

Data assimilation in the FOAM operational short-range ocean forecasting system: a description of the scheme and its impact

M. J. Martin,* A. Hines and M. J. Bell
Met Office, Exeter, UK

ABSTRACT: A detailed description of the data assimilation scheme used in the Forecasting Ocean Assimilation Model (FOAM) operational ocean forecasting system is presented. The theoretical basis for the scheme is an improved version of the analysis correction scheme, which includes information from previously assimilated data. The scheme requires the a priori specification of error covariance information for the background model field and the observations. The way in which these error covariances have been estimated is described and some examples are given. The FOAM system assimilates sea surface temperature, sea-level anomaly, temperature profile, salinity profile and sea-ice concentration data. Aspects of the scheme that are specific to each of these observation types are described.

Two sets of experiments demonstrating the impact of the data assimilation are presented. The first set are in an idealized identical-twin setting, using the $\frac{1}{9}^\circ$ -resolution North Atlantic FOAM configuration in which the state of the true ocean is assumed to be known. These experiments show that the analyses and forecasts are improved by assimilating the altimeter sea-level-anomaly data. The second set of experiments comprise data impact studies in a realistic hindcast setting. These experiments show a positive impact on the analyses from the Argo temperature- and salinity-profile data. © Crown Copyright 2007. Reproduced with the permission of the Controller of HMSO. Published by John Wiley & Sons, Ltd

KEY WORDS altimeter data assimilation; Argo data; operational oceanography

Received 12 October 2006; Revised 19 February 2007; Accepted 28 February 2007

1. Introduction

The operational Forecasting Ocean Assimilation Model (FOAM) system run at the UK Met Office is used to predict properties of the open ocean (temperature, salinity and currents) and sea ice. It consists of a suite of nested model configurations, which produce daily analyses and forecasts, out to 5 days ahead, of the ocean and sea-ice variables. FOAM is the main UK input to the Global Ocean Data Assimilation Experiment (GODAE), which aims to provide a demonstration of the international ocean-forecasting community's ability to provide information about the current and future state of the ocean. FOAM is also contributing to the Marine Environment and Security for the European Area (MERSEA) project, which will provide monitoring and forecasting services to intermediate users and policy-makers.

The model component of FOAM is based on the ocean element of the model used for climate prediction at the Hadley Centre (Gordon *et al.*, 2000). It is a z -level model, and is described in detail in (Hines *et al.*, 2007). Configurations that are currently run operationally

include a global 1° -resolution model, a $\frac{1}{3}^\circ$ -resolution model of the North Atlantic and Arctic, and a $\frac{1}{9}^\circ$ -resolution model of the North Atlantic, each of which has 20 vertical levels. Data from the operational 1° - and $\frac{1}{9}^\circ$ -resolution models are made freely available for research use in near-real time from the Live Access Server at <http://www.nerc-essc.ac.uk/godiva>.

Model forecasts require the specification of initial conditions, and the accuracy of the forecast depends to a large extent on the accuracy of the initial conditions. In order to produce initial conditions that accurately reflect the current state of the ocean, a previous model forecast is combined with observations, using a data-assimilation scheme. The scheme used in the FOAM system has been significantly upgraded since the paper of Bell *et al.* (2000). The aim of this paper is to provide a detailed description of the current method used in the FOAM system, together with some examples of its impact.

The theoretical basis of the data-assimilation scheme used in FOAM is described in Section 2, which includes an overview of the method used to produce the spatial analysis, together with a description of the method used to include time information in the analysis. Section 3 describes the way in which the scheme has been implemented in the operational FOAM system. A description

*Correspondence to: M. J. Martin, Met Office, FitzRoy Road, Exeter, Devon EX1 3PB, UK. E-mail: matthew.martin@metoffice.gov.uk

of the error covariances used in the scheme, together with specific issues associated with each of the observation types that are assimilated, is given. The way in which the updates to the model are implemented, and issues relating to systematic errors, are also discussed. The impact of the data-assimilation scheme in both identical-twin experiments and real-world experiments is described in Section 4, together with some observational impact studies. Section 5 provides a summary and a discussion of planned upgrades to the system.

