Methods in Oceanography 10 (2014) 21–43

Contents lists available at ScienceDirect

Methods in Oceanography

journal homepage: www.elsevier.com/locate/mio

Full length article

Oceanographic pursuit: Networked control of multiple vehicles tracking dynamic ocean features

Brooks Reed [a,b], Franz Hover [a,*]

[a] Massachusetts Institute of Technology, Department of Mechanical Engineering, Cambridge, MA, 02139, USA

[b] MIT-WHOI Joint Program, Woods Hole Oceanographic Institution, Woods Hole, MA, USA

ARTICLE INFO

Article history:
Received 2 February 2014
Received in revised form
24 April 2014
Accepted 8 May 2014
Available online 29 May 2014

Keywords:
Autonomous underwater vehicles
Collaborative control
Feature tracking
Ensemble forecasts
Linearization
System identification

ABSTRACT

We present an integrated framework for joint estimation and pursuit of dynamic features in the ocean, over large spatial scales and with multiple collaborating vehicles relying on limited communications. Our approach uses ocean model predictions to design closed-loop networked control at short time scales, and the primary innovation is to represent model uncertainty via a projection of ensemble forecasts into local linearized vehicle coordinates. Based on this projection, we identify a stochastic linear time-invariant model for estimation and control design. The methodology accurately decomposes spatial and temporal variations, exploits coupling between sites along the feature, and allows for advanced methods in communication-constrained control. Simulations with three example datasets successfully demonstrate the proof-of-concept.

© 2014 The Authors. Published by Elsevier B.V.
This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/3.0/).

1. Introduction

The behavior of ocean fronts and similar structures such as plumes and filaments has long been of interest to oceanographers (Gangopadhyay and Robinson, 2002; Ferrari, 2011). Recent measurements in a front off Japan have revealed sub-mesoscale structure that figures unexpectedly large in the

* Corresponding author. Tel.: +1 617 253 6762.
E-mail addresses: brooksr8@mit.edu (B. Reed), hover@mit.edu (F. Hover).

http://dx.doi.org/10.1016/j.mio.2014.05.001
2211-1220/© 2014 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/3.0/).

energy balance (D'Asaro et al., 2011). Fronts and plumes are implicated in foundational work on Lagrangian coherent structures (Olascoaga and Haller, 2012), and can show dramatic physical, chemical, and biological variability that is critical to understanding ocean–atmospheric coupling, ecological systems, and pollution (Camilli et al., 2010; Farrell et al., 2005).

Despite continual advances in modeling of complex natural processes, ocean features at the mesoscale and smaller remain challenging (Canuto and Dubovikov, 2005; Hanna and Yang, 2001), and hence have emerged as a primary focus area for mobile sensing systems. Here, progress has been rapid, e.g. Fiorelli et al. (2006) and Wang et al. (2009). Zhang et al. (2012b,a) carried out at-sea experiments where measurements both drove trajectory decisions and triggered collection of large samples. A single vehicle has successfully tracked a plankton bloom (Godin et al., 2011), while a coordinated approach for a similar problem using a drifter and vehicle has been studied in Das et al. (2012) and Graham et al. (2013). A collaborative control technique for tracking Lagrangian coherent structures is presented in Michini et al. (2014), and a distributed approach for plume and thermocline tracking is considered in Petillo et al. (2012). Supporting all these developments, basic water properties are routinely measured today from mobile robots, while sophisticated chemical and biological analyses *in situ* are becoming mature technologies, for example DNA probes (Scholin et al., 2006) and mass spectrometers (Camilli et al., 2010). In turn, ocean modeling is becoming integrated with real-time sampling systems, e.g., Willcox et al. (2001), Haley et al. (2009) and Smith et al. (2010), and is increasingly taking on multi-disciplinary aspects (Stow et al., 2009). Non-cooperative path-planning based on current forecasts has been studied extensively, for example by Smith et al. (2010) and Lolla et al. (2012). Ocean model uncertainty predictions and communication constraints, however, have not been a central focus in these works.

Already exploited regularly in the terrestrial and air domains, networks of mobile agents are an attractive means for tracking and pursuit of dynamic processes over mixed spatial scales (Dunbabin and Marques, 2012), although wireless communication inevitably brings fundamental challenges in control (Baillieul and Antsaklis, 2007). Underwater, wireless communication over distances beyond about one hundred meters is made almost exclusively via acoustics, which suffer packet losses caused by ambient noise, multipath, and other environmental conditions (Heidemann et al., 2012). This packet loss, combined with low data rates and long delays, has limited the use of acoustic communications in high-performance, real-time tasks. Our own experience with acoustic modems (Reed et al., 2013) strongly suggests that control system design should encompass communication limits from the beginning. To this end, there has been considerable recent work in the field of control under communication constraints. Constructive results exist for lossy estimation (Sinopoli et al., 2004; Gupta et al., 2007), lossy commands (Quevedo et al., 2011) and H_∞ sampled-data control (Lall and Dullerud, 2001). We extended the work of Imer et al. (2006) to the case of independent multi-channel packet losses (Reed and Hover, 2013); Imer's dynamic programming approach results in a highly tractable recursion. These principled methods for networked control design, however, usually require linear time-invariant (LTI) system representations.

In this paper, we combine the themes above to focus on tracking and pursuit of dynamic ocean fronts by multiple unmanned vehicles, posing the problem in such a way as to accommodate the most promising developments in communication-constrained feedback control. As diagrammed in Fig. 1, our approach fits as an intermediary between high-bandwidth vehicle flight control (at the seconds time scale) and lower-frequency procedures in numerical ocean modeling, assimilation, and adaptive sampling. Notably, we are using linearization for a completely different purpose here than the norm in physical oceanography, where it has helped characterize instability and maximum sensitivity directions through adjoint models (Moore et al., 2004). As described in full below, our approach explicitly leverages *ocean forecast ensembles, a projection onto vehicle coordinates, and stochastic system identification*, yielding a dynamics description that is directly suitable for control system design. These elements enable a *reactive control methodology for dynamically sampling the ocean*, that may surpass approaches used today. Looking forward, we hope that our framework may provide a basis for tradeoff studies in designing complex deployments.

We describe the overall technical framework in Section 2, with some additional background comments on forecasting and linearization. Projection is detailed in Section 3, and the integration of projection, system identification, estimation and control in Section 4. Projection and identification

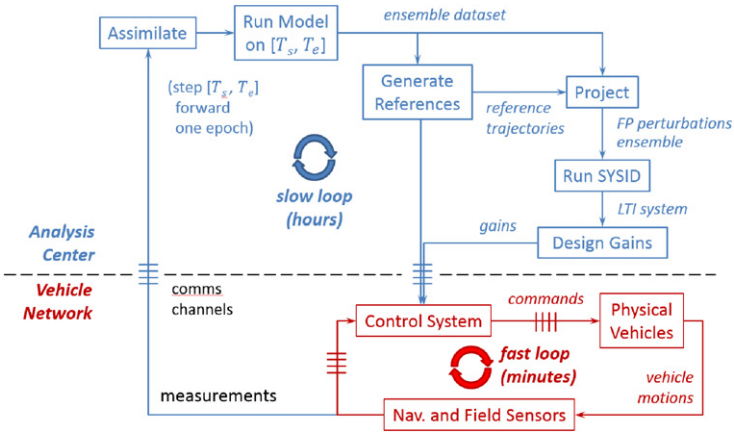

Fig. 1. High-level overview of the ocean front tracking system. Inputs on the Analysis Center side include human decisions and other data sources not available to the vehicle network; these may or may not be embedded in the projector. Additional inputs on the Vehicle Network side include channel losses, sensor noise, and physical disturbances.

outcomes with test and model data are described in Section 5, and Section 6 has control simulations. Section 7 concludes.

2. Technical setup

2.1. Scope of the field operation

We consider a dynamic scalar field ϕ in two dimensions, although there is nothing inherent to our methodology that prohibits three dimensions, or a vector field. ϕ is a function of space and time, but we will usually drop these arguments. Two key assumptions are that the instantaneous field contains areas of spatial gradient, which induce a favorable measurement gain for data-driven servoing, and that the operation of vehicles on gradients is desirable. The second point is certainly true if we want to track a front, and characterize the water properties close to it. Beyond following gradients as a primary mission objective, there is also a broader scenario in which one set of vehicles might have the task of monitoring feature boundaries – characterized by a threshold value and a gradient – while another set operates within the feature, where structure is much harder to detect.

We focus on scalar field measurements that can be adequately sampled by autonomous vehicles while they are in motion. This could include standard quantities like temperature and salinity, or more advanced chemical and biological measurements. Our general framework can be expanded to fit another setting as well, of vehicles taking "metameasurements". For example, if the feature dynamics is slow enough compared to vehicle speeds, local surveys within a time step could make local estimates of the gradient, or be used to enhance spatial resolution by visiting multiple nearby FPs in one time step.

2.2. Known and unknown frontal points

In loose terms, a front in the field variable at a given time is an elongated region of high gradient magnitude; in the plane, a front would comprise a tightly-packed set of contour lines for field variable $\phi(t)$. A *projector* expresses such a front in terms of desirable locations for vehicles – which we term *frontal points* (FPs) $\boldsymbol{p}(t)$ – and as the physical front evolves over time, so do the frontal points. Taking ϕ and $\boldsymbol{p}$ to denote a finite time series of the field and a set of FPs, respectively, our overall strategy in this paper is to develop *a process that quickly generates acceptable* $\boldsymbol{p}$ *from given* ϕ, and separately *a process by which an actual vehicle group pursues a physical realization of* $\boldsymbol{p}$. The second process perhaps

Table 1
Summary of general constraints for projector $\mathbb{D}$. $\mathbb{P}_6$ is desirable, but not required.

$\mathbb{P}_1$	General desirability of gradient areas
$\mathbb{P}_2$	Feasibility
$\mathbb{P}_3$	Accuracy and strength of reference gradients
$\mathbb{P}_4$	Slowly varying reference FPs and gradients
$\mathbb{P}_5$	Consistency of major features
$\mathbb{P}_6$	Coupling among perturbations

has a conceptual jump in the sense that we are asserting a physical $\boldsymbol{p}$ to go with a physical ϕ, and of course neither is fully known.

For a given ϕ, frontal points are defined by the projector $\mathbb{D}$:

$$\boldsymbol{p} = \mathbb{D}(\phi).$$

$\mathbb{D}$ can comprise algorithms as well as human input; it encodes any explicit dependence on time, and initial conditions for $\boldsymbol{p}$. $\mathbb{D}$ need not be causal, and as a function could be surjective or injective or both: different field evolutions ϕ could easily lead to the same desired vehicle trajectory, while there also exist useful vehicle trajectories that do not depend on ϕ at all. Most importantly, $\mathbb{D}$ has full access to ϕ and is thus *omniscient*.

Example 1. A causal, omniscient projector for N vehicles might have the following key operations at each time step:

1. Given $\phi(t)$, choose a level Φ and calculate regions of the physical space that satisfy $\phi(t) = \Phi$.
2. On this subset of interest, propagate the N FPs, $\boldsymbol{p}(t)$, with constraints on spacing and on speed.

This example captures the paradigm we explore in computations, but it is by no means the only one possible. For instance, instead of choosing a specific level set of $\phi(t)$, one could just as easily identify areas of high gradients for locating the FPs.

