ELSEVIER

Volume 28 Number 19 15 November 2008 ISSN 0278-4343

CONTINENTAL SHELF RESEARCH

Research Papers

www.elsevier.com/locate/csr

Continental Shelf Research 28 (2008) 2644–2658

Contents lists available at ScienceDirect

Continental Shelf Research

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

Optimizing fixed observational assets in a coastal observatory

Sergey Frolov [a,b,*], António Baptista [a], Michael Wilkin [a]

[a] NSF Science and Technology Center for Coastal Margin Observation and Prediction, Oregon Health and Science University, Beaverton, Oregon 97006, USA
[b] Now with Monterey Bay Aquarium Research Institute, Moss Landing, California, USA

ARTICLE INFO

Article history:
Received 26 March 2008
Received in revised form
20 August 2008
Accepted 22 August 2008
Available online 4 October 2008

Keywords:
Columbia River
Estuary
Plume
Optimal observational array
Coastal observatory

ABSTRACT

Proliferation of coastal observatories necessitates an objective approach to managing of observational assets. In this article, we used our experience in the coastal observatory for the Columbia River estuary and plume to identify and address common problems in managing of fixed observational assets, such as salinity, temperature, and water level sensors attached to pilings and moorings. Specifically, we addressed the following problems: assessing the quality of an existing array, adding stations to an existing array, removing stations from an existing array, validating an array design, and targeting of an array toward data assimilation or monitoring.

Our analysis was based on a combination of methods from oceanographic and statistical literature, mainly on the statistical machinery of the best linear unbiased estimator. The key information required for our analysis was the covariance structure for a field of interest, which was computed from the output of assimilated and non-assimilated models of the Columbia River estuary and plume. The network optimization experiments in the Columbia River estuary and plume proved to be successful, largely withstanding the scrutiny of sensitivity and validation studies, and hence providing valuable insight into optimization and operation of the existing observational network. Our success in the Columbia River estuary and plume suggest that algorithms for optimal placement of sensors are reaching maturity and are likely to play a significant role in the design of emerging ocean observatories, such as the United State's ocean observation initiative (OOI) and integrated ocean observing system (IOOS) observatories, and smaller regional observatories.

1. Introduction

Integrated ocean observatories are emerging as the backbone for future scientific exploration and science-based management of coastal resources. As an example, the US government is embarking on contraction of two national ocean observatories: the ocean observation initiative (OOI) (National Research Council, 2003), with the goal of advancing the scientific understanding of the ocean, and the integrated ocean observing system (IOOS) (US Commission on Ocean Policy, 2004), with the goal of supporting scientifically informed management of the coastal ocean. Although the design-objectives of the two observatories are different, they both are expected to involve sensors deployed on fixed and mobile observational platforms and, hence, are expected to face common design challenges. One such challenge is an optimal co-placement of fixed sensors; that is, a placement that will maximize the amount of collected information, maximize the synergy with existing observational networks, and minimize the development and operational costs for an observatory.

A possible strategy for the optimal sensors placement can be found using methods from the theory of optimal experiment design (Fedorov, 1972; Silvey, 1980; Pukelsheim, 1993)—a subfield of formal statistics. In the past, many of these methods were applied to design observational arrays in oceanography (Bretherton et al., 1976; Hackert et al., 1998; Sakov and Oke, 2008) and meteorology (Berliner et al., 1999; Bishop et al., 2000). However, the applications to design of coastal observatories are still few (She et al., 2006). Recent advances in the circulation modeling of the coastal ocean show great promise in providing information necessary for extending the applications of optimal design methods to design of coastal observatories. In this article, we studied how outputs of a circulation model and a data assimilation system can be used to find an optimal placement of sensors in a coastal observatory for the Columbia River (CR) estuary and plume.

The coastal observatory for the CR estuary and plume provides a great example of a challenging dynamical environment, where an increasing number of observational and modeling assets are emerging, hence necessitating an objective approach for optimal

* Corresponding author. Tel.: +1 831 775 1960; fax: +1 831 775 1646.
E-mail address: frolovs@mbari.org (S. Frolov).

0278-4343/$ - see front matter © 2008 Elsevier Ltd. All rights reserved.
doi:10.1016/j.csr.2008.08.009

placement of new and existing sensors. The recent observational assets in the CR estuary and plume include several multi-annual observational arrays, such as the CORIE observational arrays for the estuary and plume (Baptista, 2006), the RISE observational arrays (Hickey, 2004) for the CR plume, and the array of tide gages and atmospheric buoy operated by National Atmospheric and Oceanic Administration. Complementary to these observational assets, several realistic models were recently developed for the CR estuary and plume (Baptista et al., 2005; Frolov, 2007; MacCready et al., 2007).

To understand how much of the information about the variability of the CR estuary and plume is already known from the existing observational and modeling assets, and how much remains to be learned by deploying new observational assets, we identified and addressed six specific questions regarding optimal placement of fixed observational assets in the CR estuary and plume:

(1) How informative is the existing observational networks about the variability of salinity, temperature, and water levels in the CR estuary and plume?
(2) How robust is the existing observational network in the CR estuary and plume to data dropouts?
(3) Which sensors in the existing CR network collect redundant information and, hence, can be removed from the network?
(4) Where in the CR estuary and plume should new sensors be placed?
(5) How the CR network can be optimized for monitoring or for data assimilation?
(6) How the proposed optimal designs of the observational network in the CR can be validated?

Many of these questions are representative of other observatories with fixed observational assets and, hence, can be extended to optimal placement of fixed stations in the national OOI and IOOS observatories.

We report the results of our findings in the following six sections of this paper. We first present the background on the observatory for the CR estuary and plume, which served as test-bed for this study (Section 2). We then review the methodology of optimal experiment design, specifically the statistical machinery of the best linear unbiased estimator (BLUE) (Ripley, 1987), the exchange-type optimization algorithm (Fedorov, 1994), and the array modes of Bennett (Bennett, 1985, 1992). We apply our optimal design algorithms in Section 4, where we assess the quality of the exiting observational array in the CR and suggest strategies for further improvement of the array. Finally, we verify the validity of the proposed array-design methods in Section 5. Conclusions follow in Section 6.

2. Observatory for the CR estuary and plume

The test-bed for our network optimization experiments was CORIE—an integrated ocean observatory for the CR estuary and plume. The CR estuary, a classic river-dominated estuary, is a highly energetic and dynamic system that responds quickly to changes in ocean tides, regulated river discharge, and coastal winds. The tidal dynamic in the CR is significant (tidal amplitudes of up to 3.6 m) and is driven by the nonlinear interaction of astronomic tides with complicated bathymetry, non-stationary river discharge, and coastal wind (Jay and Flinchem, 1997). Compressed and often stratified, the estuarine circulation in the CR is subject to extreme variations in salinity intrusion and stratification regimes (Jay and Smith, 1990; Baptista et al., 2005). The CR plume is a dominant hydrographic feature on the US west coast that plays an important role in the transport of dissolved and particulate matter for hundreds of kilometers along and across the continental shelf (Barnes et al., 1972; Grimes and Kingsford, 1996). The far field of the CR plume has a predominant orientation towards north in fall and winter, and towards southwest in spring and summer (Hickey et al., 1998). The near field of the CR plume has a shorter response time (hours to days) and can react quickly to changes in the local wind (Garcia-Berdeal et al., 2002; Baptista et al., 2005; Hickey et al., 2005).

The CORIE observatory studies the dynamics of the CR estuary using a real-time observation network (Baptista, 2006) and a modeling system for 3D baroclinic circulation of the CR estuary

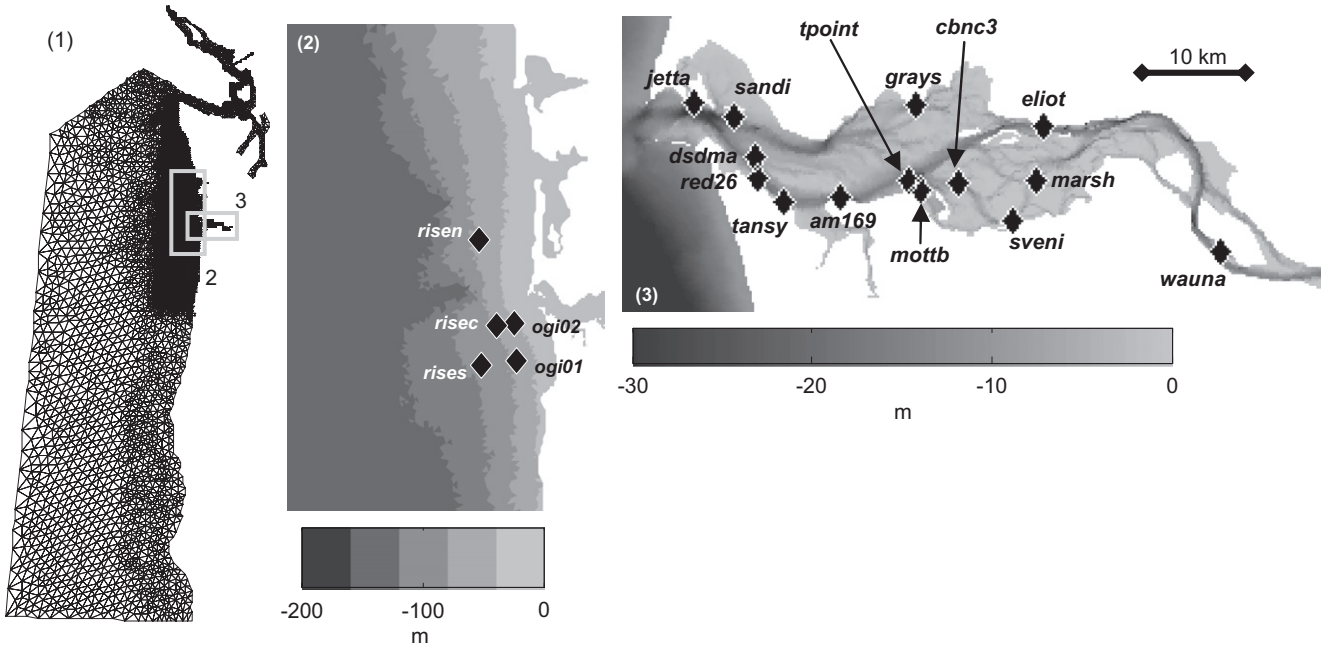

Fig. 1. (1) The computational grid. (2 and 3) The map of observation stations in the CR plume (2) and estuary (3). The map of stations in panels 2 and 3 is overlaid over the bathymetry. Stations *risen*, *risec*, and *rises* were operated by the RISE project (www.ocean.washington.edu/rise/). Station *tpoin* and *wauna* were operated by National Oceanographic and Atmospheric Agency.

and the adjacent ocean (Baptista et al., 2005). CORIE data and modeling products (Baptista et al., 2005; Baptista, 2006) are used to support the research and development of novel modeling techniques (Zhang et al., 2004; Zhang and Baptista, 2008), fisheries research (Bottom et al., 2005; Burla et al., 2007; Burla et al., accepted), and the CR ecosystem management (USACE, 2001). At the time of our experiments, CORIE observatory had three major components:

(1) A real-time observation system that included a network of fixed stations and research cruises. Fixed stations utilized in our experiments in Sections 4 and 5 are shown in Fig. 1. At each station, an array of sensors measured a variable combination of parameters, including water level, salinity, temperature, velocity, wind temperature, and wind velocity. Stations *risen*, *risec*, and *rises* in Fig. 1 were operated by the RISE project (Hickey, 2004), and stations *tpoint* and *wauna* in Fig. 1 were operated by National Atmospheric and Oceanic Administration.
(2) A modeling system for simulating 3D baroclinic circulation of the CR estuary, plume, and adjacent coastal ocean. The model simulations used in our experiments were generated using an Eulerian–Lagrangian model SELFE (Zhang and Baptista, 2008) and were forced with a combination of realistic atmospheric, ocean, and river forcings.
(3) A data assimilation system (Frolov, 2007), which was based on a recently developed, reduced-dimension Kalman filter (Frolov, 2007; Lu et al., 2007)—a fast, nonlinear extension to the classical Kalman filter.

