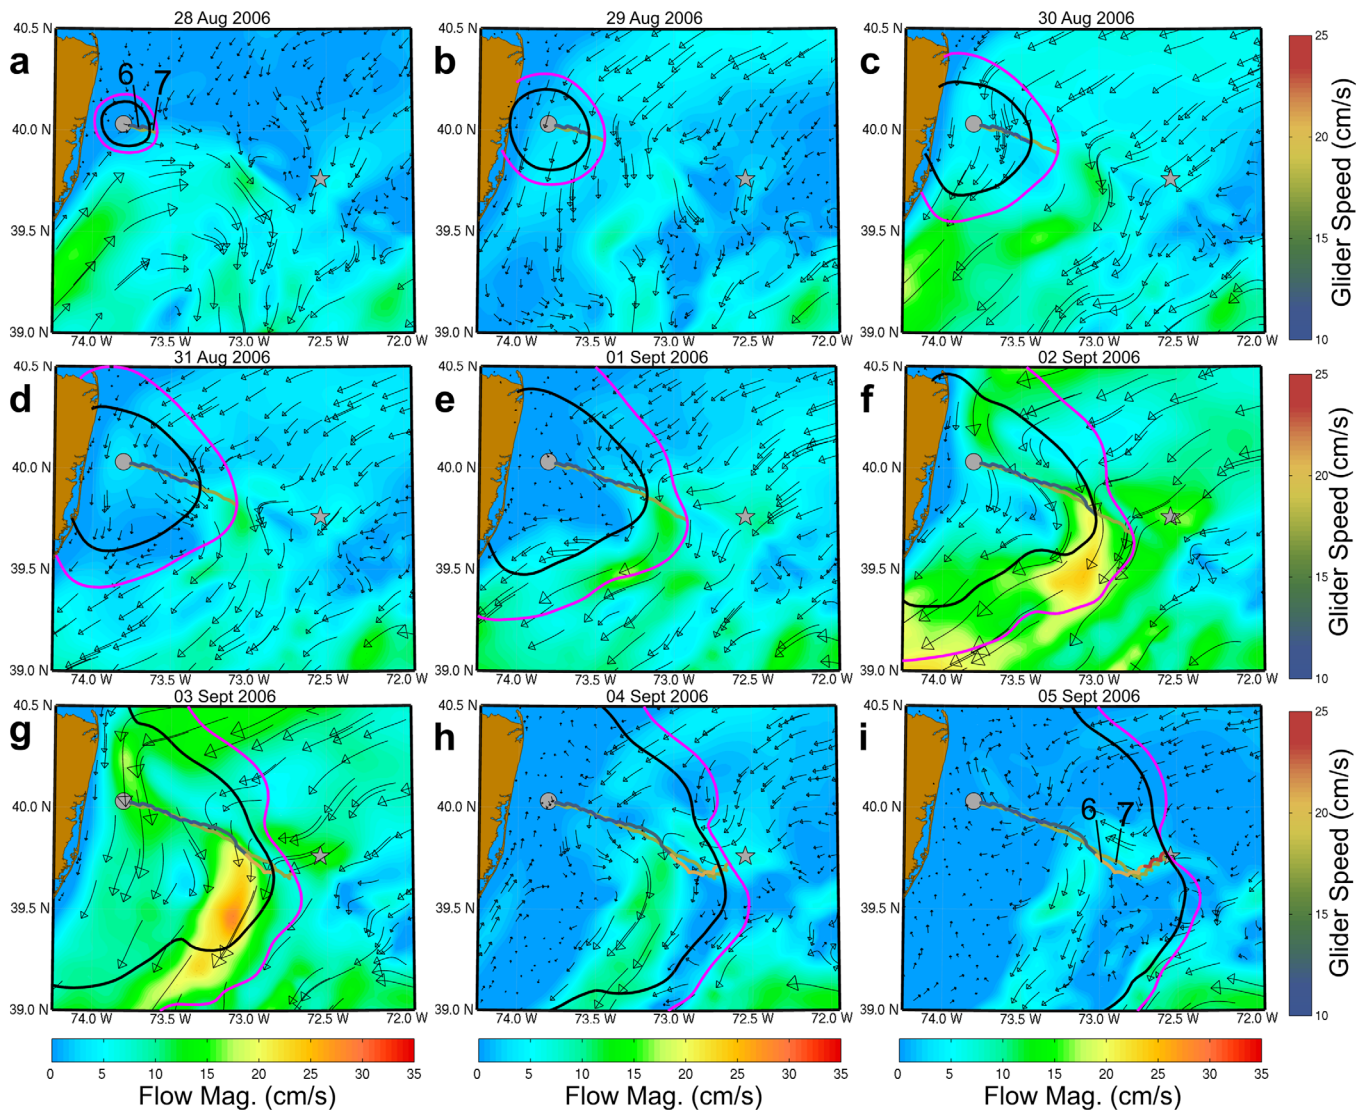


Figure 8. Paths and reachable fronts for the energy-optimal glider (#6) and time-optimal constant-speed glider (#7) that reach the target on 02 UTC 6 September 2006. Each figure shows a color plot of the daily average of the depth-averaged horizontal currents experienced by the gliders, with vectors indicating the local flow direction (color axis in cm/s at the bottom row). These ocean flows are overlaid with the paths and reachable fronts for gliders #6 and #7, from the start date until the end of the day shown on each figure. The paths are colored by their instantaneous relative vehicle-speeds (color axis in cm/s to the right column). Glider #6 is able to utilize the ocean response to the tropical storm Ernesto from 2 to 4 September. Glider #7 does not catch this weak flow. Toward the end of the mission, both gliders are affected predominantly by the tidal flow (see Figure 9). The energy-optimal glider employs a higher speed $F(t)$ in this region and is thus able to spend less energy overall and still reach the target at the same time as the time-optimal constant-speed glider.

energy savings would be larger for less stringent performance criteria such as, e.g., a glider moving in a straight line to the target (not shown here).

We first study the overall trend of the energy savings achieved. In Figure 7a, we highlighted the energy-utilized and arrival time of constant-speed gliders by coloring them with their speed. The thick black line at the bottom of the cloud of gray points is the energy-time characteristic of the energy-optimal paths. The gap between these two lines indicates the energy savings achieved, when compared with the energy required by constant-speed gliders reaching the target at the same time. Hence, these two lines can be considered as bounds of energy consumption.

In Figure 7a, the energy-optimal glider that arrives on 02 UTC 6 September 2006 is #6 and the constant-speed glider that arrives at the same time is #7. Glider #6 consumes $\approx 26\%$ less energy than glider #7. In that figure, the energy-optimal gliders #1 to #5 that correspond to other arrival times consume between 10 and 20% less energy than their constant-speed time-optimal glider benchmarks. To study how the ocean

