21. Science of Autonomy: Time-Optimal Path Planning and Adaptive Sampling for Swarms of Ocean Vehicles

Pierre F.J. Lermusiaux, Tapovan Lolla, Patrick J. Haley Jr., Konuralp Yigit, Mattheus P. Ueckermann, Thomas Sondergaard, Wayne G. Leslie

The science of autonomy is the systematic development of fundamental knowledge about autonomous decision making and task completing in the form of testable autonomous methods, models and systems. In ocean applications, it involves varied disciplines that are not often connected. However, marine autonomy applications are rapidly growing, both in numbers and in complexity. This new paradigm in ocean science and operations motivates the need to carry out interdisciplinary research in the science of autonomy. This chapter reviews some recent results and research directions in time-optimal path planning and optimal adaptive sampling. The aim is to set a basis for a large number of vehicles forming heterogeneous and collaborative underwater swarms that are smart, i. e., knowledgeable about the predicted environment and their uncertainties, and about the predicted effects of autonomous sensing on future operations. The methodologies are generic and applicable to any swarm that moves and senses dynamic environmental fields. However, our focus is underwater path planning and adaptive sampling with a range of vehicles such as autonomous underwater vehicles (AUVs), gliders, ships or remote sensing platforms.

Ocean observing is the process of collecting ocean measurements to provide the most useful information about the ocean systems and dynamics of interest. Such marine sensing is expensive and challenging. In the ocean, not everything can be measured on a sustained basis. Ocean modeling is the process of developing and utilizing theoretical and computational models for the understanding and prediction of ocean dynamics. Similar to ocean observing, ocean modeling is challenging, and not every naval and environmental process can be modeled exactly. Stochastic ocean modeling explicitly acknowledges these uncertainties in modeling and thus augments deterministic predictions with probabilistic information. Data assimilation is the process of quantitatively estimating dynamically evolving fields by combining information from observations with those predicted by models, hence limiting uncertainties. With stochastic predictive capabilities and data assimilation, one can estimate future conditions and their uncertainties, and make informed decisions and complete desired tasks. Systems that integrate some of these activities without continuous human guidance are often referred to as autonomous systems.

The science of autonomy is the systematic development of fundamental knowledge about autonomous decision making and task completing in the form of testable autonomous methods, models, and systems. In ocean applications, it involves varied disciplines that are not often connected. For example, ocean dynamics, sensing, robotics, communications, predictions, and uncertainty, as well as varied types of human–computer–vehicle interactions are inputs to ocean autonomy. There are two important ocean modeling feedbacks to autonomy. The first one is the use of models to predict optimal

paths for swarms of ocean vehicles and the second is the use of models to predict the sensing expected to be most useful. Prediction of optimal paths for ocean vehicles is referred to here as *path planning*. This search for optimality often focuses on time-optimal or energy-optimal paths. The prediction of optimal sensing strategies is referred to as *adaptive sampling*. We note that in both cases, the prediction of paths can itself be improved adaptively in the sense that observations collected along the path can be utilized to update the predictions. Such updates can be completed onboard vehicles in the form of *onboard routing* or by central/hierarchical computational entities for *high-level routing*.

Recent autonomous marine activities have started a revolution that was possibly first imagined by *Henry Stommel* in his now famous article, *The Slocum mission* [21.1]. Today, the number of autonomous platforms used in semi-coordinated sea operations can be larger than 10, and this number is increasing. This new paradigm in ocean science and operations motivates the need to carry out interdisciplinary research in the science of autonomy. This chapter reviews some recent fundamental results and research directions in time-optimal path planning and optimal adaptive sampling, setting a basis for a large number of vehicles forming heterogeneous and collaborative underwater swarms that are smart, i. e., knowledgeable about the predicted environment and their uncertainties, and about the predicted effects of autonomous sensing on future operations. The methodologies are generic and applicable to any swarm that moves and senses dynamic environmental fields. However, the focus of this chapter is underwater path planning and adaptive sampling with a range of vehicles such as AUVs, gliders, ships, or remote sensing platforms.

When compared to related control and robotics problems of large dimensions, the differences with the present focus on ocean applications are that:

- Naval platforms are heterogeneous and their data are gappy but multivariate.
- Marine fields can strongly affect vehicles' operations, are dynamic on multiple-scales, and have very large dimensions, but are predictable to some degree.
- The measurements to be collected affect these future predictions.

Therefore, there are feedbacks between optimal planning, sampling and predicting, in time and space, and across variables. The field of ocean autonomy thus uses guidance from ocean flow and acoustic propagation modeling, information theory, dynamical system theory, uncertainty prediction, decision-making under uncertainty, machine learning, bio-inspired algorithms, and distributed computing. In all cases, our focus here is to predict global dynamic swarm patterns and perform high-level optimization. This is not the detailed fully-resolved control of a single robot. In our focus, a global objective function defines the optimal dynamic and collaborative autonomy. In general, objective functions depend on the predicted environment, on the predicted values and positions of the expected measurements, and on the feedbacks between measurements and predicted dynamics.

In what follows, we focus on the science of autonomy results obtained by our multidisciplinary simulation, estimation, and assimilation systems (MSEAS) group and recent collaborators. We first outline and present a subset of our results on time-optimal path planning for swarms of ocean vehicles (Sect. 21.1). This is followed by an overview of some of our results in adaptive sampling (Sect. 21.2). Our objective is to outline overall schemes and approaches, and to illustrate their applications in idealized and realistic ocean conditions. For detailed descriptions of the algorithms and methodologies we developed and utilized, we refer to references cited in each of these sections. For general reviews on oceanic path planning, we refer to [21.2] and [21.3], and for general reviews on oceanic adaptive sampling, to [21.4–8], and [21.9]. Brief conclusions and outlook are given in Sect. 21.3.

21.1 Time-Optimal Path Planning for Swarms of Ocean Vehicles

Our new level-set based method for time-optimal path planning of vehicle swarms in time-varying uncertain velocity fields (ocean currents) was developed and described in [21.2, 3, 10–14]. To forecast the time-optimal trajectories of large numbers of robotic vehicles navigating in strong and dynamic flows, our approach first evolves the reachable sets from the starting points. A reachable set is the set of points that can be visited by a vehicle at a given time. The boundary of such a set is called the reachability front. Our approach keeps track of the reachability front at all times, so that we can determine when the front first reaches the end point. A path traced by a point on the reachability front that first reaches the end point is an optimal path that we wish to compute.

We obtained a modified level set (21.1) for ϕ^o, the scalar field that governs the evolution of reachability fronts [21.3, 10, 11, 14], linking it to the Hamilton–

Jacobi equation governing the minimum-time navigation problem in dynamic flows. The reachability fronts are evolved from the vehicle start points y_s until they reach the desired end points y_f, combining nominal vehicle motions due to steering and advection by the dynamic flow environment, as shown by (21.1). An optimal trajectory $X_P^*(y_s, t)$ and the corresponding vehicle heading directions are then extracted from the evolution of the reachability fronts. This is done by solving the backtracking (21.2) where the solution ϕ^o is differentiable, starting from an end point y_f and integrating backward in time to the start point y_s

$$\frac{\partial \phi^o}{\partial t} + F|\nabla \phi^o| + v \cdot \nabla \phi^o = 0\,, \tag{21.1}$$

$$\frac{\mathrm{d}\, X_P^*(y_s, t)}{\mathrm{d}t} = -v(X_P^*, t) - F\frac{\nabla \phi^o(X_P^*, t)}{|\nabla \phi^o(X_P^*, t)|} \quad \text{with } X_P^*(y_s, T^*(y_f)) = y_f\,. \tag{21.2}$$

We also extended the methodology to the case of coordinated motions of swarms of vehicles and to the case of uncertain stochastic flows. Coordinated motion in this context refers to maintaining swarm formations of specific geometrical shapes by the vehicles, regardless of the ocean currents. Accounting for uncertainties in the predicted ocean fields allows us to yield the paths that best compromise between least uncertainty and least total travel time.

Our approach is interdisciplinary: It is inspired by ideas in fluid and ocean sciences, computational sciences, numerical methods, and uncertainty quantification, and applies them to path planning, which has roots in robotics and optimal control. The methodology is based on solving (stochastic) partial differential equations. In this sense, it is rigorous and obviates the need for heuristics. To compute the level set evolution, i. e., the viscosity solution to the Hamilton–Jacobi equation, we either use our finite-volume code for partial differential equations (PDE) (with varied orders, limiters, etc.) or narrow-band level set and fast marching methods for efficient computations. For the uncertain flow case, we employ our stochastic PDEs solvers [21.15]. In all cases, the computational cost increases only geometrically with the number of grid points (and not with the path length). We refer to the above-mentioned references for descriptions of numerical algorithms. Next, we illustrate results under diverse fluid and ocean conditions, as well as diverse vehicle behaviors.

21.1.1 Canonical Steady Flows

We first evaluated our level-set-based methodology by computing time-optimal paths in canonical steady ocean flow features. Such problems include, for example, determining how to time-optimally enter or leave an ocean eddy, cross an ocean jet, or maneuver in a meander/wave field [21.13]. The results can be very useful as rules-of-thumb for time-optimal operations.

To exemplify such canonical steady flow problems, we consider idealized 2-D-in-space (x, y) non-dimensional *highway* scenarios (2-D: two-dimensional), as illustrated in Fig. 21.1 [21.13]. In the first case (Fig. 21.1a), the vehicle is deployed from $(0, -0.4)$, with the goal of reaching the point at $(0, 0.8)$. The speed of vehicle is 1. The flow consists of two different jets, both flowing in the positive y direction. The first jet (narrow band with black arrows on the left) has a higher speed, $v_1 = 2$, than the vehicle. The second jet (narrow band with red arrows on the right) has a constant speed equal to that of the vehicle, $v_2 = 1$. From (Fig. 21.1a), the optimal path takes advantage of the higher jet flow on the left to reach the goal point. In the second case (Fig. 21.1b), the vehicle is deployed from $(-0.6, -0.6)$ and the goal point is at $(0.6, 0.6)$. The speed of the vehicle is 1. In this case, we have two different jets flowing in opposite directions, both with a speed of 1.5 (faster than the vehicle). The first jet (upper narrow band with black arrows) flows in the negative x direction. The second jet (lower narrow band with red arrows) flows in the positive x direction. As can be seen from (Fig. 21.1b), the vehicle rides the first jet to an upstream position of the second jet in order to compensate for the unfavorable flow of the lower jet. Then, the vehicle crosses the second jet as quickly as possible to reach its goal. In another *highway* test (Fig. 21.1c), the vehicle is deployed from $(-0.6, 0.6)$ and the goal point is at $(-0.6, -0.6)$. The speed of the vehicle is 1. In this case, we have two different jets flowing with a speed of 1.5 (faster than the vehicle). The first jet (upper narrow band with red arrows) flows in the positive y direction. Therefore, it is infeasible to use any trajectory crossing this jet flow. The second jet (lower narrow band with black arrows) flows in the positive x direction. As can be seen from (Fig. 21.1c), the vehicle first skirts the upper flow on the right-hand side. Then it moves to an upstream position of the second jet and crosses it as quickly as possible to reach the goal.

