




ELSEVIER

Contents lists available at ScienceDirect

Dynamics of Atmospheres and Oceans

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



Assessing 4D-VAR for dynamical mapping of coastal high-frequency radar in San Diego

I. Hoteit^{a,*}, B. Cornuelle^a, S.Y. Kim^a, G. Forget^b, A. Köhl^c, E. Terrill^a

^a Scripps Institution of Oceanography, La Jolla, CA, USA

^b Massachusetts Institute of Technology, Boston, USA

^c Institut fuer Meereskunde, Hamburg, Germany

ARTICLE INFO

Article history:

Received 6 June 2008

Received in revised form 14 November 2008

Accepted 24 November 2008

Available online 27 December 2008

Keywords:

San Diego

Coastal modeling

HF Radar

4D-VAR

Predictability

ABSTRACT

The problem of dynamically mapping high-frequency (HF) radar radial velocity observations is investigated using a three-dimensional hydrodynamic model of the San Diego coastal region and an adjoint-based assimilation method. The HF radar provides near-real-time radial velocities from three sites covering the region offshore of San Diego Bay. The hydrodynamical model is the Massachusetts Institute of Technology general circulation model (MITgcm) with 1 km horizontal resolution and 40 vertical layers. The domain is centered on Point Loma, extending 117 km offshore and 120 km alongshore. The reference run (before adjustment) is initialized from a single profile of T and S and is forced with wind data from a single shore station and with zero heat and fresh water fluxes. The adjoint of the model is used to adjust initial temperature, salinity, and velocity, hourly temperature, salinity and horizontal velocities at the open boundaries, and hourly surface fluxes of momentum, heat and freshwater so that the model reproduces hourly HF radar radial velocity observations. Results from a small number of experiments suggest that the adjoint method can be successfully used over 10-day windows at coastal model resolution. It produces a dynamically consistent model run that fits HF radar data with errors near the specified uncertainties. In a test of the forecasting capability of the San Diego model after adjustment, the forecast skill was shown to exceed persistence for up to 20 h.

© 2008 Elsevier B.V. All rights reserved.

* Corresponding author at: Scripps Institution of Oceanography, 9500 Gilman Dr, Room 410 0230, La Jolla, CA 92093, USA
Tel.: +1 858 822 5021; fax: +1 858 534 9820.

E-mail address: ihoteit@ucsd.edu (I. Hoteit).

1. Introduction

One of the most important environmental problems for San Diego has been water quality in the ocean near the U.S./Mexico border. The Tijuana River plume during the wet season often contains high levels of contaminants from the upstream watershed, and poses human health risks which threatens the local economies. The implementation of strict water quality regulations and the development of rapid indicators would benefit the San Diego coastal ecosystems and the populations they support. The scarcity of observations of coastal ecosystems and the limited real-time data telemetry, assimilation, and analysis are impediments to the development of a predictive understanding of environmental variability and change in California's coastal waters. To improve the understanding in the near shore processes and the interactions between environmental variables in San Diego coastal waters, systematic coastal observations with sufficient duration and spatial extent have been requested. The goal is to enhance analysis and forecasting capabilities, and to respond to the environmental issues appropriately. Surface current observations by high-frequency radars are maturing as an oceanographic observational tool in support of these goals. In September 2002, three radars were deployed in the San Diego region at Point Loma, Border Field State Park located on the U.S.–Mexico, and an offshore station located on the Coronado Islands (as shown in Fig. 1).

Surface current measurements by high-frequency (3–30 MHz) radars are based on the interpretation of the back-scattered signals from surface gravity waves using shore-based antennae. The systems provide estimates of surface current radial velocities on polar coordinate grid points centered on each antenna location (De Paolo and Terrill, 2007). The observations are available in near real time with continuous temporal and broad spatial coverage. However, the signal processing for extracting surface currents from the back-scattered radar signals is complicated, and the maps of ocean surface currents can have missing observations. For instance, the application of the Multiple Signal Characterization (MUSIC) algorithm on the measured Doppler spectrum does not provide a solution for all bearing angles (De Paolo and Terrill, 2007). The estimate of current vectors along the baseline between two radars where the measurements of radial velocities are nearly aligned, suffers from poor geometrical dilution of precision (GDOP), and some processing methods produce spurious current vectors in these regions. The region with radial velocities crossing at angles less than 15–20 degrees between two radars is also commonly considered to produce unusable current vectors.

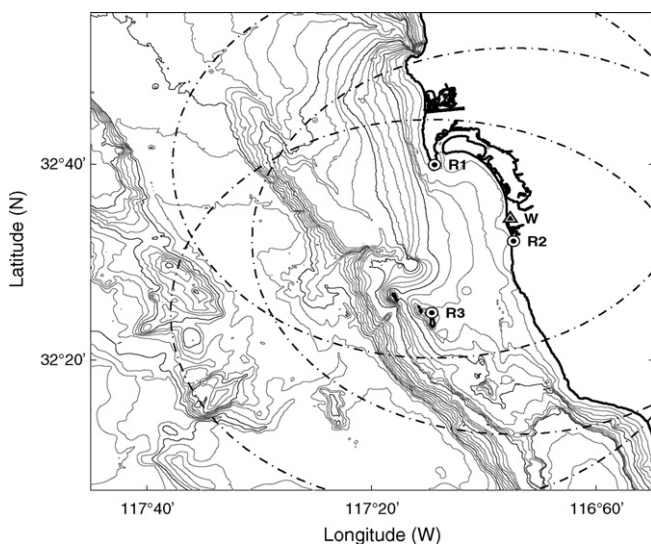


Fig. 1. The San Diego region observations of surface currents and wind. The high-frequency radar sites are Point Loma (R1), Border Park (R2), and Coronado Island (R3). The dashed-dot circles denote the coverage area by three high-frequency radars (R1, R2, and R3), respectively. The Tijuana wind station (W) is located near the Tijuana River valley.

Many applications of surface currents require time- and space-continuous data, so gappy observations must be interpolated to a regularly spaced product. Several approaches have been proposed to estimate vector velocity fields from radial velocities data measured by multiple antenna installations and to fill in the missing data points in time and space. Unweighted least squares fitting (Lipa and Barrick, 1983) has been widely used, but this often assumes discontinuous correlation functions for radial velocities and can yield spurious values along baselines between stations. Objective analysis, also called optimal interpolation (OI), techniques are now more often used. These are implemented using covariances made from normal mode analysis (Lipphardt et al., 2000), open boundary (OB) modal analysis (Kaplan and Lekein, 2007), Empirical Orthogonal Function (EOF) analysis (Beckers and Rixen, 2003; Alvera-Azcarate et al., 2005), and using idealized or smoothed observed covariances (Davis, 1985; Kim et al., 2007, 2008).

The approaches mentioned above are statistical techniques. Their performance depends on the accuracy of the covariance used for interpolating the data in space and in time. Because it is difficult to estimate covariances, the interpolation and extrapolation capabilities of these methods can be limited. Moreover, present mapping techniques often do not make full use of dynamical information for a given data set. Through dynamical linkages, current measurements may provide information about other variables and about past and future times. For instance, surface current measurements can be combined with a numerical model that describes the dynamics of the system to improve the three-dimensional picture of the ocean and to provide corrections to the wind stress fields, as suggested by Oke et al. (2002) and Paduan and Shulman (2004). Combining numerical models and data to make an optimal analysis, often for forecasts, is referred to as data assimilation.

Data assimilation methods are now widely used in oceanography to combine model outputs and data according to their respective accuracy (Ghil and Malanotte-Rizzoli, 1991; Wunsch, 1996). These methods provide dynamically based interpolation to map the data and to spread the information from measurements to non-observed variables in space and in time. By construction, they also readily provide a dynamical forecasting system that is continuously updated by the data. The advanced data assimilation schemes used for large-scale ocean general circulation models (e.g. see Stammer et al., 2002; Fukumori et al., 1999) have been less popular in high resolution limited area models, perhaps due to limited computational resources and difficulties associated with the estimation of strong, fine scale variability of coastal areas. Consequently, data assimilation schemes for high resolution limited area models were very often based on empirical approaches, e.g. Breivik and Stra (2001), Oke et al. (2002), De Mey and Benkiran (2002) and Kurapov et al. (2005). Continuous progress in computing capabilities are beginning to enable advanced four-dimensional data assimilation techniques, such as Ensemble Kalman Filter (EnKF) methods, e.g. Hoteit et al. (2005a), Zhang et al. (2006), Korres et al. (2007), and four-dimensional variational (4D-VAR) methods, e.g. Taillandier et al. (2004), Lea et al. (2006), Di Lorenzo et al. (2007), in high resolution coastal models.

Although radar observational systems are potentially one of the most important data sets for coastal ocean state estimation, the assimilation of the radar data is still evolving. Comparatively simple techniques, mainly based on nudging (e.g. Wilkin et al., 2005) and OI (e.g. Breivik and Stra, 2001; Oke et al., 2002; Paduan and Shulman, 2004), have been used to combine high-frequency (HF) radar observations with primitive equation ocean models. More recently, 3D-VAR methods have been used for operational assimilation of the HF radar network in Southern California (Li et al., 2008).