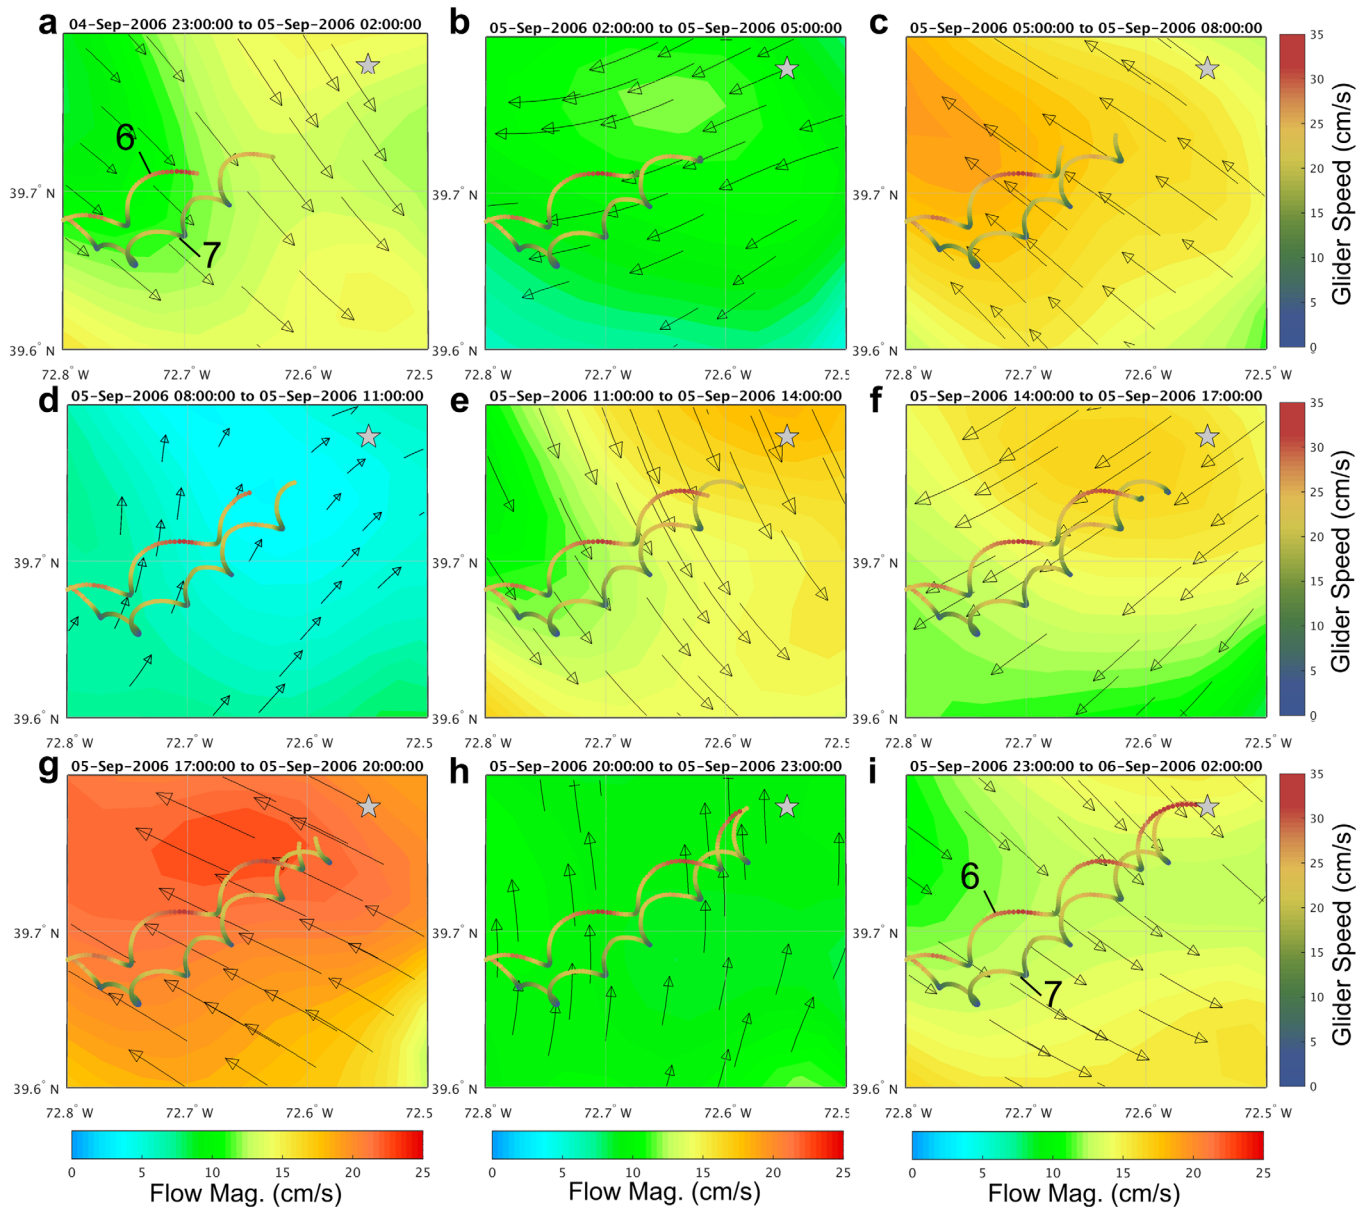


Figure 9. Effect of tides on the paths of the energy-optimal glider (#6) and the time-optimal constant-speed glider (#7) in the last 27 h: The nine figures show 3 h average of the vertically averaged horizontal flows (bottom color axis, in cm/s) from 4 September 23 UTC to 6 September 02 UTC, with vectors indicating flow direction. Each figure is a magnified region close to the target. Overlaid on the flow are the paths of gliders #6 and #7 until the end time of each figure, colored with their total *effective vehicle-speed* (right color axis, in cm/s). Both gliders execute a spiral pattern.

flows affect the energy-optimal paths, we look at the following three figures. Figure 8 shows how the reachable front (i.e., the zero level-set contour or exterior boundary of the growing reachable set) and optimal paths evolve in time, overlaid on a color plot of the daily averaged horizontal flows experienced by the gliders. Figure 9 shows the time evolution of paths in 3 h intervals during the last 27 h of the mission where wind activity is minimal and flows are mostly tidal (section 3). Figure 10 shows instantaneous and 24 h moving average of the effective vehicle-speeds, allowing a quantitative understanding of the effects of the multiple flow timescales.

In Figure 8, the nine panels correspond to the nine days from 00 UTC 28 August 2006 to 02 UTC 6 September 2006. Figures 8a–8i depict the paths of the energy-optimal glider #6 and constant-speed glider #7 from the start time until the end of the day shown in the panel. The paths are colored by the instantaneous vehicle-speed at that location and are overlaid on the daily average of the depth-averaged (400 m or near

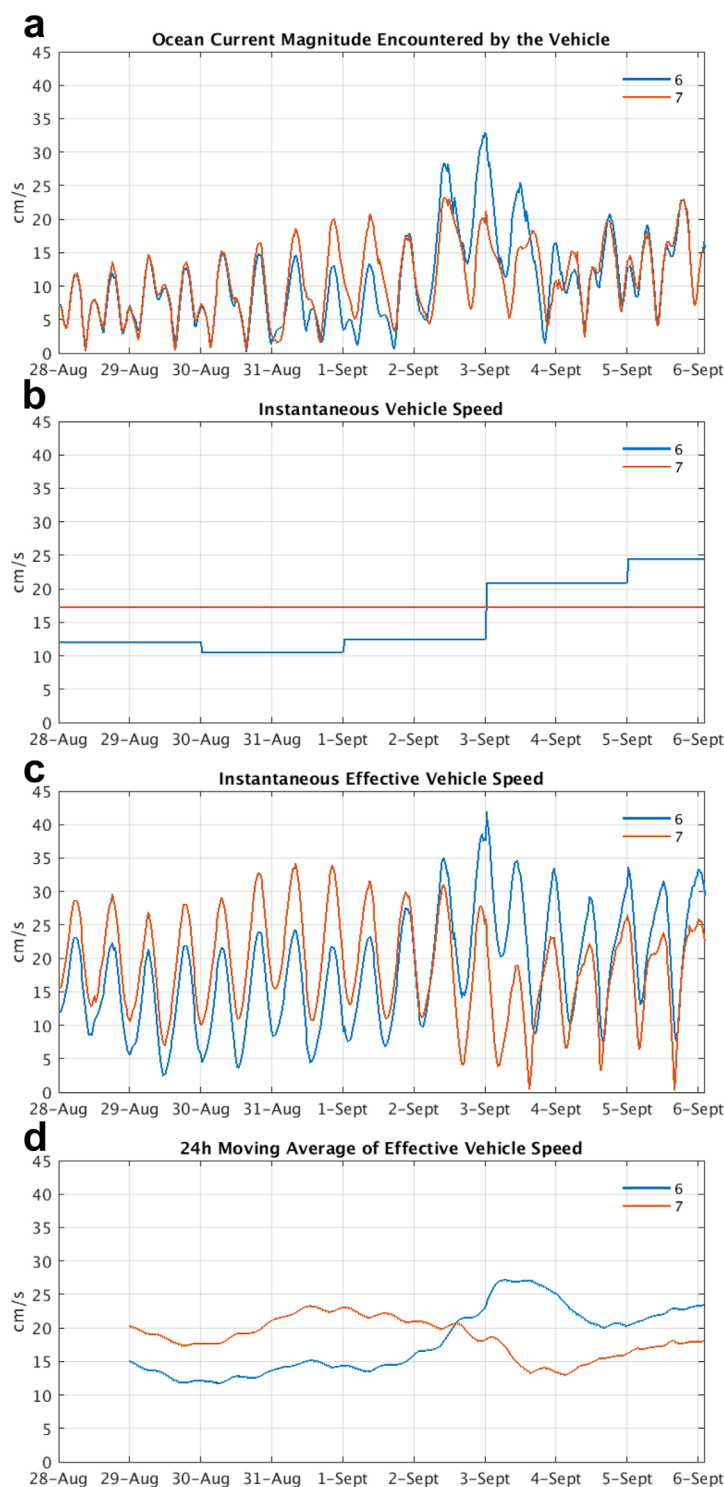


Figure 10. Time series of the (a) ocean current magnitude, (b) instantaneous vehicle-speed, (c) effective vehicle-speed, and (d) 24 h moving average of effective vehicle-speed for gliders #6 and #7. Glider #6 travels at a lower speed than glider #7 initially, and thereby catches a strong favorable flow as a response of storm Ernesto from 2 to 4 September (#6 in Figure 10a). Glider #7 does not experience this flow. On 3 September, glider #6 switches to a (b) higher speed and travels at an effective speed of up to 40 cm/s. From 4 September onward, when strong tidal flows dominate, glider #6 maintains higher relative speeds and its effective speed goes only as low as 10 cm/s, whereas glider #7 has periods with a (c) near-zero effective speed. Figures 10c and 10d together show the multiscale aspects of the effective vehicle-velocity. Instantaneous effective speeds have a large intraday variability showing the effect of tides while the 24 h moving average shows the effect of longer timescale circulations (order of days).