An acceptable $\boldsymbol{p}$ satisfies a number of constraints, $\boldsymbol{p} \in \mathbb{P} = \mathbb{P}_1 \cap \mathbb{P}_2 \cap \cdots$. Already we know from statements so far that:

$\mathbb{P}_1$: *$\boldsymbol{p}$ occurs in gradient areas, and observations near $\boldsymbol{p}$ are desirable.*

$\mathbb{P}_2$: *$\boldsymbol{p}$ is feasible for the vehicles and their control system, in terms of maximum transit speed, maneuvering, closed-loop bandwidth, collision avoidance, and communication or other operational constraints.*

To support closed-loop field operations, below we will add to $\mathbb{P}$'s specification, as summarized in Table 1.

Consider now the *physical instance* of the field, ϕ_o; the "o" subscript is used from here on to indicate the physical instance. $\mathbb{D}$ does not apply because we do not know ϕ_o. Yet desirable FPs $\boldsymbol{p}_o$ still exist, and an estimate of them plays out according to

$$\hat{\boldsymbol{p}}_o^t = \mathbb{E}_o(\boldsymbol{z}_o^t),$$

where $\mathbb{E}_o$ is a causal, actuated estimator, and sparse measurements $\boldsymbol{z}_o$ are its driver. The "t" superscript indicates times up to t. Underlying $\boldsymbol{z}_o$ is the physical instance itself, ϕ_o, and a set of physical vehicles and their control system(s), process and sensor noise channels, and so on. $\mathbb{E}_o$ is understood to embed all available information but it has a structure different from $\mathbb{D}$: it does not know the true field ϕ_o, and thus cannot carry out the first step in Example 1. Our goal then is to exploit the information available at start of mission to ensure that $\boldsymbol{p}_o \in \mathbb{P}$, and that $\hat{\boldsymbol{p}}_0 \approx \boldsymbol{p}_o$ through a sampling, estimation and control strategy. These are achieved through both $\mathbb{D}$ and $\mathbb{E}_o$.

2.3. Designing $\mathbb{D}$ via stochastic prediction

Without prior information about ϕ_o the construction from $\mathbb{E}_o$ would be extremely weak. We can use a stochastic model of the ocean process to strengthen it. Consider a set of specific *predicted* instances (realizations) of the field, indexed $i \in \{1, \ldots, M\}$. The associated set of frontal points is generated by the same omniscient projector as above: $\boldsymbol{p}_i = \mathbb{D}(\phi_i)$. Since ϕ_i is in hand, the $\boldsymbol{p}_i$ defined by $\mathbb{D}$ is easy to create and visualize. Thus a predicted instance provides a direct assessment of $\mathbb{D}$, and an ensemble of such instances allows us to design $\mathbb{D}$ such that the constraints in Table 1 are satisfied.

Let the distribution of the true field be $h(\phi_o)$, and make the assumption that $h(\phi_o) = h(\phi_i)$, i.e., there are no systematic errors in the field prediction. Formally, this assertion rests on Leith's landmark paper in climate studies, that by the ergodic hypothesis good performance in an ensemble occurs when the ensemble distribution matches the true climate distribution (Leith, 1974). This assumption underlies virtually all modern weather prediction procedures, and in our context it implies that $h(\boldsymbol{p}_o) = h(\boldsymbol{p}_i)$, since $\mathbb{D}$ is agnostic on the input data. Then $\text{prob}(\boldsymbol{p}_i \in \mathbb{P}) = \text{prob}(\boldsymbol{p}_o \in \mathbb{P})$, and consequently $\mathbb{D}$ designed via $\boldsymbol{p}_i$ achieves $\text{prob}(\boldsymbol{p}_o \in \mathbb{P}) \to 1$ if the design process fully models random $\boldsymbol{p}$, and if $\mathbb{D}$ yields a high fraction of successful trajectories. In other words, the probability that the physical instance FPs $\boldsymbol{p}_o$ will satisfy the constraints increases as the projector takes more information about the field into account.

From the stochastic prediction, we define a reference field $\bar{\phi}$ and an associated set of reference FPs $\bar{\boldsymbol{p}}$. One choice, described in Section 3.1, is to set $\bar{\phi}$ as the ensemble mean, with reference FPs equally spaced along a level set. Note that although $\bar{\boldsymbol{p}}$ does depend on a rule set, we do not write $\bar{\boldsymbol{p}} = \mathbb{D}(\bar{\phi})$, because $\mathbb{D}$ operates on instances only, and actually contains both $\bar{\boldsymbol{p}}$ and $\bar{\phi}$ implicitly. From the stochastic prediction we also define a reference gradient vector $\bar{\boldsymbol{g}}$ to go with $\bar{\boldsymbol{p}}$; both $\bar{\boldsymbol{p}}$ and $\bar{\boldsymbol{g}}$ can change over time.

2.4. Geometry of $\mathbb{E}_o$

Assuming the desirability ($\mathbb{P}_1$) and feasibility ($\mathbb{P}_2$) conditions are met by $\boldsymbol{p}_o$, the next question is how to make $\hat{\boldsymbol{p}}_o \approx \boldsymbol{p}_o$. In this and the next subsection, we consider geometric aspects, from which conditions $\mathbb{P}_3$, $\mathbb{P}_4$, and $\mathbb{P}_5$ will follow. As noted, $\mathbb{E}_o$ includes the vehicle system, and hence following $\boldsymbol{p}_o$ is a complicated function of many parameters as well as random processes. At each time step, however, $\mathbb{E}_o$ can be thought of as having a geometric part that establishes $\hat{\boldsymbol{p}}'_o \approx \boldsymbol{p}_o$ through algebraic constraints only, and a feedback and estimation part that in turn establishes $\hat{\boldsymbol{p}}_o \approx \hat{\boldsymbol{p}}'_o$, when the vehicles stay close by. This second approximation is essentially set by the closed-loop performance of the integrated system, the topic of a later section.

Regarding the first approximation, suppose for the moment that there are no constraints from the vehicle system, i.e., that all measurements and communications are perfect, and the vehicles' physical motion is unconstrained. All that is left in $\mathbb{E}_o$ is the causal mapping from perfect measurements at vehicle locations $\boldsymbol{q}$ to the FP estimates $\hat{\boldsymbol{p}}'_o$. Let us restrict our attention to a single FP and a single time step.

The geometry of $\mathbb{E}_o$ is built on one of two simple linearization models. Consider Fig. 2(a), which illustrates the first. Suppose that p_o instantaneously satisfies the level condition $\phi(p_o) = \Phi$ plus constraints $p_o \in \mathcal{C} = \mathcal{C}_1 \cap \mathcal{C}_2$.[1] To resolve ambiguities in a complicated field $\mathcal{C}_1$ restricts each p_o to the neighborhood of the corresponding reference point $\bar{p}$; for uniqueness, $\mathcal{C}_2$ is a line that sufficiently constrains p_o along the level set. The intuition is that if this FP coincides with a strong gradient, and if tracking errors are small, then a vehicle located at q would be able to directly servo to p_o, by measuring $[\phi(q) - \Phi]$ and enforcing $q \in \mathcal{C}_2$. Note that $\hat{p}'_o$ as drawn is not at p_o because the estimator does not know the true level set.

Each linearization at a site in $\mathbb{E}_o$ uses the nominal gradient $\bar{g}$, which is supposed to apply in the neighborhood of p_o. This is a strong assumption, because gradient variability can be quite large. A

[1] As a slight abuse of notation, we will use p_o and q to indicate 2-vectors of coordinates in $\mathbb{R}^2$; boldface will be used to refer to sets of FPs and vehicles, respectively.

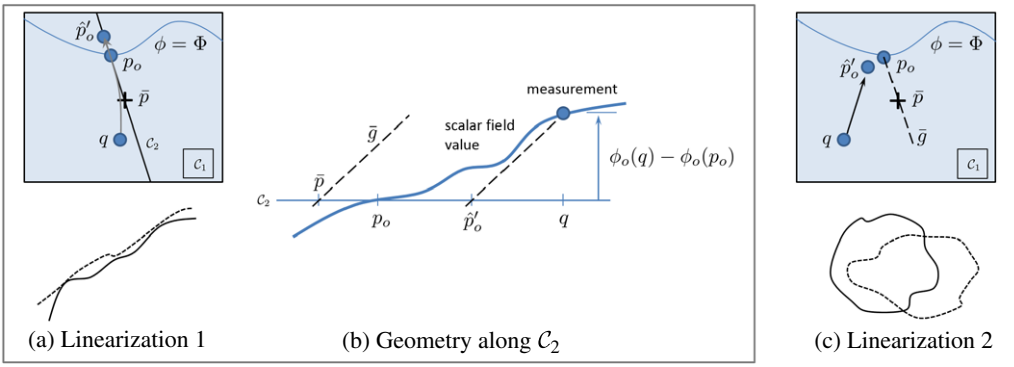

Fig. 2. Spatial linearization geometry. The pursuit control system should drive a vehicle at q to the estimated frontal point $\hat{p}'_o$. The upper boxes of (a) and (c) show two linearizations for $\mathbb{E}_o$, for a given time instant and FP; the left one (a) has a line constraint $\mathcal{C}_2$ and is our main focus in computations. The right one (c) has a reference gradient imposed and thus requires some additional constraint. The bottom sketches of (a) and (c) illustrate the perturbation type for which each linearization is suited; the solid line is the reference feature and the dotted line is a physical instance. Diagram (b) shows geometry along $\mathcal{C}_2$ (*e.g.* the reference gradient direction) for the first linearization, illustrating how $\bar{g}$ establishes $\hat{p}'_o$.

robust stochastic prediction serves as an adequate indicator, however, and our simple approach of using $\bar{\boldsymbol{p}}$ and $\bar{\boldsymbol{g}}$ directly is effective for the datasets described in Section 5.

Some simple manipulations give a key analytic relationship, illustrated in Fig. 2(b):

$$\phi_o(q) - \phi_o(p_o) = \bar{g}^T(q - \hat{p}'_o) = \int_{p_o}^{q} \nabla\phi_o(s)^T ds, \text{ to yield}$$

$$\int_{p_o}^{q} (\nabla\phi_o(s) - \bar{g})^T ds = \bar{g}^T(p_o - \hat{p}'_o), \tag{1}$$

where s denotes any path in $\mathbb{R}^2$ from p_o to q. The formula is exact, and says that the error in extrapolating a measurement at q to the anticipated measurement at (unknown) p_o is equal to the projection of the reference gradient onto the error in FP position. It offers two crucial insights. First, with all of q, p_o and $\hat{p}'_o$ assumed to lie within $\mathcal{C}_1$, for the error to be small we need $\nabla\phi_o$ to be near $\bar{g}$ in $\mathcal{C}_1$, at least in the average sense and at least in the direction of $q - p_o$. Given that q is on $\mathcal{C}_2$, gradient error that is orthogonal to $q - p_o$ has no bearing on the FP error. Second, stronger $|\bar{g}|$ makes for a smaller error in $\bar{g}$'s direction. Error orthogonal to $\bar{g}$ is not controlled by the integral (nor by $|\bar{g}|$), and this is why we require the constraint $\mathcal{C}_2$ or coupling in our models.

Example 2. To illustrate specific $\mathbb{D}$ and $\mathbb{E}_o$, suppose there are two vehicles and that $\mathbb{D}$ codes for one vehicle to follow the northernmost extreme of the level set $\phi_i = \Phi$, and the other vehicle to follow the southernmost extreme. This is an easy program to write if ϕ_i is known. Assume $\bar{\boldsymbol{p}}$ is representative of each instance; then $\mathbb{E}_o$ fixes the east–west coordinate for each vehicle according to that of $\bar{\boldsymbol{p}}$ – this is the constraint $\mathcal{C}_2$ at each site – and drives each vehicle north or south depending on the error signal $[\phi_o(q) - \Phi]$ and some control law.