These methods use a prior estimate of the forecast error covariance matrix that can be difficult to obtain in practice. They tend to match the data very well, but do not produce a completely dynamically consistent interpolation of the data in space and time. Kurapov et al. (2003) used a simplified model with representer-based 4D-VAR (Bennett, 2002) to assimilate HF radar data for studying the M2 internal tides. The goal of this study is to use a high resolution configuration of the Massachusetts Institute of Technology general circulation model (MITgcm) and its adjoint to see if there exists a dynamically consistent model evolution that matches high-frequency radar data collected near the San Diego coast. The MITgcm/adjoint ocean data assimilation system adjusts initial conditions, boundary conditions, and atmospheric forcing fields in order to improve the model match to the observations. This could provide an efficient and dynamically consistent approach to correct errors in the wind stress, which can be a significant source of error in coastal models (Paduan and Shulman, 2004). It has not yet been tested for assimilation of hourly ocean data or at 1 km resolution. The predictability of the system and

the impact of the high-frequency radar data on the coastal ocean circulation and the forecasting skill of the model are also studied.

The paper is organized as follows. Sections 2 and 3, respectively, describe the dynamical model and the radar radial velocity observations. The characteristics of the assimilation system, including the specified uncertainties for the observations and the control variables, is presented in Section 4. The assimilation system is evaluated and the state estimates analyzed and discussed in Section 5. Concluding remarks and a general discussion are provided in Section 6.

2. The dynamical model

The MITgcm (Marshall et al., 1997) solves the primitive equations on a sphere under the Boussinesq approximation to the Navier–Stokes equations. The equations are written in z -coordinates and discretized using the centered second order finite differences approximation in a staggered “Arakawa C-grid”. The numerical code is designed to enable the construction of its adjoint using the TAF (Transformation of Algorithms in Fortran) automatic differentiation tool (Giering and Thomas, 1998). The MITgcm/adjoint assimilation system was used by the ECCO (Estimating the Circulation and the Climate of the Ocean) consortium to produce the first global state estimation (Stammer et al., 2002). The system was then successfully applied in regional eddy-permitting configurations (Hoteit et al., 2005b; Gebbie et al., 2006; Hoteit et al., 2008). The system has not previously been used for coastal applications.

The model domain covers the San Diego coast region and is centered on Point Loma as shown in Fig. 2. It extends 117 km offshore and 120 km alongshore with the southeast corner at 32° N and 243.1° E. The maximum bottom depth is 2000 m. The model was integrated on a 1×1 km grid with 40 vertical levels. The spacing of the vertical levels is 4 m at the surface and increases logarithmically to 10 m at 90 m depth, 100 m at 850 m depth, and 150 m at 2000 m depth. The time step is 90 s. The model operates in hydrostatic mode with an implicit free surface. No-slip conditions are imposed at the lateral boundaries while the bottom friction is quadratic with a drag coefficient equal to 0.001. The sub-grid scale physics use diffusive operators of second order in the horizontal and the vertical. Horizontal mixing coefficients for the diffusive and viscous operators are $40 \text{ m}^2 \text{ s}^{-1}$ and $80 \text{ m}^2 \text{ s}^{-1}$, respectively. Vertical diffusivity and viscosity are parametrized by Laplacian mixing with background values $1 \times 10^{-6} \text{ m}^2 \text{ s}^{-1}$ and $1 \times 10^{-4} \text{ m}^2 \text{ s}^{-1}$, respectively. In the surface mixed layer, the eddy coefficients are parametrized by the K-profile parametrization (KPP) model (Large et al., 1994).

Open boundaries are set in the south, north, and west. They are implemented as in Zhang and Marotzke (1999). The fields within a 10 km buffer region near the boundaries are relaxed to values




Fig. 2. Bathymetry of the model domain (m).

specified on the boundaries with timescale 1 day on the boundary and 7 days at the interior edge. For the reference run, temperature and salinity at the boundaries are specified to be horizontally uniform, taken from a single observed profile and the horizontal components of the velocity are set to zero. The reference model run is forced uniformly in space with hourly winds from a single shore station and with no heat or salinity fluxes. The model starts from rest and the initial temperature and salinity are horizontally uniform, using the same single profile. The reference model was started from rest with no spin-up because of the weak wind forcing and the low expected skill of the single shore station wind. The reference run will therefore have minimal structure to contaminate the final estimates. The skill of the reference run is indeed poor in the comparisons with data to be shown later.

3. High-frequency radar data

Measurements of the radial components of surface currents from an array of high-frequency (25 MHz) direction-finding-style systems (Codar Ocean Sensors, Palo Alto, CA) deployed in the San Diego region since September 2002 were used to constrain the surface velocity fields of the San Diego MITgcm configuration. Radial speed measurements are related to the zonal U and meridional V components of the surface velocity fields according to

$$R_{\text{CE:ITALIC}} = U \cos(\theta) + V \sin(\theta) \quad (1)$$

where θ is the radial angle. Individual radar sites are located at Point Loma, Border Field State Park at the U.S.–Mexico border, and an offshore station located on the Coronado Islands as shown in Fig. 1. The Doppler shifts of back-scattered radio signals from surface gravity waves are used to estimate the current component in the direction of the radial vector originating at the radar site. Radial velocities are observed on a cylindrical grid with roughly 1 km spacing at 20 km range. Tidal movements were removed from the time series of observations using a least squares fit to four diurnal and four semi-diurnal tidal lines over a 1-year period. The hourly count of available radial velocities at each radar site between November 1, 2003 and January 31, 2004 is shown in Fig. 3. Since at least two radial velocities from different sites are required for a vector solution, more accurate estimates of the surface currents are expected in the regions with overlapping radar range cells from multiple sites.

Quality control was applied to the radar observations to remove outliers (values in the tail of the probability distribution function). A 2-day running mean of the time series is subtracted in order to remove the local ocean's response to wind events and to produce a high-pass filtered time series. These anomaly time series are then filtered with a moving 5-day window in which the deviation from the




Fig. 3. Hourly count of data from the three radars deployed at Point Loma, Border Park, and Coronado Island. Assimilation experiments were performed over 10-day periods in December.

5-day running mean is compared to the standard deviation in the same period. Anomalies outside 5 standard deviations are flagged as outliers. 5 standard deviations were chosen based on the observed probability distribution of radial velocities and visual checks of the outliers. The hourly data during 5 days provide enough realizations to detect statistically anomalous data.

The wind has a spectral peak at 1 cycle per day, which is very close to the local inertial frequency, so the two are not distinguishable over a 10-day period. The HF radar velocities have a surprisingly small (roughly 3%) contribution from the inertial frequency, perhaps because of the complicated coastline and topography.

4. The assimilation procedure

To fit the MITgcm model trajectory to the HF radar data the so-called “strong constraint” formulation of the 4D-VAR problem is used. This problem is posed as the minimization of a cost function measuring the difference between the model state and data over a specified period of time (hereafter referred to as assimilation period) and constrained by the nonlinear model equations subject to a set of control variables. The gradient of the cost function, which is obtained by integrating the adjoint of the tangent linear model backward in time (Le Dimet and Talagrand, 1986), is then used to determine descent directions toward the minimum in an iterative procedure. The cost function is a weighted sum of quadratic norms of model–data misfit (J_{data}) and changes to the control variables ($J_{control}$) between the initial time (t_0) and the final time (t_f) of the assimilation window,

$$J(u) = \underbrace{\sum_{t=t_0}^{t_f} [\mathbf{y}(t) - \mathbf{H}_t(\mathbf{x}(t))]^T \mathbf{R}^{-1}(t) [\mathbf{y}(t) - \mathbf{H}_t(\mathbf{x}(t))]}_{J_{data}} + \underbrace{\sum_{t=t_0}^{t_f} [\mathbf{u}(t) - \mathbf{u}^b(t)]^T \mathbf{Q}^{-1}(t) [\mathbf{u}(t) - \mathbf{u}^b(t)]}_{J_{control}} \quad (2)$$

where $\mathbf{x}(t)$ is the model state vector and $\mathbf{u}^b(t)$ a first guess (or “background”) for the control vector $\mathbf{u}(t)$ at time t . The control vector accounts for errors in the external forcing fields, initial and boundary conditions. Errors in the internal model physics were not included. The vector $\mathbf{y}(t)$ contains all observations available at time t and is related to the model state through the observational operator \mathbf{H}_t plus some observational errors $\boldsymbol{\varepsilon}(t)$, according to

$$\mathbf{y}(t) = \mathbf{H}_t(\mathbf{x}(t)) + \boldsymbol{\varepsilon}(t) \quad (3)$$

$\mathbf{R}(t)$ and $\mathbf{Q}(t)$ are the covariance matrices of observational and first-guess control uncertainties, respectively. Their inverses are used as weights in the cost function. In the formulation of (2), it is assumed that the errors in the control vector are uncorrelated with the observation errors, and that both errors are uncorrelated in time. For assimilation of radar data, the observational operator \mathbf{H}_t is linear and converts the surface velocity components of the MITgcm state vector to radial observations according to (1). The adjoint model of the MITgcm was generated using the TAF automatic differentiation tool (Giering and Thomas, 1998). The KPP parameterization was used in the forward run but the adjoint of KPP was not used because of difficulties with conditional statements. Mixing parameters generated in the forward run were used for the adjoint run. The descent directions toward the minimum are determined using the Quasi-Newton M1QN3 algorithm (Gilbert and Le Maréchal, 1989). The background state is taken as the starting point for minimization. At each step the adjoint is calculated for linear perturbations to the current state estimate, while the reference state for the control cost is not changed. This method will be called “4D-VAR” in the sequel, at the risk of confusion with incremental methods or methods that change only the initial conditions.