bottom to the surface) horizontal currents. Also shown are the reachable fronts of both gliders at the end of the day (black for glider #6 and pink for glider #7). On 28 August (Figure 8a), both gliders start sailing away from the start point off New Jersey coast. Both reachable fronts are roughly circular, with a slight elongation due to the weak offshore-ward shelf circulation on that day. On 29 and 30 August (Figures 8b and 8c), southward coastal flows of 5–10 cm/s establish and affect both reachable sets. On 30 August, the flow is around 15 cm/s at the southernmost part of both reachable sets. On 31 August and 1 September (Figures 8d and 8e), the energy-optimal glider is in a region with a favorable flow, but the optimal constant-speed glider is caught in a flow that advects it southward. This effect can be seen by the shape of the reachable fronts at the end of 1 September. On 2 and 3 September (Figures 8f and 8g), the gliders experience the ocean response to Storm Ernesto. The energy-optimal glider is in a region that experiences maximum favorable flow making it travel a longer distance to the southeast (toward the target). The constant-speed glider is advected south (off target). Moreover, it travels a shorter total distance on these 2 days. After the storm response, on 4 and 5 September (Figures 8h and 8i), an adverse circulation hinders the growth of the reachable sets north of the target. Since the paths of both gliders were advected south during the storm, as a result, starting on 4 September, both vehicles have to turn northeast toward the target.

The final stretches of both paths show the effect of tidal flows. Figure 9 shows these paths during the last 25 h before the target is reached. Each of the nine panels shows the 3 h average of the horizontal flow experienced by the gliders. The paths in Figures 9a–9i are from the start of the mission until the end of the 3 h window shown in that panel. The paths are colored with the total *effective vehicle-velocity* (not relative as in Figure 8).

Both the constant-speed and energy-optimal gliders execute a spiral pattern due to tides (panels of Figure 9). The flow is first to the southeast (a). Next, it is southwest (b), then northwest (c), then southwest (d), and then southeast again (e). Therefore, the gliders complete an arc through (b)–(e). Figures 9g–9i also experience similar flow patterns and paths. In Figures 9b and 9f, the flows are directly opposing the motion of the gliders to the target. The effective speeds are the lowest during these periods. Moreover, the constant-speed glider #7 executes a small loop in Figure 9f and travels briefly in a direction directly opposite to the target. The energy-optimal glider #6 has a higher speed during this period and does not loop around.

Figure 10 shows the time series of ocean current magnitude encountered by the gliders, the instantaneous relative vehicle-speeds, the instantaneous effective vehicle-speeds, and 24 h moving averages of the effective vehicle-speeds. The instantaneous effective vehicle-speed time series shows a periodic behavior matching the M2 diurnal tidal frequency. When averaged over a 24 h hour period, we see the effects of the weak flows in the region (with time scales in the order days) on the gliders.

Overall, both the gliders arrive at the same time, but the energy-optimal glider employs a lower speed for the first 6 days, and thereafter catches a helpful flow that increases its effective speed. The constant-speed glider missed the opportunity to utilize this favorable flow. This example illustrates the utility of leveraging ocean forecasts for energy-optimal path planning and demonstrates that our method is capable of rigorously incorporating this information.

4.1.2.2. Unattainable Arrival Times for Constant-Speed Gliders: Effect of Tides

Our next interest is to examine the reason for unattainable arrival times for constant-speed gliders and the wavy pattern of energy-time characteristics of the energy-optimal gliders and constant-speed gliders (see the patterns of the colored and black lines, respectively, in Figure 7a). In *Subramani and Lermusiaux* [2016], we noted that the energy utilization characteristics can be complex for dynamic flows such as the double-gyre circulation. For example, reducing the relative speed can reduce the energy usage initially, but can rapidly increase the energy usage on further reduction if the glider has to travel much longer distances (e.g., to cross the strong unfavorable wind-driven jet). In the present missions (Figure 7a), we see that the energy consumption oscillates in a 24 h period with a local energy minima and maxima that roughly matches the M2 tidal frequencies. Moreover, there are some arrival times that are unattainable as the time between the arrival of constant-speed glider #8 and #9 (Figure 7a).

Glider #8 travels at 16.51 cm/s and #9 at a marginally lower speed of 16.48 cm/s. Figure 11 shows six panels with the paths of these gliders during the 18 h period from 03 UTC 6 September 2006 to 21 UTC 6 September 2006. As before, each panel shows a 3 h average of horizontal currents encountered by the gliders and the paths are overlaid on this color map. Here the paths are colored by solid colors, gray for the glider #8 and black