We now state more specifically our two major categories of linearization:

1. The first involves perturbations relative to a reference point and its reference gradient; see Fig. 2(a). Perturbations are only allowed along the reference gradient line through the reference point. This constraint is suitable when deviations perpendicular to a level set are the dominant effect, e.g., undulations of a long front.
2. The second linearization allows unconstrained translations relative to the reference FP, but maintains the reference gradient, as in Fig. 2(c). This is useful when the shape of the tracked feature is consistent but its translation is not. In the absence of a given lateral constraint, this linearization requires coupling between FPs, which will be described below.

In this paper we give application-related studies concentrating on the first linearization. As an aside, we note that range-only tracking and pursuit of a point target can be seen as a special case of the second linearization, in which speed limits and nonholonomic constraints for the target are replacing the ocean model. That the observing vehicles have to maintain a wide aperture, and a suitable distance for sensing and communication, comprises a simple rule set within the real-time system $\mathbb{E}_o$. We have studied this scenario experimentally (Reed et al., 2013).

2.5. Gradients for $\mathbb{E}_o$

An instance ϕ_i induces p_i at each site through $\mathbb{D}$, and the associated gradient is $\nabla\phi_i(p_i)$. If this gradient is close to the reference gradient $\bar{g}$, then we can use the latter in real-time operation. Desirable properties for $\mathbb{P}$ in this regard are:

> $\mathbb{P}_{3a}$: $\|\bar{g}\|_2 \geq \kappa$ *for each FP. The reference gradient is strong, improving SNR of a noisy field measurement.*[2]
>
> $\mathbb{P}_{3b}$: $\|\nabla\phi_i(p_i + \Delta) - \bar{g}\|_2 \leq \epsilon\kappa$ *for each FP, where* $\Delta \in \mathbb{R}^2 : \|\Delta\|_2 \leq r$ *is a positioning error representing a neighborhood of radius* r *around* p_i*, and* $0 < \epsilon \ll 1$*. The instance gradient is close to the reference and is spatially robust.*

Both of these properties conform to minimizing the error $\hat{p}'_o - p_o$ in Eq. (1), and also relate to control system performance, as discussed in Section 6.2.1. Two additional issues arise. First, $\bar{\boldsymbol{p}}$ and $\bar{\boldsymbol{g}}$ can vary over the prediction horizon. Consistent with control systems practice, we require that these variations are slow enough that they do not interfere with the closed-loop system near its break frequency:

> $\mathbb{P}_4$: $\bar{p}$ *and* $\bar{g}$ *for each FP are slowly-varying relative to the control system.*

Second, despite best efforts to corral the frontal points through advantageous trajectories, they may still fail to capture the dynamic behavior of interest. For example, suppose that a simple, positive-sloping front like that shown in Fig. 2(b) develops a local minimum in the vicinity of $\bar{p}$, leading to an "S"-shaped slice. We have observed this and other failure mechanisms in our work with ocean model data, and they lead to a general constraint on feature consistency, which is intrinsically related to the gradients:

> $\mathbb{P}_5$: *An instance* $\boldsymbol{p}_i$ *stays on the major features of interest.*

2.6. Dynamics linearization

Our stochastic understanding of the field is a reference trajectory $\bar{\phi}$ and a limited variation of it ϕ_o. This applies to the FPs as well, with $\bar{\boldsymbol{p}}$ and $\boldsymbol{p}_o$. For the purpose of designing and implementing an observation system, we now assume that the perturbations $\boldsymbol{p}_o - \bar{\boldsymbol{p}}$ have a locally linear dynamic behavior, driven by both known inputs (e.g., wind forcing if the estimate is accurate) and unknown (long-term nonlinear behavior of ocean flows). Linear dynamics brings access to mature and scalable multi-variable identification, analysis and synthesis tools, but the assertion is clearly a tradeoff against the descriptive power of fully nonlinear modeling. We think the tradeoff worthwhile since communication constraints and navigation are preeminent problems in application. That said, any linearization of a nonlinear process, and certainly one that projects onto a low-dimensional space, has limitations. Our approach is advocated only for persistent features within the ability of the network and vehicle system to pursue.

The physics of ocean processes introduce spatial and dynamic structure into the field, which may be reflected in the motions of the frontal points. Such physical aspects include the length scales of turbulence vs. smoothness in diffusion, and geostrophic or tidal forcing at large scales. Establishing

[2] A loose definition of SNR is gradient magnitude multiplied by the standard deviation of the navigation noise, divided by the standard deviation of the scalar field measurement noise.

coupling between the frontal points is useful since an integrated control and estimation strategy would enable coupled vehicles to perform better than vehicles making decisions based only on local information. To emphasize this point, we observe that in general the Kalman filter systematically reduces error covariance trace as the coupling between subsystems is increased, with all other parameters held fixed. This trend holds up to the point where the system's frequency content is beyond the capability of the filter. Our control simulation results in Section 6.2.2 similarly show that coupling improves performance; see the "Lower Bound" controller. On the other hand, a controller that leverages uncoupled dynamics still has its place in scenarios where communication is difficult, or sensor noise is low; see the "Loners" controller. Thus, we set a *desirable but not required* constraint for the projector:

$\mathbb{P}_6$: *An instance* $\boldsymbol{p}_i$ *exhibits coupling.*

To establish coupling, we apply system identification techniques to the stochastic frontal points perturbations; more details are given in Section 4.2. This method may be criticized first because it is equivalent to claiming that the coupled dynamic behavior of the ocean field can be expressed as a model of very low order; and second because stochastic identification (in which the system is driven by unknown forcing) is difficult enough even for illustrative problems in the literature. At the same time, linear analysis has been extremely useful in elucidating fundamental behavior of ocean systems, e.g., Thompson and Battisti (2000) and Huck and Vallis (2001), and indeed a full linear tangent dynamic behavior has been developed in Moore et al. (2004), and implemented in the widely-used *ROMS* modeling system. Along the lines of our current objectives, linear estimation of ocean systems using sensor networks has been considered before by Zhang et al. (2009).

2.7. Modeling framework and other assumptions

Ensemble description of uncertainty: The quality of predictions is of course a perennial concern in modeling any stochastic, nonlinear process (Palmer, 2000; Rixen et al., 2009). An ensemble of model runs with variable forcing and initial conditions is a popular means for describing uncertainty, and this will be our language in the rest of the paper. More generally, however, any description of uncertainty could be accommodated insofar as it allows projection into the vehicle space. Uncertainty in prediction arises primarily from initial condition errors, modeling errors, and unknown disturbances; we shall assume that all three elements are unbiased in each epoch.

Limited time scales and renewal: No climate nor ocean model enjoys sustained accuracy as the prediction horizon lengthens. For employing an ocean model over long time scales, we assume a data assimilation schedule as found in numerical weather prediction; predictions are made for a given forward horizon, on which we immediately carry out our entire design process, and implement it. When a new prediction becomes available, the process can be run again. An alternative information paradigm would be to employ repeated sections from longer simulations of cyclic processes, e.g., tidal flows in a coastal area.

Timing: The integrated observation system has several time scales. Following Fig. 1, we define an epoch as the period between T_s and T_e over which a given model prediction (and the associated ensemble data) is valid, and the pursuit task is undertaken. For our examples in stochastic identification and controller performance, we study a single epoch. The model prediction for each epoch consists of T unit time steps (the resolution of the model forecast), and thus, $T_e - T_s = T$. The integrated observation and control system, *i.e.* all sensing, communications, and actions in the feedback loop, operates at its own time step, δt. In the current study, we use $\delta t = 1$, so that the networked control system matches the model time step. This is not far from reality, since the temporal resolution of ocean models is often of order one hour or shorter, while the command and control cycle of a multi-vehicle system – at least over large scales and if multi-hop acoustics are involved – could be many minutes. The coincidence of time steps is a convenient choice for this paper, but not required.

Centralized control architecture: For the multi-vehicle front tracking problem, we assume a centralized approach for the design stage, i.e., the upper half of Fig. 1, with measurement packets sent to a fusion center, and control commands sent back out to the vehicles. Many multi-vehicle systems

deployed at sea today have a similarly centralized architecture (Ramp et al., 2009; Schneider and Schmidt, 2010). At the same time, there is no reason that the multi-vehicle control system (in the lower right of Fig. 1) has to be centralized.

Breakdown: The integrated feedback system may fail for many reasons, some of which we allude to above, and illustrate below. In general, however, a failure does not compromise the basic abilities of individual vehicles, sensors, or communication system—a great many problems will reside in the modeling and the FPs. This suggests that on a failure of the integrated system, we can still operate assets in the field, and in many cases recover the survey and group capabilities that are available today. Recovery after breakdown is related to the problem of initial detection and convergence to the feature.

Navigation and jump aspects: Although we can accommodate certain operational considerations in our unified statement, for the most part we will consider navigation to be a separate problem, except as described through a standard LTI model with disturbances and sensor noise. Additionally, ocean model assimilation systems often generate renewed updates at intervals of several hours. Thus, while the stochastic identification procedure gives an LTI model over the prediction horizon (the full T_s to T_e), in practice a rolling horizon approach would be used, where system matrices will change at every prediction update. This implies a jump linear system, which requires special treatment in control design. Operationally, there are also issues of vehicles entering and leaving the fleet (for example to charge batteries), which require a jump system approach.

3. Details of the projector

3.1. Fully constrained perturbations

We first describe instance frontal points constrained to the reference gradient direction, passing through the reference frontal point. This model is suitable when perturbations are primarily in the gradient direction, as in when following moderate undulations perpendicular to a long front. Tracking the lateral motion of a cable or chain is a useful analogy for coupled behavior. For M prediction instances (realizations) and N vehicle sites, the projection problem on $r \in \mathbb{R}^2, t \in \{T_s, T_s+1, \ldots, T_e\}$, is:

given	$\phi_i(r,t)$	(scalar field instance) $i \in \{1,\ldots,M\}$
find	$\bar{\phi}(r,t)$	(reference field)
	$\Phi_j(t)$	(target level for ϕ at each vehicle) $j \in \{1,\ldots,N\}$
	$\bar{p}_j(t)$	(reference FP for each vehicle) $j \in \{1,\ldots,N\}$
	$\tilde{p}_{i,j}(t)$	(scalar FP perturbations) $i \in \{1,\ldots,M\}, j \in \{1,\ldots,N\}$
such that	(dropping the time argument and the ranges for i and j)	

(i) $\bar{g}_n(r) := \nabla\bar{\phi}(r)/\|\nabla\bar{\phi}(r)\|_2$ (definition of normalized reference gradient vector)
(ii) $p_{i,j} := \bar{p}_j + \bar{g}_n(\bar{p}_j)\tilde{p}_{i,j}$ (definition of constrained frontal point instance)
(iii) $\bar{\phi}(\bar{p}_j) = \Phi_j$ (reference frontal point identity)
(iv) $\phi_i(p_{i,j}) = \Phi_j$ (frontal point identity per instance)
(v) $p_{i,j} \in \mathbb{P}$ (see Table 1).

Selection of the baseline field $\bar{\phi}$ is a matter of user choice, and could be as simple as the mean field. The target level Φ_j for a given vehicle j is possibly time varying, and defines both the reference FP as well as the instance FP, (constraints (iii) and (iv)). Constraint (ii) establishes that instance FPs lie on a line passing through the reference FP, and in the reference gradient direction.