4.1. Control variables

The control variables are temperature (T), salinity (S), and horizontal velocities fields (U , V) at the initial time and at the open boundaries, and the atmospheric forcing fluxes: zonal and meridional wind stress and heat and fresh water fluxes. A novelty of the present approach is to allow frequent,

high resolution adjustments to air-sea fluxes, which can largely determine surface velocities in the real ocean. The term representing the constraints on the control variables in (2) has the form

$$\begin{aligned} \mathbf{J}_{\text{control}} = & [\mathbf{x}_0 - \mathbf{x}_0^b]^T \mathbf{Q}_0^{-1} [\mathbf{x}_0 - \mathbf{x}_0^b] + \sum_{t=t_0}^{t_f} [\mathbf{x}_{ob}(t) - \mathbf{x}_{ob}^b(t)]^T \mathbf{Q}_{ob}^{-1} [\mathbf{x}_{ob}(t) - \mathbf{x}_{ob}^b(t)] \\ & + \sum_{t=t_0}^{t_f} [\mathbf{f}(t) - \mathbf{f}^b(t)]^T \mathbf{Q}_f^{-1} [\mathbf{f}(t) - \mathbf{f}^b(t)] \end{aligned} \quad (4)$$

where \mathbf{x}_0 , $\mathbf{f}(t)$, and $\mathbf{x}_{ob}(t)$ represent, respectively, the initial conditions, the forcing fields, and the open boundaries conditions, and \mathbf{Q}_0 , \mathbf{Q}_f , and \mathbf{Q}_{ob} the associated error covariances.

The model sea surface height is sensitive to the barotropic component of the normal velocities at the OB. This makes the estimation of velocities at the OB poorly conditioned when simultaneously adjusted with other control variables. To avoid large contrasts in sensitivities, which can severely slow gradient descent optimization (Zupanski, 1996), the expected uncertainty variance of the barotropic component of the normal velocities at the OB was set to $2.5 \times 10^{-3} \text{ m s}^{-1}$ compared to 20 cm s^{-1} at the surface decreasing to 2 cm s^{-1} at the bottom for the baroclinic component. This corresponds to a root-mean-square transport across the open boundaries of roughly 0.03 Sverdrup for the entire domain. This was achieved using a decomposition of the normal velocity into barotropic and baroclinic modes as described by Hoteit et al. (2005b). In the MITgcm, the boundary conditions require the complete specification of the state U , V , S and T . Accordingly, four separate penalty terms were added to the cost function; one for each state variable. The cost function term for the adjustment of the open boundaries then reads

$$\begin{aligned} \mathbf{J}_{ob} = & \sum_{t=t_0}^{t_f} \{ [T_{ob}(t) - T_{ob}^b(t)]^T \mathbf{Q}_{T_{ob}}^{-1} [T_{ob}(t) - T_{ob}^b(t)] + [S_{ob}(t) - S_{ob}^b(t)]^T \mathbf{Q}_{S_{ob}}^{-1} [S_{ob}(t) - S_{ob}^b(t)] \\ & + [\tilde{U}_{ob}(t) - \tilde{U}_{ob}^b(t)]^T \mathbf{Q}_{\tilde{U}_{ob}}^{-1} [\tilde{U}_{ob}(t) - \tilde{U}_{ob}^b(t)] + [\tilde{V}_{ob}(t) - \tilde{V}_{ob}^b(t)]^T \mathbf{Q}_{\tilde{V}_{ob}}^{-1} [\tilde{U}_{ob}(t) - \tilde{U}_{ob}^b(t)] \\ & + [\bar{V}_{ob}(t) - \bar{V}_{ob}^b(t)]^T \mathbf{Q}_{\bar{V}_{ob}}^{-1} [\bar{V}_{ob}(t) - \bar{V}_{ob}^b(t)] + [\bar{V}_{ob}(t) - \bar{V}_{ob}^b(t)]^T \mathbf{Q}_{\bar{V}_{ob}}^{-1} [\bar{V}_{ob}(t) - \bar{V}_{ob}^b(t)] \} \end{aligned} \quad (5)$$

where $\mathbf{x}_{T_{ob}}(t)$, $\mathbf{x}_{S_{ob}}(t)$, $\mathbf{x}_{U_{ob}}(t)$, and $\mathbf{x}_{V_{ob}}(t)$ represents the temperature, salinity, and zonal and meridional velocities ($\tilde{\cdot}$ denotes the barotropic component and $\bar{\cdot}$ the baroclinic components) at the open boundaries, respectively.

4.2. Uncertainties, smoothness and dynamical constraints

The specification of the covariance matrices of the uncertainties on the control variables \mathbf{Q} is important for the performance of any 4D-VAR system as they determine the magnitude and the length scales of the adjustments to the control variables (Bonekamp et al., 2001). These matrices have to be non-diagonal to enforce smoothness in control vector adjustments. Lacking sufficient information to fully determine error covariances from observations, rough estimates are used. The correlations in \mathbf{Q} for the initial conditions and the forcing fields were, respectively, specified to be approximately two- and three-dimensional Gaussian functions. Because of the large number of elements, the diffusion operator approach (Weaver et al., 2003) is used to implement these non-diagonal Gaussian covariances in space. The horizontal and vertical covariance length scales are set to 7 km and 40 m, respectively. Adjustments at scales larger than the grid-scale are thus favored, but the grid-scale is small enough to capture fine structures that may be missing in the forcing fields (Bonekamp et al., 2001). Although these scales are expected to be inhomogeneous in the real ocean, especially in the vertical, constant values are used here for simplicity of interpretation and the lack of a good prior. The forcing errors may also be dominated by the diurnal peak, but we chose to assume no time correlation for the forcing errors for simplicity.

The uncertainties (i.e. square-root of variance in \mathbf{Q}) are set to 0.025 N m^{-2} for the wind stress, 60 W m^{-2} for the net heat flux, and $2 \times 10^{-8} \text{ m year}^{-1}$ for the net fresh water flux. For three-dimensional controls, the uncertainty is set to vary with depth: between 0.05° C at the bottom and 0.5° C at the surface for T , 0.01 PSU and 0.1 PSU for S , and 2 cm s^{-1} and 20 cm s^{-1} for the horizontal velocities U and V , respectively. These values have been chosen to obtain an appropriate relative weighting of the control and observation cost function terms.

Although only univariate correlations were modeled for the control variables in the present system, several constraints were added to the cost function in order to penalize dynamically unbalanced adjusted fields that may cause deterioration in the final state estimate. For instance, the time-mean vertical velocity over the assimilation period was constrained to be zero with a standard deviation of 10^{-5} m s^{-1} in order to reduce the spurious effect of model adjustments to changes in the initial temperature and salinity as suggested by [Stammer et al. \(2002\)](#). Similarly, deviations from the thermal wind balance on the boundary was penalized in a “soft constraint” approach to bring the boundaries closer to geostrophic balance as suggested by [Gebbie et al. \(2006\)](#). The penalty was zero over the first 40 m linearly increasing to a maximum of 0.01 m s^{-1} between layers at 100 m. The contribution of these two constraints to the total cost function was not significant. This does not mean that the constraints were irrelevant, because they may contribute to the gradient of the cost function, but separate contributions were not evaluated. The constraints are a standard part of the assimilation procedure and were retained because they provide an easy measure of plausibility of the solution (e.g. high cost would mean imbalance). The boundary velocities are required to have zero net mass flux.

The error covariance matrix for the observations \mathbf{R} combines estimates of errors in the measurements and errors due to processes unresolved by the model, mainly internal waves and residual tides. These are sometimes called representational errors and are more correctly modeled as errors in the dynamics ([Bennett, 2002](#)). There is no prior estimate for the spatial and temporal correlation of these errors between observations. In contrast to the control fields, on which smoothing operators can be applied because they are complete and regular, enforcing non-diagonal covariances for irregularly spaced observations is more complex. Although the errors are expected to vary in time and space, for simplicity, the errors in the observations of radial velocity are assumed to be uniform at 5 cm s^{-1} , and temporally and spatially uncorrelated so that \mathbf{R} is diagonal.

5. Evaluation of the assimilation system

To evaluate the performance of the assimilation system and to assess the relevance of 4D-VAR to dynamically map HF radial velocity data, two overlapping 10-day periods are considered – starting December 1, 2003 and December 5, 2003, respectively. This was the beginning of a period of good coverage from the three radars, and winter has less complicated stratification than summer. The results were compared over the overlapping period to check the continuity of the estimates and the sensitivity of the wind adjustments to the assimilation period. The state estimates were used in forecast experiments for further testing.