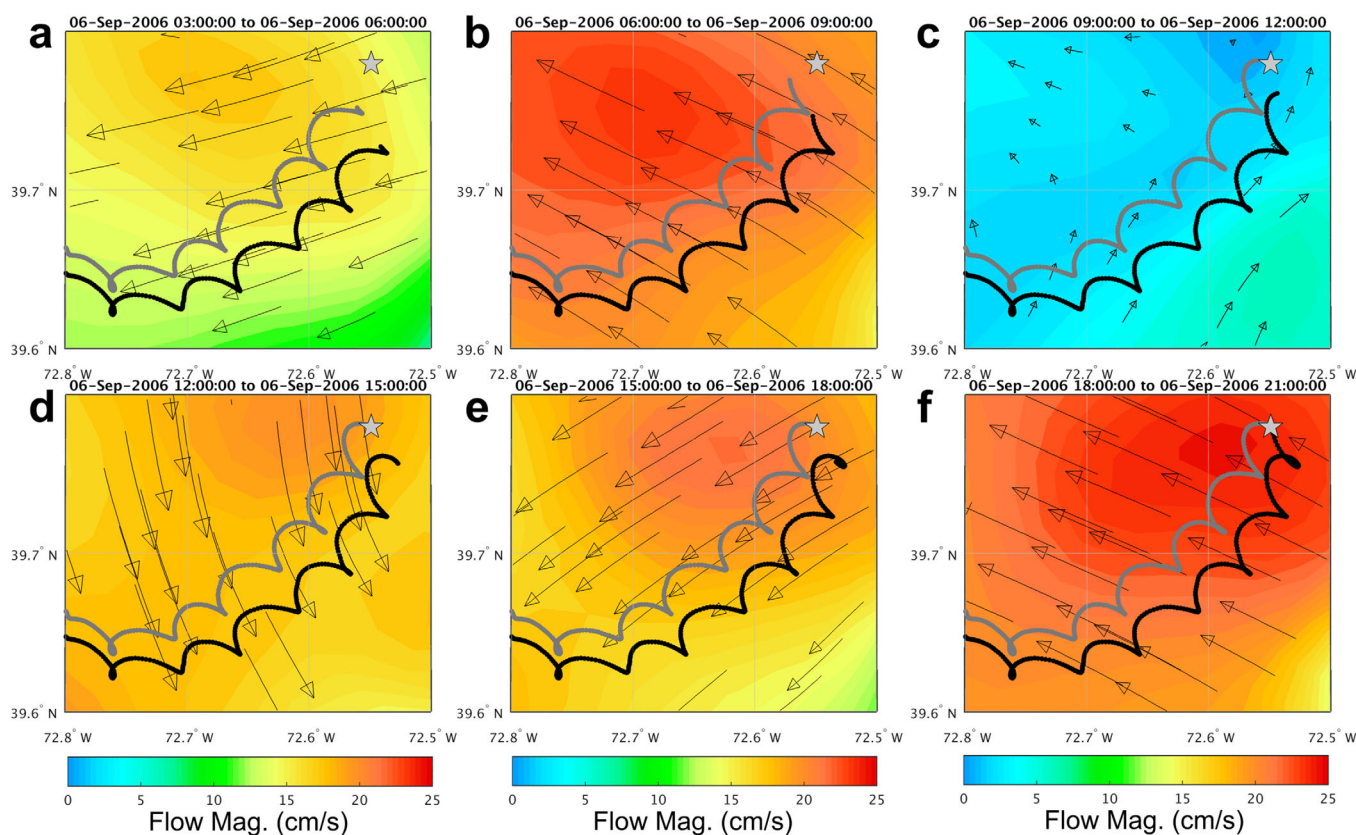


Figure 11. Effect of tides on the paths of two constant-speed gliders with speeds 16.51 cm/s (#8 in gray) and 16.48 cm/s (#9 in black): The six figures show 3 h time-averages of the vertically averaged horizontal flows (color axis in cm/s to the bottom), from 03 UTC 6 September 2006 to 21 UTC 6 September 2006, with vectors indicating the flow direction. Each figure is a magnified region close to the target. Overlaid on the flows are the paths of gliders #8 in gray and #9 in black, until the end time of each figure. Even though glider #9 is only marginally slower, it takes nine more hours and consumes more energy than glider #8. The energy requirement of both the gliders is given in Figure 7a.

for glider #9. The tidal flows are such that glider #9 has to travel for nine more hours and is caught in a particularly strong opposing tidal flow. Therefore, it utilizes a higher energy than glider #8. Moreover, the glider speeds and currents are such that there is no speed for which gliders can reach the target in the intervening 9 h period of, therefore, unattainable arrival times. Obviously, ocean engineering and forecasting cannot yet control vehicle-speeds and predict flows at four digits of accuracy. The point here is to showcase a well-known fact due to advection: who misses the bus for a few seconds (e.g., because of not running fast enough) may both be hours late and use a lot more energy to reach the destination.

4.2. Crossing the Shelfbreak Front From the Coast of New Jersey

For our second set of start and end points, we compute the energy-optimal paths for gliders navigating from (39.4°N, 74.1°W) to (38.7°N, 72.5°W) (see darker gray points in Figure 3). These gliders have to cross the shelfbreak front. We employ the same stochastic classes of vehicle-speeds and parameters as for the first set of missions (section 4.1).

Figure 12a shows the distribution of optimal arrival time and corresponding energy utilization. As previously mentioned, each gray dot is a sample of vehicle-speed time series and each colored dot a constant-speed glider (color scale is speed). Six arrival times are chosen to exhibit the energy-optimal paths that reach the target at those times and marked #1–#6 in Figure 12a. Figure 12b shows these paths colored by their instantaneous vehicle-speeds (right color axis, in cm/s). In this case, the energy savings achieved by energy-optimal gliders is between 7 and 10% when compared to constant-speed gliders that reach the target in the same arrival time. The gliders that arrive after 5 September (#3–#6) are affected by the tropical storm Ernesto and advected southward when compared to gliders that arrive before 5 September (#1 and #2). For the former gliders (#3–#6), after the favorable advection by the storm response, they first sail northeastward to reach the shelfbreak frontal jet upstream. They then sail in the jet, some further

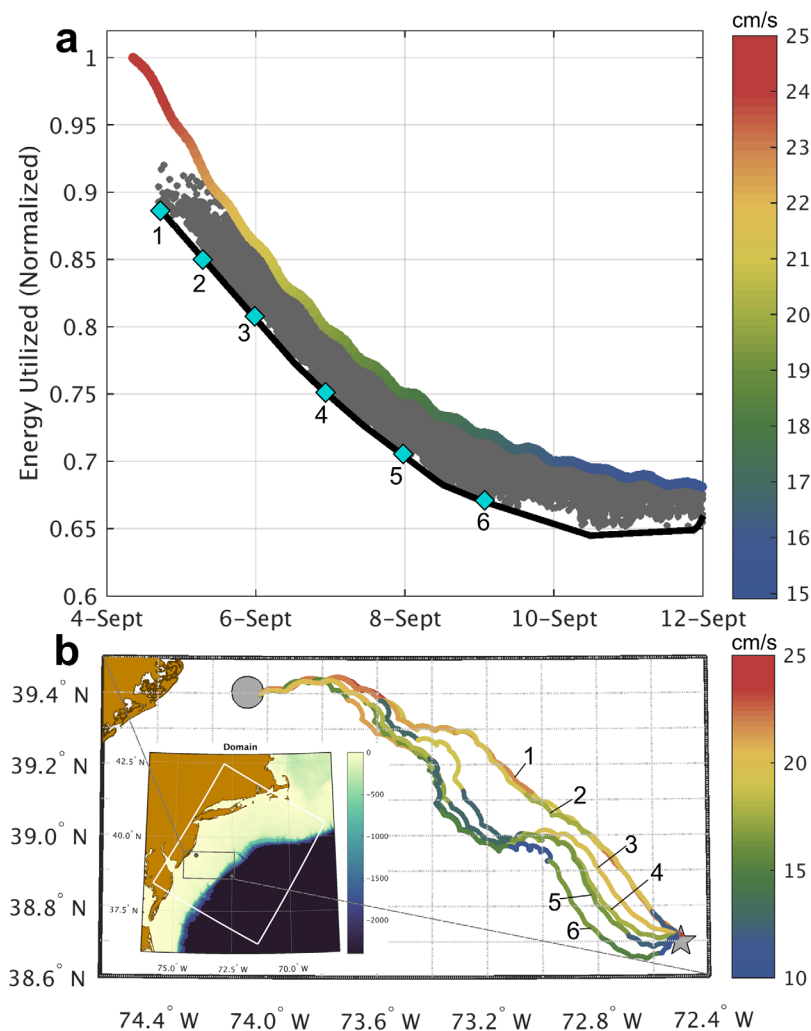


Figure 12. As Figure 7, but for energy-optimal path planning of gliders crossing the shelfbreak front from the coast of New Jersey. Longer duration gliders (#3–#6) stay on the shelf during the ocean response to storm Ernesto and are advected southward by its ocean response. Gliders #1 and #2, however, cross this region where the storm’s response is intensified before it does intensify; they are thus not much affected by the storm.

south than the end point, to finally sail northeastward again and reach the target. This latter pattern of using and crossing the jet is a realistic embodiment of the idealized energy-optimal crossing of a canonical steady front studied in *Subramani and Lermusiaux* [2016]. Considering tidal effects, they are less significant than in the first set of missions (Figure 8) since gliders sail further, southwestward and off the shelf. Nonetheless, tides induce energy-usage oscillations (Figure 12a) and spiral features in the paths (Figure 12b), both of which are more prominent on the tidally dominated shelf (see Figure 5) than beyond the shelfbreak. We note that Figure 12a shows samples for arrival times beyond 9 September until 12 September. We did not show however these winds and flows in Figures 4 and 6, respectively, because for these extra 3 days (9 September to 12 September) there is no strong flow/winds in the areas that affect our level-sets. The paths we show are also only until 9 September; the other paths that reach the target after 9 September have similar characteristics (hence are not shown).

4.3. Adaptive Replanning, Optimal Sampling, and Model Uncertainty

4.3.1. Adaptive replanning

The computational efficiency of the stochastic DO level-set optimization method for open-loop energy-optimal path planning allows its use for dynamic replanning. For example, as the glider relays its measurements

about local currents and other physical-bio-geo-chemical properties, they can be assimilated into the ocean model to obtain an updated data-assimilative forecast. Replanning can then be performed remotely with the updated flow fields or even on board if such computations are feasible. Similarly, if gliders or other vehicles collect data about the local currents, they could quickly reestimate new flow fields and our methodology could then be reapplied directly on board as well [Subramani and Lermusiaux, 2016]. These adaptive variants of our planning methodology are not illustrated in the present work, but our framework can be used to perform such replanning and onboard routing in the future. In fact, with such replanning and computational efficiency, the open-loop control solution is repetitively computed and the resulting adaptive control tends to the closed-loop control solution [Lermusiaux *et al.*, 2016]. For a discussion of computational costs, see Subramani and Lermusiaux [2016, Appendix A2]. Typically, for our coastal applications, the computational costs of integrating the stochastic DO level-set equations and the ocean model are of the same order. They allow similar real-time predictions and replanning.