2. Theoretical framework of the data-assimilation scheme

A large number of methods for combining model and observational data are described in the literature. These can be broadly split into three approaches: variational methods, such as 3D-Var and 4D-Var (Lorenc, 1986), based on minimization of a cost function that measures the differences between the model and the observations; various levels of approximation to the extended Kalman filter, generally known as sequential schemes (Daley, 1991); and ensemble-based schemes such as the ensemble Kalman filter (Evensen, 1994). Each of these approaches has its advantages and disadvantages with respect to the level of approximations made, complexity and computational cost.

In an operational system such as FOAM, there are various constraints that influence the type of data assimilation scheme that can be used. The main constraints are the amount of time and computer power available to run the scheme, together with the requirement that the system be robust. Other issues that influence the type of scheme used include the nonlinear nature of the ocean at the high resolution used in FOAM, and the number of observations available for assimilation on a daily basis. All of these constraints, together with the fact that we only have a limited knowledge of the errors in the observations and model, have led to the development of a relatively simple methodology for use in the FOAM system. It is based on an 'optimal interpolation'-type scheme that approximates the Kalman filter (where the error covariance matrices are not evolved with time), as described in Section 2.1. Some aspects of more complicated schemes (such as 4D-Var) are also included in order to extract the maximum amount of information from the available observations and model (see Section 2.2).

2.1. Spatial analysis

The generalized optimal-interpolation analysis equation at time t_k can be written as:

$$\mathbf{x}_k^a = \mathbf{x}_k^f + \mathbf{B}\mathbf{H}^T(\mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R})^{-1}(\mathbf{y}_k - h(\mathbf{x}_k)), \quad (1)$$

where \mathbf{x}_k is a vector containing the state variables at each grid point, \mathbf{y}_k is a vector containing the observations, h is the observation operator (which converts variables from

model space to observation space), \mathbf{H} is the Jacobian of the observation operator, \mathbf{B} is the model forecast error covariance, and \mathbf{R} is the observation error covariance; superscript 'a' indicates an analysis and superscript 'f' indicates a forecast. The notation is based on that of (Ide *et al.*, 1997).

The assimilation scheme used in the FOAM system is based on the analysis correction (AC) method introduced by Lorenc *et al.* (1991). This method provides an efficient means of calculating the solution to Equation (1) using an iterative procedure. The iteration is initialized by setting $\mathbf{x}_k^0 = \mathbf{x}_k^f$ and $\mathbf{y}_k^0 = \mathbf{y}_k$, and the ν th iterate is obtained using the equations:

$$\left. \begin{aligned} \mathbf{x}_k^{\nu+1} &= \mathbf{x}_k^\nu + \mathbf{B}\mathbf{H}^T\mathbf{Q}(\mathbf{y}_k^\nu - h(\mathbf{x}_k^\nu)) \\ \mathbf{y}_k^{\nu+1} &= \mathbf{y}_k^\nu - \mathbf{R}\mathbf{Q}(\mathbf{y}_k^\nu - h(\mathbf{x}_k^\nu)) \end{aligned} \right\}, \quad (2)$$

for $\nu = 0, 1, \dots, N - 1$. The analysis is then set equal to the last iterate: $\mathbf{x}_k^a = \mathbf{x}_k^N$. The component of the matrix \mathbf{Q} in the i th column and j th row is defined as

$$\mathbf{Q}_{ij} = \delta_{ij} \left\{ \sum_{j=1}^I (\mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R})_{ij} \right\}^{-1}, \quad (3)$$

where

$$\delta_{ij} = \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases},$$

and I is the number of observations.

Bratseth (1986) and Lorenc *et al.* (1991) have shown that provided the \mathbf{x}_k^ν and \mathbf{y}_k^ν obtained using Equation (2) converge, they converge respectively to the analysis given by Equation (1) and to the 'observation' values $h(\mathbf{x}_k^a)$ corresponding to that analysis. A fixed number of 10 iterations are used in the FOAM system. This number was determined by performing some sensitivity experiments in which the number of iterations was varied and the results compared with independent *in situ* temperature data. Fewer iterations than this resulted in a less accurate analysis, and there was little gain in accuracy from increasing the number of iterations. However, the number of iterations will depend on various aspects, including the model resolution and the data density, and should be recalibrated when any of these aspects are changed.