Single adjoint model integrations will first be examined, in which the adjoint model is integrated forced by HF radar data at the end of the 10-day period only. Sensitivity maps from the adjoint model provide information on the backward propagation in time of surface radial velocity information, tracking the influence of the ocean state and air-sea fluxes upon final surface velocities. The maps quantify the time evolution of the sensitivity of the observations to the model state, initial conditions, forcing, and open boundaries. The adjoint showed no fast-growing sensitivities over the 10-day intervals, so no extra damping was necessary in the optimization ([Köhl and Willebrand, 2002](#); [Hoteit et al., 2005b](#)).

In the assimilation experiments, the model was adjusted to reproduce the radial velocities measured hourly by the three radars on the San Diego coast by iteratively adjusting the model initial conditions, the forcing, and open boundaries every hour. After the assimilation, three purely forward model runs were carried out separately using the assimilation adjustments to the initial conditions, open boundaries, and forcing fields. These runs were used to study the contributions of the different types of control variables to the improvement in the data fitting. Several forecasting experiments were

then conducted starting from the estimated state by integrating the model beyond the assimilation window.

The evaluation of the assimilation system focused on (i) the cost function in space and time to evaluate the behavior of the 4D-VAR, (ii) the quality of the mapped surface currents to assess the accuracy of dynamically mapping radial velocity data, (iii) the reliability of the wind adjustments which provide information on the model dynamics and on the information content of the ocean current observations for the winds, and (iv) the impact of assimilating radial velocities on non-observed ocean variables, particularly at depth. The adjustments to the wind fields reflect the information in the ocean observations about the atmospheric state. They could be analyzed statistically to explore the possibility of improving the wind stress fields through the assimilation of HF radar data.

The total number of iterations used for the optimization of the cost function with the M1QN3 algorithm varies for each assimilation experiment, but does not exceed 30 iterations. The optimization was stopped when the rate of cost function decrease was relatively small (less than 1%) and the differences between the model and the observations were reduced to less than 1.5 times the guessed levels of expected root-mean-square misfit for the data. This was done to avoid iterating with low efficacy and to prevent excess structure in the state and control estimates since \mathbf{Q} and \mathbf{R} are not perfectly known. In all cases the misfits were reduced sufficiently before the cost function decrease slowed to the stopping level.

5.1. Adjoint sensitivities

One of the interesting questions raised by assimilation of surface velocities in a coastal region is how sensitivity spreads from the surface to the rest of the domain. The sensitivity propagates as free waves in addition to advection by existing flows. In this case the flows are slow, and internal waves, Rossby waves, and topographically trapped waves, such as coastally trapped waves (Brink, 1991), carry the information. Fig. 4 shows the gradient of the reference model cost function contribution from the current observations in the final hour of the 10-day window with respect to the temperature field going backward in time. The misfits between the radial speeds from the reference model run and the HF radar observations are immediately (after one backward time step) sensitive to the temperature over nearly the entire water depth. Subsurface density perturbations produce surface pressure gradients through the pressure solver, and the gradients accelerate surface currents according to the momentum equations, eventually reaching thermal balance at the equilibrium. As the integration of the adjoint model proceeds backward in time, the sensitivity spreads out horizontally away from the observed region and shows obvious interactions with the bottom, showing the form of both internal waves and bottom-intensified waves. The misfit between the observations and the reference model run is large-scale, so the cost function sensitivity approaches a regional average of velocity sensitivity. The sensitivity to zonal and meridional velocities look qualitatively similar and are not shown.

To further investigate the spatial structure of the spreading sensitivity, Fig. 5 shows the sensitivity of the same cost function (the radial speed misfits at the final hour (day 10)) to temperature at 250 m depth propagated backwards over 1 day. The initial sensitivity is located directly under the points with misfits in the radial velocities, and spreads to the west and south, while also becoming smoother as small-scale structures are dissipated. After slightly more than a day, significant sensitivity to the boundary conditions in the west and south is seen, as expected. There is also significant propagation along the coast, as expected from theory (Brink, 1991) and studies of remote forcing (Pringle and Riser, 2003). The sensitivity fields gradually damp out as time passes, and no large sensitivities due to nonlinearities are seen.

The adjoint shows that the surface data are sensitive to the three-dimensional structure of velocity and density, although there is significant ambiguity without more observations. For example, because there are no observations of temperature or salinity, the velocity observations only constrain density gradients, and the assimilation cannot distinguish between temperature and salinity perturbations. The temperature and salinity sensitivities are identical in form with opposite sign and are scaled by the ratio of the corresponding expansion coefficients. In addition, there is ambiguity between initial conditions and forcing that must be mitigated by specifying prior weights for the controls. Fig. 6

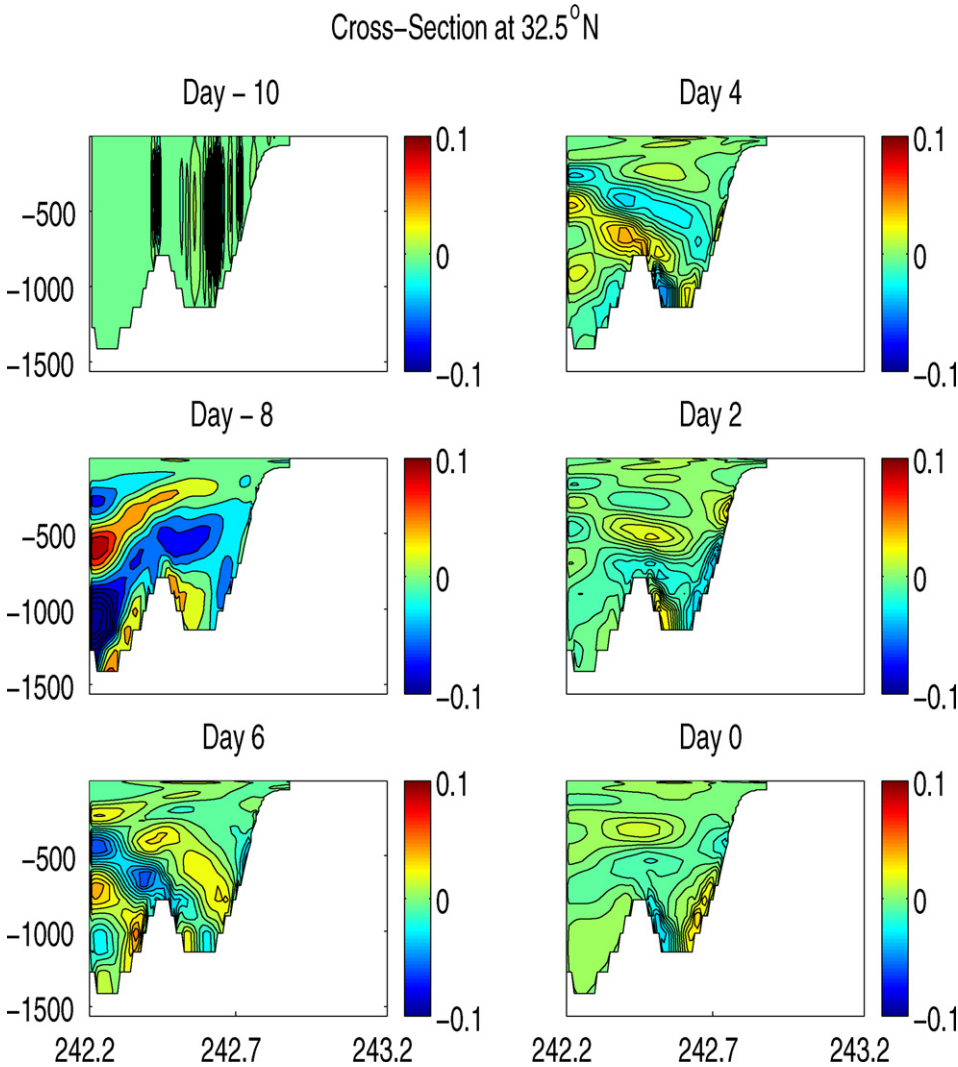


Fig. 4. Time evolution (backward) of the adjoint model (forced by misfits between the model and observations at the end of the assimilation window only) for a cross-section of the temperature at 32.5°N. Snapshots are shown every 2 days starting from the end of the first backward time step.

shows the time history of the sensitivity of the same cost function (data misfits on the last hour of the assimilation period) to the wind forcing along a meridional section at 242.5°E, showing the inertial waves which dominate the response to the wind at short timescales but damp out after a few days. The local inertial period is 22.3 h at 32.5°N. The sensitivity is large and positive at the time of the measurements, then oscillates with the inertial frequency and decays as the time lag increases. In other words, for similar effects on the final hour misfits, wind stress one-half inertial period before the observations should be in the opposite direction to wind at the time of the observations or an inertial period before due to the rotation of the driven currents. Wind 3 days or more before the observations has little effect.




Fig. 5. Time evolution (backward) of the adjoint model (forced by misfits between the model and observations at the end of the assimilation window only) for the temperature at 250 m depth. Snapshots are shown every 5 h starting from the end of the first backward time step and ending at 24 h.

5.1.1. Cost function and data fitting

The assimilation was performed in 10-day windows, the first starting on December 1, 2003. The total cost functions vs. optimization iteration count is shown in Fig. 7 for the first assimilation window. The decrease in the total cost function is greatest during the early iterations and slows nearly monotonically with iteration to less than 1% at iteration 22 where the iteration is terminated. The total cost function decreases from 11.5×10^5 to about 4.1×10^5 (a 65% decrease) representing a 70% decrease in the squared weighted misfit with observations, since the contribution of the control variables to the final cost is about 5% of the total.