We show how this construction allows for an explicit estimate of the distance between the vehicle and the frontal point, from measurement of the field variable. Let $\tilde{\phi}_i := \phi_i - \bar{\phi}$. For an instance in the

field (again dropping the time argument), invoking a Taylor series gives

$$\begin{aligned}\phi_i(p_{i,j}) &= \bar{\phi}(p_{i,j}) + \tilde{\phi}_i(p_{i,j}) \\ &= \bar{\phi}(\bar{p}_j + \bar{g}_n(\bar{p}_j)\tilde{p}_{\,i,j}) + \tilde{\phi}_i(\bar{p}_j + \bar{g}_n(\bar{p}_j)\tilde{p}_{\,i,j}) \\ &= \bar{\phi}(\bar{p}_j) + \|\nabla\bar{\phi}(\bar{p}_j)\|_2\ \tilde{p}_{\,i,j} + \tilde{\phi}_i(\bar{p}_j) + \text{h.o.t.}\end{aligned}$$

"h.o.t." indicates higher-order terms, which we drop, and the second term here results from the simple fact that $\nabla\bar{\phi}(\bar{p}_j)$ is parallel to $\bar{g}_n(\bar{p}_j)$. Now let $q_{\,i,j}$ be the location of vehicle j; with $q_{\,i,j} := \bar{p}_j + \bar{g}_n(\bar{p}_j)\tilde{q}_{i,j}$, the same expansion above gives

$$\phi_i(q_{\,i,j}) \approx \bar{\phi}(\bar{p}_j) + \|\nabla\bar{\phi}(\bar{p}_j)\|_2\ \tilde{q}_{i,j} + \tilde{\phi}_i(\bar{p}_j).$$

Finally, define the noisy physical measurement, expressed as a deviation from the reference value:

$$\begin{aligned}z_j^{\phi} &= \left[\phi_i(q_{o,j}) + \nu_j^{\phi}\right] - \Phi_j \\ &\approx \|\nabla\bar{\phi}(\bar{p}_j)\|_2(\tilde{q}_{o,j} - \tilde{p}_{o,j}) + \nu_j^{\phi},\end{aligned} \tag{2}$$

where ν_j^{ϕ} is the measurement noise of the scalar field sensor on vehicle j. The distance between the true frontal point and the vehicle is thus captured in the field measurement.

3.2. Unconstrained translation

Now we look briefly at the more general case, in which the instance FPs do not have to remain on the reference gradient line through the reference FP; they can translate arbitrarily. Such a projector could be used to follow a consistently-shaped structure moving through space, for example by placing FPs along fixed directions from its centroid. The problem statement of the previous section is modified only slightly. Scalar $\tilde{p}_{i,j}$ is changed to a two-vector, and constraints (i) and (ii) are replaced with $p_{i,j} = \bar{p}_j + \tilde{p}_{i,j}$. We then simply write the vector form of the earlier measurement equation:

$$z_j^{\phi} \approx [\nabla\bar{\phi}(\bar{p}_j)]^T(q_{o,j} - p_{o,j}) + \nu_j^{\phi}. \tag{3}$$

It is important to note that this approach, by Eq. (1), cannot follow $\boldsymbol{p}_o$ in the absence of coupling. The target pursuit analogue is illustrative. A vehicle monitoring target range in the N–S direction plus a vehicle monitoring range in the E–W direction are clearly sufficient to pursue the target if they communicate. Implicit at each tracker is the fact that the range gradient points exactly away from the target, and this is parallel to $q - p$. The explicit model coupling is simply that the two trackers are seeing the same target, or $p_1 = p_2$.

4. System integration steps

As we discuss integration of the system described so far, one should bear in mind that beginning with the system identification step, all measurements and state variables now relate to *perturbations.*

4.1. Implementing projection

As noted earlier, projection can be made generic or problem-specific; the main objective is to provide oceanographers with a good balance of automation and interface. We will take here the key ideas from Section 2, and the notation of Section 3, to describe the projection used in our example cases.

First, we specify the reference field $\bar{\phi}$ as the ensemble mean. The set of interest on $\bar{\phi}$ is taken as a front defined by the scalar field value Φ, constant through the epoch; we set $\Phi_j = \Phi$ for all the FPs j. We look through the ensemble data by hand (as an oceanographer might), and select confined areas (the $\mathcal{C}_1$ box from Section 2.2) for each FP at each time step, containing desirable parts of the Φ-level set. When selecting the $\mathcal{C}_1$ box for each FP, we consider realistic vehicle motions ($\mathbb{P}_2$), strong

gradients ($\mathbb{P}_3$), stable reference features ($\mathbb{P}_4$), and feature clarity ($\mathbb{P}_5$). $\mathbb{P}_2$ focuses on satisfying vehicle speed constraints, for which we limit the distance any reference FP can move between time steps: $|\bar{p}_j(t+1)-\bar{p}_j(t)| \leq K_u u_{\max}, \forall j \in \{1,\ldots,N\}, t \in \{T_s,\ldots,T_e-1\}$, where $u_{\max}$ is the vehicle maximum speed, and $0 < K_u < 1$ is a user parameter. Actual vehicles have to follow dynamic perturbations, not only the reference, so K_u is typically well below one. This approach for $\mathbb{P}_2$ can be extended to more complex and practical constraints specific to a given deployment, such as expected distance traveled, mission time, energy, and so on.

To pick the specific $\bar{\boldsymbol{p}}_j$ at each time step, we developed an interface in which the user draws rays across the reference front within $\mathcal{C}_{1,j}$, and the intersections then define $\bar{\boldsymbol{p}}_j$. The reference gradient directions $\nabla\bar{\phi}(\bar{p}_j)$ are computed in the same process, defining a set of N line constraints $\mathcal{C}_{2,j}$. We next project instance variations of the front onto $\mathcal{C}_{2,j}$, using Algorithm 1, giving scalar FP perturbations $\tilde{p}_{i,j}$. The next two paragraphs detail variations of this algorithm that account for grid interpolation and complicated local contours.

Algorithm 1 FP Projector for constrained linearization

Require: $\phi_i(t)$ for $i \in \{1,\ldots,M\}$, and
$[\bar{p}_j(t), \mathcal{C}_1(t), \mathcal{C}_{2,j}(t)]$ for $j \in \{1,\ldots,N\}$, all for $t \in \{1,\ldots,T\}$
for all epoch time steps $t = 1,\ldots,T$ **do**
for all instances $i \in \{1,\ldots,M\}$ **do**
for all FPs $j \in \{1,\ldots,N\}$ **do**
$\mathbb{S}_{0,j} \leftarrow \Phi$-contour of field ϕ_i in $\mathcal{C}_{1,j}(t)$
$p_{i,j}(t) \leftarrow$ unique intersection of $\mathbb{S}_{0,j}$ and $\mathcal{C}_{2,j}(t)$
$\tilde{p}_{i,j}(t) \leftarrow p_{i,j}(t) - \bar{p}_j(t)$
end for
end for
end for

Computed contours consist of vertices and straight segments, so ensuring smooth interpolation to find the intersection of a local contour $\mathbb{S}_0$ and a line $\mathcal{C}_2$ takes some care. We choose a number of points on $\mathbb{S}_0$ in the immediate neighborhood of $\bar{p}_j$, compute the signed distance of each point to $\mathcal{C}_2$, and record the closest point on either side of $\mathcal{C}_2$. $p_{i,j}$ is then computed as the intersection of $\mathcal{C}_2$ with the line connecting these two closest points. This neighborhood around $\bar{p}_j$ is unique to each FP and can be chosen tighter than $\mathcal{C}_1$ if needed—this helps to ensure uniqueness of frontal points in each instance.

The instance front is not always well-behaved. Eddies and similar fluctuations occur, and these create complicated $\mathbb{S}_0$, which can induce multiple intersections with $\mathcal{C}_2$. We attempt to smooth these out by point selection based on the median filter, well-known in signal and image processing for removing outliers (Gallagher and Wise, 1981). First, find a set of n tight contours for ϕ-levels very close to Φ; we refer to this as a contour family $\mathbb{S}(\supset\mathbb{S}_0)$. For a given frontal point, compute the intersection of $\mathcal{C}_{2,j}$ with each member of $\mathbb{S}$; there may be several intersections for a given member, but we only keep the one closest to $\bar{p}_j$. Compute the signed distance to $\bar{p}_j$ for each of these n closest intersections, and then take the median across $\mathbb{S}$ as $\tilde{p}_{i,j}$. In the extreme case where $\mathbb{S} \cap \mathcal{C}_{2,j} = \emptyset$ (that is, no crossing contours are available in the area of interest), we expand the local search neighborhood, and with it $\mathbb{S}$, enough to find at least one intersection, so the FP can be placed. For our computations, this overall approach has proven very effective for limiting sudden jumps in $\tilde{p}_{o,j}$.

In some instances the front breaks down to a point that the real-time system $\mathbb{E}_o$ simply cannot follow, and so creates major outliers in the FPs. This is a failure to satisfy $\mathbb{P}_5$. Linear system identification does not handle outliers well, and if needed we can concede to *exclude outlier realizations from system identification*. If the fraction of excluded trials in the ensemble is high, however, we clearly cannot expect the real-time system to do well. Practically speaking, any operational vehicle system should be able to switch to a "recovery" mode if such conditions are encountered in a physical experiment; we expect that even if the oceanographic pursuit mission fails, the ensemble can still aid failure-detection algorithms.

Once the population of FP perturbations, $\tilde{p}_{i,j}(t)$ has been created, it should to be reviewed to ensure that all (or an acceptable fraction) are in $\mathbb{P}$. $\mathbb{P}_2$ is again notable because it relates to vehicle control, speed and maneuvering capabilities. For example, after running the projector, we check the vehicle speeds: $|p_{i,j}(t+1) - p_{i,j}(t)| \leq u_{\max}$. Corrections in the process are likely to be made at the level of the reference FPs, since Algorithm 1 even with its variations has few tuning parameters.

4.2. Stochastic identification of instance frontal points

Linear system identification (Ljung, 1999) – the determination of an LTI dynamical system model to explain measurements – is highly effective when the input signals are known. For single-input, single-output systems, strong results are regularly produced by robust time- and frequency-domain approaches; for multi-input, multi-output systems, subspace identification algorithms (Van Overschee and De Moor, 1996) are widely used. The stochastic analogue, where we wish to explain the data as the output of a coupled system driven by noise inputs, is much less developed. In fact, aside from the two key references above, there seem to be few technical results on the topic in the last fifteen years. Intuitively, one can appreciate that the difficulty of stochastic identification reflects a tenuous optimization problem—the raw constraints are noisy output traces, the belief that the driving signals are Gaussian, and a system order specification. System order selection is aided by computing singular values of the Hankel matrix on the input data; this is very similar to what is done for balanced realization and model reduction (Moore, 1978). We describe specific choices for model order selection in Section 5.

Given the order, stochastic subspace methods involve first estimating a sequence of *states* from the data; this is in strong contrast to the input–output viewpoint of classical system identification. The states are interpreted as the output of Kalman filter predictors, and they can be obtained through QR factorization and singular value decomposition (SVD). Once these states are identified, a linear least-squares regression yields the unknown system matrices. The specific numerical algorithm we use is Matlab's `n4sid` (Van Overschee and De Moor, 1996), providing a model in innovations form. In our case, the outputs are the linearized perturbations $\tilde{\boldsymbol{p}}$ (with t denoting the time step):

$$\begin{aligned}\boldsymbol{x}_p(t+1) &= A_p\boldsymbol{x}_p(t) + K\boldsymbol{e}(t)\\ \tilde{\boldsymbol{p}}(t) &= C_p\boldsymbol{x}_p(t) + \boldsymbol{e}(t),\end{aligned}$$

and A_p, K, C_p, and $\mathrm{Cov}(\boldsymbol{e}) \triangleq R_e$, are generated by the algorithm. We reiterate that $\tilde{\boldsymbol{p}}$ is available only for the model ensemble, not for the physical instance.