21.1.2 Time-Dependent 2-D Flows

Time-dependent flows are now considered. To evaluate robustness, a large number of 2-D flows were utilized, including flows at the exit of a strait, flows behind islands, and wind-driven flows. Such flows have also been utilized to evaluate our approach for the path planning of swarms of vehicles [21.3, 10–12, 14].

Next, we present three examples of time-optimal path planning for swarms of vehicles in the context of

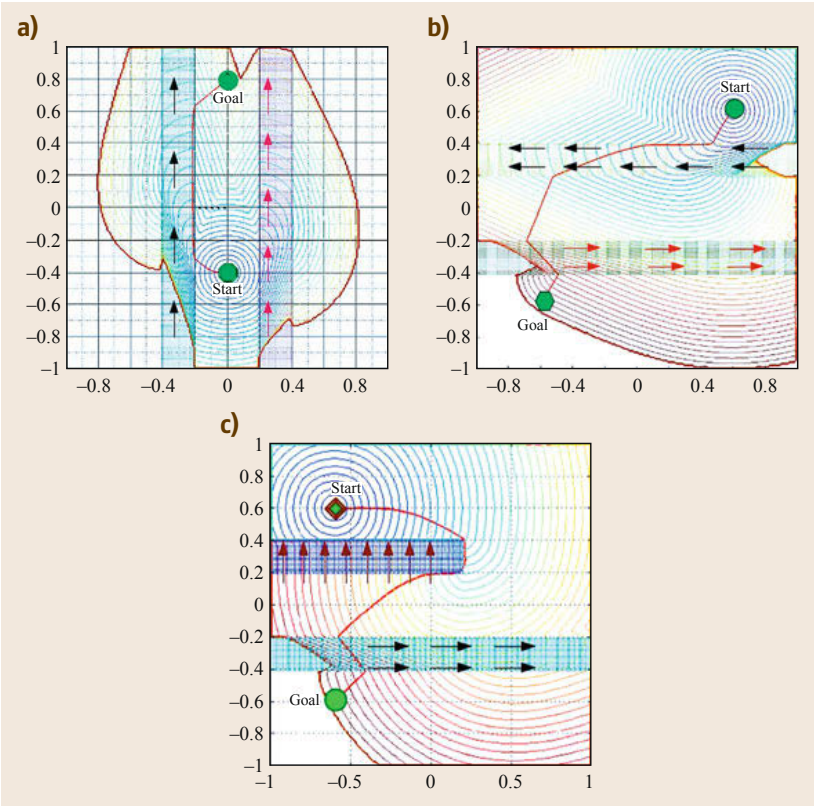

Fig. 21.1a–c Highway test cases. **(a)** 2 ideal jets. Faster jet on left (*black arrows*). Level-set optimal path rides the faster jet. **(b)** Two ideal jets in opposite directions. Level-set optimal path first rides top jet to get to an upstream position of opposing lower jet. It then crosses the opposing jet as quickly as possible to reach goal. **(c)** This case includes one truncated jet (at twice the vehicle velocity) that is effectively an impenetrable obstacle. The second jet is a transverse jet, similar to the bottom jet in **(b)**. Here the level-set optimal path moves around the truncated jet to reach an upstream position of lower jet. It then crosses the lower jet as quickly as possible to reach the goal (after [21.13])

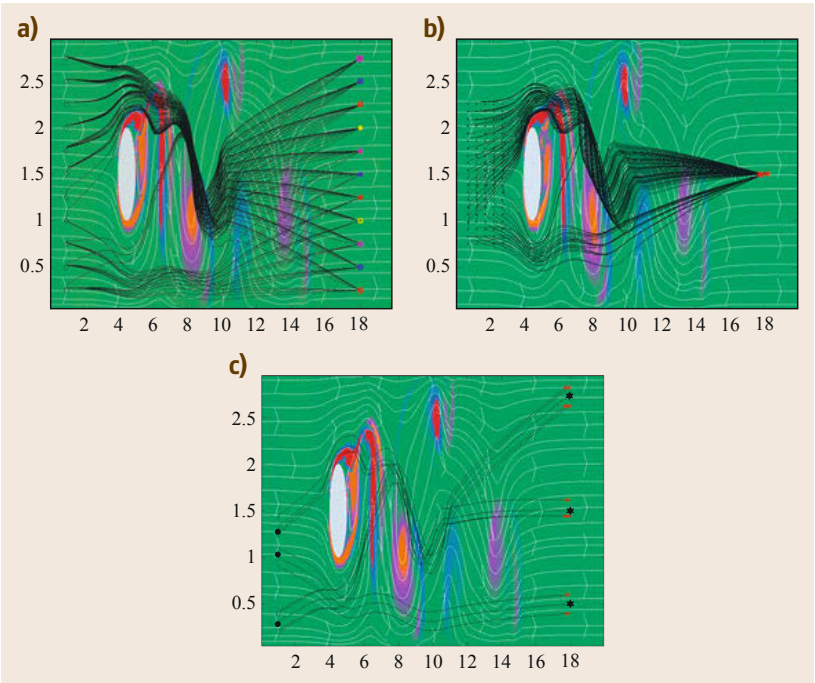

Fig. 21.2a–c Vorticity of the flow past an island (*gray oval*) overlaid with white streamlines. **(a)** *All-to-all broadcast* with swarms while crossing strong eddy field in optimal time – 11 swarms of 11 members each, avoid island, utilize eddies/wakes, reform swarms, and communicate. **(b)** *All back to ship* optimal recovery of large swarm formation of vehicles – formation of 100 vehicles to return to ship for pick up, avoid island, utilize eddies/wakes to return in shortest time. **(c)** *Maintain swarm formation* with coordination based on local gradients or dynamics – 3 swarms of 2–4 members (*line*, *triangle*, *square*), avoid island, shortest time (after [21.10, 12, 14])

flows past an island. The circular island is an obstacle to both the flow and to the vehicles. The flow is fast enough to cause the shedding of vortices off the leeward side of the island.

The first swarm example is an *all-to-all broadcast* to swarm leaders, in which 11 swarms of 11 members each optimally cross the strong eddy field in the wake of the island while exchanging information with each other:

Specifically, one member of each swarm switches to another swarm and all subsequently reform the swarms in fastest time at 11 target locations (Fig. 21.2a). In other words, we start with 11 swarms of 11 vehicles each at the 11 starting points upstream of the island (shown by open circles). Each of the initial swarms is to break up and distribute one member to each of the 11 target points downstream of the island (shown by colored markers), resulting in 11 new swarms, each made up from 1 member of each of the original 11 swarms. The goal is to achieve the new configuration in minimum time, communicating data to each swarm leader, while avoiding the island. The optimal paths computed by our level-set-based method are the black curves in (Fig. 21.2a). We see three classes of paths. Most paths go above the island and get a boost from the first shed eddy. A second group goes below the island and intercepts the eddy at a later point. The final grouping remains below the island and avoids the eddy.

The second example is an *all-back-to-ship* fastest-time recovery scenario (Fig. 21.2b). A formation of 100 vehicles starts upstream of the island (black dots). The goal is to find the quickest paths to bring all the vehicles to a single pick-up point (18, 1.5) downstream of the island, utilizing wakes and eddies while avoiding the island (simulating the recovery of a swarm of AUVs by a single ship). The optimal set of paths are found by the level-set-based method and are shown as black curves in Fig. 21.2b. Roughly three quarters of the vehicles are sent over the island to take full advantage of the first eddy. Most of the remaining vehicles avoid the eddy by passing under the island, although some go under the island and catch the eddy later.

21.1.3 Maintain Swarms Formations

Building on the above, we created algorithms for *maintaining swarm formations*. The first example consists of three swarms of two to four members each, organized in the shape of a line, triangle or square, coordinating their paths based on gradients or dynamics (Fig. 21.2c). For the algorithm illustrated, the coordination is based on local gradients (Fig. 21.2c), using a local control algorithm.

Here, the goal is to bring each of the three swarms from their start locations to their target locations in minimum time while maintaining their formation and avoiding the island. We utilize our level-set-based method to compute the optimal paths for the center of mass of each swarm (black curves starting from black dot and finishing at black star). Around each center-of-mass-path we construct the swarm paths by using the headings for the optimal center-of-mass-path and adding a velocity component at each step to maintain each vehicle position relative to the center of mass. To prevent the vehicle paths from diverging, we allow the center of mass to move at a speed lower than the maximum vehicle speed (i. e., the center of mass moves with a velocity that is a function of the instantaneous flow).

A drawback of this local control approach is that the allowable separation between vehicles depends on the flow dynamics. The distance between vehicles has to be small if the flow speed is highly variable in space (large flow gradients). The smoother the flow in space, the larger the spacing between vehicles can be. In [21.12], we develop a more powerful methodology based on local level-sets which allows maintaining complex shapes in large flow gradients; a time optimal path is computed for the group leader and finite-time local level sets are computed continuously in time for the other vehicles in the swarm. The latter provide reachable sets for each of these vehicles for each finite-time horizon (which could be a time step) and allow a subsequent shape optimization algorithm to find and sustain an optimal swarm formation.

21.1.4 Forbidden Regions

Forbidden regions are areas which are denied to the vehicles (for safety reasons or security restrictions) but do not affect the flow field [21.3]. In the example that follows, this constraint information is provided to the vehicles ahead of time, just as restricted regions or forecasts of hazardous weather can be provided to an aircraft. In these conditions, the goal is still to compute time-optimal paths, but under the additional constraint of avoiding the provided or forecast forbidden regions.