A histogram with bars representing the contribution of the individual cost terms for each of the control variables terms to the final cost function is shown in Fig. 8. Large bars mean large adjustments to the control variable. The control term costs are zero at the start of the assimilation, by construction, and the initial vertical velocity drift term and the deviations from geostrophic balance at the open




Fig. 6. Time evolution (backward) of the adjoint model (forced by misfits between the model and observations at the end of the assimilation window only) for a zonal section of the zonal and meridional wind stress at 242.5°E.

boundaries are non-zero but less than 0.1% at the start and end of the iteration. The most important adjustments are made to the winds. This is expected (as shown below) because wind is the main source of energy for surface velocities, and because of the weak sensitivity of the surface currents a few days later to the initial conditions. Cost terms for the adjustments of the open boundaries are tiny, but the time series of their cost for the final iteration (not shown) shows important changes, suggesting more




Fig. 7. Decrease of the total cost function with iterations of the optimization. Note that the lower limit of the cost function axes is not zero.




Fig. 8. Individual cost function contribution for each control term after 22 optimization iterations.

contributions during certain events. The controls remain less than the specified variance at all times. Adjustments to the heat and fresh water fluxes are not significant and remain much smaller than the specified variance, confirming weak contributions to the decrease in the total cost function.

The contributions to the observation cost function in 13×15 km boxes summed over the entire interval are shown in Fig. 9 for several combinations of adjusted and reference control parameters. The values were divided by the cost function terms before assimilation, so that a value of 1 in a given box would mean that the assimilation did not improve the model fit to the data compared to the reference run. Plots are shown for the adjusted controls, and for three model runs using only adjustments to the (i) initial conditions, (ii) open boundaries, and (iii) forcing fields. The model–data misfit is clearly reduced over most of the domain using the adjusted controls. The experiment with only forcing adjustments seems to mainly improve the model–data agreement in the center of the domain, with less impact near the coast and the open boundaries. The wind adjustments provide about 65% of the improvement of the model fit to the data. The contribution of the initial conditions to the improvement is less significant and is estimated at about 17% of the overall decrease in the total cost function. Its distribution is uniform over the entire domain. The contribution of the open boundaries is about 20%, slightly larger than that of the initial conditions, with most influence coming from the western boundary. This is contrary to expectations that coastally trapped waves should propagate in from the south delivering the effects of remote forcing, but it may be due to the small time range covered. Because the separate cost function contributions of the different controls approximately sum to the total decrease suggests that the model response to each of type of control is approximately linear within the assimilation window.

The data cost term is plotted in Fig. 10 as a function of time for different iterations (left panel) and for the model runs but using the control adjustments separately as described in the previous section (right panel). The terms were normalized by the number of variables in the sum, so a value of 1 means that the data misfit variance matched the specified observational errors. Because it is difficult to know a priori what misfits are expected between the observations and model, the plausible error level of 5




Fig. 9. Spatial distribution of the data cost term in the cost function over 72 boxes as obtained from the model run using the adjusted controls after 22 optimization iterations (top left) and from three model runs, respectively, only using the control adjustments from the assimilation of (i) initial conditions, (ii) open boundaries, and (iii) forcing fields. The plotted values were normalized by the cost term from the reference run, so a value of less than 1 indicates that the assimilation improves the model fit with the data. Zero values indicate regions with no observations.




Fig. 10. *Left panel:* Time evolution of the radar data cost function term for the reference run and for assimilated runs after 5, 10 and 22 iterations. *Right panel:* Time evolution of the radar data cost function term from the assimilated run after 22 iterations compared to model runs with only a single type of control. Curves represent runs using control adjustments for (i) initial conditions, (ii) open boundaries, and (iii) forcing. The terms were normalized by the number of observations, so a value of 1 indicates that the model fits the data within the specified uncertainties.

cm s^{-1} was underestimated in retrospect. The unreliability of the prior error estimates means that the Chi-squared test is not applicable to judge the fit as suggested by Bennett (2002).

In the left panel, the model/data misfit is significantly improved over the entire time window after the first few iterations, suggesting good controllability of the system. The reference model error variance shows quasi-semi-diurnal peaks. These are partially due to diurnal variability in the observations, which have only had narrow tidal lines removed, and partially due to the diurnal and semi-diurnal wind forcing of the reference run, which showed similar quasi-periodic variance peaks. Tidal lines were not removed from the winds, which are not expected to have large tidally driven components. The assimilation greatly reduces these misfit variance peaks. Later iterations have reduced impact, but still contribute nearly equally over the whole window. Overall, the assimilation brings the root-mean-square (rms) model–data misfit to about 6 cm s^{-1} over the entire assimilation window, about 1.2 times the guessed observational error of 5 cm s^{-1} , and consistent with other estimates of errors (Kim et al., 2008). The right panel of Fig. 10 shows that the initial conditions affect the ocean state in the early part of the assimilation period with almost no contribution after approximately 18 h. The adjustments to the forcing fields alone show control of the surface velocities by the wind over the entire assimilation period. Adjustments to the open boundaries seem to only contribute during certain events and show an impact mainly after about 100 h suggesting a time lag for the signal to propagate into the interior where observations exist.

5.1.2. Estimated solution

Fig. 11 shows the estimated surface current and SSH fields at 2-day intervals during the first assimilation period. There is partial visual correlation between surface height gradients and velocity that will be explored in more detail. The fields and the wind stress evolve relatively slowly since the snapshots sampled the same phase of the diurnal and inertial cycles. The wind adjustments are more concentrated over the observed area and are weaker toward the end of the assimilation period because the late wind adjustments can only affect relatively few observations before the assimilation window cut off.

A depth-longitude cross-section of the estimated state at 32.5° E is shown in Fig. 12 for the temperature, salinity and horizontal velocities at December 6, 2003. The flow is mostly in geostrophic balance and shows a baroclinic structure that is visible also in temperature but appears more pronounced in the velocities. The estimated velocities show decreasing length scales, both vertically and horizontally, near the coast where boundary influences become more important.

5.2. Validation

The cost function discussion above focused on the reproduction of the observed radial surface currents. Assimilated model output may reproduce the observations but show distorted dynamical balances due to weighting choices or model error. To investigate this, several checks on the assimilated result were performed. First, state and control estimates at the same time from different but overlapping assimilation intervals were compared for consistency. Second, the momentum balances in the adjusted runs were investigated for consistency with expected coastal dynamics. Finally, prediction experiments were performed to test the stability of the model state evolution and to cross-validate with independent (future) observations.

5.2.1. Sensitivity to the assimilation period

A comparison of surface velocity, SSH, and adjusted winds is shown in Fig. 13 for December 7 from two assimilation experiments, started on December 1 (top) and 6 (bottom). The estimates from the two periods are qualitatively similar, providing reassurance that the solution is not greatly changed by changing the assimilation interval. Some difference between the two is expected due to the ambiguity of estimating a three-dimensional time-evolving field from surface velocity data only. Although the surface velocity fields are constrained to be similar by the observations, the SSH and forcing are diagnostic of the momentum balance, and are subject to the ambiguity. An ensemble of assimilation solutions from different intervals and with differing (within uncertainty) weights could be used to estimate error bars.