Several specific challenges that practitioners of stochastic algorithms face include high sensitivity to pre-filtering of data, and, perhaps more seriously, the requirement of long signals. Indeed, guarantees of optimality and asymptotic unbiasedness only exist for an infinite number of samples, and statistical analysis with finite sample length remains an open problem (Van Overschee and De Moor, 1996). In practice `n4sid` can generate useful results with short sequences, and this would be a necessity for oceanographic pursuit, where typically few, if any, repeated events are observed. One could conjecture that from a system identification point of view, longer data traces are desirable even if the latter portions of them are useless for forecasting; this question is beyond our scope.

`n4sid` includes a few user parameters that can be tweaked to improve performance: SVD weighting, forward prediction horizon, and number of past inputs used for prediction. In both Van Overschee and De Moor (1996) and Ljung (1999), however, it is made clear that optimal choices for most of these settings are still open subjects of research.

4.3. Connecting with the vehicle system

We take in our framework a linear time-invariant model of vehicle system behavior; see Kinsey et al. (2006) for a recent review of underwater vehicle navigation systems that support control. Our formulation is developed for a *group* of vehicles, which may or may not share physical disturbances and navigation aspects. This aggregate vehicle system is described by the state-space dynamics matrix A_q, gain B_q, and output C_q; the vehicles' process noise vector $\boldsymbol{w}_q$ has covariance Q_q, and measurement

noise ν_q has covariance R_q. A_q and B_q are block diagonal, since the vehicles have no coupling except through the control and possibly the disturbances. The integrated open-loop system is:

$$\begin{Bmatrix} \boldsymbol{x}_p(t+1) \\ \boldsymbol{x}_q(t+1) \end{Bmatrix} = \left[\begin{array}{c|c} A_p & 0 \\ \hline 0 & A_q \end{array}\right] \begin{Bmatrix} \boldsymbol{x}_p(t) \\ \boldsymbol{x}_q(t) \end{Bmatrix} + \left[\begin{array}{c} 0 \\ \hline B_q \end{array}\right] \boldsymbol{u}(k) + \left[\begin{array}{c|c} K & 0 \\ \hline 0 & I \end{array}\right] \begin{Bmatrix} \boldsymbol{e}(t) \\ \boldsymbol{w}_q(g) \end{Bmatrix}. \tag{4}$$

As written, dynamics of the frontal point perturbations and vehicle perturbations are decoupled, giving the macro block-diagonal structure shown. Clearly $\boldsymbol{x}_p$ is affected by neither $\boldsymbol{x}_q$ nor $\boldsymbol{u}$, but is $\boldsymbol{x}_q$ affected by $\boldsymbol{x}_p$? Although in principle one could establish such coupling, the fact that A_p is developed from stochastic identification, whereas A_q is not, implies that simulation of vehicles and their control system would have to be added into $\mathbb{D}$. This turn toward a self-referential projector is left for future work. Off-diagonal blocks in the process noise gain matrix can be left at zero by the same rationale.

The output equations are where the oceanographic and the vehicle systems most strongly interact:

$$\boldsymbol{z}(t) = \begin{Bmatrix} \boldsymbol{z}^{\phi}(t) \\ \tilde{\boldsymbol{q}}(t) \end{Bmatrix} = \left[\begin{array}{c|c} -GC_p & GC_q \\ \hline 0 & C_q \end{array}\right] \begin{Bmatrix} \boldsymbol{x}_p(t) \\ \boldsymbol{x}_q(t) \end{Bmatrix} + \left[\begin{array}{c|c|c} -G & I & 0 \\ \hline 0 & 0 & I \end{array}\right] \begin{Bmatrix} \boldsymbol{e}(t) \\ \nu^{\phi}(t) \\ \nu_q(t) \end{Bmatrix}, \tag{5}$$

with $G = \mathrm{diag}(\|\nabla\bar{\phi}(\bar{p}_1)\|_2, \ldots, \|\nabla\bar{\phi}(\bar{p}_N)\|_2)$, and it is assumed that G varies slowly enough not to interfere with the feedback controller ($\mathbb{P}_4$). Note that the disturbance $\boldsymbol{e}(t)$ enters the output equation, in accordance with the innovations form. Its inclusion highlights a subtlety in Eq. (2); whereas $\tilde{q}$ represents a physical vehicle perturbation via a strictly proper system, $\tilde{p}$ does not. $\tilde{p}$ is the output of our projection algorithm with no low-pass constraints at all, and as defined in the stochastic identification it has "jitter".

4.4. Estimation of the integrated system

The aggregate process noise vector (dropping the time argument) is $\boldsymbol{w} = [(K\boldsymbol{e})^T,\ (\boldsymbol{w}_q)^T]^T$ and the total measurement noise vector is $\nu = [(-G\boldsymbol{e} + \nu^{\phi})^T,\ (\nu_q)^T]^T$. Process noise for the FPs may be correlated with the vehicles' process noise, for example due to currents; this is captured by the matrix Q_{eq}. We assume, however, that beyond $\boldsymbol{e}$, no process noise is correlated with any sensor noise. For use in the generalized Kalman filter (Simon, 2006), expanding out the expectations gives:

$$\mathbb{E}\left[[\boldsymbol{w}^T, \nu^T][\boldsymbol{w}^T, \nu^T]^T\right] = \left[\begin{array}{cc|cc} KR_eK^T & Q_{eq} & -KR_eG^T & 0 \\ * & Q_q & 0 & 0 \\ \hline * & * & (GR_eG^T + R_{\phi}) & 0 \\ * & * & * & R_q \end{array}\right]. \tag{6}$$

5. Examples of projection and identification

We now present three projection and stochastic identification examples, for the fully constrained linearization of Section 3.1. The first case is an LTI chained mass system, for validation. Next, we study a double gyre model, simulated using a finite-volume Navier–Stokes solver (Ueckermann and Lermusiaux, 2012; Sapsis and Lermusiaux, 2012). The double gyre is a canonical fluid mechanics problem, highly nonlinear and unstable; while these factors would seem to position it poorly for system identification, there is also a dominant wave-like behavior that greatly contributes to coupling, and to a locally linear behavior. The double gyre is a generic and dimensionless scenario with few physical parameters, and thus useful for benchmarking. The third dataset is substantially more difficult and realistic, focusing on three-dimensional flows north and east of Taiwan. This particular set was part of a larger study (Gawarkiewicz et al., 2011) emphasizing prediction and uncertainty for an ocean setting with complex multiscale dynamics, from internal tides and waves to large-scale currents. One of the main features noted was sporadic intrusions by the Kuroshio Current into the so-called "Cold Dome". These elements make the Taiwan front a challenging case for oceanographic pursuit.

Table 2
Parameter values used for chained mass system.

k_n	k_g	b_n	b_g	$w_{1,\ldots,7}$	w_8
0.02	0.048	0.3	0.01	0.13	0.26

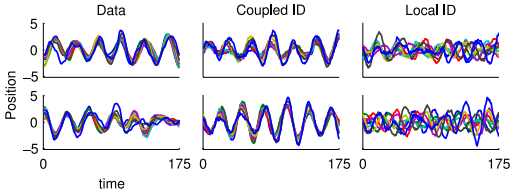

(a) Time series of selected chained mass perturbations. The eight sites are shown on each plot. Left: two selected instances of original simulated data. Middle: simulations using coupled system identification. Right: simulations using local system identification.

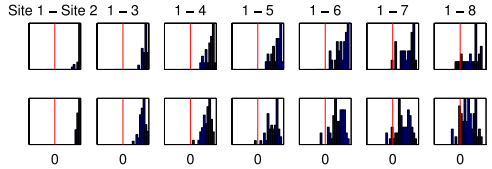

(b) Top: histograms of FP correlation coefficients from original chained mass simulated data (25 instances). Bottom: from identified model simulations (50 instances). Only correlations between Site 1 and the others are shown.

Fig. 3. Chained mass system identification results.

5.1. Identification of an LTI chained mass

The chained mass system has eight masses arranged nominally along a line and undergoing lateral perturbations. Each mass is tied to ground lightly by a spring k_g and damper b_g, connected to adjacent neighbors with springs k_n and dampers b_n, and forced by zero-mean white noise of variance w_j. The parameters are tuned to give magnitudes and frequencies similar to those in the double gyre FP perturbations; see Table 2. The initial condition for each point is $p_j(0) = 0.1 \cdot \sin\left(j\frac{2\pi}{N-1}\right) + 0.5\eta,\ \forall j \in \{1, \ldots, N\}$, where $\eta = \mathcal{N}(0, 1)$.

With the `n4sid` algorithm, we used a forward prediction horizon of four, and four past inputs. We chose a model order twice the number of sites; this is of course the order of the generating system, but also consistent with our choices later. Fig. 3a compares time series for two "data" instances (left) to two time series constructed by stochastic simulation of the identified system and its noise statistics (middle). The "data" perturbations show a dominant resonant mode and tight coupling in phase; this strong coupling is supported by the correlation coefficients shown in Fig. 3b. Time series and correlations from the identified model show virtually the same properties. The Hankel singular values of the input data (Fig. 4 (left)), however, show that the stochastic ID does not exactly recover the original system. We also ran a local second-order identification for each mass. This is shown in Fig. 3a (right), where all relative phase information has been lost. This model is what we use below in a "loners" control scheme made up of non-communicating vehicles.

5.2. Identification of a double gyre front

The feature of interest in the double gyre model is a vorticity contour; although vorticity is not a typical scalar measurement taken at sea, turbulence scales and shear stress are. We study a

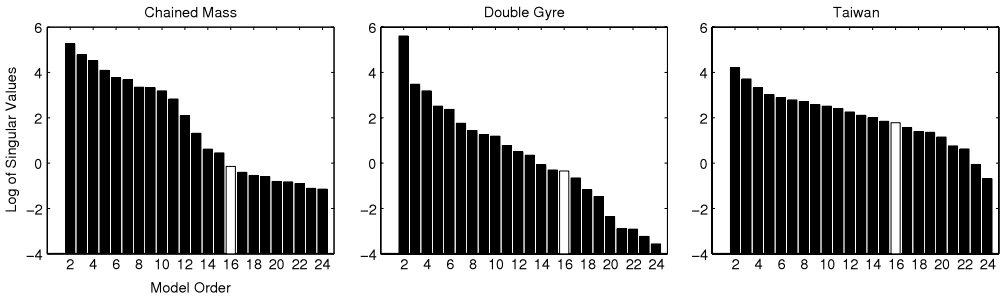

Fig. 4. Hankel singular values for identification data. White bars indicate the model order chosen for each case study.

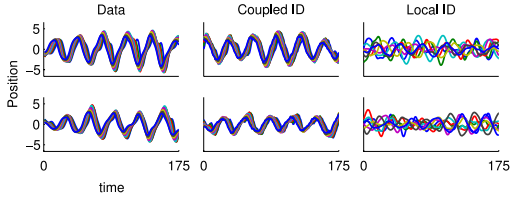

(a) Time series data, same format as Fig. 3a.