Results of recent hindcast simulations (Baptista et al., 2005; Baptista, 2006) showed good predictive skill in both the CR estuary and the plume, even under such challenging conditions as high river discharge and fast-changing winds. Specifically, water levels and salinities were represented robustly in the estuary across spatial (from channels to tidal flats) and time (from tidal to inter-annual) scales. In the plume, models showed reliable skill in representing fronts and direction of the plume (as compared to remote sensing and field data). Recent work (Frolov, 2007) on application of data assimilation to the CR estuary and plume showed that data assimilation was able to further improve simulated water levels, salinity, and temperature in the CR estuary and plume.

3. Optimal experiment design methods

In the existing oceanographic and meteorological literature, the methods of optimal experiment design fall into two large categories: (1) methods based on a framework of statistical experiment design (Fedorov, 1972; Silvey, 1980; Pukelsheim, 1993) and (2) methods based on a framework of adjoint sensitivity fields (Langland and Rohaly, 1996; Baker and Daley, 2000). The seminal article by Berliner (Berliner et al., 1999) reconciled the two frameworks and identified some theoretical limitations to using the adjoint sensitivity fields for guiding placement of adaptive measurements. From the two methodological frameworks, we used the framework of statistical experiment design to find optimal placement of sensors in the CR estuary and plume, since the statistical framework is closely aligned with the framework of Kalman filter, the data assimilation scheme of choice in the CORIE observatory. An additional consideration for choosing the statistical framework was the ease with which it can be applied to other computational domains and variables, and the independence of the framework from many details of a forward model and a data assimilation system.

To implement the framework of statistical experiment design in the CR estuary and plume, we used the following methodological choices and approximations:

- To evaluate how the location of a potential measurement contributed to reducing the uncertainty about the field of interest (e.g. salinity or errors in salinity), we employed the statistical machinery of the BLUE estimator (Ripley, 1987).
- To characterize the covariance structure for the field of interest, required by the BLUE estimator, we used a realistic, stationary, low-rank approximation to the covariance of state and error fields. The state covariance was computed from the output of long hindcast simulations of the CR estuary and plume (Baptista et al., 2005; Baptista, 2006; Frolov, 2007). The error covariance was approximated using the temporal average of error covariance computed by the data assimilation system for CR estuary and plume (Frolov, 2007).
- To compare alternative designs of observational arrays, we used two different cost functions: the trace (mean square error (MSE)) and the determinant (volume) of the posterior covariance (Fedorov, 1994).
- To find optimal array configuration, we used several modified versions of the exchange-type optimization algorithm (Fedorov, 1994).
- To quantify the statistical redundancy of the existing observational array, we analyzed the eigen-spectrum of the observation covariance matrix, using analysis similar to the array modes of (Bennett, 1985, 1992).

3.1. Best linear unbiased estimator

To evaluate how the location of a potential measurement contributed to reducing the uncertainty about the field of interest (e.g. salinity or errors in simulated salinity), we employed the statistical machinery of the BLUE estimator (Ripley, 1987). The BLUE estimator—also known as the Gauss–Markov theorem, kriging, and the objective analysis method—was introduced as early as (Gandin, 1963) in meteorology and (Bretherton et al., 1976) in oceanography. In the following description of the BLUE estimator, we adopted the exposition of the method from Fedorov (1994).

Consider a discretized random field $x = x(\xi,k)$ where subscript ξ is the spatial index on the computational grid and k is the index of the discrete time. An example of such field is the field of simulated salinities or errors in simulated salinity. The statistics of the field x is approximated using a stationary mean $\bar{x}(\xi) = E[x(\xi,k)]$ and a stationary covariance matrix $\mathbf{C}(\xi,\xi') = \mathrm{cov}(x(\xi), x(\xi'))$. In geosciences, we are often interested in predicting the field values

$$x_p(k) = \mathbf{H}_p x(\xi,k) \tag{1}$$

at spatial locations $\xi_p \in \xi$, given sparse, noisy observations of the field

$$y(k) = \mathbf{H}_y x(\xi,k) + \varepsilon \tag{2}$$

where $\mathbf{H_p}$ and $\mathbf{H_y}$ are selection or interpolation operators for prediction and observation points, and $\varepsilon \sim N(0,\sigma^2\mathbf{I})$ is the Gaussian observation noise. For example, in the CR estuary, we are interested in predicting the values of salinity inside of the estuary $x_p(k)$, given sparse *in situ* measurements of salinity $y(k)$.

The solution to this prediction (interpolation) problem in Eq. (1) is the well-known BLUE estimator, which is defined as

$$x_p(k) = \bar{x}_p + \mathbf{C}_{py}\mathbf{C}_{yy}^{-1}(y(k) - \mathbf{H}_y\bar{x}) \tag{3}$$

$$\mathbf{D}_{pp|y} = \mathbf{C}_{pp} - \mathbf{C}_{py}\mathbf{C}_{yy}^{-1}\mathbf{C}_{yp} \tag{4}$$

where $\bar{x}$ is the mean state, $\mathbf{C}_{pp}$ is the prior covariance, and $\mathbf{D}_{pp|y}$ is the posterior covariance conditioned on observations y. The cross-covariance $\mathbf{C}_{py}$ and the observations covariance $\mathbf{C}_{yy}$ in (Eqs. (3) and (4)) are defined as

$$\mathbf{C}_{py} = \mathbf{H}_p \mathbf{C} \mathbf{H}_y^T$$
$$\mathbf{C}_{yy} = (\mathbf{H}_y \mathbf{C} \mathbf{H}_y^T + \sigma^2 \mathbf{I}) \quad (5)$$

The diagonal elements of the prior and posterior covariance matrices $\mathbf{C}_{pp}(\xi,\xi)$ and $\mathbf{D}_{pp|y}(\xi,\xi)$ in Eq. (4) are the prior and the posterior variance of the field x_p, which characterize our uncertainty about the field values before and after the measurements are taken, e.g. the lower variance indicates a higher certainty about the field. The off-diagonal elements of the covariance matrices $\mathbf{C}_{pp}(\xi,\xi')$ and $\mathbf{D}_{pp|y}(\xi,\xi')$ indicate how strongly the field values are correlated for any two spatial locations ξ and ξ'. For example, the stronger the cross correlations, the fewer measurements will be required to achieve small posterior uncertainty $\mathbf{D}_{pp|y}$.

Eq. (4) of the BLUE estimator shows that additional information from observations reduces the posterior uncertainty $\mathbf{D}_{pp|y}$, as compared to the initial (prior) uncertainty $\mathbf{C}_{pp}$. Eqs. (4) and (5) also show that for linear selection operators $\mathbf{H_y}$ and $\mathbf{H_p}$, the posterior uncertainty $\mathbf{D}_{pp|y}$ only depends on the locations of observation and prediction points and does not depend on the value of the field at these locations.

The BLUE algorithm in Eqs. (3)–(5) can be easily extended to prediction of parameters that are integrals or nonlinear functions of the fields $x(\xi,k)$, such as the volume of the CR plume or the maximum length of the daily salinity intrusion. This extended algorithm uses the same Eqs. (3)–(5), but defined for the augmented vectors $x^a(k)$

$$x^a(k) = \begin{bmatrix} h_1(x(\xi,k))/\omega_1 \\ \vdots \\ h_i(x(\xi,k))/\omega_i \\ x(\xi,k) \end{bmatrix}$$

where h_1–h_i are i nonlinear functions of the field $x(\xi,k)$ that one wants to predict, and ω_1–ω_i are the normalization weights that normalize each of the functions h_1–h_i to have unit variance. The field values $x(\xi,k)$ remain non-normalized, as they are not part of the prediction vector x_p. In the case when it is desirable to normalize field values $x(\xi,k)$, e.g. to improve conditioning of the covariance or when $x(\xi,k)$ is a part of the prediction vector x_p, one should also not forget to normalize the observational error variances σ^2 in Eq. (5).

3.2. Optimization criteria

In optimal experiment design, we are interested in selecting, based on some criterion of optimality, the best observation strategy from a set of all possible strategies. One such criterion of optimality is the size of the posterior covariance $\mathbf{D}_{pp|y}$ (Eq. (4)), i.e., the smaller posterior covariance indicates a higher certainty about the field. However, the set of covariance matrices does not posses a natural order; hence, a cost function $J(\mathbf{D}_{pp|y})$ needs to be defined in order to map a set of posterior covariance matrices to the ordered set of real numbers. In our experiments, we used two such cost functions: the minimum MSE and the minimum determinant (DET) of the posterior covariance $\mathbf{D}_{pp|y}$.

The MSE cost function, for a fixed number of predictions points x_p, can be measured with the trace of the posterior covariance $\mathbf{D}_{pp|y}$

$$J(\mathbf{D}_{pp|y}) = \text{trace}(\mathbf{D}_{pp|y}) = \sum \lambda_i \quad (6)$$

where trace(·) is the trace of a matrix, and λ_i is the ith eigen-value of the covariance matrix $\mathbf{D}_{pp|y}$. The trace of the matrix $\mathbf{D}_{pp|y}$ is defined as a summation of the diagonal elements of the matrix $\mathbf{D}_{pp|y}$, which is equivalent to the sum of the posterior variances (posterior uncertainties) at all prediction points x_p. Eq. (6) shows that minimizing for the trace (MSE) criterion is the same as minimizing for the average eigen-value of the matrix $\mathbf{D}_{pp|y}$.

The DET cost function measures the determinant det(·) of the posterior covariance

$$J(\mathbf{D}_{pp|y}) = \det(\mathbf{D}_{pp|y}) = \prod \lambda_i \quad (7)$$

where λ_i is the ith eigen-value of the covariance matrix $\mathbf{D}_{pp|y}$. The DET criterion is equivalent to the product of the eigen-values, hence minimizing for the DET criterion is equivalent to minimizing for the smallest volume of the covariance $\mathbf{D}_{pp|y}$.

The advantage of using the MSE over the DET criterion is that MSE is the same criterion as used by the Kalman filter and the variational data assimilation (Bennett, 2002). The drawback of the MSE criterion is the higher computational cost of computing Eq. (6) than Eq. (7), since it is possible to evaluate Eq. (7) in the usually smaller space of observations y rather than the space of prediction points x_p.