4.3.2. Optimal Data Collection and Adaptive Sampling

Here we directly forecast energy-optimal paths to target locations that are decided or predicted a priori. Two examples of such missions where minimizing energy is critical include deploying a glider to reach a specific area in the ocean for subsequent data collection within that area, or recovering a glider at a specific location after its mission is completed. Another example where our theory and method directly applies consists of specifying an ordered set of optimal sampling locations that need to be sampled within a certain maximum time and to utilize our methodology to compute a path that reaches all of these sampling locations in the required order, minimizing energy within the allowed time. Of course, sampling missions where there is a preferred track or line for the glider to follow are not directly amenable to any energy-optimization: in such missions, the sensing vehicle must sample along that line and it is unlikely that an energy-optimal path would exactly be that line. What our reachable-set approach can do however is forecast if this sampling along a line is at all feasible or not [Lolla *et al.*, 2014a]. The approach can also remain compatible with optimal sampling goals if it is first used to compute energy-optimal paths for the prechosen set of sampling stations and then augmented with a method that determines if the energy-optimal paths in between station locations are still representative of the properties to be measured. If yes, the energy-optimal path could be used; if not, a path that is not energy-optimal should be used. For example, if the energy-optimal path catches a current (e.g., the Gulf Stream) to save energy, the resulting data sampled could have an undesired bias. For additional discussions on optimal sampling including adaptive sampling with or without time-optimality constraints, we refer for example to Lolla [2016] and Lermusiaux [2007], respectively.

4.3.3. Model Uncertainty

In the present work we considered the ocean forecasts to be deterministic. When model forecast uncertainties are available, they can nonetheless be rigorously incorporated into our optimal planning. In that case, the result is a probability of energy-optimal paths and a risk formulation corresponding to the uncertainty in the ocean forecast. Such planning in stochastic flow predictions is discussed in our other works [Wei, 2015; Lermusiaux *et al.*, 2016].

5. Conclusion

We demonstrated that fundamental PDEs can be used for energy-optimal path planning in the coastal ocean, integrating multiscale ocean forecasts with stochastic level-set optimization. Our focus was on glider-missions in the Middle-Atlantic Bight and shelfbreak front region for two sets of start and end points. We hindcast energy-optimal paths from among all time-optimal paths, for a range of arrival times during 28 August 2006 to 8 September 2006. The methodology—stochastic dynamically orthogonal (DO) level-set optimization—has three stages: (i) data-driven ocean forecasting to obtain realistic flow fields, (ii) stochastic DO level-set integration to compute the optimal arrival times and total energy utilized for a stochastic class of vehicle-speeds, and (iii) optimization to identify the energy-optimal vehicle-speeds, paths, and headings. In the first stage, the currents were hindcast using the data-assimilative MSEAS modeling system. Primitive equations were integrated with a nonlinear free-surface, realistic bathymetry, and wind and tidal forcing, over implicit two-way nested computational domains. Varied observations were used for initialization and assimilation, including data collected during the real-time SW06 and AWACS-06 experiments. In the second stage, we solved the stochastic DO level-set PDE that governs the time-optimal reachable sets for vehicles

that navigate with stochastic relative speeds in the ocean reanalysis. We then computed the joint probability density function (pdf) of total energy utilized, optimal arrival time, and vehicle-speeds. In the third and final stage, we selected the vehicle-speeds that minimize energy from the marginal energy-vehicle-speed pdf, and computed corresponding headings and paths by solving a particle backtracking equation. This completed the planning.

We studied energy-optimal results for two sets of missions. One set started from the coast of New Jersey to reach the Hudson River Canyon. The other started off the coast of Atlantic City, New Jersey, to sail beyond the shelfbreak in the AWACS/SW06 region. For the first set, we illustrated energy-optimal paths for a range of arrival times. We focused on the 02 UTC 6 September 2006 arrival and analyzed how currents advect the corresponding energy-optimal glider. We benchmarked this glider with a time-optimal constant-speed glider that arrived at the same time and found that the energy-optimal glider required $\sim 26\%$ less energy. This energy-optimal glider sailed slower at first so as to later ride the favorable ocean response to the tropical storm Ernesto. As a result, it was able to catch up with the constant-speed glider while critically using less energy overall. We also revealed and studied the effect of tides on the energy-usage and paths of optimal gliders, including the resulting spiral path patterns. We explained how tidal flows (or other strong currents) can lead to unattainable arrival times for constant-speed gliders. For the second set of missions, further south and beyond the shelfbreak, we also studied energy-optimal paths for a range of arrival times. These missions again displayed effects of storm Ernesto and of the tides on the shelf. Critically, they showcased how the shelfbreak front affects the energy-optimal paths, revealing that the shelfbreak jet is at first order a realistic embodiment of the canonical jet [Subramani and Lermusiaux, 2016].

Overall, the stochastic DO level-set PDEs derived from first principles provide a promising capability for prolonging the endurance of marine vehicles, by taking advantage of the ocean's dynamic environment. Particularly noteworthy is the computational efficiency: the DO level-set PDEs for $O(10^5)$ vehicle-speed time series were here 3–4 orders of magnitude faster than traditional Monte-Carlo schemes and were as fast or faster than solving the ocean primitive equations. Moreover, the unique combination of environmental forecasting with the fundamental level-set PDEs provide a template for energy-optimal path planning for drones or land robots navigating in dynamic air and land environments.

In the future, the approach can be extended to predict time- and energy-optimal paths in uncertain ocean currents [e.g., Wei, 2015; Lermusiaux et al., 2016]. Schemes can also be developed for energy-optimal path planning with onboard routing [Wang et al., 2009], with coordination for swarms of vehicles [Lolla et al., 2015] and with other ocean platforms such as kayaks [Xu et al., 2008] and ships [Mannarini et al., 2016; Mirabito et al., 2017]. Other avenues for research include energy-optimal adaptive sampling [Leonard et al., 2010] and adaptive modeling [Lermusiaux, 2007; Lermusiaux et al., 2007]. Energy-optimal planning for real gliders or propelled AUVs can also be completed in real-time, similarly to our recent real-time at-sea demonstration of time-optimal path planning and forecasting with AUVs [Subramani et al., 2017; Edwards et al., 2017].

Appendix A: A Stochastic Dynamically Orthogonal Level-Set Optimization for Energy-Optimal Path Planning

In this appendix, we first provide the stochastic DO level-set equations (box 2b in Figure 2). Finally, we describe the algorithm for the stochastic DO level-set optimization including the switch-sampling procedure and computational costs. Details of this randomized algorithm are in Subramani and Lermusiaux [2016]. At the end of the Appendix, we provide a list of the acronyms and notation employed in the manuscript.

A1. Stochastic DO Level-Set Equations

To solve the S-PDE equation (3), we employ a dynamic stochastic model-order reduction, the Dynamically Orthogonal decomposition. The DO method achieves model-order reduction by exploiting the nonlinearities of the flow $\mathbf{v}(\mathbf{x}, t)$, which tend to concentrate the responses of the scalar level-set field (ϕ) into specific dynamic patterns. The reduced DO level-set equations are computationally less expensive to solve. Next, we outline the derivation of the DO level-set equations, as a summary of Subramani and Lermusiaux [2016].