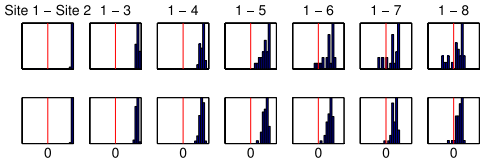

(b) Correlation coefficients, same format as Fig. 3b.

Fig. 5. Double gyre system identification results.

176-timestep period that takes place when the gyre system is just starting to become unstable, soon after which all structure is lost. Eight reference FP trajectories were designed, maintaining equal spacing along the reference front. We then ran the projector on 25 randomly-selected ensemble members, and performed stochastic identification. The perturbations were lightly low-pass filtered before being passed into `n4sid`, so as to reduce grid effects. The Hankel singular values shown in the middle plot of Fig. 4 have only one clear break point, at a model order of *two*. This choice would allow only for one oscillator, with differences between the behavior at each site merely representing phase shifts. The actual response is somewhat more complicated, and we chose a model order of sixteen. A higher number than this would risk overfitting, while sixteen can be justified at least from a comparison point of view: the uncoupled dynamics is naturally described with eight independent pairs of states. We used a forward prediction horizon of four, and four past inputs in `n4sid`.

As for the chained mass, in Fig. 5a we compare perturbations for two instances (left) to two time series constructed by stochastic simulation of the identified system and its noise statistics (middle). The instance fluctuations are oscillatory and coordinated, with a traveling wave characteristic; amplitude, phase relations, and frequency content are all well-captured in the identified model. Correlation coefficients shown in Fig. 5b confirm the strong coupling.

5.3. Identification of a front off Taiwan

The Taiwan ensemble has fifty instances, the time step is three hours, the run is 68 steps long (about eight days), and the model grid is 4.5 km. The feature of interest here is a persistent temperature

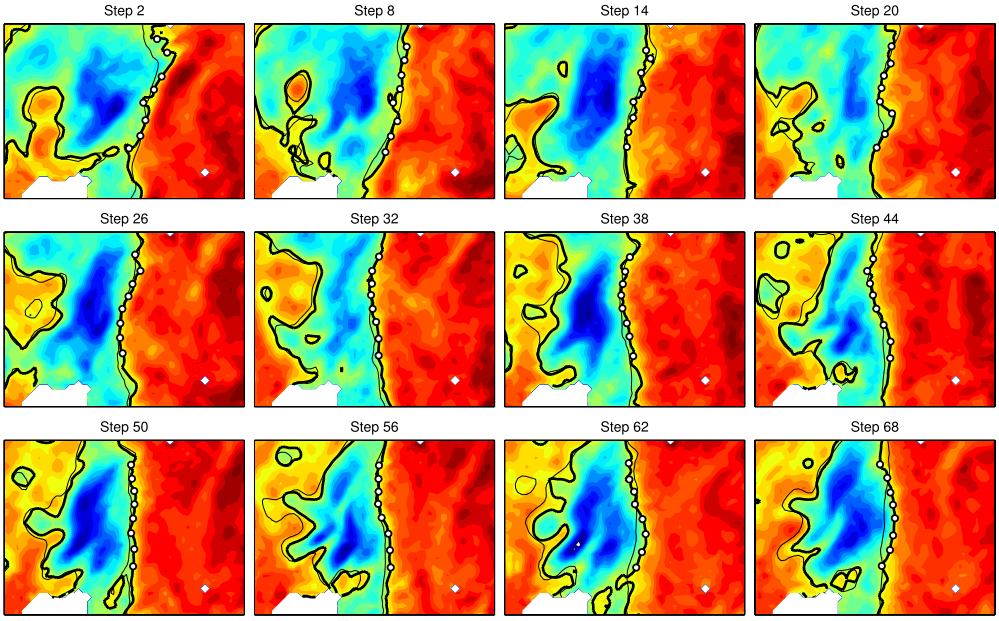

Fig. 6. Snapshots over time of a single instance of the Taiwan dataset; the northern edge of Taiwan is visible at the bottom of each frame. This instance was used for projection and system identification. The reference contour is a thin black line, and the true contour is a thick black line. True FPs are white with black outlines. The domain of each box is 248 × 180 km. The time between snapshots is 18 h, which makes out-of-phase semi-diurnal (and diurnal) internal tidal effects visible.

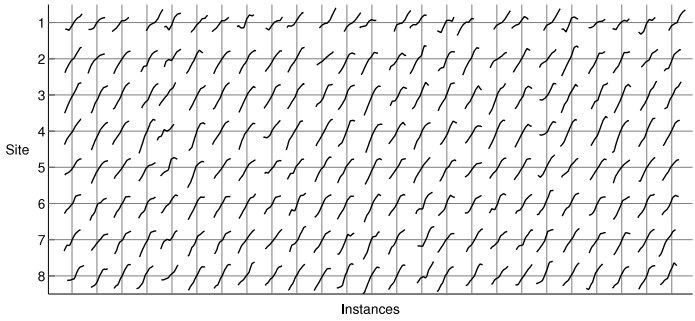

Fig. 7. Slices of temperature field along constraint $\mathcal{C}_2$ for each FP at time step 30. Twenty-five ensemble instances used for projection and identification are shown. The horizontal axis represents distance along $\mathcal{C}_2$ and the vertical axis represents temperature. Slices are positioned horizontally relative to the reference front, which is the origin of each subplot.

front at 50 m depth. Compared with the double gyre case above, this front exhibits much smaller perturbations relative to motion of the reference, and far less structure. We followed the same general procedure in the projector, picking a temperature level and a physical area where gradients are strong and the front stays well-formed. Snapshots over time of a single realization are shown in Fig. 6.

To illustrate the key aspects of this scenario pertinent to identification and oceanographic pursuit, Fig. 7 shows slices of the temperature field along the reference gradient cut, and centered at each reference FP. These slices at time step 30 are from twenty-five instances used for projection and identification. Most of them have a clean albeit nonlinear shape, confirming definition of the front and strong gradients. Some sharp corners are visible, caused by interpolation on the model grid.

Hankel singular values on the right side of Fig. 4 show no break point for any order; this flatter shape indicates the difficulty and lack of structure in the dataset. We chose a system order of sixteen,

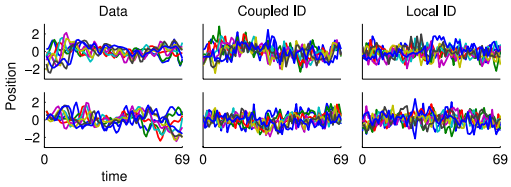

(a) Time series data, same format as Fig. 3a.

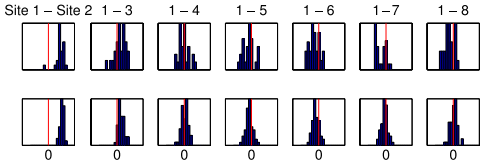

(b) Correlation coefficients, same format as Fig. 3b.

Fig. 8. Taiwan system identification results.

for consistency with twice the number of FPs. With the `n4sid` algorithm, we used a forward prediction horizon of six, and three past inputs. Time series are plotted in Fig. 8a, and these confirm reduced coupling. Yet there are still meaningful correlations between neighboring points, as seen in Fig. 8b. These correlations are replicated in the identified system, and we will show in the next section that even this modest level brings a benefit to communication-constrained control.

6. Closed-loop control examples

For each of the three scenarios above, our controller design uses an LTI system model found by identification on a subset of ensemble instances; the design is evaluated on separate subsets. We emphasize that these cases are considered only proof of concept; the settings and operational parameters used do not necessarily reflect real conditions.

6.1. Explanation of controllers

As stated in the Introduction, one of our main objectives is to develop controllers suited for communication constraints; one of the most difficult is packet loss. We define four vehicle communication models:

- **NC**: no communication between vehicles for the purpose of oceanographic pursuit.
- **PC**: perfect communication (lossless and instantaneous).
- **IL**: independent losses. In a cycle, vehicles report measurements instantaneously to a fusion center, which then sends out a command set instantaneously. There are thus N inbound sensor packets and N outbound sensor packets. Losses in the sensor packets are binary, and described as a set of independent Bernoulli processes, having success probabilities $[\beta_1, \ldots, \beta_N]$. Losses in the command packets are similarly described by the Bernoulli parameters $[\alpha_1, \ldots, \alpha_N]$.
- **TL**: total losses. In a cycle, vehicles report measurements instantaneously to a fusion center, which then sends out a command set instantaneously. The sensor packets all succeed with probability β, or they all fail. The command packets all succeed with probability α, or they all fail. These total failures on the command and the sensor sides are independent.

This setup ignores the role of communication delays, interference, and scheduling—aspects that, like packet loss, can be handled rigorously from an LTI framework. We compare five controllers, each

subject to vehicle navigation noise and physical disturbances:

- "Non-reacting" uses local linear quadratic Gaussian controllers (LQG, comprising a Kalman filter (KF) for state estimation and a linear quadratic regulator (LQR) for full-state feedback) to place vehicles at the reference FPs. The vehicles do not alter their trajectories based on field measurements, and thus do not have to communicate: **NC**. The local model from stochastic identification is used for estimation with a KF at each vehicle. This approach gives a practical upper bound on real-time estimation error.
- "Loners" uses a set of N independent LQG controllers to servo to the estimated front. Each vehicle has only a local model derived from the stochastic identification: **NC**.
- "Naïve" applies a standard LQG controller design, given a fully coupled model and assuming **PC**. The simulation uses **IL**, while the KF takes the standard approach for handling missed measurements.
- "All-or-none" uses the dynamic programming procedure of Imer et al. (2006), on a fully coupled model and where **TL** is assumed. In simulation, we impose **IL**. The control design uses the means of the α and β vectors that make up the simulation **IL** model. A regular missed-measurement KF applies, but the **TL** scalar α invokes an adjustment to the prior due to the uncertainty of the control action (Garone et al., 2010).
- "Lower Bound" applies a standard LQG controller design, given a fully coupled model and assuming **PC**. The simulation is **PC** also, and hence this controller is expected to be the best.

When control packets are not received, vehicles stay in place. The control objective is to minimize the squared positioning error over time, for all sites/vehicles: $\Sigma_{j=1}^{N}\Sigma_{t=1}^{T}\left(\tilde{q}_{o,j}(t)-\tilde{p}_{o,j}(t)\right)^2$. For LQR design, we weight positioning error ten times more heavily than control actions.

6.2. Results

6.2.1. Gain margins

We return briefly to the gradient condition $\mathbb{P}_{3b}$: $\|\nabla\phi_i(p_i+\Delta)-\bar{g}\|_2 \leq \epsilon\kappa$ for each FP, where $0<\epsilon\ll 1$. We consider Δ as a positioning error constrained to be along $\mathcal{C}_2$. ϵ is related explicitly to gain margin, a standard means for quantifying robustness in the closed loop, while κ ultimately limits the bandwidth of the closed-loop system. The first ten instances of the Taiwan data (from Fig. 7) are re-plotted in Fig. 9, now showing more specifically the slices (thick line) overlaid on top of twice the reference gradient slope (thin line), and centered at the instance FP. Differences between the reference gradient (not shown) and the instance slice can be interpreted as sensor gain variations from the point of view of an estimator. Although we use a time-varying KF, it is noteworthy that the steady-state KF and the LQR each have a guaranteed per-channel gain margin of zero to two (Safonov and Athans, 1977) and thus the thin diagonal line and the horizontal define a [0,2]-sector that is highly desirable.[3] We see that all of these instantaneous slices fall within (or very close to within) the sector, at least near the origin.