3.3. Optimization algorithm

To find the minimum of the cost function in Eqs. (6) and (7), we defined the following optimization problem:

$$\mathbf{Y}_o = \underset{\mathbf{Y}_o \in \mathbf{Y}_f}{\text{argmin}} \left[J(\mathbf{D}_{pp|\mathbf{Y}_o}) \right] \quad (8)$$

where $\mathbf{Y}_f = \{y_1,\ldots,y_n\}$ is the set of all possible observations and $\mathbf{Y}_o \in \mathbf{Y}_f$ is the set of selected optimal observations that minimizes cost function $J(\mathbf{D}_{pp|y})$. To search for the minimum of the optimization problem in Eq. (8), we used several variants of the exchange-type algorithm. The basic exchange-type algorithm, as described in (Fedorov, 1994), iterates over the following two steps:

1. *Delete step*: from the set $\mathbf{Y}_o$ of candidate stations delete observation y_o^- that contributes the least to minimizing the optimization criteria $J(\mathbf{D}_{pp|y})$

$$y_o^- = \underset{y_o^- \in \mathbf{Y}_o}{\text{argmin}} \left[J(\mathbf{D}_{pp|\mathbf{Y}_o - y_o^-}) \right] \quad (9)$$

2. *Add step*: from the set of all feasible observations $\mathbf{Y}_f$ add observation y_l^+ that contributes the most to minimizing the optimization criteria $J(\mathbf{D}_{pp|y})$

$$y_o^+ = \underset{y_o^+ \in \{\mathbf{Y}_f - \mathbf{Y}_o\}}{\text{argmin}} \left[J(\mathbf{D}_{pp|\mathbf{Y}_o - y_o^- + y_o^+}) \right] \quad (10)$$

Iteration over the delete and add steps continues until $J(\mathbf{D}_{pp|y})$ decreases. After each successful iteration, the active set of optimal observations is updated as following:

$$\mathbf{Y}_o = \mathbf{Y}_o - y_o^- + y_o^+$$

We implemented the delete and the add steps in Eqs. (9) and (10) differently for the MSE and the DET criteria. For the MSE criteria, we implemented the optimization algorithm as a direct search through the set of all active observations $\mathbf{Y}_o$ (the delete step) and through the set of all feasible observations $\mathbf{Y}_f$–$\mathbf{Y}_o$ (the add step). For the DET criteria, we used a computationally more efficient algorithm (Fedorov, 1994) that operates in a smaller space of observations y. Specifically, Fedorov (Fedorov, 1994) showed that the minimum of det($\mathbf{D}_{pp|y}$) is equivalent to the maximum of det($\mathbf{C}_{yy}$). That is, the minimum volume of $\mathbf{D}_{pp|y}$ (the posterior uncertainty) is achieved at the maximum volume of

$\mathbf{C}_{yy}$ (the observation covariance). Computing $\det(\mathbf{C}_{yy})$ is usually computationally more efficient than computing $\det(\mathbf{D}_{pp|y})$.

The modified algorithm for the DET cost function uses the following delete and add steps to find a minimum of $J = \det(\mathbf{D}_{pp|y})$:

$$\text{Delete} : y_o^- = \underset{y_o^- \in \mathbf{Y}_o}{arg\max} \left[c_{yy}^{-1} \right] \tag{11}$$

$$\text{Add} : y_o^+ = \underset{y_o^+ \in (\mathbf{Y}_f - \mathbf{Y}_o)}{arg\min} \left[c_{yy}^{-1} \right] \tag{12}$$

where ${c_{yy}}^{-1}$ are the diagonal elements of matrix ${\mathbf{C}_{yy}}^{-1}$, and ${y_o}^-$ and ${y_o}^+$ are the observations that correspond to the largest (${y_o}^-$) and the smallest (${y_o}^+$) value in the vector ${c_{yy}}^{-1}$.

Although the exchange-type algorithm (Eqs. (9) and (10)) is guaranteed to converge to a fixed point, this fixed point is not guaranteed to be the global minimum of the optimization problem (Eq. (8)) (Fedorov, 1994). To avoid convergence to a suboptimal local minimum, we restarted the optimization algorithm using multiple random configurations of the initial observation network. Alternatively, we often reduced the exchange-type algorithm to the pure add- or delete-algorithm, both of which have a unique, but suboptimal minimum. In Section 5.1, we compare the optimization results using all three modifications of the exchange algorithm.

3.4. Choices of the prior covariance

To define our prior knowledge about the field of interest, we used two distinct choices of the prior covariance ($\mathbf{C}$): the state and the error covariance. The state covariance characterizes the uncertainty about the current state of the system with respect to the mean state. The error covariance characterizes the uncertainty about the accuracy of the model prediction. The observation locations optimized using the state covariance are optimal for reconstructing the state of the system using model-derived EOFs, that is, for monitoring of the system state. The observation locations that are optimized using the error covariance are optimal for reconstructing the error field of the model, that is, for improving the model simulation through a data assimilation procedure.

To characterize the covariance structure for the state variables, such as salinity and temperature in the CR estuary, we used an empirical covariance

$$\mathbf{C}^{state} = E[(x - \bar{x})(x - \bar{x})^T] \tag{13}$$

where the samples of states x were approximated by model states extracted from a long hindcast database. The primary choice for computing the state covariance matrix $\mathbf{C}^{state}$ was the CORIE hindcast database DB14 that included both the CR estuary and plume and that did not use data assimilation. For sensitivity experiments in Section 5.1, we computed the state covariance based on two additional datasets: the non-assimilated DB16 and the assimilated DB14. The design of DB16 was similar to DB14; however, the numerical grid of DB16 had a higher resolution and was restricted to the CR estuary. The data assimilation method was described in a separate publication (Frolov, 2007).

To characterize the covariance structure of model errors, such as errors in simulated salinity and temperature, we used a scaled time-average of the forecast error covariance matrices

$$\mathbf{C}^{error} = \alpha \frac{1}{n} \sum_{k=1}^{n} \mathbf{P}_{xx}^{f}(k) \tag{14}$$

where $\mathbf{P}_{xx}^{f}(k)$ is a forecast error covariance computed by the data assimilation system described in (Frolov, 2007), n is the number of averaged time steps, and α is the inflation factor that accounts for a possible inconsistency between the predicted and the true size of the forecast error covariance. The inflation factor α, which establishes a correct balance between the model error and the observation error in the Eq. (5), was computed using the following formula:

$$\alpha = [\text{trace}(\mathbf{C}_{yy}^{observed}) - \text{trace}(\mathbf{R})]/\text{trace}(\mathbf{C}_{yy}^{predicted})$$

$$\mathbf{C}_{yy}^{predicted} = \mathbf{H}_y \left[\frac{1}{n} \sum_{k=1}^{n} \mathbf{P}_{xx}^{f}(k) \right] \mathbf{H}_y^T \tag{15}$$

where ${\mathbf{C}_{yy}}^{observed}$ and ${\mathbf{C}_{yy}}^{predicted}$ are the observed and the predicted error covariance at the locations of existing observation stations, and $\mathbf{R} = \sigma^2\mathbf{I}$ is the covariance of observation errors.

Since the size of the state and error covariance (Eqs. (13) and (14)) can be quite large ($\sim 10^6 \times 10^6$), we operated with a factorized form of these covariance matrixes

$$\mathbf{C}_{pp} = \mathbf{H_p} \mathbf{C} \mathbf{H_p}^T = \mathbf{H_p} \mathbf{L} \mathbf{S} \mathbf{L}^T \mathbf{H_p}^T \tag{16}$$

where $\mathbf{H_p}$ is the selection operator for the prediction points, $\mathbf{L}$ is the matrix of leading eigen-vectors for a covariance matrix $\mathbf{C}$, and $\mathbf{S}$ is the diagonal matrix of the eigen-values. The eigen-spectrum $\mathbf{S}$ was truncated to improve computational performance. Typically, we retained more than ~95% of the variance, with the exact truncation value determined through a series of sensitivity studies (not presented here). The efficient implementation of the BLUE formulas (Eq. (3)) with the factorized covariance matrix (Eq. (16)) is detailed in Appendix 1.

3.5. Advantages and limitations of the proposed method

There are several advantages to using the optimal experiment design method described in Sections 3.1–3.4. The method is generic, e.g. we used it for placing salinity sensors in the CR estuary and plume, but similar methods were also used for finding optimal sensor-locations in many open ocean (Bretherton et al., 1976; Hackert et al., 1998; Sakov and Oke, 2008) and meteorological (Berliner et al., 1999; Bishop et al., 2000) applications. The method is closely related to the theory of data assimilation, since the MSE cost function (Eq. (6)) that we used in our experiments is the same cost function as used by most sequential and variational data assimilation algorithms (Bennett, 2002). The method does not require development of an adjoint model. In fact, it can be used without any model input in cases where high-resolution observational data are available (Sakov and Oke, 2008).

The two main limitations of the method stem from the assumption of stationary statistics and from the choice of the optimization criteria. It is likely that the stationary (averaged in time) covariance matrix is well-suited for optimizing fixed observational assets, since averaging in time is likely to provide an adequate proxy for large-scale correlations of persistent features that can be monitored well with a fixed observational array. However, it is also likely that the stationary covariance matrix may not suffice for adaptive sampling of a fast-moving feature like a plume front. In the case of adaptive measurements, the fine-scale correlations associated with fronts are likely to be non-stationary and will require an accurate prediction of their structure, for example using a data assimilation system. For more references on adaptive measurements see (Berliner et al., 1999; Daescu and Carmichael, 2003; Ogren et al., 2004; Bertozzi et al., 2005; Leonard et al., 2006).

The second limitation is rooted in the optimization criteria. Both the MSE and the DET criteria optimize measurements for state estimation; that is, for reconstruction of the field given sparse measurements. However, the observations can be targeted

for applications other than the state estimation, for example for detecting extreme events or for tracking the location of a maximum concentration. In these cases, observations optimized using the MSE and the DET criteria may not be optimal, and using alternative optimization criteria can be prudent (Fedorov, 1994; Berliner et al., 1999; Chang and Tseng, 1999).

4. Design of an optimal observational array in the CR estuary and plume

To design an optimal array for the CR estuary and plume, we applied the optimal experiment design methods from Section 3 to quantify how informative is the current observational array and to identify the strategies further improvement of the array. The evaluation of the existing array (discussed in Section 4.1) suggested two such strategies: (1) removing redundant stations (discussed in Section 4.2) and (2) adding new salinity sensors in the estuary and plume (discussed in Section 4.3). In most of our experiments, we compared the array designs based on two choices of the covariance matrix: the state and the error covariance, hence comparing sensor placements optimized for monitoring of the CR estuary and plume or optimized for improved data assimilation. The suggested array designs were validated using data assimilation experiments, presented in Section 5.