To derive the DO level-set equations, we first introduce a generalized dynamic Karhunen-Loeve decomposition (i.e., a DO decomposition) of the level-set field ($\phi(\mathbf{x}, t; \omega)$) and the stochastic vehicle-speed ($F(t; \omega)$),

$$\phi(\mathbf{x}, t; \omega) = \bar{\phi}(\mathbf{x}, t) + \sum_{i=0}^{n_{s,\phi}} Y_i(t; \omega) \tilde{\phi}_i(\mathbf{x}, t), \tag{A1a}$$

$$F(t; \omega) = \bar{F}(t) + \sum_{i=0}^{n_{s,F}} z_i(t; \omega) \tilde{F}_i(t), \tag{A1b}$$

into the stochastic level-set equation (3) to obtain the DO-expanded equation

$$\begin{aligned} \frac{\partial}{\partial t} \left(\bar{\phi}(\mathbf{x}, t) + \sum_{i=0}^{n_{s,\phi}} Y_i(t; \omega) \tilde{\phi}_i(\mathbf{x}, t) \right) &= \left(\bar{F}(t) + \sum_{i=0}^{n_{s,F}} z_i(t; \omega) \tilde{F}_i(t) \right) \\ \left| \nabla \left(\bar{\phi}(\mathbf{x}, t) + \sum_{i=0}^{n_{s,\phi}} Y_i(t; \omega) \tilde{\phi}_i(\mathbf{x}, t) \right) \right| &\mathbf{v}(\mathbf{x}, t) \cdot \nabla \left(\bar{\phi}(\mathbf{x}, t) + \sum_{i=0}^{n_{s,\phi}} Y_i(t; \omega) \tilde{\phi}_i(\mathbf{x}, t) \right). \end{aligned} \tag{A2}$$

Here $\bar{\phi}(\mathbf{x}, t)$ is the mean, $\tilde{\phi}_i(\mathbf{x}, t)$ s are the spatial modes, and $Y_i(t; \omega)$ s are the stochastic coefficients of the scalar level-set field $\phi(\mathbf{x}, t; \omega)$. Similarly, $\bar{F}(t)$ is the mean, $\tilde{F}_i(t)$ is the spatial-mode, and $z_i(t; \omega)$ is the stochastic coefficient of the vehicle-speed $F(t; \omega)$. By definition, the stochastic coefficients are zero-mean stochastic processes. Next, we apply the expectation operator on equation (A2) to obtain the mean PDE, we multiply equation (A2) with the coefficients and apply the expectation operator to obtain the mode PDEs, and we project the equation (A2) on the modes to obtain the coefficient ODEs. The equations thus obtained—the stochastic DO level-set equations—are as follows:

$$\frac{\partial \bar{\phi}}{\partial t} = -(\bar{F} \mathbb{E}[\gamma] + \mathbb{E}[z\gamma] \bar{F}) - \mathbf{v} \cdot \nabla \bar{\phi}, \tag{A3}$$

$$\frac{dY_i}{dt} = -(\bar{F}(\gamma - \mathbb{E}[\gamma]) + \tilde{F}(z\gamma - \mathbb{E}[z\gamma]) + Y_k \mathbf{v} \cdot \nabla \tilde{\phi}_k, \tilde{\phi}_i), \tag{A4}$$

$$\begin{aligned} \frac{\partial \tilde{\phi}_i}{\partial t} &= -C_{Y_i Y_i}^{-1} (\bar{F} \mathbb{E}[Y_i \gamma] + \tilde{F} \mathbb{E}[z Y_i \gamma]) + \mathbf{v} \cdot \nabla \tilde{\phi}_i, \\ &- \langle -C_{Y_i Y_i}^{-1} (\bar{F} \mathbb{E}[Y_j \gamma] + \tilde{F} \mathbb{E}[z Y_j \gamma]) + \mathbf{v} \cdot \nabla \tilde{\phi}_i, \tilde{\phi}_n \rangle \tilde{\phi}_n, \end{aligned} \tag{A5}$$

where we have dropped the parenthesis denoting spatial, temporal, and stochastic indices for ease of notation. Also, $\gamma \equiv |\nabla \phi|$, \mathbb{E} is the expectation operator and the spatial inner product over a domain Ω for any two fields $u(\mathbf{x}, t)$ and $v(\mathbf{x}, t)$ in Ω is defined as

$$\langle u(\mathbf{x}, t), v(\mathbf{x}, t) \rangle = \int_{\Omega} u(\mathbf{x}, t)^T v(\mathbf{x}, t) d\mathbf{x}. \tag{A6}$$

A key step in deriving the DO equations is handling the non-polynomial nonlinearity γ . The above equations (A1a) and (A1b) do not assume a DO decomposition for this γ norm. Other equivalent DO equations which consider an explicit DO decomposition for γ have also been derived in *Subramani and Lermusiaux* [2016], but not utilized in the present paper.

A2. Algorithm and Computational Cost

The algorithm is listed in Table A1. It consists of three core successive steps, the computation of the velocity field from an ocean modeling system; stochastic simulation for a given function space of vehicle-speeds; and optimization for energy-optimal paths. The second and third core steps can be iterated if the function space of vehicle-speeds is resampled or improved by (machine) learning. For the stochastic simulation, the stochastic class $F(t; \omega)$ is represented as realizations $F_{DO}(t; r)$, where $r = 1 \dots n_r$ is the realization index and the subscript denotes that it is a DO decomposition. For initialization of the DO representation of the level-set field $\phi(\mathbf{x}, t; \omega)$, several choices exist and all commonly lead to the same overall results [Subramani and Lermusiaux, 2016]. Here we choose the initial modes as the left singular vectors of the matrix of level-set fields saved at every time step from a very short deterministic simulation of the fastest $F_{DO}(t; r)$. This choice ensures that all other level-set evolutions are within that largest reachability set. The stochastic simulation steps compute the energy

Table A1. Stochastic DO Level-Set Optimization: Algorithm

Ocean Modeling System

1. Obtain the velocity-field forecast from a realistic simulation of the region of interest for the planning horizon using an ocean modeling system.

Stochastic DO Level-Set Simulation

2. Sample $r=1, \dots, n_r$ realizations of relative vehicle-speeds for the planning horizon from a class of stochastic processes and its DO (KL) representation $F_{DO}(\cdot; r) = \bar{F}(\cdot) + z(\cdot, r)\bar{F}$.
3. Solve the stochastic DO level-set equations ((A1a)a) and (A1b), with n_r realizations and $n_{s,\phi}$ modes to compute the n_r optimal arrival times at the target $T(\mathbf{x}_f; F_{DO}(\cdot; r))$ corresponding to each relative-speed realization $F_{DO}(\cdot; r)$.
4. Compute the energy utilization for each realization using $E(r) = \int_0^T(\mathbf{x}_f; F_{DO}(\cdot; r))p(t; F_{DO}(\cdot; r))dt$.

Optimization

5. From the stochastic DO level-set simulation, identify all realizations that arrive at the target in each of the queried arrival time windows.
6. For each of these subsets of realizations, select the realization that utilizes the minimum energy, and compute its path using the backtracking equation (2).

Iterate

7. If needed, resample the function space $F(\cdot; \omega)$ or expand this function space by (machine) learning

distribution $E(r) = \int_0^{T(\mathbf{x}_f; F_{DO}(\cdot; r))} p(t; r)dt$ for all realizations. To find energy-optimal vehicle-speeds, we perform the optimization $F_{DO}^*(t) = \arg \min_{F_{DO}(\cdot; r)} E(r)$. The result is energy-optimal paths among all time-optimal paths for vehicles navigating in the coastal ocean.

A2.1. Switch-Sampling Procedure

To obtain realizations $F(t; r)$, the stochastic class $F(t; \omega)$ is sampled by *switch sampling* [Subramani and Lermusiaux, 2016], a structured approach to sample the large combinatorial space of discretized vehicle-speed time series. For l levels of discretized speed and $n_{sw,F}$ discrete times, the total size of this combinatorial space is $l^{n_{sw,F}}$. In *switch sampling*, we use a coarser resolution in both discrete speeds and switching times than in direct sampling, but a randomization is employed within this coarser resolution that aims to remedy for this. We are aided by the fact that the structured and relatively strong environmental flows are such that ranges of speeds lead to very similar behavior: these ranges do not need to be sampled with refined grids, the transitions due to flow forcing is what needs to be captured. Since a structured approach is taken to sample the large combinatorial space, with enough samples, the method can adaptively find the energy-optimal speed time series and thus paths. Figure A1 shows four samples obtained using this procedure for




Figure A1. Four realizations sampled using the *switch-sampling* scheme from a stochastic class where vehicle-speeds *switch* every 48 h.

a 16 day duration.

A2.2. Computational Cost

The computational cost of the direct Monte Carlo simulation is n_r times the cost of one deterministic solution of the level-set PDE equation (1b). The DO method solves a mean PDE and $n_{s,\phi}$ mode PDEs and n_r coefficient ODEs. The mean and mode PDEs are similar to the level-set PDE and the coefficient ODEs are cheaper than the PDE-solve. Hence, the net cost is of the order of $(1+n_{s,\phi})$ solutions of the level-set PDE. Since in our applications $(1+n_{s,\phi}) \ll n_r$, the DO method is 3–4 orders of magnitude cheaper than MC.