For example, consider a uniform barotropic jet exiting a strait or estuary (sudden expansion or 2-D coastal flow), as illustrated in Fig. 21.3. The non-dimensional speed of the vehicle is 0.5, while the largest current speed is 2.5 (Fig. 21.3a,b). The width of the inlet is one third of the total width of the channel. In the simulation (Fig. 21.3c), nine vehicles are released from the lower edge of the inlet (black dot) and are required to achieve a triangular formation (colored dots). The unconstrained time-optimal paths constructed from the level-set-based algorithm are shown by colored lines. The vehicle heading for the tip of the triangle rides along the central jet, while the vehicles targeting the four endpoints closest to the inlet use the upper and lower recirculation eddies. Overall, we can see from Fig. 21.3c that the algorithm correctly predicts the shapes of the optimal paths.

In the companion case (Fig. 21.3d), two forbidden regions are added (gray shaded area), which block seven of the nine paths in the free run (Fig. 21.3c). The

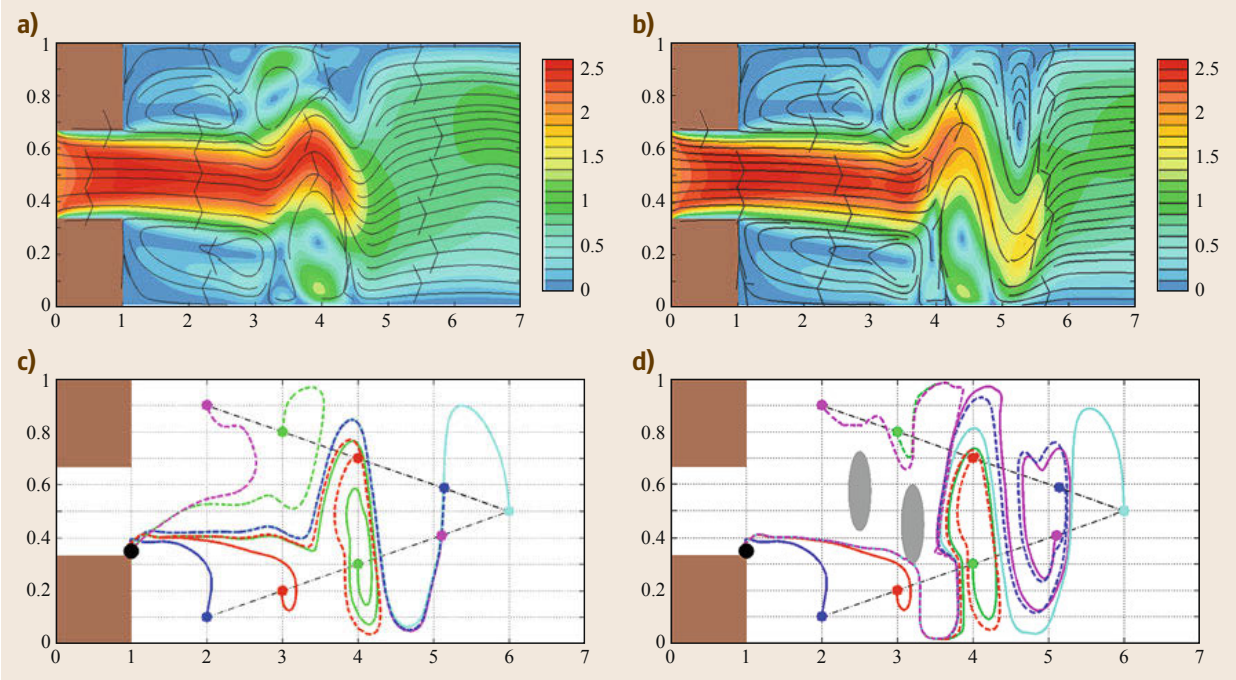

Fig. 21.3a–d Time-optimal paths without **(c)** and with **(d)** forbidden regions. **(a,b)** Snapshots of flow field for a jet exiting a strait or estuary (sudden expansion flow) **(a)** At time of initial vehicle deployment and **(b)** at final time of vehicle maneuvers in **(c)**. **(c,d)** Optimal vehicle paths for nine vehicles deployed from a single point (*black dot*) in the flow field of **(a,b)**. Results for two situations: **(c)** no constraints or forbidden regions. Vehicle paths take full advantage of the evolving jets and eddies to reach their final positions (*colored dots*) in the shortest times. **(d)** Two forbidden regions. Vehicles are denied access to the *gray shaded* regions. The level set algorithm provides seven new time optimal paths for the paths computed in **(c)** that are blocked while correctly leaving unchanged the two paths that are not blocked (after [21.3])

new optimal paths for these seven vehicles all ride the lower edge of the main jet, just skirting the bottom of the second forbidden region. They then ride down one eddy and up an adjoining eddy (Fig. 21.3a), to rejoin the main jet behind the forbidden regions. The two paths from Fig. 21.3c that did not pass through the forbidden areas remain unaffected.

We note that forbidden regions can be learned by vehicles as they progress, using onboard routing (Sect. 21.2.1). In that case, data about forbidden regions are collected by the vehicle along its optimal forecast route. As these new data are utilized, optimal paths are re-computed, leading to a new optimal forecast route for the remaining travel plan.

21.1.5 Uncertain Flow Fields

It is well known that ocean predictions are limited in accuracy because of the predictive capabilities of present modeling systems [21.16, 17] and also due to fundamental limits of predictability for nonlinear dynamics. As a result, the practical horizon for skillful forecasts of time-optimal paths is also limited. In addition, some paths will be more uncertain than others, for example, some fast paths may be very uncertain, while some somewhat slower paths may be very certain. The path planner can thus aim to find a compromise between time optimality and certainty of arrival.

To illustrate this compromise, we consider the case of time-optimal path planning in a 2-D stochastic wind-driven double-gyre flow with large uncertainties (Fig. 21.4), [21.10]. This is a strong flow field in which the uncertainties are predicted using dynamically orthogonal (DO) stochastic PDEs [21.15, 18, 19]. Questions that we consider include:

- What are the effects of flow uncertainties on the time optimal path?
- Which path should be selected?
- Is there a path which is both fastest and least uncertain (most robust or most likely)?

For the chosen start and end points, the shortest-distance path is against the jet flow that is stronger than the vehicle speed. We thus expect at least a bi-modal behavior for the probability density functions (PDFs) and

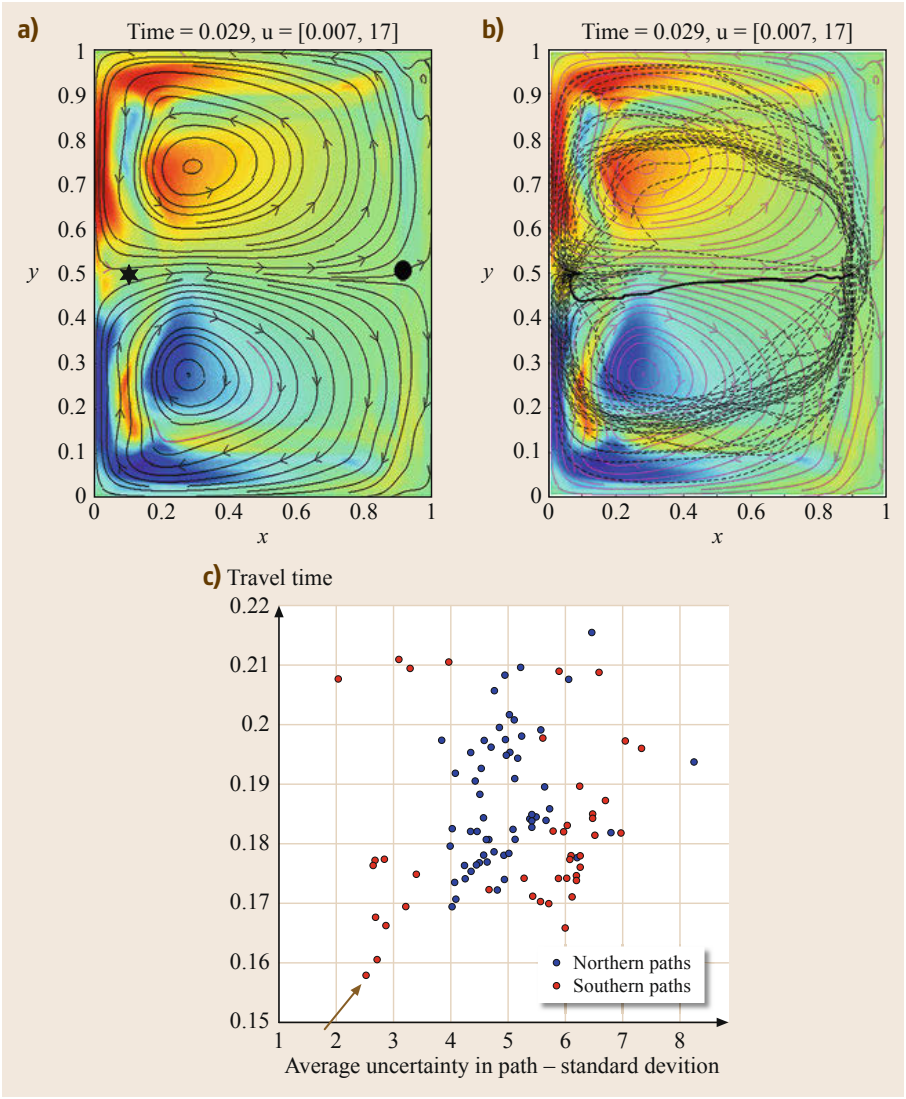

Fig. 21.4a–c Path planning in uncertain flow. **(a)** One realization of strong wind-driven, double-gyre flow. Vehicle starting point indicated by *filled circle*. End point indicated by *black star*. Wind forcing is stochastic. **(b)** Ensemble of optimal paths for different winds shown by *dashed lines*. Mean path shown by *solid line*. The mean path is unrealizable; the ensemble of paths is close to a bi-modal distribution (either following northern or southern gyre). **(c)** Travel time vs. uncertainty of ensemble of gyre paths. Most probably path follows southern gyre (*green arrow*)

sensitivity to initial conditions. We use the DO equations to quantify the uncertainty and the level-set-based path-planning (21.1)–(21.2) to solve for optimal paths for a number of realizations of the DO simulation. As is shown in Fig. 21.4b, as expected, we find that the mean path is not realizable as the paths generated are close to a bi-modal distribution (the flow field PDF leads to a PDF for the paths). For each path, we integrate the standard deviation of the flow velocity along the path, which provides a measure of the extent of *uncertainty* along the path. We then plot the travel time versus this uncertainty (Fig. 21.4c), providing a Pareto-like quantity. This plot indicates that a fast and least uncertain path is a path to the south. We note that this preferred south path arises mainly because the initial mean and PDF conditions are not symmetric with respect to the shortest-distance (horizontal) path between the start and end points.