4.1. Evaluating the existing CORIE array

4.1.1. How representative is existing array of the state and error variability?

To determine whether the existing array was representative of state and error variability in the CR estuary, we used the R^2 statistics computed using the following formula:

$$R^2 = 100 \times \left(1 - \frac{\mathrm{trace}(\mathbf{D}_{pp|y})}{\mathrm{trace}(\mathbf{C}_{pp})}\right) \tag{17}$$

where $\mathbf{C}_{pp}$ and $\mathbf{D}_{pp|y}$ are the prior and the posterior covariance matrixes, and trace$(\cdot)$ is the trace of the matrix. The R^2 statistics measures the percentage of the prior variance (uncertainty), explained by the observations. For example, an R^2 of 0% indicates uninformative measurements that cannot reduce the prior uncertainty about the field of interest. An R^2 of 100% indicates highly informative measurements that can be used to reconstruct the field of interest in its entirety, short of inaccuracies in the prior covariance $\mathbf{C}$.

The computed R^2 values for the CORIE array are presented in Fig. 2. The R^2 values for the state (panel 1) and the error (panel 2) covariances are plotted against the number of sensors retained in the array, with sensors ordered from the most important (first sensor) to the least important (last sensor). The order of sensors was determined using the sensor-reduction sequences discussed in Section 4.2. Fig. 2 shows that the salinity, temperature, and water level arrays were highly informative of the state variability (R^2 of $>90\%$), but less informative of the error variability in the estuary (as low as 62% for salinity errors). Fig. 2 suggests that a possible strategy for extending of the CORIE array in the estuary is to invest in additional salinity sensors, since the current array can explain only ~62% of the salinity error variance, but more than 80% of the temperature and water level error variance.

4.1.2. How redundant is the existing array?

To determine whether the existing salinity array in the CR estuary provided redundant information, we analyzed the eigen-spectrum of the observation covariance matrix $\mathbf{C}_{yy}$ (Eq. (5)), which was computed using actual observations for the state covariance or model-data misfits for the error covariance. Our eigen-spectrum analysis was similar to the array modes of Bennett (Bennett, 1985, 1992) and relied on the analysis of the truncation error e_i:

$$e_i = 1 - \sum_{k=1}^{i} \lambda_k \Big/ \sum_{k=1}^{k_{\max}} \lambda_k \tag{18}$$

where λ_k is the kth eigen-value of the observation covariance matrix $\mathbf{C}_{yy}$, k_{max} is the number of all eigen-values, and i is the number of retained eigen-values.

The results of our experiments are presented in Fig. 3 showing truncation error e_i (Eq. (18)) for salinity (panel a), temperature (panel b), and water level (panel c) arrays plotted against the number of retained eigen-modes λ_k. From our experiments, we found that the existing arrays were slightly redundant. For example, from a total of 15 salinity sensors, 99% of the variance was captured by 8 sensors for the state covariance and by 10 sensors for the error covariance. The comparison of the truncation errors between the state and the error covariance in Fig. 3 shows that more sensors were required for data assimilation than for monitoring of the CR estuary, suggesting that the error covariance had shorter correlations scales than the state covariance.

4.1.3. Can the redundancy of the existing array offset gaps in data coverage?

The analyses presented in Figs. 2 and 3 assumed that all sensors in the existing CORIE network were active and reporting data. However, our experience with the operation of the CORIE network is to the contrary, showing that some of the sensors in the network are not reporting data at any given time, due to a combination of biofouling, routine maintenance, and dropouts in telemetry. To study the robustness of the existing CORIE array to gaps in data coverage, we show in Fig. 4 how the changing

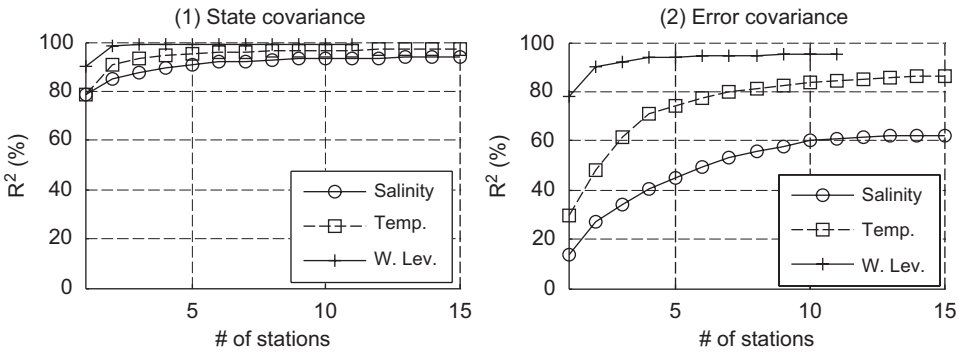

Fig. 2. The R^2 values for the state (1) and the error (2) covariance as a function of the number of sensors retained in the existing observation array.

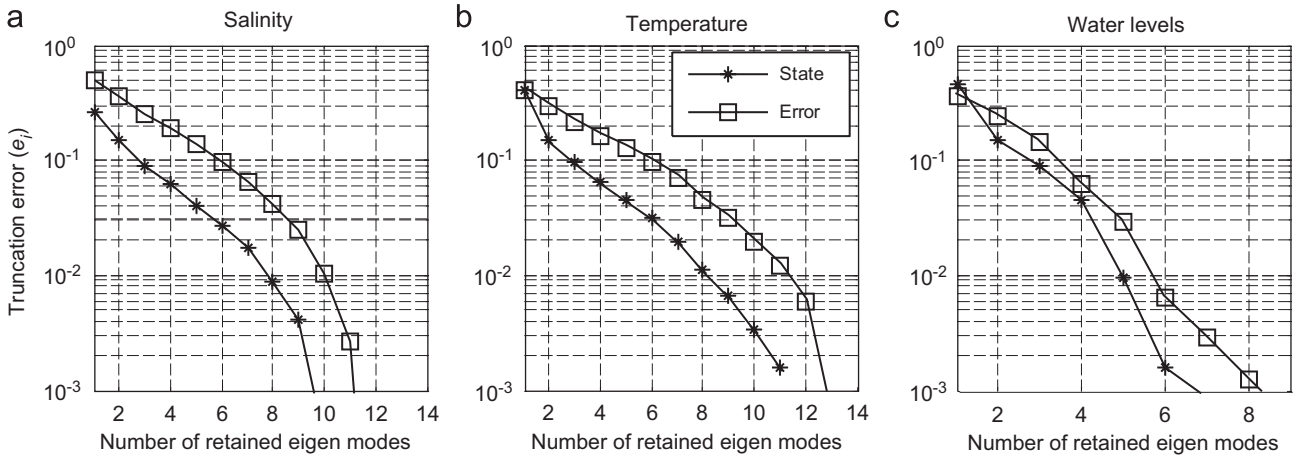

Fig. 3. Eigen-spectrum of the observation covariance C_{yy} for the existing array of salinity (a), temperature (b), and water level (c) sensors in the CR estuary. The truncation errors e_i (Eq. (18)) are plotted against the number of retained eigen-modes. The observation covariance for the state (line marked with +) was computed using actual measurements and for the error (line marked with *squares*) was computed using errors for the DB14 prediction.

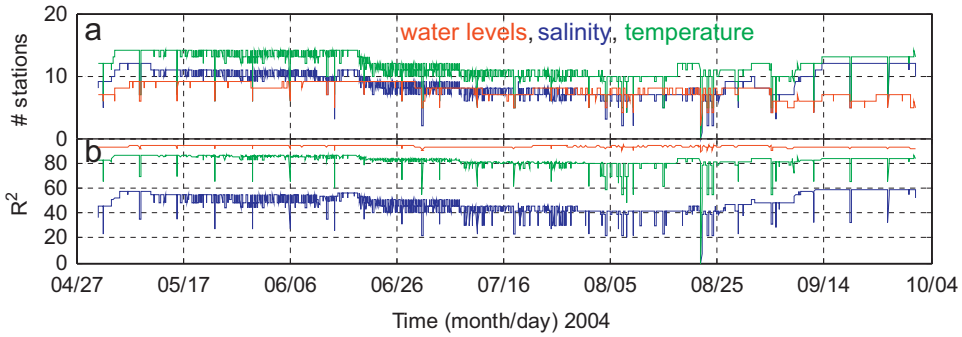

Fig. 4. (a) Number of available water level (red), salinity (blue), and temperature (green) sensors. (b) Time series of the R^2 values for the same three categories of sensors.

number of available sensors in the CORIE array (panel a) affected the ability of the array to measure the statistics of the salinity, water level, and temperature errors in the estuary (shown using R^2 statistics in panel b). The results in Fig. 4 show that the water level and the temperature arrays were redundant enough to compensate for gaps in observational data. For example, the R^2 for the water level array dropped only slightly (from 93% to 90%) as the number of water level sensors dropped from 9 to 5. In contrast, the salinity array was not redundant enough to compensate for gaps in observational data. For example, as the number of sensors dropped from 12 to 6, the R^2 dropped from 57% to 30%. It is likely that this increased sensitivity of the salinity array was due to a preferential biofouling of salinity sensors in the lower CR estuary. For example, in July–August of 2004 only 2 out of 10 salinity sensors located in the lower CR estuary remained active and were not biofouled.

4.2. *Optimizing operations of the CORIE array by removing redundant sensors from the network*

The analysis of the existing array, presenting in Section 4.1, showed that the existing CORIE array was redundant, yet this redundancy was not sufficient to offset the gaps in observational data, suggesting an opportunity for the improvement of the CORIE array. This opportunity is to remove redundant stations in the CORIE array, hence freeing the limited resources of the field crew, and allowing more frequent maintenance of the remaining salinity sensors in the estuary. More frequent maintenance is likely to prevent biofouling of sensors—a key contributor to the reduced coverage of the salinity array in the summer of 2004.

To determine which salinity sensors in the CORIE array provided redundant information, we applied the deleting procedure (Eq. (9)), which sequentially removed redundant sensors from the salinity array in the CR estuary. Table 1 presents the order in which the salinity sensors were removed and the corresponding R^2 values for the truncated arrays. In general, the suggested order in Table 1 was logical. For example, sensors that were removed first, such as *marsh*, *eliot*, and *cbnc*3, were located upriver and saw salt only occasionally. In contrast, sensors that were removed last, such as *sandi*, *dsdma*, and *jetta*, were located in the mouth of the CR estuary and saw large variation in salinity at each tidal cycle.

Analysis of R^2 values for the truncated arrays in Table 1 suggests that many sensors in the array were equally important. For example, it took removing four sensors in the error-based array (Table 1 columns IV and V) to change the value of R^2 from 62% to 61%, which suggests that these four sensors had equally small effect on the value of R^2, and hence, the order in which these sensors were removed can be easily altered. To determine whether there was a single most important or least important salinity sensor in the CR estuary, we conducted two additional experiments described below.

To determine which single salinity sensor was most important, we computed R^2 (Eq. (17)) values for arrays that had only one active salinity sensor in the CR estuary. Based on our experiments,

Table 1
Suggested order in which salinity sensors may be removed from the existing CR array.

Order	State		Error	
	Salinity sensor	R^2(%)	Salinity sensor	R^2(%)
I	II	II	IV	V
	all 15 sensors	94	all 15 sensors	62
1	eliot	94	marsh	62
2	marsh	94	eliot	62
3	cbnc3	94	cbnc3	62
4	mottb	93	red26 (mid.)	61
5	tansy	93	am169 (bot.)	60
6	am169 (mid.)	93	tansy	58
7	red26 (mid.)	93	am169 (mid.)	56
8	dsdma	92	mottb	53
9	am169 (top)	92	red26 (bot.)	49
10	sandi	91	sandi	45
11	red26 (bot.)	89	red26 (top)	40
12	am169 (bot.)	87	grays	34
13	grays	85	jetta	27
14	jetta	79	dsdma	14
15	red26 (top)	0	am169 (top)	0

R^2 is computed after the sensor is removed from the network.