6.2.2. Controller performance

For our simulations, we assume that all of the vehicles have identical dynamics and sensors. The scalar field measurement noise covariance R_ϕ is the most sensitive parameter, and our results compare estimation performance varying R_ϕ, with other parameters fixed. For all cases, we simulate control on instances separate from those used for projection and identification, and use $\tilde{p}_{o,j}$ generated by the omniscient projector for ground-truth. We use one hundred simulation instances for the chained mass, and twenty-five instances each for the double gyre and Taiwan datasets. Vehicle noise parameters are set as $R_q = 0.01$ and $Q_q = 0.01$. These are chosen to roughly describe a physical scenario, but in reality would depend on the specific vehicles and environmental conditions. We consider vehicle noise only along $\mathcal{C}_2$. Packet success probabilities are set as $\bar{\alpha} = [0.7, 0.8, 0.6, 0.9, 0.7, 0.8, 0.6, 0.9]$,

[3] Despite its wide use in applications, the generic LQG does not have guaranteed stability margins (Doyle, 1978).

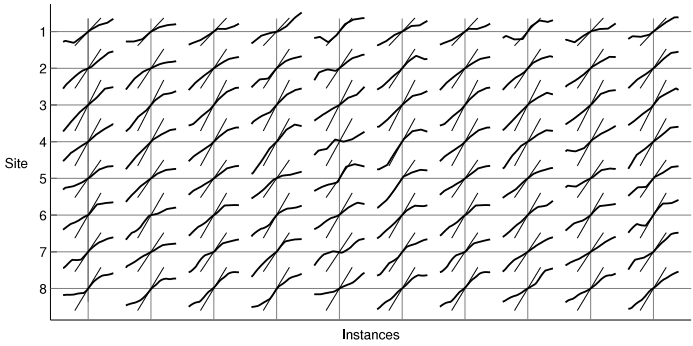

Fig. 9. Ten instances of Taiwan temperature slices (thick line) along $\mathcal{C}_2$, overlaid on top of twice the reference gradient (thin line). The horizontal axis represents distance along $\mathcal{C}_2$, centered at the instance FP, and the vertical axis represents temperature. The [0,2] sector defines a region of stability for the Kalman filter.

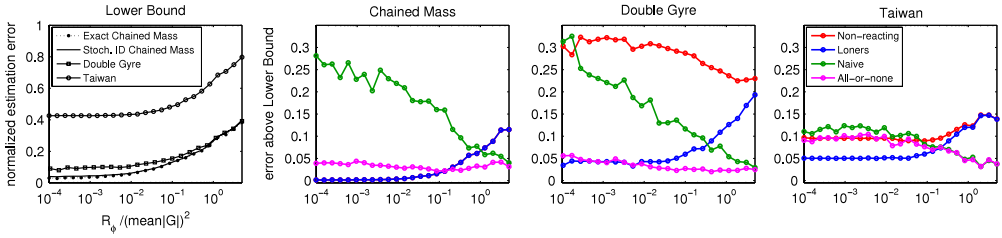

Fig. 10. Summarized estimation performance. Estimation error is the RMS over time of $(\hat{\tilde{p}} - \tilde{p})$, averaged across points and realizations. **Left plot:** performance of Lower Bound in each dataset, as a function of R_ϕ. For the chained mass, the "Exact" controller is designed using the true system and noise. All other controllers use the outcome of stochastic identification. **Right three plots:** difference in estimation error relative to Lower Bound for all controllers in each dataset.

$\bar{\beta} = [0.6, 0.9, 0.7, 0.6, 0.8, 0.9, 0.7, 0.8]$.[4] From examination of empirical transfer functions, we verified that the control system bandwidth is sufficient to track the dominant motion of the feature.

We show results comparing performance with varying R_ϕ in Fig. 10. R_ϕ in the plots is scaled with the mean gradient squared, so as to capture the importance of noise relative to the ambient signal.[5] We compute the RMS (over time) estimation error $(\hat{\tilde{p}} - \tilde{p})$, averaged across sites and across simulation instances, and this error is normalized by the same statistic on the perturbations used for system identification, in order to compare results across datasets. The leftmost plot of Fig. 10 shows the performance of Lower Bound. Major differences underscore the relative difficulties mentioned earlier, and in particular Taiwan has less structure in its perturbations and considerably more noise, caused both by grid effects as well as erratic projections in the noisy field. At the other extreme, Lower Bound for the chained mass, designed using the exact system model, performs slightly better than Lower Bound based on the identification—demonstrating that the control system resulting from identification is successful. Lower Bound for the chained mass does far better than in either the double gyre or Taiwan cases, largely because the chained mass has an ideal projector and exact gradients, such that estimation performance is independent of positioning.

In the right plots of Fig. 10, we plot differences between the estimation error of each control method and that of Lower Bound. We will describe a few key aspects: positioning, coupling, and packet losses.

[4] In other work, we have investigated a mixed-loss controller that explicitly designs for **IL** (Reed and Hover, 2013); however, mixed-loss control performs similarly to "All-or-none" with this set of loss probabilities.

[5] We have fixed R_q, so that these figures have essentially 1/SNR on the horizontal axis.

Non-reacting and Loners use the same local model for estimation; the difference is that Loners positions vehicles reactively to pursue the front. Non-reacting and Loners are identical in the chained mass case, but with the ocean model datasets, estimation becomes more difficult as positioning worsens in the nonlinear field (see Fig. 2(b)). This fact yields a dramatic improvement in estimation for Loners over Non-reacting in the double gyre, where slices are smooth but slope deforms with distance from the front. In the Taiwan dataset, Loners also outperform Non-reacting, but to a lesser extent and mostly at low noise values. The Taiwan front has a more variable shape and shallower slope, along with more pronounced grid effects. For every scenario, the bandwidth of the estimator decreases with large R_ϕ—the entire system becomes sluggish, vehicles move less, and thus the differences between Loners and Non-reacting disappear.

The role of coupling can be understood by comparing performance of Loners and Lower Bound. In the chained mass model, the improvement of Lower Bound over Loners increases significantly from small values as R_ϕ gets larger, reflecting clearly that coupling can offset the bad effects of sensor noise. In the two ocean model datasets, Lower Bound outperforms Loners even at very low R_ϕ, which may seem surprising but is due to the nonlinear shapes of the slices. The KF does the best it can with the given noise level, but the slice nonlinearities invariably create a "pseudo-noise", so the filter is mistuned. Loners and Lower Bound both suffer, but Lower Bound again has the advantage of coupling.

A more nuanced effect of coupling relates to packet loss, and an interesting comparison can be seen between All-or-none and Loners. Communication losses hurt estimation because of missed measurements, but also hurt positioning directly through dropped control commands. At low levels of sensor noise, Loners outperform the controllers that are subject to packet loss since uncertain communication deteriorates the benefits of coupling. At larger R_ϕ, above a crossover, Loners cannot succeed with only local modeling, and All-or-none prevails, even with lossy communication.

Comparing now the two controllers subject to packet loss, Naïve and All-or-none use the same coupled model for estimation, and are subject to the same stochastic packet loss sequence in each simulation. All-or-none, however, explicitly takes loss into account in both controller design and in the uncertainty of the KF prior. This is a strong stabilizing effect, as Naïve shows much more sensitivity to sequences of packet losses; its errors can be very small or very large, giving rise to the jumps up and down in Fig. 10. Naïve performs poorly especially at low R_ϕ, the error converging to that of All-or-none as R_ϕ increases. This convergence simply reflects decreased estimator bandwidth.

The estimation performance over time from one example site and realization for each dataset is shown in Fig. 11. The double gyre and chained mass cases show all methods clearly tracking the dominant oscillations, although Loners and Naïve often exhibit more over- or undershoot than the coupled methods. In the double gyre, estimation errors of Non-reacting are clear, especially near the peaks of perturbation magnitude. The Taiwan case is drastically different, again showing the difficulty of this dataset. None of the controllers is able to track the high-frequency variations, but Loners and Non-reacting are much more sluggish than the coupled methods. These time series also illustrate system startup (in the first ten or so samples), an important factor to consider in real-world operations. We have had no difficulties in initializing to the reference FPs in any of our test cases.

7. Conclusion

We have articulated an integrated framework for dynamically sampling the ocean using a group of communicating mobile agents. Our new concept is that locally linear behavior of an ocean process admits strong estimation and control techniques on short time scales; this will allow multiple cooperating vehicles to decompose spatial and temporal variations in the field, and actively follow a dynamic feature of interest. The stochastic dynamical model supporting our controller design is created via a projection from an ocean forecast ensemble into succinct vehicle coordinates, and this is the main innovation of our work. We have demonstrated that control and estimation designs resulting from these identified models are successful, in studies with three example datasets.

Local linearization of ocean model simulations is a critical element of the procedure, that we address first with the frontal points concept and then with subspace identification methods. Neither of these is easy. Frontal point generation has an implicit ergodicity assumption, and requires domain expertise as implied in the constraints $\mathbb{P}$. At the same time, subspace identification, even for systems

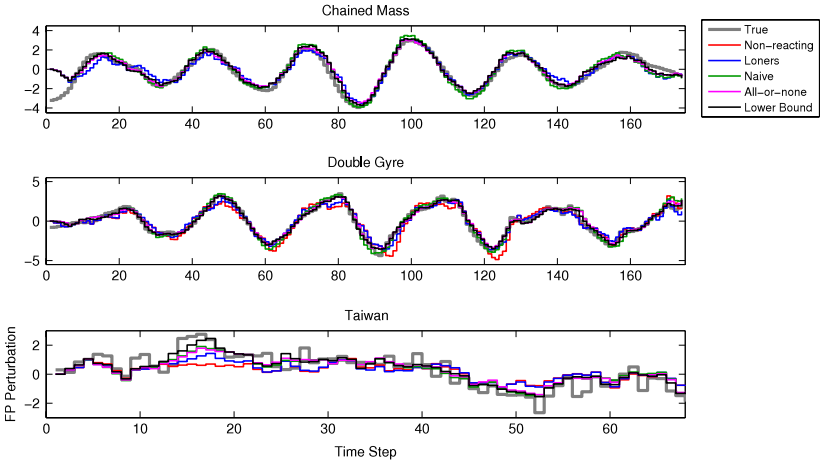

Fig. 11. Estimation performance for one example site and realization. These cases are in the region where All-or-none outperforms Loners: $R_\phi/(\text{mean}(|G|))^2 = 0.82$.

which are known to be linear, is difficult, and we have as well restrictions on record length, plus gridding effects. Many specializations and improvements can be made to the preliminary methods we have described.

We firmly believe that linear stochastic models, however, are key to cogent analysis and design procedures when multiple vehicles have to operate with realistic navigation and communication limits. The LTI model allows for classical and scalable multivariable estimation and control, as well as rigorous contemporary approaches for lossy communications. In the control simulation results presented, we can clearly see the importance of actively pursuing the front as opposed to passive tracking, and of global model-based estimation. Moreover, these outcomes illustrate performance of a pursuit system in different conditions, illuminating an interesting trade-off space for designing deployments. An important dimension in it is the number of vehicles and frontal points compared to the length of the feature, as this affects both spatial resolution of the reconstruction as well as the expected level of coupling between sites. These factors need to be weighed in turn against the costs and capabilities of the individual vehicles. For example, if highly accurate sensors are available, it may not be necessary to set up communication at all for the purposes of oceanographic pursuit. On the other hand, vehicles with less expensive, lower-quality sensors can be deployed in larger numbers and will likely benefit from a coupled model and hence collaboration. More broadly, heterogeneous sensor and vehicle networks are increasingly likely in practice. The integrated framework we have developed should support strong trade-off studies along these lines.