21.1.6 Realistic Ocean Conditions: Three-Dimensional Multi-Scale Flows

Our level-set-based method has also been employed for path planning in realistic (data-driven) three-dimensional (3-D) ocean flow fields [21.13], as illustrated in Fig. 21.5. In this example, the MSEAS ocean prediction system provides the three-dimensional currents. The simulations are realistic (data-driven) ocean reanalyses based on the real-time exercises that occurred during August–September 2006 in the New Jersey Shelf/Hudson Canyon region [21.20, 21], specifically the autonomous wide aperture cluster for surveillance (AWACS) and Shallow Water-06 (SW06) exercises. The simulations employ nonlinear free surface primitive equations with fully implicit two-way nesting (Fig. 21.5a,b) and tides. The larger domain (Fig. 21.5a)

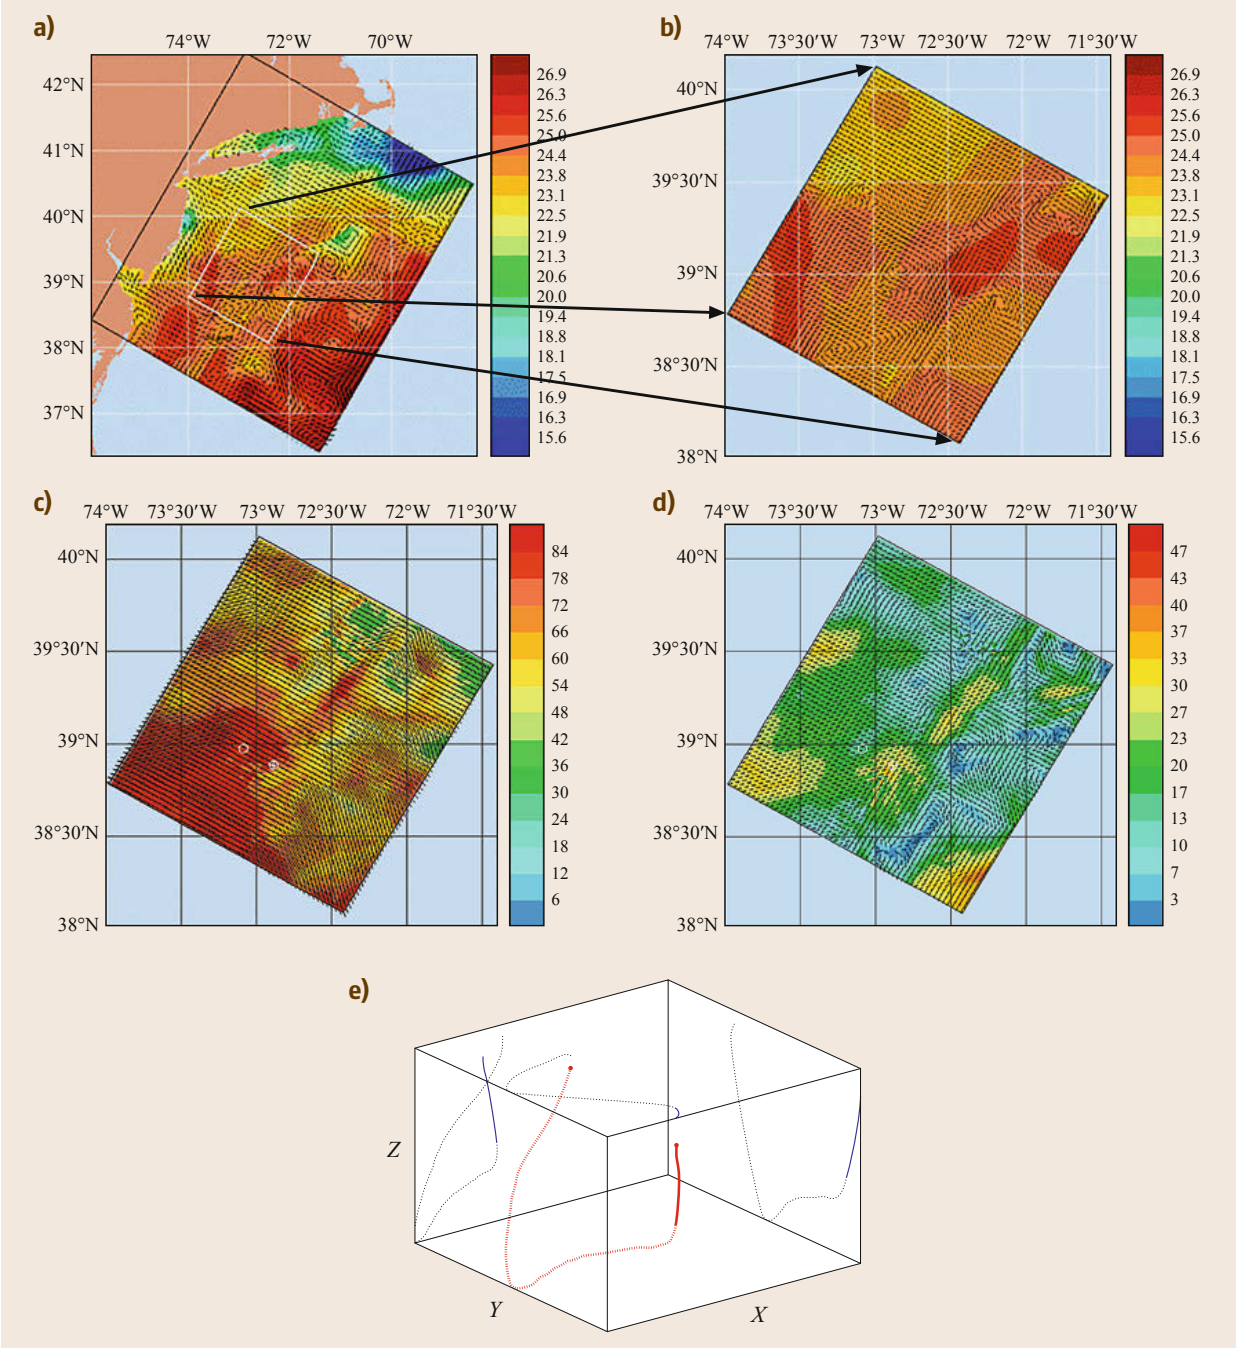

Fig. 21.5a–e Optimal path planning in 3-D with time for *fully four-dimensional (4-D) realistic ocean flows*. **(a–b)** Surface temperature and velocity from fully implicit nested SW06 simulation. **(c)** Surface flow during path planning from open circle to circle+cross. Note the strong flow opposing path. **(d)** 80 m flow during path planning. Weaker flow with some flow into goal from NE. **(e)** 3-D optimal path initially dives deep to weaker flows. The final approach to goal is from NE along the favorable flow (after [21.13])

is a 522×447 domain with 3 km resolution and 30 vertical levels, while the smaller domain (Fig. 21.5b) is a 172×155 km domain, with 1 km resolution and 30 vertical levels. The ocean initial conditions utilize real data gathered by gliders, AUVs, ship deployed conductivity temperature depth (CTD) and historical data taken from the World Ocean Database, Gulf Stream feature analyses, National Marine Fisheries Service, etc. The duration for this real-time simulation was 43.5 days. Here, we focus on the period during which Tropical

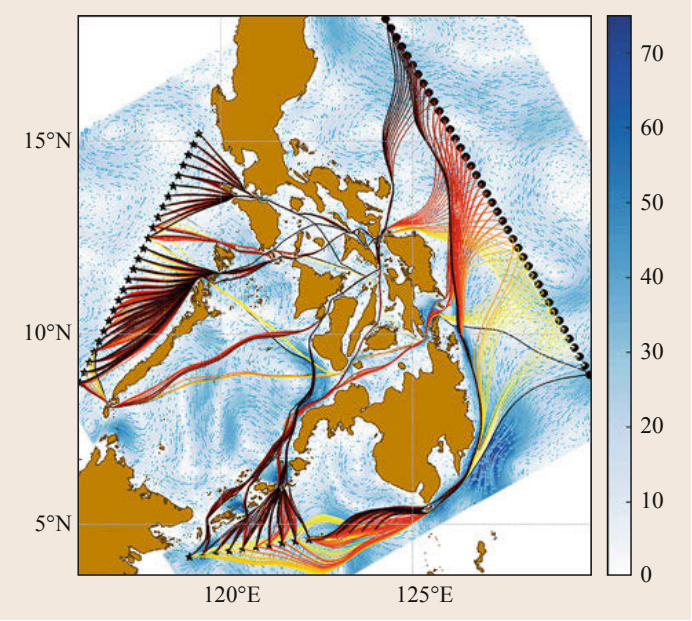

Fig. 21.6 *Time-optimal path planning in the Philippines Archipelago for vehicle swarms*: time-optimal paths of 1600 vehicles deployed from 40 different release points in the Pacific (marked by *circles*, in the east of the picture). The vehicles cross the Philippines Archipelago from east to west in the fastest time. Most of the paths are indirect as they take advantage of currents and other multi-scale ocean features ►

Storm Ernesto passed by the region (late August to early September 2006).

In order to simulate realistic and most practical conditions, we set our vehicle speed to $0.25\,\mathrm{m\,s^{-1}}$, which is often a maximum speed for most currently operated underwater gliders. The vehicle is deployed at the ocean surface (white open circle in Fig. 21.5c) and needs to travel to a second, offshore surface location (white circle with cross, Fig. 21.5c). During the path planning exercise, Tropical Storm Ernesto generates a strong, opposing, onshore current of around $0.8\,\mathrm{m\,s^{-1}}$ (Fig. 21.5c). Therefore, the glider cannot use the straightforward surface trajectory from the starting location to its offshore goal. However, the level-set algorithm finds a time-optimal 3-D path (Fig. 21.5e) in which the glider initially dives and *meanders* in 3-D to take advantage of the weaker flow field in the deeper ocean region (Fig. 21.5d).