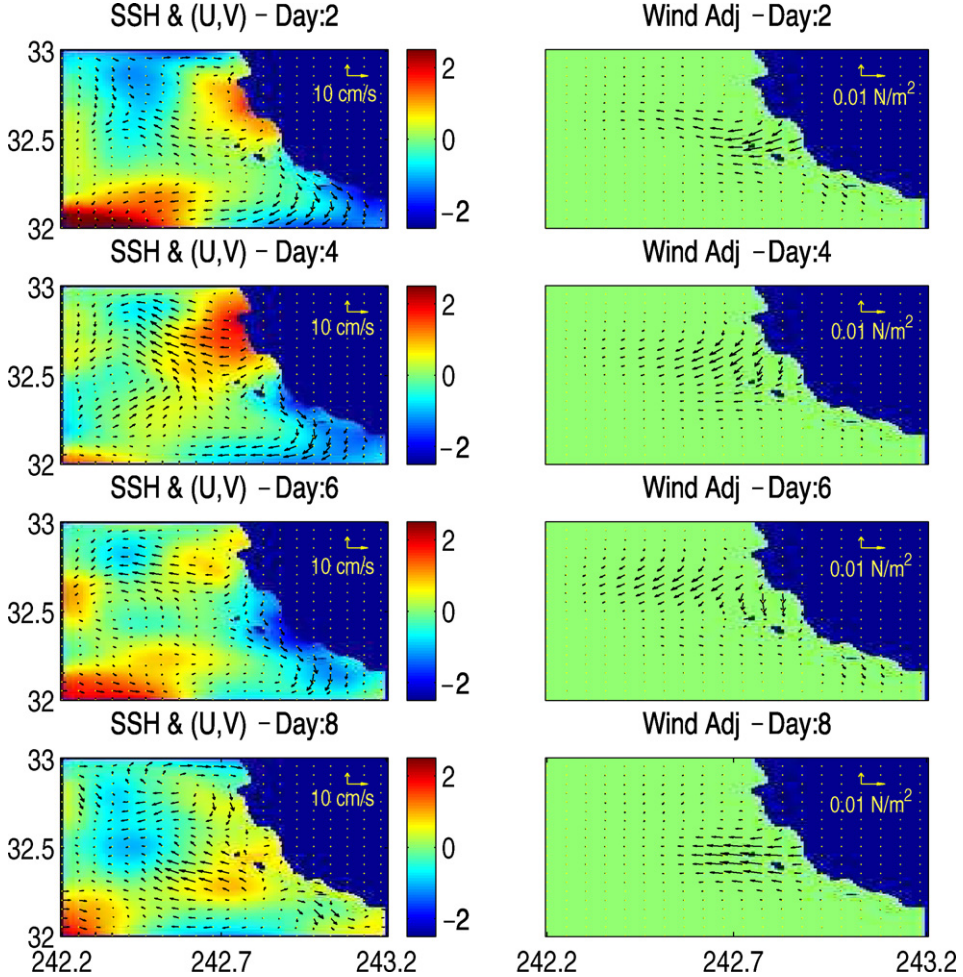


Fig. 11. Left column: Sea surface height (cm) and surface velocities (cm s^{-1}) estimated by the assimilation for days 2, 4, 6, and 8. Right column: Adjustments to the winds from the assimilation (N m^{-2}) on the same days.

5.2.2. Momentum balance

An example of daily mean momentum balances in the adjusted runs is shown in Fig. 14 at the surface and 20 m depth. The deterministic terms in the momentum equation are calculated from the 1-day mean model fields in the center (day 6) of the assimilation window, with the difference between the Coriolis term and the right-hand side (RHS) of the equation due primarily to the frictional and turbulent stress divergences. Showing deterministic terms only:

$$-fv = -\frac{\partial u}{\partial t} - \mathbf{u} \cdot \nabla u + \frac{\partial \phi}{\partial x} + \frac{\partial}{\partial z} \left(K_v \frac{\partial u}{\partial z} \right) = Rhs_x + \frac{\partial}{\partial z} \left(K_v \frac{\partial u}{\partial z} \right) \quad (6)$$

and

$$fu = -\frac{\partial v}{\partial t} - \mathbf{u} \cdot \nabla v - \frac{\partial \phi}{\partial y} + \frac{\partial}{\partial z} \left(K_v \frac{\partial v}{\partial z} \right) = Rhs_y + \frac{\partial}{\partial z} \left(K_v \frac{\partial v}{\partial z} \right) \quad (7)$$

where ϕ is the pressure divided by density, \mathbf{u} is the 3D velocity, f is the Coriolis parameter, and K_v is the vertical viscosity parameter. The momentum flux divergence due to friction is the difference




Fig. 12. Depth-longitude cross-section at 32.5° N of the assimilation solution for temperature, salinity, zonal and meridional velocities on December 6, 2003.

between the stress at the top and bottom of each layer. For the surface layer, the stress at the top is the wind stress while the stress at the bottom of the top layer has a deterministic component proportional to the velocity shear (set by the viscosity parameter via KPP) and stochastic component from stresses due to turbulence resolved in the model, which is expected to be small in the 1-day average.

Selected zonal (left) and meridional (right) acceleration terms (m s^{-2}) for the 1-day mean of the adjusted run are plotted in Fig. 14 as a function of latitude at a longitude 242.4° E in the middle of the model domain. The horizontal viscosity terms are negligible in these balances. The time derivatives (cyan curves) are important even in the average, while the nonlinear terms (advection of velocity gradients) are less than 2% of the balance and were not plotted separately. The mismatch between the Coriolis force (green line in Fig. 14) and the total of all non-friction terms on the right-hand side (red line in figure) are mainly due to viscosity and Reynolds stresses. So the difference between the green and red curves shows the viscous forcing terms. The distance between the red and blue shows the size of the ageostrophic component. Both surface and 23 m budgets show a geostrophic balance to within about 20%, but the frictional fluxes are more important at the surface than at depth where the red (RHS) and green (Coriolis) curves are much closer together. This illustrates the fall-off of wind-derived viscous forcing with depth. The momentum balance is in agreement with expectations from other analysis (e.g. Oke et al., 2002), although the strength of the wind driving is lower. The balances serve both as an example of the diagnoses possible in a dynamically consistent analysis and as a check on the state estimate.




Fig. 13. Left column: Sea surface height (cm) and surface velocities (cm s^{-1}) on December 7, 2003 from two 10-day assimilation experiments starting on December 1, 2003 and December 6, 2003, respectively. Right column: The adjustments to the winds (N m^{-2}) from the same experiments. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

5.2.3. Prediction experiments

Forecasts from the adjusted model state provide cross-validation of the state estimate by comparing with observations not used in the assimilation. Given enough control parameters, a model with incorrect dynamics can be made to fit observations, but the forecast from the adjusted model state will reveal the model error. A related problem is over-fitting of the observations by allowing strong and/or small-scale control adjustments. In the case of coastal radar networks, forecast skill shows the ability of the model dynamics to correctly fill gaps in the observations and to reject noise and artifacts from signal processing.

Four 1-day pseudo-forecast experiments were run at days 6, 7, and 11 of the first 10-day assimilation window and for day 11 of the second assimilation period. The model was started from the adjusted state from the assimilation and run forward with the original forcing and boundary conditions. Because the observed winds are used (although from single location), these are not true forecasts, but they serve to evaluate the skill of the estimated initial conditions. Similarly, the top two panels are “controls” showing free runs inside the assimilation window, so that the model state was partly determined by the later observations. The errors in these cases are due to the incomplete convergence of the assimilation

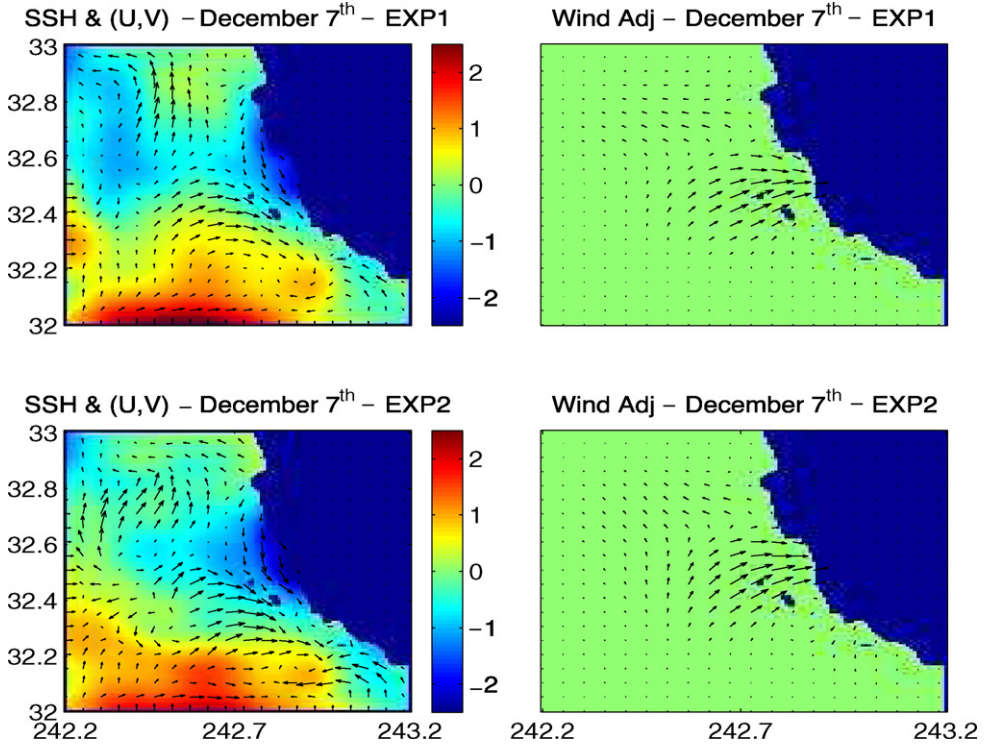


Fig. 14. Selected zonal (left) and meridional (right) momentum equation terms (ms^{-2}) for the model averaged over day 6 in the first assimilation experiment at the surface (top) and at 23 m depth (bottom). The curves are for a meridional section at longitude 242.4° E. The blue curve is the dynamic pressure gradient, the green curve is Coriolis force, the red is the sum of the non-friction terms balancing Coriolis force, and cyan is the negative of the time derivative of velocity. All terms have been multiplied by 10^7 to adjust the axes. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

and to the boundary and forcing errors and so provide a reference for the forecasts outside the assimilation windows. The goal is to highlight the forecast errors caused by errors in the starting model state. The estimated surface radial velocities were differenced from the observations and the square-root of the weighted sum of squared misfits (the cost function) was computed for the dynamical forecast. For comparison, the cost function was also computed for the reference model run over the same period, for a persistence forecast where the model state was held constant, and for a “climatology” forecast where the forecast velocities were zero. Fig. 15 shows the time evolution of the cost function over a day for the four experiments. Inside the assimilation window (top), the persistence forecast becomes worse than the zero and reference runs after about 5 h, while the error in the dynamical forecast from the assimilated state remains good for 20 h, which compares with the time of influence of the initial condition shown in Fig. 10. This prediction skill shows the effects of the subsurface structure since the low skill of the reference run is a reflection of the insufficient forcing information. The lower panels show comparisons to independent observations. The lower left panel is not very conclusive, since although the dynamical forecast outperforms persistence for hours 10 through 16, the total signal is low. The lower right panel shows a short-term jump in the error of the dynamical forecast but it eventually beats persistence. The larger errors in the lower right panel reflect the less complete convergence of the second state estimate (19 iterations). Some of the misfits are due to unpredictable wind changes, as highlighted by the top panels, and an average over many more forecast experiments would be needed to make a reliable estimate of the prediction skill. These figures do suggest that the model adds skill in reproducing independent observations, i.e. to cross-validate the estimated state. The 10-day assim-