The cost of one deterministic solution is $O(n_x n_y)$ and

the cost of one DO solution is $O((1+n_{s,\phi})n_x n_y)$. The cost of integrating the primitive equations by the MSEAS modeling system is $O(4n_x n_y n_z)$. In the present case $n_x=175$, $n_y=150$, $n_z=100$, $n_{s,\phi}=120$. For the current implementation, the CPU time requirement in a single CPU with four cores for integrating the PE for 1 day with two-way nesting is ~ 1.5 h and the DO level-set equations is also ~ 1.5 h. Further implementation improvements can be achieved, for example, by using narrowband approaches, parallel computing, and/or by reducing the level-set computation to a smaller domain estimated by the maximum distance a glider may travel for the planning horizon.

Notation

Scalar

$i \in \mathbb{N}$	stochastic subspace index.
$F(\bullet) \in \mathbb{R}$	vehicle-speed functional.
$\bar{F}(\bullet) \in \mathbb{R}$	mean vehicle-speed functional.
$\tilde{F}(\bullet) \in \mathbb{R}$	DO mode (spatial) of vehicle-speed functional.
$n_x, n_y, n_z \in \mathbb{N}$	number of grid points in the x , y , and z directions, respectively.
$n_g \in \mathbb{N}$	total number of grid points in the horizontal direction.
$n_{s,\phi} \in \mathbb{N}$	dimension of the stochastic subspace of level-set.
$n_{s,F} \in \mathbb{N}$	dimension of the stochastic subspace (in time) of vehicle-speed functional.
$Y_i \in \mathbb{R}$	random variable describing the pdf of the orthonormal level-set(ϕ) modes $\tilde{\phi}_i$.
$z \in \mathbb{R}$	random variable describing the pdf of the vehicle-speed F .
$n_r \in \mathbb{N}$	number of Monte Carlo realizations of the level-sets (and vehicle-speeds).
$n_{sw,F} \in \mathbb{N}$	number of switches made in the vehicle-speed throughout the duration of the simulation for the switch-sampling algorithm.
$l \in \mathbb{N}$	number of discrete levels into which the vehicle-speed spread is divided for the switch-sampling algorithm.
$r \in \mathbb{N}$	realization index.
$T \in \mathbb{R}^+$	optimal arrival time random variable.

Vector

$\phi \in \mathbb{R}^{n_g}$	level-set field.
$\bar{\phi} \in \mathbb{R}^{n_g}$	mean level-set field.
$\tilde{\phi}_i \in \mathbb{R}^{n_g}$	DO mode i of ϕ : Dynamically orthogonal basis for the stochastic subspace of ϕ .
$\gamma \in \mathbb{R}^{n_g}$	$ \nabla\phi $.

Acronyms

DO	Dynamically Orthogonal.
PDE	Partial Differential Equation.
MC	Monte Carlo.

References

- Alvarez, A., A. Caiti, and R. Onken (2004), Evolutionary path planning for autonomous underwater vehicles in a variable ocean, *IEEE J. Oceanic Eng.*, 29(2), 418–429.
- Athans, M., and P. Falb (2007), *Optimal Control: An Introduction to the Theory and Its Applications*, Dover Books Eng. Ser., New York.
- Bellingham, J. G., and K. Rajan (2007), Robotics in remote and hostile environments, *Science*, 318(5853), 1098–1102.
- Chapman, N. R., and J. F. Lynch (2010), Special issue on the 2006 shallow water experiment, *IEEE J. Oceanic Eng.*, 35(1), 1–2.
- Edwards, J., J. Smith, A. Girard, D. Wickman, D. N. Subramani, C. S. Kulkarni, J. P. J. Haley, C. Mirabito, S. Jana, and P. F. J. Lermusiaux (2017), Data-driven learning and modeling of AUV operational characteristics for optimal path planning, in *Oceans '17 MTS/IEEE Conference*, in press.
- Egbert, G. D., and S. Y. Erofeeva (2002), Efficient inverse modeling of barotropic ocean tides, *J. Atmos. Oceanic Technol.*, 19(2), 183–204.
- Gangopadhyay, A., A. R. Robinson, P. J. Haley Jr., W. G. Leslie, C. J. Lozano, J. J. Bisagni, and Z. Yu (2003), Feature-oriented regional modeling and simulations in the Gulf of Maine and Georges Bank, *Cont. Shelf Res.*, 23(3), 317–353.
- Haley, P. J., Jr., and P. F. J. Lermusiaux (2010), Multiscale two-way embedding schemes for free-surface primitive equations in the “Multidisciplinary Simulation, Estimation and Assimilation System,” *Ocean Dyn.*, 60(6), 1497–1537, doi:10.1007/s10236-010-0349-4.
- Haley, P. J., Jr., A. Agarwal, and P. F. J. Lermusiaux (2015a), Optimizing velocities and transports for complex coastal regions and archipelagos, *Ocean Modell.*, 89, 1–28, doi:10.1016/j.ocemod.2015.02.005.

Acknowledgments

We thank the members of our MSEAS group particularly T. Lolla for useful discussions. We are grateful to the Office of Naval Research (ONR) for research support under grants N00014-09-1-0676 (Science of Autonomy - A-MISSION), N00014-14-1-0476 (Science of Autonomy - LEARNS), and N00014-15-1-2616 (NASCar-OPS) to the Massachusetts Institute of Technology (MIT). We also thank the MIT-Tata Center Program for the Fellowship support of DNS. The MSEAS SW06-AWACS-06 ocean reanalyses employed were made possible by research support to MIT from ONR under grant N00014-11-1-0701 (MURI-IODA) and the National Science Foundation under grant OCE-1061160. We thank all of our SW06 and AWACS-06 colleagues, especially G. Gawarkiewicz, P. Abbot, and T. Duda for their ocean data and M. Taylor and J. Hare for their NMFS survey data. We also thank J. Evans, S. Glenn, and J. Wilkin for their real-time WRF atmospheric fluxes and the FNMOC teams for their own products. Data used in this paper are available at <http://mseas.mit.edu/Research/IODA/index.html>. We are grateful to the Editor and the anonymous reviewers for their useful feedback. D.N.S. thanks the MIT Writing and Communication Center for their advice.