Table 2
Results of the most important and the least important salinity sensor experiment.

Salinity sensor	Most important (%)		Least important (%)	
	State	Error	State	Error
I	II	III	IV	V
jetta	74.1	12.0	−0.7	−3.9
sandi	73.2	12.1	−1.1	−4.6
red26 (top)	78.9	11.2	−0.6	−2.2
red26 (mid)	74.4	8.7	−0.2	−0.5
red26 (bot)	72.0	7.8	−0.3	−0.7
dsdma	77.5	13.5	−0.6	−4.0
tansy	74.4	9.4	−0.2	−2.1
am169 (top)	64.3	13.7	−0.4	−3.6
am169 (mid)	58.5	8.9	−0.2	−0.7
am169 (bot)	51.3	7.8	−0.2	−0.7
mottb	15.9	6.9	−0.1	−1.7
cbnc3	25.5	8.3	−0.1	−0.4
grays	27.4	9.1	−0.5	−1.6
eliot	0.0	0.2	0.0	0.0
marsh	0.0	0.1	0.0	0.0

shown in Table 2 columns II and III, we could not locate a single most important salinity sensor. Instead, we found that most salinity sensors were equally important in the CR estuary, since each had similar value of R^2. For example, there were at least seven sensors, each of which was able to explain between 70% and 80% of the state variability. We also found that a single sensor can explain a much higher percentage of the state variance (as much as 79%) than the error variance (no more than 14%).

To determine which single salinity sensor was least important, we computed the R^2 (Eq. (17)) values for arrays with one of the salinity sensors in Table 2 removed. The results of the experiment, columns IV and V, suggest that, while the impact of an individual sensor varied, removing a single sensor had very little impact on the R^2 of the remaining sensors. For example, removing a single salinity sensor led to a decrease in R^2 of no more than 1.1% for the state and 4.6% for the error variance.

4.3. Extending the existing CORIE array with new sensors

To find optimal placement of new salinity sensors in the CR estuary and plume, we reduced the exchange-type algorithm to several sequential add-steps (Eq. (10)). We considered two realistic scenarios: adding one salinity sensor in the CR estuary and re-organizing the existing salinity sensors in the CR plume. In the estuary, our optimization procedure placed the new salinity sensor in the North Channel, leading to an improved R^2 for the salinity errors (from 62% to 69%). In the plume, the optimization procedure placed two sensors north and one station south of the CR mouth. The three optimally placed salinity sensor in the CR plume were more informative (R^2 of 23% for salinity errors) than the five historical sensors; the R^2 value for the historical plume sensors *risen*, *risec*, *rises*, *ogi*01, and *ogi*02, shown in Fig. 1, was 18%. The comparisons of the arrays optimized using the error and the state covariance showed that both arrays had similar configuration in the estuary and plume. The detailed description of the logistical constraints and the experimental results follow.

To characterize some of the logistical constraints that are likely to arise in the deployment of new salinity sensors in the CR estuary and plume, we added the following constraints to the optimization problem (Eq. (8)). We constrained the locations of feasible salinity sensors to areas deeper than 10 m in the plume and deeper than 4 m in the estuary. We constrained the domain of interest, defined by the operator $\mathbf{H}_p$ in Eq. (1), to the CR plume for placing the plume sensors, and to the CR estuary for placing the estuarine sensors. We further constrained our search by deploying candidate sensors at the depth of the highest variability at each horizontal location, hence reducing the three-dimensional (3D) optimization problem to a two-dimensional (2D) one. The experiments in this section were for May–September 2004.

4.3.1. Adding salinity sensor in the CR estuary

The results of the array-optimization in the CR estuary are presented in Fig. 5. The configurations for both the state (panels *s*.1, *s*.2, and *s*.3) and the error covariance (panels *e*.1, *e*.2, and *e*.3) are presented. The top and the middle panels depict the local uncertainty reduction R^2_{local} (from Eq. (19) below) before (panels *s*.1 and *e*.1) and after (panels *s*.2 and *e*.2) the new salinity sensor was added. The bottom panels (*s*.3 and *e*.3) depict the map of the cost function that shows the added value of placing the next salinity sensor at one of the logistically feasible locations. The added value was characterized by an increase in the global R^2 value (Eq. (17)).

The local uncertainty reduction values R^2_{local} on panels *e*.1-2 and *s*.1-2 were computed for each spatial location $x(i)$ using the following formula:

$$R^2_{local}(\xi) = 100 \times \left(1 - \frac{\mathbf{D}_{pp|y}(\xi,\xi)}{\mathbf{C}_{pp}(\xi,\xi)}\right) \tag{19}$$

where R^2_{local} is the local uncertainty reduction, and $\mathbf{C}_{pp}(\xi,\xi)$ and $\mathbf{D}_{pp|y}(\xi,\xi)$ are the prior and the posterior variances at the spatial location with the index ξ. In Fig. 5, the 3D fields of R^2_{local} are reduced to 2D maps by showing, for each horizontal location, which percentage of total variability in the vertical water column was explained by the existing measurements.

Results in Fig. 5 show that both the state- and the error-based configurations agreed on placing the next salinity sensor in the CR estuary close to the river end of the North Channel. Both optimal locations were in the shallower water (4.9 m depth), ~300 m away from each other. The maps of the cost function in panels *s*.3 and *e*.3 showed that most of the location in the North Channel of the CR estuary were equally beneficial, hence provided additional

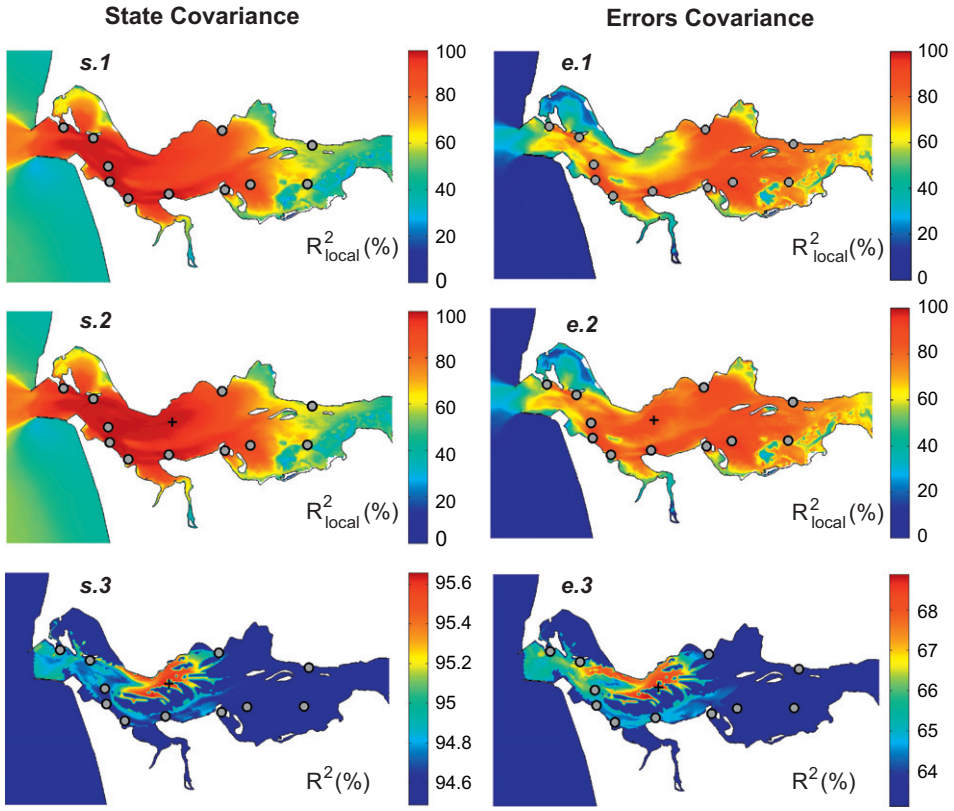

Fig. 5. Optimal location of the next estuarine sensor (marked with a diamond), optimized using the state (panels *s*.1, *s*.2, and *s*.3) and the error (panels *e*.1, *e*.2, and *e*.3) covariance. Top panels display local uncertainty reduction (R^2_{local} in %) for the existing salinity sensors (marked with circles), middle panels display R^2_{local} for augmented array (new sensor marked with a diamond), and bottom panels display the added value of placing the next sensor in one of the feasible sensor-locations. All panels display maximum-over-depth quantities.

information for the operational staff of the CORIE observatory. Placing the next salinity sensor in the North Channel of the CR estuary contributed slightly to uncertainty reduction for the state-based array, improving R^2 from 96% to 97%. However, the new salinity sensor had a higher impact on the error-based array, improving R^2 from 62% to 69%. Although these increases in the average R^2 were modest, the local R^2 in the North Channel of the estuary increased substantially, from ~50% to >80% in Fig. 5, panel *e*.2.

4.3.2. *Adding salinity sensors in the CR plume*

The results of the array-optimization in the CR plume are presented in Fig. 6: panels s.1–s.4 for the state covariance and panels e.1–e.4 for the error covariance. The panels in the top row display the prior salinity variance (panels s.1 and e.1) and the local uncertainty reduction R^2_{local} (panels s.2 and e.2) before any salinity sensors were added in the CR plume. The panels below the top row display the added value of placing new salinity sensors in the CR plume. Panels s.3 and e.3 depict the map of the cost function that shows the added value of placing the next salinity sensor at one of the feasible sensor-locations in the CR plume. The added value was characterized by the increase in the global R^2 value (Eq. (17)). Panels s.4 and e.4 show the increase in local uncertainty reduction R^2_{local} (Eq. (19)) after the new sensor was added in the CR plume. Panels 3 and 4 show three consecutive iterations of the algorithm, each sequentially adding another new salinity sensor in the CR plume.

Comparison of the optimal designs in for the state (panels s.1–s.4) and the error covariance (panels e.1–e.4) in Fig. 6 shows a remarkable agreement between the designs. Both designs place two salinity sensors north and one salinity sensor south of the CR mouth. These configurations contrast the historic array designs in the CR plume that favored sampling in southern and central plume; for an example of historical array, see three RISE and two OGI sensor on Fig. 1.2. A possible explanation for this difference is that historic arrays were biased towards sampling of the plume's southern branch—the dominant orientation of the CR plume in summer. In contrast to the historic designs, the proposed optimal designs in Fig. 6 stress the importance of placing multiple sensors in the northern plume in order to determine the extent of the plume's northern propagation.

To compare the optimality of the proposed salinity-array designs in the CR plume with that of the historic arrays, we list the computed R^2 values (Eq. (17)) for the historic and the proposed arrays in Table 3. Table 3 shows that the proposed arrays had consistently higher R^2 values than the historic arrays with comparable number of moorings. For example, the error-based array with three optimally placed moorings had higher R^2 value (24%) than the RISE array with three moorings (17%) and the combined RISE+OGI array with five moorings (20%).