21.1.7 Realistic Ocean Conditions: Swarms in Multi-Scale Flows with Complex Geometries

Our level-set-based methodology has been parallelized to efficiently plan optimal paths for swarms of vehicles in realistic multi-scale ocean flows with complex geometries. To illustrate this, we consider tidal-to-mesoscale ocean simulations for the Philippines Archipelago [21.22]. In this example (Fig. 21.6), there are 40 different swarms of vehicles, each released at a different location in the Pacific (east of the Archipelago). Each swarm and location consists of 40 different vehicles, one vehicle per swarm is the leader and aims to reach one of the end points in the west while the other vehicles aim to reach another of the 39 end points. When all vehicles reach their final destination, vehicles in the reformed-swarms are close to each other and can communicate underwater to exchange all needed information from their previous swarms. In other words, at the final time, an information exchange among all swarms can occur, in the form of an *all-to-all broadcast* (as in Fig. 21.2, but for realistic complex flows). The key is that each of the vehicles is navigated in fastest time to its specific end point, to the west (South China Sea) and the south (Sulawesi Sea) of the Archipelago. To do so, each vehicle optimally combines the following behaviors for time-optimality:

- Accounting for the wind events and the corresponding strong currents
- Waiting for favorable strong tidal flows in the many straits
- Taking the major currents in the region, following their meanders and eddies
- Looping in rings as needed or avoiding them all together, depending on their final destinations and variability of the 3-D ocean flows.

Figure 21.6 shows the fastest paths for all the vehicles overlaid on the map of the Philippines [21.14]. In this example, paths for 1600 vehicles have been planned by efficiently parallelizing the level-set-based algorithm. As observed, none of the vehicle trajectories is a direct path to its target. Vehicles clearly utilize ocean currents and multi-scale features to reduce their travel time.

21.2 Adaptive Sampling for Swarms of Ocean Vehicles

With our collaborators, we have developed a range of adaptive sampling schemes and applied them in varied coastal regions. In what follows, we start by outlining some of the schemes that were first applied in real

ocean situations, as well as the results of these schemes (Sect. 21.2.1). The constraints of the realistic ocean forecasts and the computational capabilities of the last decade led to a range of approximations necessary to issue optimal sampling paths forecasts in real time. After these results, we have started new theoretical research towards schemes and methods that are both rigorous and applicable to real ocean systems. This progress is summarized in Sect. 21.2.2.

21.2.1 Early Results: Approximate Schemes for Realistic Applications

During the last decade, we have developed and applied a range of adaptive sampling and onboard routing schemes for optimal sensing in varied regions of the world's oceans. They include:

- Adaptive sampling via error subspace statistical estimation (ESSE) with nonlinear predictions of error reductions [21.5, 23]
- Mixed integer linear programming (MILP) for optimal-sampling path planning [21.24]
- Nonlinear optimal-sampling path planning using genetic algorithms [21.25]
- Dynamic programming and onboard routing for optimal-sampling path planning [21.26]
- Command and control of surface kayaks over the Web, directly read from model instructions [21.27].

These schemes and results are outlined next.

Real-Time Adaptive Sampling via ESSE

The results of the ESSE adaptive sampling scheme [21.20, 28] utilized in real time are illustrated in Fig. 21.7. For data assimilation, the ESSE scheme integrates an ensemble of nonlinear simulations and assimilates data by Kalman updates in the adaptive error subspace obtained from this ensemble. For adaptive sampling, the same is carried out, but in a forecast sense and for many candidate sampling plans, since the goal is the nonlinear prediction of the ideal future sampling. Specifically, a set of candidate sampling plans (data locations and sensor types) was chosen a priori to reduce computational costs and directly satisfy operational constraints. For each plan, a nonlinear ESSE prediction and data assimilation was then carried out. This estimated the error reduction for each plan and so determined the candidate plan with the best impact.

Adaptive Sampling with Routing Schemes

Predictive adaptive sampling and onboard adaptive routing have been combined for thermocline tracking and adaptive sampling for acoustic fields with AUVs [21.26], as illustrated in Fig. 21.8. The method consists of two complementary steps. First, data assimilative environmental and acoustic propagation ESSE ensemble modeling provides input to a dynamic scheme that computes parameter values for autonomous sensing behaviors, with the goal of optimally reducing the forecast acoustic uncertainties. Behavioral parameters include, for example, the dive angle or the times to switch depths. Second, these parameters for the autonomous sensing behavior are refined onboard the vehicle in realtime, in response to the data sampled.

A related approach consists of computing the high-level routing remotely (in a centralized computer). A vehicle is then guided by downloading remotely-computed optimal waypoints from a website. During

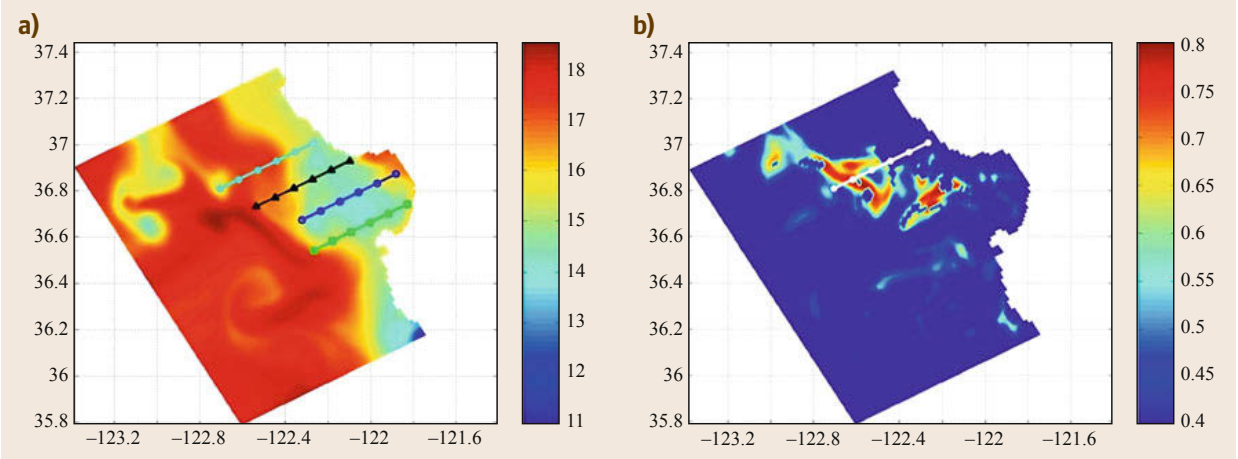

Fig. 21.7a,b Adaptive sampling via ESSE, for an 24–27 August sample forecast. **(a)** Central forecast of the surface T for August 26, overlaid with the four candidate tracks. **(b)** Predicted optimal track for August 26 and its relative error reduction for August 27 (after [21.5]). The question answered is *which of the four sampling track tomorrow will reduce uncertainties the most the day after tomorrow?*

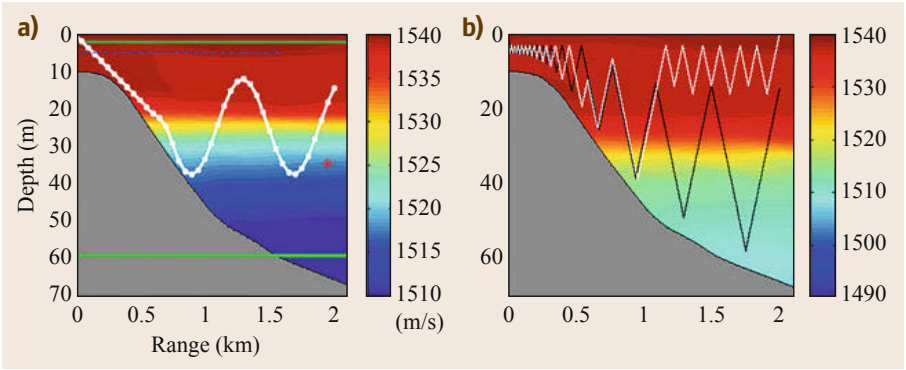

Fig. 21.8a,b Simulated adaptive thermocline tracking using yo-yo adaptive sampling and onboard routing, overlaid on sound–speed ($m\,s^{-1}$). **(a)** AUV path; *green lines* are the minimum and maximum depths allowed. **(b)** The optimized AUV automatically captures the *afternoon effect* on the surface thermocline, i. e., the warming of the upper ocean layers due to the strong daylight sun (after [21.26])

the Persistent Littoral Undersea Surveillance Network 2007 exercise (PN07) in Dabob Bay [21.27, 29]; such a website was used to optimally guide kayaks at sea, directly based on data-assimilative ocean model predictions of the optimal sampling plans, without humans in the loop. Computers ran the ocean model forecasts and adaptive sampling forecasts and then wrote the waypoints file on a website that was read by a kayak, all without human input. This was a technical achievement compared to other larger and more extensive exercises, e.g., Monterey Bay 06 [21.8, 9, 30], where humans were intermediaries between ocean models and vehicles, e.g., entering waypoints or interpreting ocean model forecasts and then selecting waypoints.

Path Planning for Fixed-Uncertainty Optimal Sensing

Path planning for fixed-uncertainty optimal sampling aims to answer the following type of questions: Assuming the error field is not dynamic during the sampling, which sampling path is optimal? This approach often assumes simple updates of field estimates due to data but more importantly neglects how dynamics over time evolves future observations, evolution which affects the truly optimal adaptive sampling. The planning then amounts to the optimization of a fixed objective function whose solution often amounts to visiting the locations of highest uncertainty. Here, we outline two schemes for such planning. The first one solves the optimization using MILP [21.23, 24, 31], as shown in Fig. 21.9. The solution is obtained by a *branch and bound algorithm*, which provides exact and globally optimal solutions. The path is segmented by waypoints, and the goal is to solve for their coordinates. The number of waypoints is a parameter fixed as a function of vehicle range and grid dimensions. An extension of this approach for time-dependent fields has been obtained [21.32].