- Haley, P. J., Jr., P. F. J. Lermusiaux, and S. Kelly (2015b), Multiscale dynamic relaxation and sponge open boundary conditions, MSEAS report, Dep. of Mech. Eng., Mass. Inst. of Technol., Cambridge.
- He, R., and J. L. Wilkin (2006), Barotropic tides on the southeast new England shelf: A view from a hybrid data assimilative modeling approach, *J. Geophys. Res.*, *111*, C08002, doi:10.1029/2005JC003254.
- Hsieh, M. A., E. Forgoston, T. W. Mather, and I. B. Schwartz (2012), Robotic manifold tracking of coherent structures in flows, in *IEEE International Conference on Robotics and Automation*, pp. 4242–4247, IEEE, St. Paul, Minn.
- Inanc, T., S. C. Shadden, and J. E. Marsden (2005), Optimal trajectory generation in ocean flows, in *Proceedings of American Control Conference*, vol. 1, pp. 674–679, IEEE, Portland, Oreg.
- Keen, T. R., and S. M. Glenn (1995), A coupled hydrodynamic-bottom boundary layer model of storm and tidal flow in the Middle Atlantic bight of North America, *J. Phys. Oceanogr.*, *25*(3), 391–406.
- Kelly, S. M., and P. F. J. Lermusiaux (2016), Internal-tide interactions with gulf stream and Middle Atlantic bight shelfbreak front, *J. Geophys. Res. Oceans*, *121*, 6271–6294, doi:10.1002/2016JC011639.
- Kohut, J. T., S. M. Glenn, and R. J. Chant (2004), Seasonal current variability on the New Jersey inner shelf, *J. Geophys. Res.*, *109*, C07S07, doi:10.1029/2003JC001963.
- Kruger, D., R. Stolkin, A. Blum, and J. Briganti (2007), Optimal AUV path planning for extended missions in complex, fast-flowing estuarine environments, in *2007 IEEE International Conference on Robotics and Automation*, pp. 4265–4270, IEEE, Roma, Italy.
- Latombe, J.-C. (1991), *Robot Motion Planning*, Kluwer Acad., New York.
- Lentz, S. J. (2008), Observations and a model of the mean circulation over the Middle Atlantic Bight continental shelf, *J. Phys. Oceanogr.*, *38*(6), 1203–1221.
- Leonard, N. E., D. A. Paley, R. E. Davis, D. M. Fratantoni, F. Lekien, and F. Zhang (2010), Coordinated control of an underwater glider fleet in an adaptive ocean sampling field experiment in monterey bay, *J. Field Rob.*, *27*(6), 718–740.
- Lermusiaux, P. F. J. (1999), Data assimilation via error subspace statistical estimation, part II: Mid-Atlantic Bight shelfbreak front simulations, and ESSE validation, *Mon. Weather Rev.*, *127*(7), 1408–1432, doi:10.1175/1520-0493(1999)127<1408:DAVESS>2.0.CO;2.
- Lermusiaux, P. F. J. (2002), On the mapping of multivariate geophysical fields: Sensitivities to size, scales, and dynamics, *J. Atmos. Oceanic Technol.*, *19*(10), 1602–1637.
- Lermusiaux, P. F. J. (2007), Adaptive modeling, adaptive data assimilation and adaptive sampling, *Physica D*, *230*(1), 172–196.
- Lermusiaux, P. F. J., P. J. Haley Jr., W. G. Leslie, O. Logutov, and A. R. Robinson (2006), Autonomous wide aperture cluster for surveillance (AWACS): Adaptive sampling and search using predictive models with coupled data assimilation and feedback. [Available at http://mseas.mit.edu/archive/AWACS/index_AWACS.html]
- Lermusiaux, P. F. J., P. J. Haley Jr., and N. K. Yilmaz (2007), Environmental prediction, path planning and adaptive sampling: sensing and modeling for efficient ocean monitoring, management and pollution control, *Sea Technol.*, *48*(9), 35–38.
- Lermusiaux, P. F. J., T. Lolla, P. J. Haley Jr., K. Yigit, M. P. Ueckermann, T. Sondergaard, and W. G. Leslie (2016), Science of autonomy: Time-optimal path planning and adaptive sampling for swarms of ocean vehicles, in *Springer Handbook of Ocean Engineering: Autonomous Ocean Vehicles, Subsystems and Control*, edited by T. Curtin, chap. 21, pp. 481–498, Springer, New York, doi:10.1007/978-3-319-16649-0_21.
- Lin, Y.-T., A. E. Newhall, T. F. Duda, P. F. J. Lermusiaux, and P. J. Haley (2010), Merging multiple-partial-depth data time series using objective empirical orthogonal function fitting, *IEEE J. Oceanic Eng.*, *35*(4), 710–721, doi:10.1109/JOE.2010.2052875.
- Logutov, O. G. (2008), A multigrid methodology for assimilation of measurements into regional tidal models, *Ocean Dyn.*, *58*(5–6), 441–460, doi:10.1007/s10236-008-0163-4.
- Logutov, O. G., and P. F. J. Lermusiaux (2008), Inverse barotropic tidal estimation for regional ocean applications, *Ocean Modell.*, *25*(1–2), 17–34, doi:10.1016/j.ocemod.2008.06.004.
- Lolla, S. V. T. (2016), Path planning and adaptive sampling in the coastal ocean, PhD thesis, Dep. of Mech. Eng., Mass. Inst. Technol., Cambridge.
- Lolla, T., and P. F. J. Lermusiaux (2016), A forward reachability equation for minimum-time path planning in strong dynamic flows, *MIT MSEAS Report*.
- Lolla, T., P. F. J. Lermusiaux, M. P. Ueckermann, and P. J. Haley Jr. (2014a), Time-optimal path planning in dynamic flows using level set equations: Theory and schemes, *Ocean Dyn.*, *64*(10), 1373–1397, doi:10.1007/s10236-014-0757-y.
- Lolla, T., P. J. Haley Jr., and P. F. J. Lermusiaux (2014b), Time-optimal path planning in dynamic flows using level set equations: Realistic applications, *Ocean Dyn.*, *64*(10), 1399–1417, doi:10.1007/s10236-014-0760-3.
- Lolla, T., P. J. Haley Jr., and P. F. J. Lermusiaux (2015), Path planning in multiscale ocean flows: Coordination and dynamic obstacles, *Ocean Modell.*, *94*, 46–66, doi:10.1016/j.ocemod.2015.07.013.
- Mannarini, G., N. Pardini, G. Coppini, P. Oddo, and A. Iafrazi (2016), VISIR-I: small vessels – least-time nautical routes using wave forecasts, *Geosci. Model Dev.*, *9*, 1597–1625, doi:10.5194/gmd-9-1597-2016.
- Michini, M., M. A. Hsieh, E. Forgoston, and I. B. Schwartz (2014), Robotic tracking of coherent structures in flows, *IEEE Trans. Rob.*, *30*(3), 593–603, doi:10.1109/TRO.2013.2295655.
- Mirabito, C., et al. (2017), Autonomy for surface ship interception, in *Oceans '17 MTS/IEEE Conference*, in press.
- Newhall, A. E., et al. (2007), Acoustic and oceanographic observations and configuration information for the WHOI moorings from the SW06 experiment, *Tech. Rep. WHOI-2007-04*, Woods Hole Oceanogr. Inst., Woods Hole, Mass.
- Ramp, S. R., et al. (2009), Preparing to predict: The second Autonomous Ocean Sampling Network (AOSN-II) experiment in the Monterey Bay, *Deep Sea Res., Part II*, *56*(3–5), 68–86, doi:10.1016/j.dsr2.2008.08.013.
- Rao, D., and S. B. Williams (2009), Large-scale path planning for underwater gliders in ocean currents, in *Proceedings of Australasian Conference on Robotics and Automation*, pp. 28–35, ARAA, Sydney, Australia.
- Rudnick, D. L., R. E. Davis, C. C. Eriksen, D. M. Fratantoni, and M. J. Perry (2004), Underwater gliders for ocean research, *Mar. Technol. Soc. J.*, *38*(2), 73–84.
- Subramani, D. N. (2014), Energy optimal path planning using stochastic dynamically orthogonal level set equations, MS thesis, Sch. of Eng., Mass. Inst. of Technol., Cambridge.
- Subramani, D. N., and P. F. J. Lermusiaux (2016), Energy-optimal path planning by stochastic dynamically orthogonal level-set optimization, *Ocean Modell.*, *100*, 57–77.
- Subramani, D. N., T. Lolla, P. J. Haley Jr., and P. F. J. Lermusiaux (2015), A stochastic optimization method for energy-based path planning, in *DyDESS 2014, LNCS*, edited by S. Ravela and A. Sandu, vol. 8964, pp. 1–12, Springer, Boston, Mass., doi:10.1007/978-3-319-25138-731.
- Subramani, D. N., P. F. J. Lermusiaux, J. P. J. Haley, C. Mirabito, S. Jana, C. S. Kulkarni, A. Girard, D. Wickman, J. Edwards, and J. Smith (2017), Time-optimal path planning: Real-time sea exercises, in *Oceans '17 MTS/IEEE Conference*, in press.

- Tang, D., et al. (2007), Shallow water06: A joint acoustic propagation/nonlinear internal wave physics experiment, *Oceanography*, 20(4), 156–167.
- Wang, D., P. F. J. Lermusiaux, P. J. Haley Jr., D. Eickstedt, W. G. Leslie, and H. Schmidt (2009), Acoustically focused adaptive sampling and on-board routing for marine rapid environmental assessment, *J. Mar. Syst.*, 78, supplement, S393–S407, doi:10.1016/j.jmarsys.2009.01.037.
- Wei, Q. J. (2015), Time-optimal path planning in uncertain flow fields using stochastic dynamically orthogonal level set equations, Bachelor's thesis, Dep. of Mech. Eng., Mass. Inst. of Technol., Cambridge.
- WHOI (2006), Shallow water experiment 2006. [Available at <http://acoustics.whoi.edu/sw06/>, (Last Accessed 24 July 2016).]
- Witt, J., and M. Dunbabin (2008), Go with the flow: Optimal AUV path planning in coastal environments, in *Proceedings of Australasian Conference on Robotics and Automation*, edited by J. Kim and R. Mahony, pp. 86–94, ARAA. Canberra, Australia.
- Xu, J., P. F. J. Lermusiaux, P. J. Haley Jr., W. G. Leslie, and O. G. Logutov (2008), Spatial and temporal variations in acoustic propagation during the PLUSNet-07 exercise in Dabob Bay, in *Proceedings of Meetings on Acoustics (POMA)*, vol. 4, p. 11, Acoust. Soc. of Am., Aix en Provence, France, doi:10.1121/1.2988093.
- Zhang, W., T. Inanc, S. Ober-Blobaum, and J. E. Marsden (2008), Optimal trajectory generation for a glider in time-varying 2D ocean flows B-spline model, in *IEEE International Conference on Robotics and Automation*, pp. 1083–1088, IEEE Pasadena, Calif.