The matrix multiplication in Equation (2) – $\mathbf{B}\mathbf{H}^T\mathbf{z}$ where $\mathbf{z} = \mathbf{Q}(\mathbf{y}_k^\nu - h(\mathbf{x}_k^\nu))$ – is the most costly aspect of this procedure. It has been shown by Lorenc (1992) that successive applications of a recursive filter on the values of \mathbf{z} (in model space) provide a good approximation to this matrix multiplication when Gaussian or second-order autoregressive (SOAR) functions are used to approximate the off-diagonal elements of \mathbf{B} . This application of the filter is used in FOAM to approximate the effect of SOAR correlation functions using the covariance information described in Section 3.1. The term $\mathbf{H}\mathbf{B}\mathbf{H}^T$ in Equation (3) is calculated by filtering the background error variances and then interpolating them to the observation locations using a bilinear interpolation.

Equation (2) is written as a fully three-dimensional analysis of the model state. In practice, the analysis is decomposed into separable horizontal and vertical parts. For *in situ* profile data, a vertical analysis is performed first, to produce a vertical profile of observational increments and their error estimates on model levels. These are then used in a set of horizontal analyses, one on each vertical level. For surface data types such as altimeter data, a horizontal analysis is performed first, and the vertical correlations are specified analytically, as described in Section 3.2. This decomposition has the disadvantage that the full three-dimensional correlations cannot be applied, although these are rarely known in practice.

2.2. Timeliness aspects

The above method is designed to work at one particular ‘analysis’ time at which the background model field is valid. The observations that have become available during the period since the last analysis are usually treated as if they are all valid at the analysis time, taking into account the time correlations of the observations (unlike in a ‘first-guess-at-appropriate-time’ scheme). However, the number of new observations that become available during one analysis period is small, and a scheme has been devised that allows the information from previous analysis cycles to be properly included in the current analysis. This timeliness aspect of the assimilation scheme is now described.

In order to take account of information from previous analyses, we first present the result that it is possible to process observations in batches using the *same* background error covariance matrix, and still obtain the optimal analysis: for two subsets of observations,

$$\left. \begin{aligned} \mathbf{y}_k &= \begin{pmatrix} \mathbf{y}_{1,k} \\ \mathbf{y}_{2,k} \end{pmatrix} \\ h(\mathbf{x}_k) &= \begin{pmatrix} h_1(\mathbf{x}_k) \\ h_2(\mathbf{x}_k) \end{pmatrix} \\ \mathbf{R} &= \begin{pmatrix} \mathbf{R}_1 & 0 \\ 0 & \mathbf{R}_2 \end{pmatrix} \end{aligned} \right\}, \quad (4)$$

it is possible to analyse each subset separately, with the second analysis using the first analysis instead of a background field, and the observational increments for the first analysis set to zero. The second analysis then produces the same values as would have been obtained if both sets of observations had been analysed together. That is, if the analysis is performed in two steps, both of which use the same background error covariance matrix –

$$\mathbf{x}_{1,k}^a = \mathbf{x}_k^f + \mathbf{B}\mathbf{H}_1^T(\mathbf{H}_1\mathbf{B}\mathbf{H}_1^T + \mathbf{R}_1)^{-1}(\mathbf{y}_{1,k} - h_1(\mathbf{x}_k^f)) \quad (5)$$

and

$$\mathbf{x}_{2,k}^a = \mathbf{x}_{1,k}^a + \mathbf{B}\mathbf{H}^T(\mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R})^{-1} \begin{pmatrix} 0 \\ \mathbf{y}_{2,k} - h_2(\mathbf{x}_{1,k}^a) \end{pmatrix} \quad (6)$$

– then provided there are no cross-correlations between the errors in the first set of observations and those in the second set, and that the observation operator is linear, the second analysis is identical to the analysis produced from Equation (1), that is,

$$\mathbf{x}_{2,k}^a \equiv \mathbf{x}_k^a \quad (7)$$

See Appendix A for a proof. This result applies to any number of subsets of observations. It also applies to observations taken over a period of time when the interpolation operator includes the evolution of the model state to the observation time (although the observation operator will not usually then be a linear operator). This is different from the method presented by Anderson and Moore (1979, pp. 142–146) and used by various people including Houtekamer and Mitchell (2001), in which the background error covariance is updated between batches of observations.