Acknowledgments

The work is supported by the Office of Naval Research, Grant N00014-09-1-0700 and the National Science Foundation, Contract CNS-1212597. We thank Prof. P. Lermusiaux for the model data used in this paper.

References

Baillieul, J., Antsaklis, P., 2007. Control and communication challenges in networked real-time systems. Proc. of the IEEE 95 (1), 9–28.

Camilli, R., Reddy, C., Yoerger, D., Van Mooy, B., Jakuba, M., Kinsey, J., McIntyre, C., Sylva, S., Maloney, J., 2010. Tracking hydrocarbon plume transport and biodegradation at Deepwater Horizon. Science 330 (6001), 201.

Canuto, V., Dubovikov, M., 2005. Modeling mesoscale eddies. Ocean Model. 8 (1), 1–30.

Das, J., Py, F., Maughan, T., O'Reilly, T., Messié, M., Ryan, J., Sukhatme, G.S., Rajan, K., 2012. Coordinated sampling of dynamic oceanographic features with underwater vehicles and drifters. Int. J. Robot. Res. 31 (5), 626–646.

D'Asaro, E., Lee, C., Rainville, L., Harcourt, R., Thomas, L., 2011. Enhanced turbulence and energy dissipation at ocean fronts. Science 332 (6027), 318–322.
Doyle, J., 1978. Guaranteed margins for LQG regulators. IEEE Trans. Automat. Control 23 (4), 756–757.
Dunbabin, M., Marques, L., 2012. Robots for environmental monitoring: significant advancements and applications. IEEE Robot. Autom. Mag. 19 (1), 24–39.
Farrell, J.A., Pang, S., Li, W., 2005. Chemical plume tracing via an autonomous underwater vehicle. IEEE J. Ocean. Eng. 30 (2), 428–442.
Ferrari, R., 2011. A frontal challenge for climate models. Science 332 (6027), 316–317.
Fiorelli, E., Leonard, N.E., Bhatta, P., Paley, D.A., Bachmayer, R., Fratantoni, D.M., 2006. Multi-AUV control and adaptive sampling in Monterey Bay. IEEE J. Ocean. Eng. 31, 935–948.
Gallagher Jr., N.C., Wise, G.L., 1981. A theoretical analysis of the properties of median filters. IEEE Trans. Acoust. Speech Signal Process. 29 (6), 1136–1141.
Gangopadhyay, A., Robinson, A.R., 2002. Feature-oriented regional modeling of oceanic fronts. Dyn. Atmos. Oceans 36 (1), 201–232.
Garone, E., Sinopoli, B., Casavola, A., 2010. LQG control over lossy TCP-like networks with probabilistic packet acknowledgements. Int. J. Syst. Control Commun. 2 (1), 55–81.
Gawarkiewicz, G., Jan, S., Lermusiaux, P.F., McClean, J.L., Centurioni, L., Taylor, K., Cornuelle, B., Duda, T.F., Wang, J., Yiing, J.Y., et al., 2011. Circulation and intrusions northeast of Taiwan: chasing and predicting uncertainty in the Cold Dome. Oceanography 24, 110–121.
Godin, M., Zhang, Y., Ryan, J., Hoover, T., Bellingham, J., 2011. Phytoplankton bloom patch center localization by the Tethys Autonomous Underwater Vehicle. In: OCEANS'11. MTS/IEEE, pp. 1–6.
Graham, R., Py, F., Das, J., Lucas, D., Maughan, T., Rajan, K., 2013. Exploring space–time tradeoffs in autonomous sampling for marine robotics. In: Experimental Robotics. In: Springer Tracts in Advanced Robotics, vol. 88. pp. 819–839.
Gupta, V., Hassibi, B., Murray, R., 2007. Optimal LQG control across packet-dropping links. Systems Control Lett. 56 (6), 439–446.
Haley Jr., P., Lermusiaux, P., Robinson, A., Leslie, W., Logoutov, O., Cossarini, G., Liang, X., Moreno, P., Ramp, S., Doyle, J., et al., 2009. Forecasting and reanalysis in the Monterey Bay/California current region for the Autonomous Ocean Sampling Network-II experiment. Deep Sea Res. Part II 56 (3), 127–148.
Hanna, S.R., Yang, R., 2001. Evaluations of mesoscale models' simulations of near-surface winds, temperature gradients, and mixing depths. J. Appl. Meteorol. 40 (6), 1095–1104.
Heidemann, J., Stojanovic, M., Zorzi, M., 2012. Underwater sensor networks: applications, advances and challenges. Phil. Trans. R. Soc. A 370 (1958), 158–175.
Huck, T., Vallis, G.K., 2001. Linear stability analysis of the three-dimensional thermally-driven ocean circulation: application to interdecadal oscillations. Tellus A 53 (4), 526–545.
Imer, O., Yüksel, S., Başar, T., 2006. Optimal control of LTI systems over unreliable communication links. Automatica 42 (9), 1429–1439.
Kinsey, J., Eustice, R., Whitcomb, L., 2006. A survey of underwater vehicle navigation: recent advances and new challenges. In: IFAC Conference of Manoeuvering and Control of Marine Craft, Lisbon, Portugal, Invited Paper.
Lall, S., Dullerud, G., 2001. An LMI solution to the robust synthesis problem for multi-rate sampled-data systems. Automatica 37 (12), 1909–1922.
Leith, C., 1974. Theoretical skill of Monte Carlo forecasts. Mon. Weather Rev. 102 (6), 409–418.
Ljung, L., 1999. System Identification: Theory for the User. Prentice Hall PTR.
Lolla, T., Ueckermann, M., Yigit, K., Haley, P., Lermusiaux, P., 2012. Path planning in time dependent flow fields using level set methods. In: IEEE International Conf. on Robotics and Automation. IEEE, pp. 166–173.
Moore, B.C., 1978. Singular value analysis of linear systems. In: IEEE Conf. Decision and Control, vol. 17, IEEE, pp. 66–73.
Moore, A.M., Arango, H.G., Di Lorenzo, E., Cornuelle, B.D., Miller, A.J., Neilson, D.J., 2004. A comprehensive ocean prediction and analysis system based on the tangent linear and adjoint of a regional ocean model. Ocean Model. 7 (1), 227–258.
Michini, M., Hsieh, M., Forgoston, E., Schwartz, I., 2014. Robotic tracking of coherent structures in flows. IEEE Trans. Robot. PP (99), 1–11. http://dx.doi.org/10.1109/TRO.2013.2295655.
Olascoaga, M.J., Haller, G., 2012. Forecasting sudden changes in environmental pollution patterns. Proc. Natl. Acad. Sci. 109 (13), 4738–4743.
Palmer, T.N., 2000. Predicting uncertainty in forecasts of weather and climate. Rep. Progr. Phys. 63 (2), 71.
Petillo, S., Schmidt, H., Balasuriya, A., 2012. Constructing a distributed AUV network for underwater plume-tracking operations. Int. J. Distrib. Sens. Netw.
Quevedo, D., Ostergaard, J., Nesic, D., 2011. Packetized predictive control of stochastic systems over bit-rate limited channels with packet loss. IEEE Trans. Automat. Control 56 (12), 2854–2868.
Ramp, S., Davis, R., Leonard, N., Shulman, I., Chao, Y., Robinson, A., Marsden, J., Lermusiaux, P., Fratantoni, D., Paduan, J., 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.
Reed, B., Hover, F., 2013. Tracking ocean fronts with multiple vehicles and mixed communication losses. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, pp. 3374–3381.
Reed, B., Leighton, J., Stojanovic, M., Hover, F., (2013). Multi-vehicle dynamic pursuit using underwater acoustics. In: International Symposium on Robotics Research, ISRR.
Rixen, M., Book, J.W., Carta, A., Grandi, V., Gualdesi, L., Stoner, R., Ranelli, P., Cavanna, A., Zanasca, P., Baldasserini, G., et al., 2009. Improved ocean prediction skill and reduced uncertainty in the coastal region from multi-model super-ensembles. J. Mar. Syst. 78, S282–S289.
Safonov, M., Athans, M., 1977. Gain and phase margin for multiloop LQG regulators. IEEE Trans. Automat. Control 22 (2), 173–179.
Sapsis, T.P., Lermusiaux, P.F., 2012. Dynamical criteria for the evolution of the stochastic dimensionality in flows with uncertainty. Physica D 241 (1), 60–76.
Schneider, T., Schmidt, H., 2010. Unified command and control for heterogeneous marine sensing networks. J. Field Robot. 27 (6), 876–889.

Scholin, C., Jensen, S., Roman, B., Massion, E., Marin, R., Preston, C., Greenfield, D., Jones, W., Wheeler, K., 2006. The Environmental Sample Processor (ESP)–an autonomous robotic device for detecting microorganisms remotely using molecular probe technology. In: OCEANS'06. MTS/IEEE, pp. 1–4.
Simon, D., 2006. Optimal State Estimation: Kalman, H infinity, and Nonlinear Approaches. Wiley-Interscience.
Sinopoli, B., Schenato, L., Franceschetti, M., Poolla, K., Jordan, M., Sastry, S., 2004. Kalman filtering with intermittent observations. IEEE Trans. Automat. Control 49 (9), 1453–1464.
Smith, R., Chao, Y., Li, P., Caron, D., Jones, B., Sukhatme, G., 2010. Planning and implementing trajectories for autonomous underwater vehicles to track evolving ocean processes based on predictions from a regional ocean model. Int. J. Robot. Res. 29 (12), 1475.
Stow, C.A., Jolliff, J., McGillicuddy Jr., D.J., Doney, S.C., Allen, J., Friedrichs, M.A., Rose, K.A., Wallhead, P., 2009. Skill assessment for coupled biological/physical models of marine systems. J. Mar. Syst. 76 (1), 4–15.
Thompson, C., Battisti, D., 2000. A linear stochastic dynamical model of ENSO. Part I: model development. J. Clim. 13 (15), 2818–2832.
Ueckermann, M.P., Lermusiaux, P.F.J., 2012. 2.29 Finite Volume MATLAB Framework Documentation. MSEAS Report-14, August 2012.
Van Overschee, P., De Moor, B., 1996. Subspace Identification for Linear Systems: Theory, Implementation, Applications. Kluwer Academic Publishers.
Wang, D., Lermusiaux, P.F., Haley, P.J., Eickstedt, D., Leslie, W.G., Schmidt, H., 2009. Acoustically focused adaptive sampling and on-board routing for marine rapid environmental assessment. J. Mar. Syst. 78, S393–S407.
Willcox, J., Bellingham, J., Zhang, Y., Baggeroer, A., 2001. Performance metrics for oceanographic surveys with autonomous underwater vehicles. IEEE J. Ocean. Eng. 26 (4), 711–725.
Zhang, Y., Godin, M., Bellingham, J., Ryan, J., 2012a. Using an autonomous underwater vehicle to track a coastal upwelling front. IEEE J. Ocean. Eng. 37, 338–347.
Zhang, H., Moura, J., Krogh, B., 2009. Dynamic field estimation using wireless sensor networks: tradeoffs between estimation error and communication cost. IEEE Trans. Signal Process. 57 (6), 2383–2395.
Zhang, Y., Ryan, J.P., Bellingham, J.G., Harvey, J.B., McEwen, R.S., 2012b. Autonomous detection and sampling of water types and fronts in a coastal upwelling system by an autonomous underwater vehicle. Limnol. Oceanogr. Methods 10, 934–951.