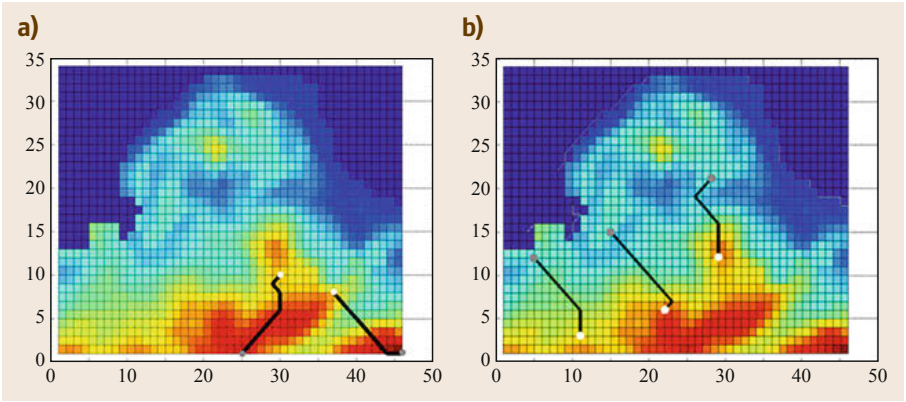

Fig. 21.9a,b Generation of sampling paths for fixed objective fields using mixed-integer programming (MIP). Fields are the ESSE T error standard deviation averaged over 0–50 m. *Gray dots* are starting points for the AUVs and *white dots* are the MIP optimal termination points. **(a)** Optimal path of two vehicles. **(b)** Optimal path of three vehicles (after [21.24]). The question answered is *assuming the error field remains constant for the next day, on which path do I send my AUVs?*

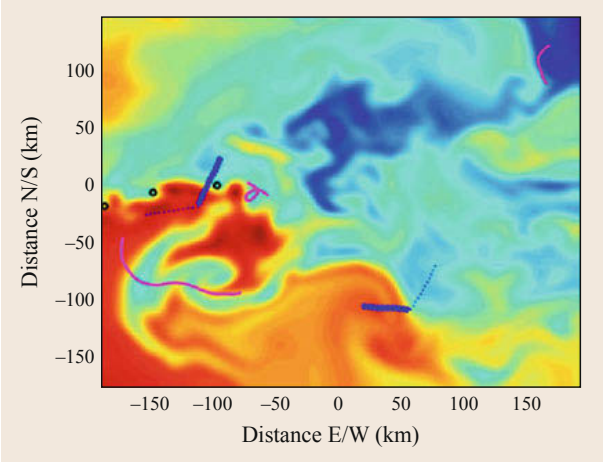

Fig. 21.10 Optimal sampling path solution computed by a genetic algorithm (GA) in the Middle Atlantic Bight and Shelfbreak Front regions (after [21.25]). The plan is computed for three moorings, three gliders, and two REMUS vehicles. The GA places all moorings (*black circles*) directly within the Shelfbreak Front, which passes over them during the simulated 5 days of deployment. The REMUS vehicles (*blue tracks*) pass over Shelfbreak events and fronts, and the gliders (*magenta tracks*) track large gradients with time. A distance-potential constraint maintains the paths of platforms apart by at least a Rossby radius (key horizontal length scale for mesoscale ocean dynamics, 15 km in this example)

The second scheme uses genetic algorithms to solve the optimization problem [21.25]; the results are illustrated in Fig. 21.10. With genetic algorithms, one can carry out nonlinear path planning for any type of cost function, which is a significant practical advantage. The weights of each term in the cost function are chosen by users. However, a disadvantage of the genetic algorithm solution is that one cannot always theoretically guarantee that it is at, or is near, the optimum. In the results illustrated, the cost function includes terms representing ocean variability, acoustic transmission loss sensitivity, and temperature uncertainty. Results illustrate that the temporal dependence of the ocean fields strongly influences the optimal location of the moorings that are fixed in space and the gliders that are slowly drifting. The remote environmental monitoring unit (REMUS) vehicles move quickly and are less sensitive to time-dependent ocean effects.

21.2.2 Recent Progress: Towards Rigorous Schemes for Realistic Applications

Our recent research in adaptive sampling has been motivated by the ESSE examples presented above. The first research theme that we investigate is to increase the number of candidate sampling paths towards functional paths covering the whole physical domain and to increase the durations of look-aheads towards durations up to the end of the experiment or the predictability limit. Ideally, there should be no limitation on the candidate paths other than the operational constraints. A second theme is to continue the utilization of nonlinear uncertainty evolution and estimation of future data impacts, but replace the ESSE ensemble approach by our DO PDEs and replace the ESSE Kalman updates by non-Gaussian updates, including Gaussian Mixture Models updates [21.33, 34]. Mixture models are semiparametric frameworks consisting of a sum (mixture) of parametric distributions whose parameters and number (total number of distributions themselves, referred to as mixture complexity) are commonly fit to an ensemble of realizations. A third theme is to extend the adaptive sampling optimality criterion to information theory (mutual information schemes) and to include nonlinear smoothing effects. In what follows, we outline a subset of results related to these research themes.

Adaptive Sampling Using Partially Observable Markov Decision Processes (POMDPs)

Computational research has begun on the use of POMDPs as a means of adaptively sampling fields of interest in the ocean. A POMDP in this case is a mathematical framework that allows one or more agents to optimize an objective function under uncertainties both in sensing and knowledge of the field dynamics. We have demonstrated its use on two simplified dynamical model testbeds. The first is a stochastic game of life where a robot aims to optimally reach an end goal while sampling minor rewards and avoiding minor losses along the way. The second is the prediction of optimal sampling paths for two vehicles operating in the idealized Lorenz-95 atmospheric model [21.35]:

1. *Stochastic game of life*: We have adapted John Conway's game of life [21.36] by including stochasticity in the governing equations, as described in [21.37]. As such, we consider this to represent a (simplified) discrete flow model. Figure 21.11 shows one time instant in that game, all actors of the game are dynamic: the black circle is the AUV; the green circle is a slowly moving global feature of high reward; blue plus signs are minor rewards; and red crosses are minor cost/loss. Both of the latter fields evolve independently according to the rules of the game. The route of the AUV (not plotted in the snapshot) is being optimally planned. The domain is discretized such that all motions are restricted to be up or down, or left or right. The AUV applies the POMDP-like algorithm, intelligently maneuvering

through the field such as to maximize its reward. Specifically, it seeks to approach the green circle, primarily sampling blue dots along its path. This test bed naturally extends to multiple vehicles. The game of life results (Fig. 21.11) may also directly apply to real ocean fields in the sense that green, blue, and red regions could be specific ocean field properties or discretized regions of instability or of interesting dynamics that one wishes to sample.

2. *Lorenz-95 model*: The system studied is the Lorenz-95 atmospheric model, which is a time-varying one-spatial dimension model, as illustrated in Fig. 21.12. The Lorenz-95 model replicates a number of phenomena observed in geophysical fluid dynamics, and as such lends itself well to this analysis. Each vehicle (i. e., adaptive observation) locally predicts future consequences of proposed sampling paths, independently of other vehicles, given current knowledge of the system. In this test bed, the adaptive DO equations are used to predict the PDF of the Lorenz-95 state and artificial potential-like functions are employed to ensure that adaptive vehicles remain sufficiently apart from each other. We implemented a POMDP-like algorithm for adaptive sampling, using a GMM-DO (GMM: Gaussian Mixture Model) filter for data assimilation [21.33, 34]. With this example, we show that adaptively sampling the field using the POMDP approach outperforms the case in which the adaptive observations remain fixed in location (see the bottom left plot of Fig. 21.12).

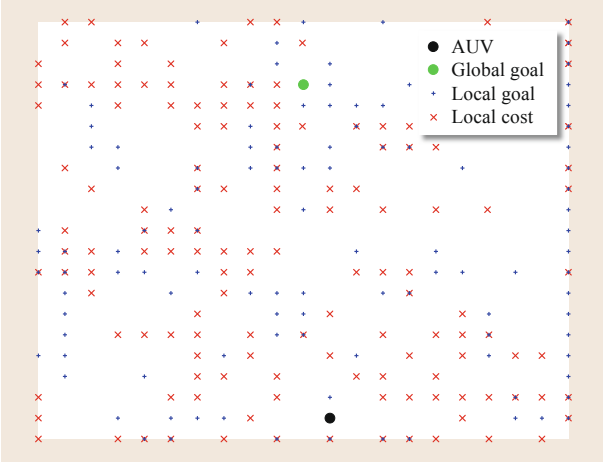

Fig. 21.11 Stochastic game of life with intelligent sampling using POMDPs. A fixed-time snapshot in the time-dependent game is shown. The route of the AUV is being optimally planned. The domain is discretized such that all motions are restricted to be up or down, or left or right. The goal for the AUV, *black circle*, is to track and reach the *green circle* as fast as possible and, if efficient along the way, sample the *blue plus signs* (e.g., food) and avoid the *red crosses* (e.g., predators), so as to increase the total reward

While POMDPs are computationally expensive, techniques are being developed that continually allow their use in domains of increasing complexity. In the future, with a detailed probabilistic description of the currents via the adaptive DO equations, one could derive novel metrics using information theoretic measures (specifically *mutual information*) to represent rewards. Our scheme may effectively be applied among multiple agents, where each agent locally applies a POMDP, communicating across a subset of its information to neighboring vehicles. Knowledge of the ocean fields, derived from the MSEAS ocean prediction model, is naturally of significant advantage. By subdividing the POMDP framework into a set of hierarchies, itself inspired by biological swarming formations, one could optimize the sampling behavior of the team of AUVs. Then, a top layer POMDP controls a number of lower layered POMDPs, themselves controlling a group of AUVs.

Adaptive Sampling with Limited-Time Look-Ahead POMDPs and ESSE in 2-D Fields

We now extend the above approach to two-spatial dimension and time-dependent problems and evaluate it on simulated ocean fields. We still employ DO equations for uncertainty prediction and a limited time look-ahead POMDP approach to select the sampling paths. However, we now use ESSE for data assimilation instead of the GMM update.

As shown in Fig. 21.13, we consider an uncertain flow behind a circular cylinder/island. The uncertainty arises due to the initial conditions and is evolved in time using the DO equations. Assuming that we have access to a single sampling vehicle (a glider), we predict a trajectory for this asset that optimally reduces the uncertainty in the estimated velocity field (horizontal components: u and v). In Fig. 21.13a, we plot the field of variance in u and v along with the vehicle trajectory obtained using a three-time-steps look-ahead approach for adaptive sampling. Figure 21.13b shows the time evolution of the variance of two stochastic DO modes. Initially, the variances of the modes are large (order 1). As the vehicle makes measurements at integer times, the variance of each of the modes is reduced. With time, the variance of the modes eventually decreases to an order of 10^{-1} for the first mode and 10^{-5} for the last mode.