5. Verification of the proposed array designs

To verify that the suggested arrays designed in Section 4 were optimal, we answered the following four questions about the validity of our array-design methods:

(1) *What were the theoretical limitations of our design algorithm*? In Section 3.5, we identify the limitations of our design algorithm, which suggested that the arrays that we designed

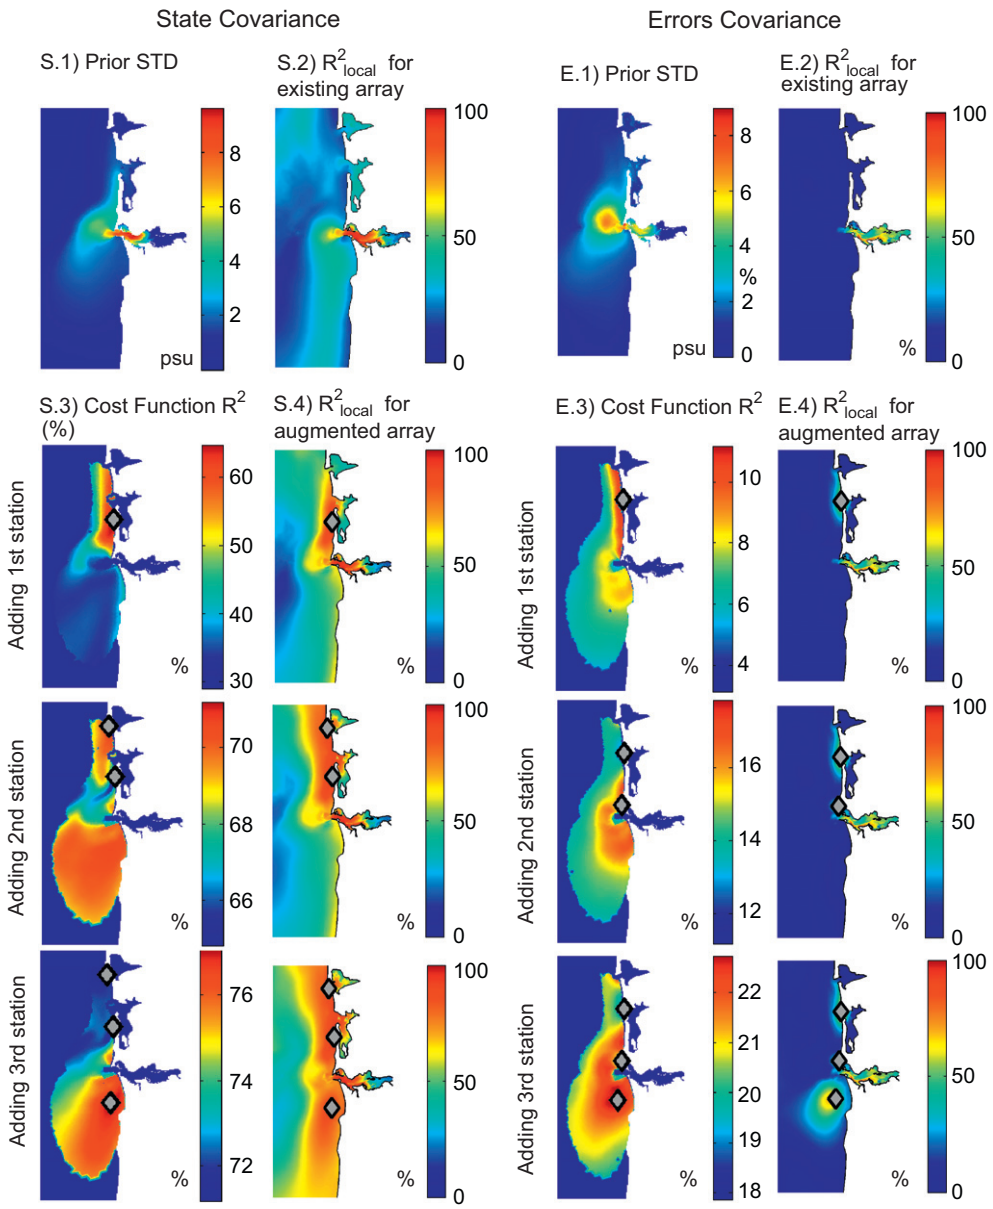

Fig. 6. Optimal locations of sensors in CR plume. Panels S.1–S.4 are based on the state covariance and Panels E.1–E4. (S.1 and E.1) Standard deviations (STD) of surface salinities (prior uncertainty). (S.2 and E.2) Local uncertainty reduction R^2_{local} (%) for the existing array; existing sensors are marked with circles. (S.3 and E.3) Value of the cost function for sensors one through three; new sensors are marked with diamonds. (S.4 and E.4) Local uncertainty reduction R^2_{local} (%) for the augmented array. All quantities were computed on the surface, where the prior uncertainty was the largest.

Table 3
Comparison of the salinity array designs in the CR plume.

Array name	Number of plume sensors	R^2 (%)	
		State	Error
Estuarine sensors	0	40	5
OGI	2	50	13
RISE	3	53	17
RISE+OGI	5	59	20
Optimal 2	2	74	19
Optimal 3	3	79	24
Optimal 5	5	84	33

Note: To make adequate comparison between the existing and the proposed optimal arrays, we only use information from the surface salinity sensors at OGI and RISE locations.

using the MSE criterion are likely to be optimal for data assimilation and for monitoring the variability of the CR estuary and plume. However, these arrays may not be optimal for detecting occurrence of extreme events, such as an event of unusually high salinity intrusion.

(2) *Was the model, used to compute the prior statistics, representative of the physical variability in the CR estuary and plume*? From the past studies (Baptista et al., 2005; Baptista, 2006; Frolov, 2007), we know that, at the time of our experiments, assimilated and non-assimilated models were capable of realistically representing variability of salinity, temperature, and water levels in the CR estuary and plume. To determine how the differences among these models influenced the optimal placement of sensors, we conducted a sensitivity study described in Section 5.1, which used four different datasets and three different lengths of the dataset.

(3) *How sensitive was the optimal sensor placement to changes in the optimization criteria and the optimization algorithm*? To determine the sensitivity of optimal sensor placement, we conducted two studies in Section 5.1 that examined the sensitivity to the choice of the optimization criteria, MSE vs. DET, and to the choice of the optimization algorithm, full exchange-type algorithm vs. add- and delete-only algorithms.

(4) *How accurately did the BLUE estimator predict the utility of alternative sensor placements*? In Section 5.2, we present the results of the cross-validation study, where we used data from the existing salinity sensors to verify the predicted utility of several alternative designs. We verified the R^2 values predicted using the BLUE algorithm using a data assimilated experiment that computed the actual value of error reduction at the locations of the validation sensors.

5.1. Sensitivity studies

To determine whether the proposed salinity arrays in Section 4 were sensitive to changes in experimental procedures, we studied the sensitivity of these optimal arrays to changes in the optimization criteria, the optimization algorithm, and the length and the type of the dataset that was used for computing of the prior covariance. From our experiments, we found that the design of these salinity arrays was strongly sensitive to the choice of the optimization criteria, but was not very sensitive to the choice of the optimization algorithm or to the choice of the dataset. To select between the arrays designed using the MSE and the DET criteria, we choose configurations optimized using the MSE criterion, since the MSE is also the criterion used by data assimilation.

5.1.1. Sensitivity to optimization criteria

To determine how the choice of the optimization criteria influences the optimal placement of salinity sensors in the CR estuary and plume, we reproduced sensor-placement experiments from Section 4 using two different optimization criteria: the MSE criterion (Eq. (6)) and the DET criterion (Eq. (7)). The reproduced experiments added a new salinity sensor in the CR estuary (shown in Fig. 7.1) and three new sensors in the CR plume (shown in Fig. 7.2). Sensor locations optimized using the MSE criterion, marked with circles in Fig. 7, differed from sensor locations optimized using the DET criterion, marked with diamonds. For example in the plume (Fig. 7.2), the MSE criterion distributed the salinity sensor throughout the CR plume, while the DET criterion clustered the salinity sensors in the high-variance area, located close to the CR mouth. We speculate, based on the limited results from Fig. 7 and on more extensive experiments not presented here, that locations selected by the DET criteria tend to cluster in the region of higher variance, while the MSE criteria tend to yield more distributed observation networks.

5.1.2. Sensitivity to optimization algorithm

To determine how sensitive the optimized locations of salinity sensors were to changes in the optimization algorithm, we compared optimization results using the full exchange-type algorithm (described in Section 3.3), with the results from the modified add-only algorithm (used in Sections 4.3) and with the delete-only algorithm (used in Section 4.2). The sensitivity experiment compared the R^2 values from 1000 random initializations of the full exchange-type algorithm (displayed as frequency bars in Fig. 8) with the deterministic results of the modified algorithms (marked with vertical lines in Fig. 8). To enable this computationally expensive experiment, we considered a small optimization problem, where the vertical locations of five salinity sensors at station *red*26 were optimized to fully represent the vertical variability of the salinity error at the location of *red*26.

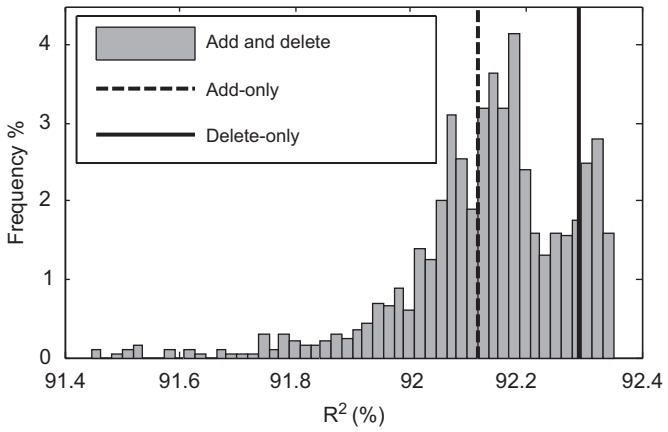

Fig. 8. Optimality of sensor placement as a function of the optimization algorithm used. The histogram shows frequency distribution of the R^2 values from 1000 randomly initialized passes through the full exchange algorithm. Vertical lines mark the R^2 values for the add-only (dashed line) and the delete-only (solid line) algorithms. Experiments were conducted for the MSE criterion, the error covariance, and a time period of $\frac{1}{2}$ years.

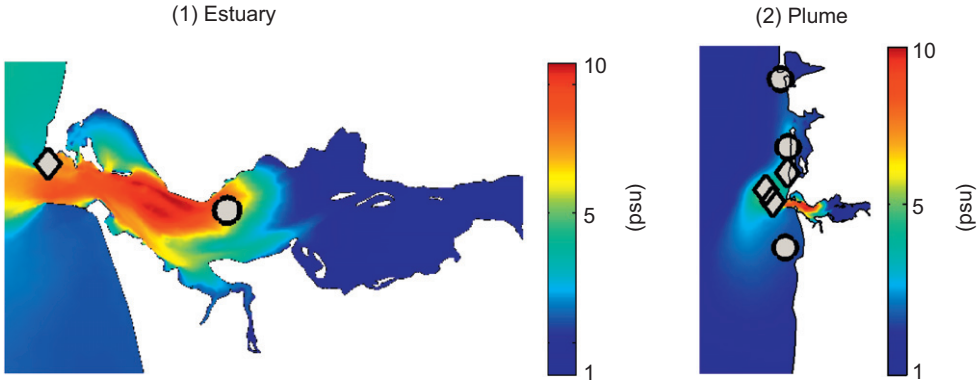

Fig. 7. Sensitivity of sensor placement to optimization criteria. Sensor locations that were located using the DET criterion are marked with diamonds and using the MSE criterion, with circles. (1) Optimal placement of the next salinity sensor in the CR estuary. (2) Optimal placement of salinity sensors in the CR plume. The locations were computed using the state covariance and are overlaying the maximum-over-depth prior uncertainty.