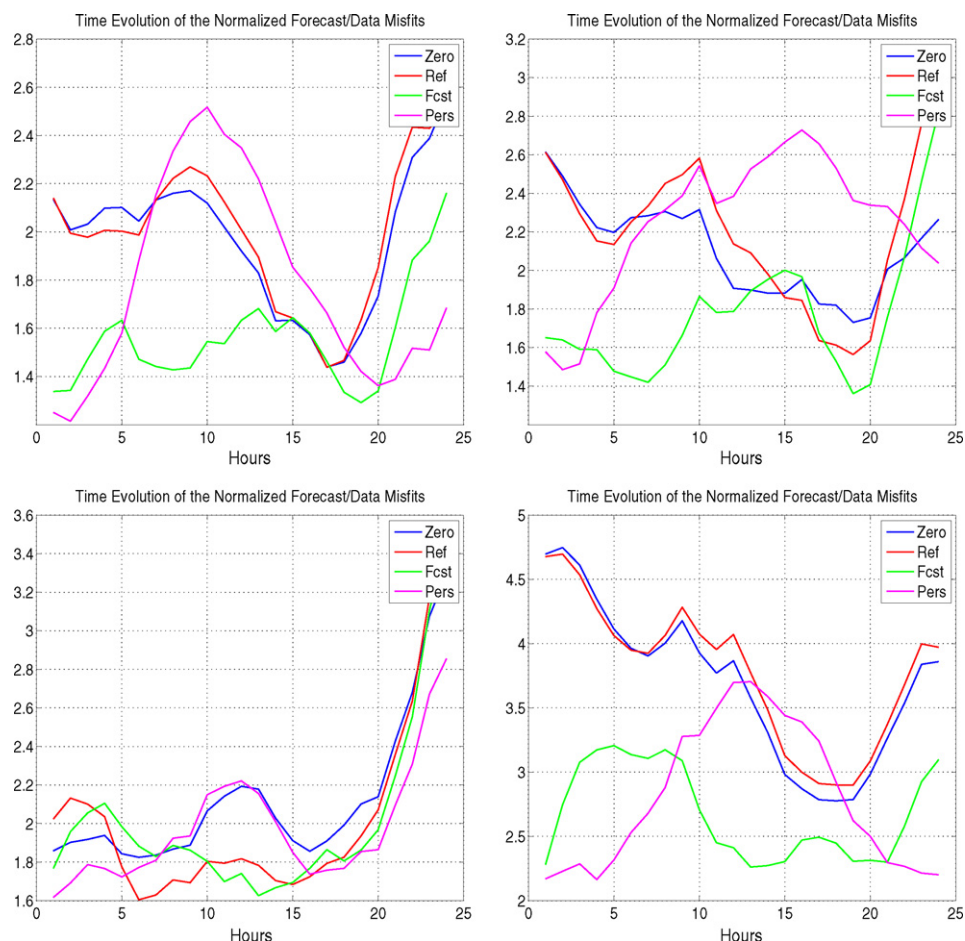


Fig. 15. Time evolution of the square-root of the normalized cost terms for four pseudo-forecast experiments, each over a 1-day period. Error for a forecast of zero is plotted in blue, error for the reference model is in red, error for a model forecast initialized from the assimilation solution is in green, and error for a forecast using the same assimilation solution but without time evolution (persistence) is in purple. The top two panels show integrations starting from the assimilated state on on December 6, 2003 (left) and December 7, 2003 (right) from the first assimilation run. The two bottom panels show integrations starting from the state estimate for December 11, 2003 from the first assimilation run (left) and from the state estimate for December 16, 2003 from the second assimilation run (right).

ilation window was chosen to try to maximize the length of the dynamically consistent periods, not for forecasting. Further experiments are needed to explore forecast performance for varying length of assimilation window. The influence of the observations seems to be limited to a few days, so a shorter assimilation window may allow a quicker and more complete fit to observations and perhaps better forecasts.

6. Discussion

A coastal 4D-VAR assimilation system was implemented to assess the possibility of dynamically mapping high-frequency radar surface radial velocity data from three sites on the coast of San Diego. The assimilation system is inherited from the Circulation and the Climate of the Ocean (ECCO) sys-

tem which uses the adjoint method to constrain the MITgcm model to available data sets over the global ocean. The present study is a first application of this system to a coastal context focusing on scales of a few kilometers. Assimilation experiments are performed over 10-day periods. An advantage of the MITgcm-ECCO system is to allow for estimation of not only initial conditions but also lateral boundary conditions and air-sea fluxes, which are expected to be important in coastal regions.

This study shows the practicality of radar velocity assimilation using the MITgcm-ECCO system. The model solutions can be closely fitted to the observations over 10 days, largely through adjusting wind stress controls. The model runs with adjusted controls provide a dynamical interpretation of the observations. It was also illustrated that dynamically mapping observations can yield skillful forecasts of surface ocean currents. The predicted surface currents showed skill over persistence for about 20 h.

The experiments presented here emphasize the controllability of fluctuations in surface currents near the coast by wind forcing. In the experimental setup, the first-guess solution was initialized with only a single temperature and salinity profile and forced by spatially constant wind fields extrapolated from a single shore station to the entire model domain. This first-guess solution showed weak variability and low skill in reproducing the observations. The variability after assimilation thus results almost entirely from observational constraints. To understand the information content of observations, their sensitivity to model state was propagated backward in time using the adjoint model and examined in detail. It became clear that wind stress adjustments were the main control in this system. The effects of all controls, (either initial conditions, boundary conditions, or air-sea fluxes) become small after about two inertial periods (i.e. days). However wind stress can be adjusted throughout a 10-day period, allowing an efficient control of the full model evolution.

A study of the underlying momentum balance shows that multiple contributions matter at the scales we consider – acceleration, pressure gradient, Coriolis force, and momentum fluxes from deterministic and statistical mixing. The solution is thus sensitive to the mixing parameterizations, which were not varied in these experiments. The assimilation of other observations, particularly subsurface temperature and salinity profiles is necessary to determine the complete model state. Surface radial velocities alone are not sufficient to completely constrain the ocean evolution, e.g. leaving ambiguity between temperature and salinity. Heat and fresh water balances over the coastal region are left for further investigation, pending inclusion of such observations.

The results from our experiments suggest that the ECCO-MITgcm system is feasible for small-scale and high-frequency applications over coastal areas. The limited set of experiments here only achieve basic tests of the assimilation system, and much remains to be done. The data base ought to be complemented with other types of observations as argued above. Extending the assimilation window beyond 10 days might also help improve some aspects of the results, which would then satisfy longer-term dynamical balances. The model dynamics may also be significantly improved compared with our test configuration. In particular, these experiments did not include tides, which is an important limitation for coastal applications. Starting from a system resolving wind effects on the circulation, including a tide model would open additional perspectives. The resulting analysis of radar observations should help to understand the trade-off between tides and wind generation of coastal fluctuations. Also, a hierarchy of nested model domains could ultimately be needed to better represent along-shore pressure gradients and the effects of remote forcing.

Further work is also needed to adjust the weights and balance the various controls, which ought to improve the dynamics of the assimilation results. This remains a challenge because errors in initial conditions, etc. are poorly known. The degree of uncertainty in wind products is especially questionable. In analyzing the assimilation results, the cost function reduction is smaller during some events with strong winds. This is presumably because the size of the adjustments to the wind are limited by the wind uncertainty range, which was set to be constant in the cost function. Choosing larger but constant errors for the winds would mean that the wind adjustments would dominate the rest of the adjustments even during periods of weak winds. This suggests the use of time-varying errors for the wind forcing which could be estimated from the observed wind statistics or might be inferred from the radar measurements. The lack of uncertainty estimates after assimilation remains an outstanding issue, which could be addressed in practice using an ensemble approach.

Acknowledgments

This work was supported by NOAA and the State of California as part of the Southern California Coastal Ocean Observing System. Computer resources came from the COMPAS facility at SIO and from the TAG-Cobalt machine through Teragrid project grant number TAG-ICE060003N.