Figure 21.13c,d shows the true solution for the velocity field and the mean of the adaptive-sampling-based estimate of that velocity field at two different times. At time $t = 6$ (Fig. 21.13c), we see that the mean of the estimated solution is still very different from

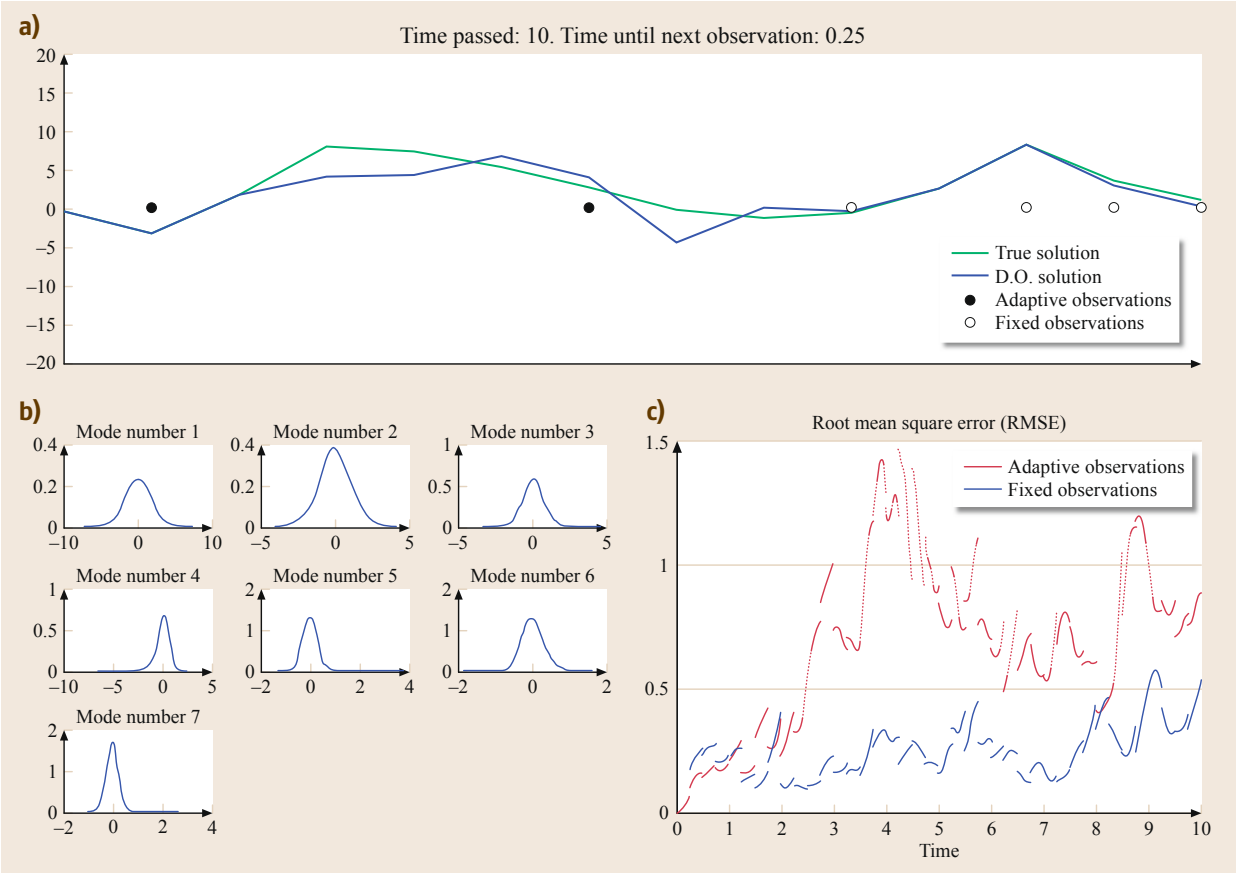

Fig. 21.12a–c Adaptive sampling one-dimensional (1-D) example: the Lorenz-95 application. The adaptive sampling is based on: i) non-Gaussian data assimilation (GMM-DO Filter) combining Gaussian Mixture Models (GMMs) with the DO equations; ii) three-step look-ahead POMDPs to decide which direction (left or right) to sample for each of the two adaptive observations; and iii) an artificial potential-like function to ensure the two adaptive observations remain separated in space. **(a)** True (unknown) and estimated solutions at final non-dimensional time 10. The latter estimated solution is obtained by adaptive sampling by GMM-DO, three look-ahead POMDPs and artificial potential-like functions. Three routine observations are fixed in space (*open circles*) while two observations have paths that are forecast by the adaptive sampling (*black circles*). **(b)** Marginal PDFs for the seven DO modes active at final time 10. **(c)** Comparison between the time-evolution of the averaged true error of the estimates obtained using adaptive sampling or not using it (i. e., assimilating data but keeping the two *black circle* observation locations fixed in place). The true error reduction obtained by predicting the best sampling location is substantial (about 100% better)

the true solution because very few measurements have been collected so far. As time progresses (Fig. 21.13d) the mean of the estimate matches quite well with the true solution, and the algorithm correctly predicts the shape, phase, and frequency of vortex shedding behind the island. Another interesting result in this example is that the vehicle automatically navigates towards the cylinder, indicating that it is the optimal location where measurements must be made to gain maximum information about the field. It also follows the locations where eddies are shed. Clearly, the dynamics of the problem, the limited-time look-ahead POMDP behavior and the assimilation scheme govern the track of the vehicle.

21.3 Conclusions and Outlook

Some of the results obtained by our MSEAS group and recent collaborators have been summarized here. We first presented a subset of our results on time-optimal path planning for swarms of ocean vehicles. We then

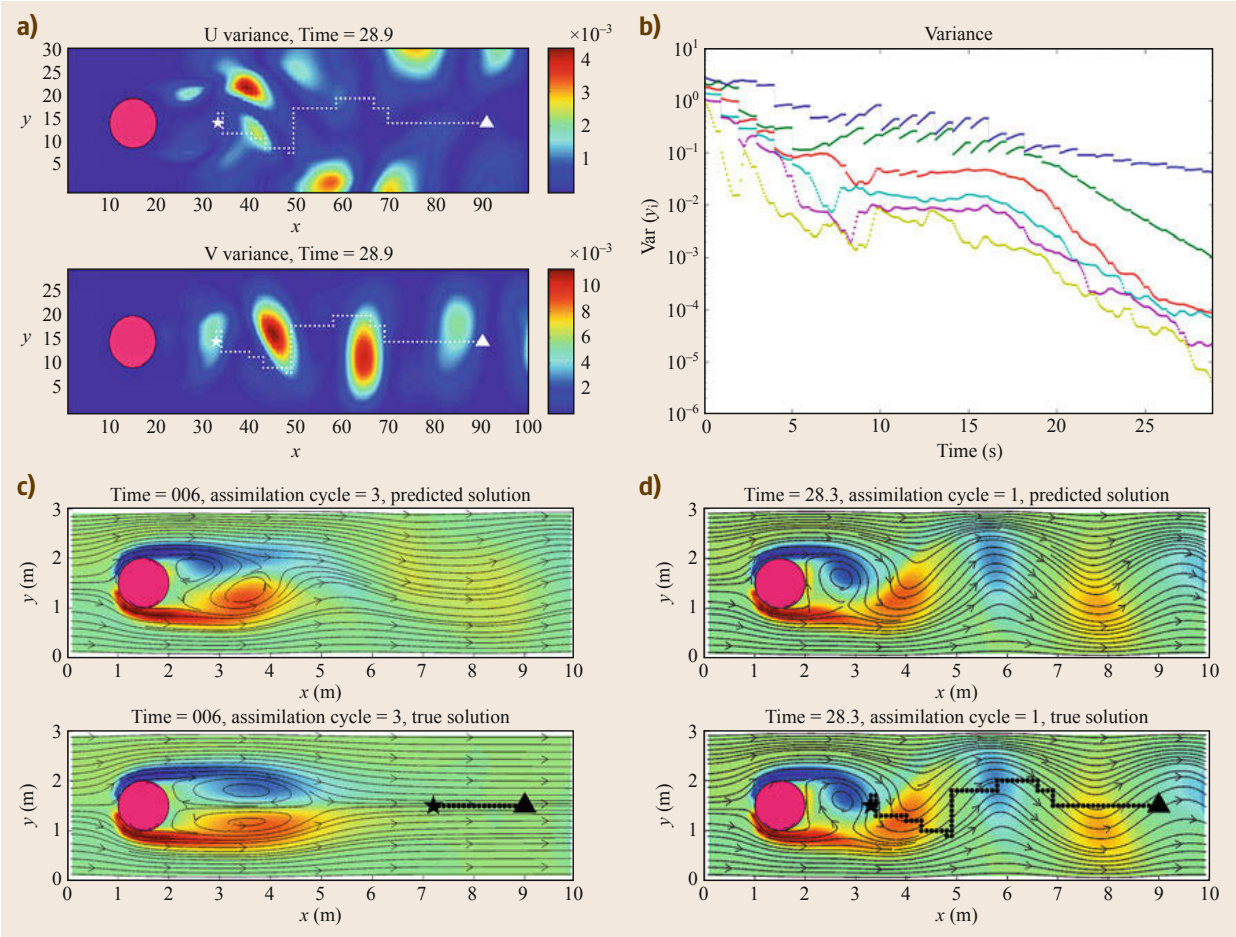

Fig. 21.13a–d *Adaptive sampling for flow past a circular island.* We predict the trajectory of a vehicle that makes measurements in the field so that the uncertainty in the field is optimally reduced. The uncertainty due to initial conditions is propagated forward in time by the DO equations and an ESSE update is used for the assimilation. **(a)** Variance in the u and v velocities of the flow field. **(b)** Reduction of variance of each of the DO modes due to the measurements made by the vehicle. **(c)** The mean of the predicted solution compared with the true velocity field. The predicted solution is significantly different from the true solution at this time. **(d)** As the vehicle measurements become more effective, the predicted solution is driven towards the true solution and eventually, both these fields match up well

reviewed some of our research in adaptive sampling. The path planning and adaptive sampling schemes and approaches presented were illustrated by their applications in idealized and realistic ocean conditions.

Of course, many more results, complementary to ours, have been obtained, both in optimal path planning and adaptive sampling. Even if we were to limit ourselves to oceanic and atmospheric applications, the depth and breadth of these complementary results are significant. For general reviews on oceanic adaptive sampling, we refer the reader to [21.4–9]. For general reviews on oceanic path planning, we refer the reader to [21.2], as well as to Chaps. 14, 19, and 20 of this Handbook.