Experimental results in Fig. 8 show that the exchange-type algorithm produced array designs with a range of R^2 values (from 91.5% to 92.4%), which in several occasions were somewhat higher than the R^2 values computed by the deterministic add-only ($R^2 = 92.1\%$) and delete-only ($R^2 = 92.3\%$) algorithms. All three optimization algorithms produced arrays with the R^2 higher than an average randomly initialized array ($R^2 = 89\%$). It is very likely that some of the arrays optimized using the exchange-type algorithm were at or close to the global optimum. These experimental results suggest that multiple initializations of the full exchange-type algorithm are likely to find better array designs. However, the add-only and the delete-only algorithms are also likely to find good approximations to the globally optimal array design.

5.1.3. Sensitivity to the length and the type of the dataset

To determine how sensitive was the location of the next salinity station in the CR estuary to the type of the dataset that was used in computing of the prior covariance **C**, we repeated the experiment from Section 4.3.1 using four different datasets: assimilated DB14, non-assimilated DB14, non-assimilated DB16, and a scaled time-average of prior error $\mathbf{C}_{error}$ (Eq. (14)). To determine the sensitivity to the length of the dataset, we repeated the same experiment using the following three lengths of the dataset: $\frac{1}{2}$ year, 1 year, and 6 years. We found that the placement of the salinity sensor in the CR estuary was not very sensitive to the type or to the length of the dataset. At most, the suggested locations for the new salinity sensors were ~1.5 km apart.

5.2. Validation experiments

To determine whether the BLUE estimator accurately predicted the utility of alternative observation strategies in the CR estuary, we used a validation procedure that leveraged existing observations of salinity in the CR estuary. From our experiments, we found that the BLUE estimator overpredicted the R^2 values by 10–20%. However, the BLUE estimator correctly predicted the relative importance of alternative salinity sensor placements, suggesting that the optimal designs of salinity arrays in Section 4 were likely correct, even though the BLUE estimates of their utility, the R^2 values, were likely too optimistic. The details of the validation strategy and the detailed analysis of the experimental results follow.

We used the following algorithm to validate the predicted utility of a candidate salinity array using a data assimilation experiment:

(1) We choose a validation station with extensive historical dataset. We used two such stations—*red*26 and *am*169—each of which had three vertically spaced salinity sensors.
(2) We predicted, using the BLUE estimator (Eqs. (3) and (4)), how assimilating data from the existing salinity sensors can reduce salinity errors at the location of the validation station. The predictions were computed for the salinity error covariance matrix (Eq. (14)) and a two-week-long period in September of 2004. We also considered several alternative configurations of the observational network. In each network, one of the existing salinity sensors was removed from the assimilation.
(3) We verified the BLUE predictions of salinity error reductions, computed in item 2 above, by assimilating salinity data from candidate arrays using the data assimilation method described in (Frolov, 2007).

Tables 4 and 5 show the experimental results for validation stations *red*26 and *am*169. The alternative arrays are ranked based on their predicted R^2 values. The arrays with the lower R^2 values indicate salinity sensors with highly valuable observations, since removing them from the assimilation reduced the data assimilation accuracy at validation station significantly.

Table 4
Predicted and actual R^2 (%) values at the validation station *red*26.

Observation array	R^2	
	Predicted	Actual
no stations	0	0
–dsdma	53	43
–tansy	68	48
–jetta	78	49
–am169 (top)	78	49
–am169 (bot)	79	49
–sandi	79	50
–cbnc3	79	50
–am169 (mid)	79	50
–eliot	79	51
all stations	79	51

Note: Minus sign '–' in the first column indicates that the station was removed from the assimilation array.

Table 5
Predicted and actual R^2 (%) values at the validation station *am*169.

Observation array	R^2	
	Predicted	Actual
no stations	0	0
–tansy	34	00
–cbnc3	44	40
–jetta	49	40
–red26 (top)	**50**	**47**
–sandi	**50**	**42**
–red26 (bot)	50	41
–dsdma	**51**	**40**
–eliot	51	41
–red26 (mid)	51	41
all stations	51	41

Note: Minus sign '–' in the first column indicates that the station was removed from the assimilation array. Bold indicates stations for which the predicted and actual rankings disagree.

Results in Tables 4 and 5 show that the BLUE estimator consistently overpredicted the R^2 values at both validation stations, by about 10–20%. For example, the BLUE estimator predicted that using all sensors (last line in Tables 4 and 5) the errors in simulated salinity will be reduced at the validation station *red*26 by 79% and at *am*169 by 51%; however, the data assimilation experiment showed that actual errors were reduced only by 51% and 41%, respectively. The results in Tables 4 and 5 also show that the predicted ranking of the alternative observation strategies was largely correct, suggesting that the salinity error covariance matrix used in the BLUE estimator captured the correlation scales in the CR estuary well. For example, the rankings agreed for all configurations at the validation station *red*26 (Table 4) and for most configurations at the validation station *am*169 (Table 5). The few candidate arrays for which the rankings disagreed are marked in bold in Table 5.

6. Summary and conclusions

Algorithms for evaluating and optimizing a fixed network of sensors were developed and tested. The developed algorithms

were based on the theory of optimal experiment design, were computationally inexpensive, and were application- and model-independent. The developed algorithms were tested in the coastal margin observatory for the CR estuary and plume.

In evaluating the existing observational array for the CR estuary, we found that the existing array of salinity sensors in the CR estuary was too sparse (capturing at most 60% of the error variability) and too prone to data dropouts (capturing as little as 20% of error variability when salinity sensors biofouled). We hypothesize that the following two strategies will improve the coverage and the robustness of the existing array:

(1) To improve the accuracy of data assimilation in the CR estuary, we suggested placing a salinity sensor in the North Channel of the CR estuary, leading to an improvement in the R^2 value for salinity errors from 62% to 69%;
(2) To reduce the maintenance costs for the salinity array in the CR estuary, we suggested a strategy for reducing the existing salinity array from 14 to 10 sensors. This strategy should free the limited field resources, lead to improve maintenance on remaining sensors, and reduce the negative impact of sensor biofouling.

To improve the historical design of the observational array in the CR plume, we suggested deploying two sensors north and one sensor south of the CR mouth. This optimized array with three sensors is predicted to be more informative of salinity errors (R^2 of 24%) than the historic RISE and OGI arrays with five sensors (R^2 of 20%);

In transitioning the developed array designs to operational staff of the CR observatory, we found that the suggested optimal observational strategies were only a part of considerations weighted by the science managers and the field staff. Alternative considerations included: logistical constraints, contractual obligations, and competing observational agendas—all of which would be impossible to express as a single cost function. However, the benefits of the proposed objective evaluation procedure were also clear. Our objective procedure allowed scientists to prove or disprove their intuitions and to identify, otherwise overlooked, flaws in the design of the CORIE observatory. For example at the time of the publication, the new station in the north channel of the estuary was under construction and there was a plan for decommissioning three out of ~14 historical stations in the CR estuary.

Our experience in the CR suggests that the following findings are likely to hold in the design of other ocean observatories:

- Popular approximate optimization algorithms (Bishop et al., 2000; Sakov and Oke, 2008; Zhang and Bellingham, 2008), based on sequential addition and deletion of candidate stations to an optimized array, can yield very close approximations to globally optimal array configurations; see Section 5.1 for more details.
- Arrays optimized for improved accuracy of ocean DA will require more stations than the arrays optimized for monitoring of the ocean, as a consequence of shorter correlation lengthscales for error fields (e.g. errors in salinity) than for state fields (e.g. salinity); see Section 4.1 for more details.
- A preliminary analysis of an array can be performed using a long historic simulation of the ocean circulation, without the need to develop a complex variational- or ensemble-based DA system; see Section 4.3 for more details. However, the availability of such DA systems expands the capabilities and the accuracy of the array design experiments.

Our success in the CR estuary and plume suggest that algorithms for optimal placement of sensors are reaching maturity and are likely to play a significant role in the design of emerging ocean observatories, such as the national OOI and IOOS observatories.

Acknowledgments

We thank Dr. Yinglong Zhang for his work on CR simulation databases. SF is grateful to the members of his thesis committee—Drs. António Baptista, Todd Leen, Robert Miller, Mike Foreman, and John Kindle—for their guidance, support, and constructive feedback.

The National Science Foundation (ACI-0121475, OCE-0424602) and National Oceanic and Atmospheric Administration (AB133F-04-CN-0033) provided financial support for this research. Any statements, opinions, findings, conclusions, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views or policies of the federal sponsors, and no official endorsement should be inferred.

Appendix 1. BLUE estimates with factorized covariance matrix

To enable a fast computation of the BLUE estimate in Eqs. (3) and (4), we used the eigen-factorization of the prior covariance $\mathbf{C}$ in Eq. (16). To efficiently compute this eigen-factorization, we used the empirical orthogonal functions (EOF) of the forward model, which were pre-computed from a long, statistically representative simulation of the CR estuary and plume. The EOF functions, described by the orthonormal operator $\prod$, capture the dominant spatial modes of the simulation and provides the time series of the EOF coefficients

$$x_s = \prod(x - \bar{x}) \tag{20}$$

which capture the temporal variability of the simulation. The EOF operator $\prod$ was computed using a memory-efficient EOF algorithm described in Frolov (2007). Following are the computational procedures that we used to compute the factorized covariance (Eq. (16)) and to compute the BLUE estimate (Eq. (3)).