References

- Alvera-Azcarate, A., Barth, A., Rixen, M., Beckers, J., 2005. Reconstruction of incomplete oceanographic data sets using empirical orthogonal functions: application to the Adriatic Sea surface temperature. *Ocean Model.* 9, 325–346.
- Beckers, J.M., Rixen, M., 2003. EOF calculations and data filling from incomplete oceanographic datasets. *J. Atmos. Ocean. Technol.* 20, 1839–1856.
- Bennett, A.F., 2002. *Inverse modeling of the ocean and atmosphere*. Cambridge University Press, 256 pp.
- Bonekamp, H., van Oldenborgh, G.J., Burgers, G., 2001. Variational assimilation of TAO and XBT data in the HOPE OGCM, adjusting the surface fluxes in the tropical ocean. *J. Geophys. Res.* 106, 16693–16709.
- Breivik, Ø., Stra, Ø., 2001. Real time assimilation of HF radar currents into a coastal ocean model. *J. Mar. Syst.* 28, 161–182.
- Brink, K.H., 1991. Coastal-trapped waves and wind-driven currents over the continental shelf. *Ann. Rev. Fluid Mech.* 23, 389–412.
- Davis, R.E., 1985. Objective mapping by least squares fitting. *J. Geophys. Res.* 90 (C7), 4773–4778.
- De Mey, P., Benkiran, M., 2002. A multivariate reduced-order optimal interpolation method and its application to the Mediterranean basin-scale circulation. In: *Ocean Forecasting, Conceptual Basis and Applications*. Springer-Verlag, Berlin/Heidelberg/New York, 472 pp.
- De Paolo, T., Terrill, E., 2007. Skill assessment of resolving ocean surface current structure using compact-antenna-style HF radar and the MUSIC direction-finding algorithm. *J. Atmos. Ocean. Technol.* 24, 1277–1300.
- Di Lorenzo, E., Moore, A.M., Arango, H.G., Cornuelle, B.D., Miller, A.J., Powell, B.S., Chua, B.S., Bennett, A.F., 2007. Weak and strong constraint data assimilation in the inverse Regional Ocean Modeling System (ROMS): development and application for a baroclinic coastal upwelling system. *Ocean Model.* 16, 160–187.
- Fukumori, I., Raghunath, R., Fu, L., Chao, Y., 1999. Assimilation of TOPEX/POSEIDON data into a global ocean circulation model: how good are the results? *J. Geophys. Res.* 104, 25647–25665.
- Gebbie, G., Heimbach, P., Wunsch, C., 2006. Strategies for nested and eddy-permitting state estimation. *J. Geophys. Res.* 111, C10073.
- Ghil, M., Malanotte-Rizzoli, P., 1991. Data assimilation in meteorology and oceanography. *Adv. Geophys.* 33, 141–266.
- Giering, R., Thomas, K., 1998. Recipes for adjoint code construction. *ACME Trans. Math. Software* 24, 437–474.
- Gilbert, J.C., Le Maréchal, C., 1989. Some numerical experiments with variable storage Quasi-Newton algorithms. *Math. Program.* 45, 407–435.
- Hoteit, I., Korres, G., Triantafyllou, G., 2005a. Comparison of extended and ensemble based Kalman filters with low and high resolution primitive equation ocean models. *Nonlinear Process. Geophys.* 12, 755–765.
- Hoteit, I., Cornuelle, B., Köhl, A., Stammer, D., 2005b. Treating strong adjoint sensitivities in tropical eddy-permitting variational data assimilation. *Q. J. R. Meteor. Soc.* 131, 3659–3682.
- Hoteit, I., Cornuelle, B., Heimbach, P., 2008. An eddy-permitting dynamically consistent hindcast of the tropical Pacific in 2000 using an adjoint based assimilation system. *J. Geophys. Res.*
- Kaplan, D.M., Lekein, F., 2007. Spatial interpolation and filtering of surface current data based on open-boundary modal analysis. *J. Geophys. Res.* 112, C12007.
- Kim, S.Y., Terrill, E.J., Cornuelle, B.D., 2007. Objectively mapping HF radar-derived surface current data using measured and idealized data covariance matrices. *J. Geophys. Res.* 112, C06021.
- Kim, S.Y., Terrill, E.J., Cornuelle, B.D., 2008. Mapping surface currents from HF radar radial velocity measurements using optimal interpolation. *J. Geophys. Res.* 113, C10023, doi:10.1029/2007JC004244.
- Köhl, A., Willebrand, J., 2002. An adjoint method for the assimilation of statistical characteristics into eddy-resolving ocean models. *Tellus* 54 (4), 406–425.
- Korres, G., Hoteit, I., Triantafyllou, G., 2007. Data assimilation into a Princeton Ocean Model of the Mediterranean Sea using advanced Kalman filters. *J. Mar. Syst.* 65, 84–104.
- Kurapov, A.L., Allen, J.S., Egbert, G.D., Miller, R.N., Kosro, P.M., Levine, M., Boyd, T., 2005. Distant effect of assimilation of moored currents into a model of coastal wind-driven circulation off Oregon. *J. Geophys. Res.* 110, C02022.
- Kurapov, A., Egbert, G., Allen, J., Miller, R., Erofeeva, S., Kosro, P., 2003. The M2 Internal Tide off Oregon: Inferences from Data Assimilation. *J. Phys. Oceanogr.* 33, 1733–1757.
- Large, W., Williams, J., Doney, S., 1994. Oceanic vertical mixing: A review and a model with a non local boundary layer parametrization. *Rev. Geophys.* 32, 363–403.
- Lea, D., Haine, T., Gasparovic, R., 2006. Observability of the Irminger Sea circulation using variational data assimilation. *Quart. J. Roy. Meteor. Soc.* 132, 1545–1576.
- Le Dimet, F.X., Talagrand, O., 1986. Variational algorithms for analysis and assimilation of meteorological observations: Theoretical aspects. *Tellus* 38, 97–110.
- Li, Z., Chao, Y., McWilliams, J.C., Ide, K., 2008. A three-dimensional variational data assimilation scheme for the Regional Ocean Modeling System: Implementation and basic experiments. *J. Geophys. Res.* 113 (C2), C05002.
- Lipa, B.J., Barrick, D.E., 1983. Least-squares methods for the extraction of surface currents from CODAR crossed-loop data: Application at ARSLOE. *IEEE J. Ocean. Eng.* 13 (2), 507–513.
- Lipphardt, B.L., Kirwan, A.D., Grosch, C.E., Lewis, J.K., Paduan Jr., J.D., 2000. Blending HF radar and model velocities in Monterey Bay through normal mode analysis. *J. Geophys. Res.* 105 (C2), 3425–3450.
- Marshall, J., Adcroft, A., Hill, C., Perelman, L., Hersey, C., 1997. A finite-volume, incompressible Navier Stokes model for studies of the ocean on parallel computers. *J. Geophys. Res.* 102, 5753–5766.

- Oke, P.R., Allen, J.S., Miller, R.N., Egbert, G.D., Kosro, P.M., 2002. Assimilation of surface velocity data into a primitive equation coastal ocean model. *J. Geophys. Res.* 107 (C9), 3122.
- Oke, P.R., Allen, J.S., Miller, R.N., Egbert, G.D., 2002. A Modeling Study of the Three-Dimensional Continental Shelf Circulation off Oregon. Part I. Dynamical Analysis. *J. Phys. Oceanogr.* 32, 1383–1403.
- Paduan, J.D., Shulman, I., 2004. HF radar data assimilation in the Monterey Bay area. *J. Geophys. Res.* 109, C07S09.
- Pringle, J.M., Riser, K., 2003. Remotely forced nearshore upwelling in Southern California. *J. Geophys. Res.* 108, 3131.
- Stammer, D., Wunsch, C., Giering, R., Eckert, C., Heimbach, P., Marotzke, J., Adcroft, A., Hill, C.N., Marshall, J., 2002. The global ocean circulation during 1992–1997, estimated from ocean observations and a general circulation model. *J. Geophys. Res.* 107 (C9), 3118.
- Taillandier, V., Echevin, V., Mortier, L., Devenon, J.L., 2004. Controlling boundary conditions with a four-dimensional variational data-assimilation method in a non-stratified open coastal model. *Ocean Dyn.* 54, 284–298.
- Weaver, A.T., Vialard, J., Anderson, D.L.T., 2003. Three- and four-dimensional variational assimilation with a general circulation model of the tropical Pacific Ocean. Part I. Formulation, internal diagnostics, and consistency checks. *Mon. Weather Rev.* 131, 1360–1378.
- Wilkin, J.L., Arango, H.G., Haidvogel, D.B., Lichtenwalner, C.S., Glenn, S.M., Hedström, K.S., 2005. A Regional Ocean Modeling System for the long-term ecosystem observatory. *J. Geophys. Res.* 110, C06S91.
- Wunsch, C., 1996. The ocean circulation inverse problem. Cambridge University Press, New York, 442 pp.
- Zhang, K., Marotzke, J., 1999. The importance of open boundary estimation for an Indian ocean GCM data synthesis. *J. Mar. Res.* 57, 305–334.
- Zhang, F., Meng, Z., Aksoy, A., 2006. Tests of an ensemble Kalman filter for mesoscale and regional-scale data assimilation. Part I. Perfect model experiments. *Mon. Weather Rev.* 134, 722–736.
- Zupanski, M., 1996. A preconditioning algorithm for four-dimensional variational data assimilation. *Mon. Weather Rev.* 124, 2562–2573.