In coming decades, in light of the growing human population on earth, the ocean, with the water and life it contains, is likely to become increasingly important. Since data collection and operations at sea are expensive, utilizing all of the knowledge we have to plan such collections and operations is logical and important. Useful knowledge here combines science and engineering, including observation, modeling, prediction, estimation, and control. The science of autonomy will necessarily play an increasing role in ocean applications. This is because sustaining healthy interactions between the ocean and humans is vital to life on earth. These interactions will have to be intelligently and autonomously managed.

References

21.1 H. Stommel: The Slocum mission, Oceanography **2**(1), 22–25 (1989)

21.2 T. Lolla: Path Planning in Time Dependent Flows Using Level Set Methods, Master's Thesis (MIT, Cambridge 2012)

21.3 T. Lolla, M.P. Ueckermann, K. Yigit, P.J. Haley Jr., P.F.J. Lermusiaux: Path planning in time dependent flow fields using level set methods, Proc. ICRA (2012) pp. 166–173

21.4 N.E. Leonard, D. Paley, F. Lekien, R. Sepulchre, D. Fratantoni, R. Davis: Collective motion, sensor networks and ocean sampling, Proc. IEEE OCEANS'07, Vol. 95 (2007) pp. 48–74

21.5 P.F.J. Lermusiaux: Adaptive modeling, adaptive data assimilation and adaptive sampling, Physica D **230**, 172–196 (2007)

21.6 N. Roy, H.-L. Choi, D. Gombos, J. Hansen, J. How, S. Park: Adaptive observation strategies for forecast error minimization, Lecture Notes Comput. Sci. **4487**, 1138–1146 (2007)

21.7 D.A. Paley, F. Zhang, N.E. Leonard: Cooperative control for ocean sampling: The glider coordinated control system, IEEE Trans. Control Syst. Technol. **16**(4), 735–744 (2008)

21.8 S.R. Ramp, R.E. Davis, N.E. Leonard, I. Shulman, Y. Chao, A.R. Robinson, J. Marsden, P.F.J. Lermusiaux, D. Fratantoni, J.D. Paduan, F. Chavez, F.L. Bahr, S. Liang, W. Leslie, Z. Li: 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 (2009)

21.9 T.B. Curtin, J.G. Bellingham: Progress toward autonomous ocean sampling networks, Deep Sea Res. Part II **56**(3–5), 62–67 (2009)

21.10 T. Lolla, P.F.J. Lermusiaux, M.P. Ueckermann, P.J. Haley Jr.: *Modified Level Set Approaches for the Planning of Time-Optimal Paths for Swarms of Ocean Vehicles*, MSEAS Report-14 (MIT, Cambridge 2012)

21.11 T. Lolla, P.F.J. Lermusiaux, M.P. Ueckermann, P.J. Haley Jr.: Time-optimal path planning in dynamic flows using level set equations: Theory and schemes, Ocean Dynamics **64**(10), 1373–1397 (2014)

21.12 T. Lolla, P.J. Haley Jr., P.F.J. Lermusiaux: Time-optimal path planning in dynamic flows using level set equations: Realistic applications, Ocean Dynamics **64**(10), 1399–1417 (2014)

21.13 K. Yigit: Path Planning Methods for Autonomous Underwater Vehicles, Master's Thesis (MIT, Cambridge 2011)

21.14 T. Lolla, P.J. Haley Jr., P.F.J. Lermusiaux: Path planning in multi-scale ocean flows: Coordination and dynamic obstacles, Ocean Dynamics **94**, 46–66 (2015)

21.15 M.P. Ueckermann, P.F.J. Lermusiaux, T.P. Sapsis: Numerical schemes for dynamically orthogonal equations of stochastic fluid and ocean flows, J. Comput. Phys. **233**, 272–294 (2013)

21.16 P.F.J. Lermusiaux, P. Malanotte-Rizzoli, D. Stammer, J. Carton, J. Cummings, A.M. Moore: Progress and prospects of U.S. data assimilation in ocean research, Oceanography **19**(1), 172–183 (2006)

21.17 P.F.J. Lermusiaux, C.-S. Chiu, G.G. Gawarkiewicz, P. Abbot, A.R. Robinson, R.N. Miller, P.J. Haley, W.G. Leslie, S.J. Majumdar, A. Pang, F. Lekien: Quantifying uncertainties in ocean predictions, Oceanography **19**(1), 92–105 (2006)

21.18 T.P. Sapsis, P.F.J. Lermusiaux: Dynamically orthogonal field equations for continuous stochastic dynamical systems, Physica D **238**, 2347–2360 (2009)

21.19 T.P. Sapsis, P.F.J. Lermusiaux: Dynamical criteria for the evolution of the stochastic dimensionality in flows with uncertainty, Physica D **241**(1), 60–76 (2012)

21.20 P.F.J. Lermusiaux: Uncertainty estimation and prediction for interdisciplinary ocean dynamics, J. Comput. Phys. **217**, 176–199 (2006)

21.21 P.J. Haley Jr., P.F.J. Lermusiaux: Multiscale two-way embedding schemes for free-surface primitive-equations in the multidisciplinary simulation, estimation and assimilation system, Ocean Dyn. **60**, 1497–1537 (2010)

21.22 P.F.J. Lermusiaux, P.J. Haley Jr., W.G. Leslie, A. Agarwal, O. Logutov, L.J. Burton: Multiscale physical and biological dynamics in the Philippines archipelago: Predictions and processes, Oceanography PhilEx Issue **24**(1), 70–89 (2011)

21.23 P.F.J. Lermusiaux, P.J. Haley Jr., N.K. Yilmaz: Environmental prediction, path planning and adaptive sampling: Sensing and modeling for efficient ocean monitoring, management and pollution control, Sea Technol. **48**(9), 35–38 (2007)

21.24 N.K. Yilmaz, C. Evangelinos, P.F.J. Lermusiaux, N. Patrikalakis: Path planning of autonomous underwater vehicles for adaptive sampling using mixed integer linear programming, IEEE-JOE **33**(4), 522–537 (2008)

21.25 K.D. Heaney, G. Gawarkiewicz, T.F. Duda, P.F.J. Lermusiaux: Non-linear optimization of autonomous undersea vehicle sampling strategies for oceanographic data-assimilation, J. Field Robotics **24**(6), 437–448 (2007)

21.26 D. Wang, P.F.J. Lermusiaux, P.J. Haley, D. Eickstedt, W.G. Leslie, H. Schmidt: Acoustically focused adaptive sampling and on-board routing for marine rapid environmental assessment, J. Mar. Syst. **78**, S393–S407 (2009)

21.27 J. Xu, P.F.J. Lermusiaux, P.J. Haley Jr., W.G. Leslie, O.G. Logutov: Spatial and temporal variations in acoustic propagation during the PLUSNet'07 exercise in Dabob Bay, Proc. POMA, Vol. 4 (2008) p. 070001

21.28 P.F.J. Lermusiaux: Evolving the subspace of the three-dimensional multiscale ocean variability: Massachusetts Bay, J. Mar. Syst. **29**, 385–422 (2001)

21.29 J. Curcio, T. Schneider, M. Benjamin, A. Patrikalakis: Autonomous surface craft provide flexibility to remote adaptive oceanographic sampling and modeling, Proc. IEEE OCEANS'08 (2008) pp. 1–7

21.30 N.E. Leonard, D.A. Paley, R.E. Davis, D.M. Fratantoni, F. Lekien, F. Zhang: Coordinated control of an underwater glider fleet in an adaptive ocean sampling field experiment in Monterey Bay, J. Field Robotics **27**(6), 718–740 (2010)

21.31 N.K. Yilmaz, C. Evangelinos, N.M. Patrikalakis, P.F.J. Lermusiaux, P.J. Haley, W.G. Leslie, A.R. Robinson, D. Wang, H. Schmidt: Path planning methods for adaptive sampling of environmental and acoustical ocean fields, Proc. OCEANS'06 (2006)

21.32 N.K. Yilmaz, P.F.J. Lermusiaux: Mixed Integer Linear Programming MILP Path Planning of AUVs for Adaptive Sampling: Real-World Simulation Results, MSEAS Tech. Rep., Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge (2012)

21.33 T. Sondergaard, P.F.J. Lermusiaux: Data assimilation with Gaussian Mixture Models using the dynamically orthogonal field equations. Part I: Theory and scheme, Mon. Weather Rev. **141**(6), 1737–1760 (2013)

21.34 T. Sondergaard, P.F.J. Lermusiaux: Data assimilation with Gaussian Mixture Models using the dynamically orthogonal field equations. Part II: Applications, Mon. Weather Rev. **141**(6), 1761–1785 (2013)

21.35 E.N. Lorenz, K.A. Emanuel: Optimal sites for supplementary weather observations: Simulation with a small model, J. Atmos. Sci. **55**(3), 399–414 (1998)

21.36 M. Gardner: The fantastic combinations of John Conway's new solitaire game *life*, Scientific American **223**, 120–123 (1970)

21.37 R.A. Monetti, E.V. Albano: Critical edge between frozen extinction and chaotic life, Physical Review E **52**(6), 5825–5831 (1995)

Acknowledgements

PFJL dedicates the results presented to Dr. Tom Curtin. Tom Curtin's vision and questions in the early 2000s on how to utilize the ocean currents as dynamic highways for autonomous vehicles were early seeds for our theoretical and applied research in marine autonomy. We are also very thankful to all members of the MSEAS group, past and present, for the many useful discussions. PFJL also thanks Drs. Tim Duda and Kevin Heaney, as well as Drs. Jim Bellingham, Russ Davis, Constantinos Evangelinos, Dave Fratantoni, Glen Gawarkiewicz, Naomi Leonard, John Leonard, Sharan Majumdar, Nick Patrikalakis, Steve Ramp, and Henrik Schmidt for their collaborations on autonomy. We are grateful to the Office of Naval Research for support under grants N00014-09-1-0676 (Science of Autonomy, A-MISSION) and N00014-12-1-0944 (ONR6.2) to the Massachusetts Institute of Technology (MIT). We especially thank Dr. Marc Steinberg, as well as Drs. T. Paluszkiewicz and S. Harper, for their support and interactions. KY is very thankful to the Turkish Navy for support. MPU and PFJL thank the Natural Sciences and Engineering Research Council (NSERC) of Canada for the postgraduate scholarship partially supporting the graduate studies and research of MPU at MIT.