7.1. Computing factorized covariance

To compute the factorized state covariance matrix $\mathbf{C}^{state} = \mathbf{LSL}^T$ for a given period of interest, we used the following three-step procedure. In the first step, we selected the pre-computed EOF coefficients x_s (Eq. (20)) that fall inside of this period of interest. In the second step, we computed the eigen-decomposition of their covariance matrix

$$\mathrm{cov}(x_s) \xrightarrow{eig} \mathbf{L}_s\mathbf{S}\mathbf{L}_s^T \tag{21}$$

where $\mathbf{S}$ is the diagonal matrix of eigen-values, and $\mathbf{L_s}$ is the matrix of eigen-vectors, defined in the EOF subspace $\prod$. In the third step, we reconstructed the reduced-space eigen-vectors $\mathbf{L}_s$ as the full-space eigen-vectors $\mathbf{L}$ using the EOF operator $\prod$

$$\mathbf{L} = \prod \mathbf{L}_s \tag{22}$$

To compute the factorized error covariance $\mathbf{C}^{error} = \mathbf{LSL}^T$, we used a similar three-step procedure as for computing the state

covariance. In the first step, we used Eq. (14) to compute a time-averaged prior covariance $\mathbf{C}_{EOF}^{error}$. Since each of the forecast error covariance matrices $\mathbf{P}_{xx}^{f}(k)$ in Eq. (14) was defined in the same low-dimensional EOF subspace, their time-average $\mathbf{C}_{EOF}^{error}$ was also defined in the same EOF subspace, characterized by the projection operator $\prod$. In the second step, we computed the eigen-decomposition of $\mathbf{C}_{EOF}^{error}$ as

$$\mathbf{C}_{EOF}^{error} \xrightarrow{eig} \mathbf{L}_{error}\mathbf{S}\mathbf{L}_{error}^{T} \tag{23}$$

In the third step, we computed the eigen-vectors $\mathbf{L}$ of the full-space error covariance $\mathbf{C}^{error}$ as:

$$\mathbf{L} = \prod \mathbf{L}_{error} \tag{24}$$

7.2. Computing the prior variance

To compute the prior variance $cpp = diag(\mathbf{C}_{pp})$ in Eqs. (6), (17), and (28), we used the following formula that avoided the unnecessary computation of cross-covariance terms

$$cpp = \sum_{i=1}^{ns}([\mathbf{H}_p\mathbf{L}(:,i)\sqrt{S(i,i)}]^2) \tag{25}$$

where cpp is a vector of prior variances, $\mathbf{L}(:,i)$ is the ith eigen-vector, $S(i,i)$ is the ith eigen-value, $[\ldots]^2$ is the element-wise squaring-operation, and ns is the number of retained eigen-vectors.

7.3. Computing posterior variance

To compute the posterior variance $dpp = diag(\mathbf{D}_{pp|y})$, without computing the unnecessary cross terms, we used the following two-step procedure. First, we computed the intermediate matrix product

$$\mathbf{A} = \mathbf{C}_{py}\mathbf{C}_{yy}^{-1} \tag{26}$$

where, from Eqs. (5) and (16), the observation covariance $\mathbf{C}_{yy}$ and the cross-covariance $\mathbf{C}_{py}$ are

$$\mathbf{C}_{py} = (\mathbf{H}_p\mathbf{L})\mathbf{S}(\mathbf{L}^T\mathbf{H}_y^T)$$
$$\mathbf{C}_{yy} = (\mathbf{H}_y\mathbf{L})\mathbf{S}(\mathbf{L}^T\mathbf{H}_y^T) + \sigma^2\mathbf{I} \tag{27}$$

In a case some of the observations were nonlinear, we linearized them around the mean state $\bar{x}$. In the second step, we compute the posterior variance dpp as

$$dpp = cpp - \sum_{i=1}^{ns}(\mathbf{A}(:,i).^*\mathbf{C}_{yp}(:,i)) \tag{28}$$

where '.*' is the element-wise multiplication of vectors, and $(:,i)$ denotes the ith column of the matrix.

References

Baker, N.L., Daley, R., 2000. Observation and background adjoint sensitivity in the adaptive observation-targeting problem. Quarterly Journal of the Royal Meteorological Society 126 (565B), 1431–1454.

Baptista, A.M., 2006. Corie: The First Decade of a Coastal Margin Collaborative Observatory, in Oceans'06 MTS/IEEE, Boston.

Baptista, A.M., Zhang, Y.-L., Chawla, A., Zulauf, M.A., Seaton, C., Myers, E.P., Kindle, J., Wilkin, M., Burla, M., Turner, P.J., 2005. A cross-scale model for 3D Baroclinic circulation in estuary-plume-shelf systems: ii. Application to the Columbia River. Continental Shelf Research 25, 935–972.

Barnes, C.A., Duxbory, A.C., Morse, B.A., 1972. Circulation and selected properties of the Columbia River effluent at sea. In: Alverson, D.L., Pruter, A.T. (Eds.), The Columbia River Estuary and Adjacent Ocean Regions: Bioenvironmental Studies. University of Washington Press, Seattle, WA., pp. 41–80.

Bennett, A.F., 1985. Array design by inverse methods. Progress in Oceanography 15, 129–156.

Bennett, A.F., 1992. Inverse Methods in Physical Oceanography. Cambridge University Press, Cambridge.

Bennett, A.F., 2002. Inverse Modeling of the Ocean and Atmosphere. Cambridge University Press, Cambridge, UK.

Berliner, M., Lu, X.-Q., Snyder, C., 1999. Statistical design of adaptive weather observations. Journal of Atmospheric Sciences 56 (15), 2536–2552.

Bertozzi, A.L., Kemp, M., Marthaler, D., 2005. Determining environmental boundaries: asynchronous communication and physical scales. Lecture Notes in Control and Information Sciences 309, 25–42.

Bishop, C.H., Etherton, B.J., Majumdar, S.J., 2000. Adaptive sampling with the ensemble transform Kalman filter. Part I: theoretical aspects. Monthly Weather Review 129, 420–436.

Bottom, D.L., Simenstad, C.A., Burke, J., Baptista, A.M., Jay, D.A., Jones, K.K., Casillas, E., Schiewe, M.H., 2005. Salmon at River's End. The Role of the Estuary in the Decline and Recovery of Columbia River Salmon. NOAA Technical Memorandum NMFS-NWFSC-68, US Department Commer.

Bretherton, F.P., Davis, R.E., Fandry, C.B., 1976. A technique for objective analysis and design of oceanographic experiments applied to mode-73. Deep-Sea Research 23, 559–582.

Burla, M., Baptista, A.M., Zhang, Y.-L., Casillas, E., Bottom, D.L., Simenstad, S.A., 2007. Salmon habitat opportunity in the columbia river estuary: modeling the physical environment to inform management decisions. In: Coastal Zone '07, Portland, OR.

Burla, M., Baptista, A.M., Casillas, E., Williams, J.G., (accepted). The influence of the columbia river plume on the survival of steelhead (Oncorhynchus Mykiss) and Chinook Salmon (O. Tshawytscha): A Numerical Exploration. Canadian Journal of Fisheries and Aquatic Sciences

Chang, N.-B., Tseng, C.C., 1999. Optimal design of a multi-pollutant air quality monitoring network in a metropolitan region using Kaohsiung, Taiwan as an example. Environmental Monitoring and Assessment 57, 121–148.

Daescu, D.N., Carmichael, G.R., 2003. An adjoint sensitivity method for the adaptive location of the observations in air quality modeling. Journal of the Atmospheric Sciences 60, 434–450.

Fedorov, V.V., 1972. Theory of Optimal Experiments. Academic Press, New York.

Fedorov, V.V., 1994. Optimal experimental design: spatial sampling. Calcutta Statistical Association Bulletin 44 (173–174).

Frolov, S., 2007. Enabling technologies for data assimilation in a coastal-margin observatory. Ph.D. thesis, Oregon Health & Science University, Environmental & Biomolecular Systems.

Gandin, L., 1963. Objective analysis of meteorological fields. Hydrometeoizdat, Leningrad.

Garcia-Berdeal, I., Hickey, B., Kawase, M., 2002. Influence of wind stress and ambient flow on a high discharge river plume. Journal of Geophysical Research 107 (C9), 3130.

Grimes, C., Kingsford, M., 1996. How do riverine plumes of different sizes influence fish larvae: do they enhance recruitment? Marine and Freshwater Research 47 (2), 191–208.

Hackert, E.C., Miller, R.N., Busalacchi, A.J., 1998. An optimized design for a moored instrument array in the tropical Atlantic ocean. Journal of Geophysical Research 103, 7491–7509.

Hickey, B., 2004. River influences on shelf ecosystems: initial impressions. In: AGU Fall meeting, San Francisco.

Hickey, B.M., Pietrafesa, L.J., Jay, D.A., Boicourt, W.C., 1998. The Columbia River plume study: subtidal variability in the velocity and salinity fields. Journal of Geophysical Research 103 (C5), 10,339–310,369.

Hickey, B., Geier, S., Kachel, N., MacFadyen, A., 2005. A bi-directional river plume: the Columbia in summer. Continental Shelf Research 25 (14), 1631–1656.

Jay, D.A., Flinchem, E.P., 1997. Interaction of fluctuating river flow with a Barotropic tide: a demonstration of wavelet tidal analysis methods. Journal of Geophysical Research 102 (C3), 5705–5720.

Jay, D., Smith, J.D., 1990. Circulation, density distribution and neap-spring transitions in the Columbia River estuary. Progress in Oceanography 25 (1–4), 81–112.

Langland, R., Rohaly, G., 1996. Adjoint-based targeting of observations for Fastex cyclones. In: Preprints, Seventh Conference on Mesoscale Processes. Amer. Meteor. Soc., Phoenix, AZ.

Leonard, N.E., Paley, D., Lekien, F., Sepulchre, R., Fratantoni, D.M., 2006. Collective motion, sensor networks and ocean sampling. In: Proceedings of the IEEE, Special Issue on Networked Control Systems.

Lu, Z., Leen, T.K., van der Merwe, R., Frolov, S., Baptista, A.M., 2007. Sequential data assimilation with a sigma-point Kalman filter on a low-dimensional manifold, technical report TR-07-001, NSF-STC for Coastal Margin Observation and Prediction, available at 〈http://www.stccmop.org/files/CMOP-TR-07-001.pdf〉.

MacCready, P., Banas, N. S., Hickey, B. H., Dever, E. P., 2007. A model study of tide- and wind-induced mixing in the Columbia River estuary and Plume. Submitted to Continental Shelf Research April 2007.

National Research Council, 2003. Enabling ocean research in the 21st century: implementation of a network of ocean observatories.

Ogren, P., Fiorelli, E., Leonard, N.E., 2004. Cooperative control of mobile sensor networks: adaptive gradient climbing in a distributed environment. IEEE Transactions on Automatic Control 49 (8), 1292–1302.

Pukelsheim, F., 1993. Optimal Design of Experiments. Wiley, New York.

Ripley, B.D., 1987. Spatial Statistics. Wiley, New York.

Sakov, P., Oke, P.R., 2008. Objective array design: application to the tropical Indian Ocean. Journal of Atmospheric and Oceanic Technology 25, 794–807.

She, J., Høyer, J.L., Larsen, J., 2006. Assessment of sea surface temperature observational networks in the Baltic sea and North sea. Journal of Marine Systems.

Silvey, S.D., 1980. Optimal Design. Chapman and Hall, London.

US Commission on Ocean Policy, 2004. An ocean blueprint for the 21st century. Final Report.

USACE, 2001. Columbia river channel improvement project: 2001 biological assessment. US Army Corps of Engineers.

Zhang, Y.-L., Baptista, A.M., Myers, E.P., 2004. A cross-scale model for 3D Baroclinic circulation in estuary-Plume-shelf systems: I. formulation and skill assessment. Continental Shelf Research 24, 2187–2214.

Zhang, Y., Bellingham, J.G., 2008. An efficient method of selecting ocean observing locations for capturing the leading modes and reconstructing the full field. Journal of Geophysical Research, 113, C04005, doi:10.1029/2007JC004327.

Zhang, Y.-L., Baptista, A.M., 2008. Selfe: A semi-implicit Eulerian–Lagrangian finite-element model for cross-scale ocean circulation. Ocean Modeling 21 (3–4), 71–96.