In order to include time dependence in the analysis without having to minimize the full 4D-Var penalty function, it would be possible to evolve the values and locations of the observations to the analysis time (although in practice the locations of the observations are not evolved). This idea can be expressed in the form of a penalty function,

$$\mathcal{J}(\mathbf{x}) = (\mathbf{x} - \mathbf{x}^f)^T \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}^f) + \sum_i (g_{k,i}(\mathbf{y}_i) - h_{k,i}(\mathbf{x}))^T \tilde{\mathbf{R}}_i^{-1}(g_{k,i}(\mathbf{y}_i) - h_{k,i}(\mathbf{x})), \quad (8)$$

where $g_{k,i}$ represents the observations made at time t_i whose values and locations have been evolved to the analysis time t_k , $h_{k,i}$ represents the interpolation of the model state to these evolved locations, and

$$\tilde{\mathbf{R}}_i = \overline{(g_{k,i}(\mathbf{y}_i) - h_{k,i}(\mathbf{x}))(g_{k,i}(\mathbf{y}_i) - h_{k,i}(\mathbf{x}))^T} \quad (9)$$

is the expected error covariance in the evolved observations at the analysis time (the overbar denoting the expectation operator).

The penalty function given in Equation (8) is not of much practical use as it stands, because it would be difficult to evolve the observations forward or backward to the analysis time sufficiently accurately, particularly if the observations are of active tracers. Moreover, the error associated with the evolution of the observation value and its location would be difficult to estimate. An alternative to minimizing Equation (8), which also makes use of the idea for sequentially updating analyses with new observations, is to derive a filter from Equation (8) and use the model M to evolve forward the best estimate

of the state while evolving the expected error covariances of previously used observations. This can be written as

$$\mathbf{x}_k^f = M(\mathbf{x}_{k-1}^a), \tag{10}$$

and

$$\mathbf{x}_k^a = \mathbf{x}_k^f + \mathbf{B}\tilde{\mathbf{H}}^T \left(\tilde{\mathbf{H}}\mathbf{B}\tilde{\mathbf{H}}^T + \tilde{\mathbf{R}} \right)^{-1} \Delta\tilde{\mathbf{y}}, \tag{11}$$

where

$$\Delta\tilde{\mathbf{y}} = \begin{pmatrix} g_{k,k-N}(\mathbf{y}_{k-N}) - h_{k,k-N}(\mathbf{x}_k^f) \\ g_{k,k-N+1}(\mathbf{y}_{k-N+1}) - h_{k,k-N+1}(\mathbf{x}_k^f) \\ \vdots \\ \mathbf{y}_k - h_{k,k}(\mathbf{x}_k^f) \end{pmatrix},$$

$$\tilde{\mathbf{R}} = \begin{pmatrix} \tilde{\mathbf{R}}_{k-N} & 0 & \cdots & 0 \\ 0 & \tilde{\mathbf{R}}_{k-N+1} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \mathbf{R}_k \end{pmatrix},$$

$$\tilde{\mathbf{H}} = \begin{pmatrix} \mathbf{H}_{k,k-N} \\ \mathbf{H}_{k,k-N+1} \\ \vdots \\ \mathbf{H}_{k,k} \end{pmatrix},$$

and N is the number of time periods retained in the analysis.

In order to properly implement the above scheme, the observation values and locations would have to be evolved from one analysis time to the next (using the operator $g_{k,k-i}$). Furthermore, the errors associated with each observation would have to include the errors associated with this evolution (as in Equation (9)). In the practical implementation used in FOAM, the model is used to evolve the background field at the location where the observation reports, and the location is not evolved. It is assumed in the scheme that the model evolves the values at previously assimilated observation locations with reasonable accuracy, and that the distance over which observations are advected during the analysis period is generally small compared with the typical separation of observations and the scales of the background errors: that is,

$$g_{k,k-i}(\mathbf{y}_{k-i}) \equiv M(h_{k-i,k-i}(\mathbf{x}_{k-i}^a)) \approx h_{k,k-i}(\mathbf{x}_k^f). \tag{12}$$

These assumptions generally hold in the ocean, but in some instances they may have limitations, particularly where there is rapid evolution of fields. Given the above approximations, Equation (11) can be written as:

$$\mathbf{x}_k^a = \mathbf{x}_k^f + \mathbf{B}\tilde{\mathbf{H}}^T \left(\tilde{\mathbf{H}}\mathbf{B}\tilde{\mathbf{H}}^T + \tilde{\mathbf{R}} \right)^{-1} \begin{pmatrix} 0 \\ 0 \\ \vdots \\ \mathbf{y}_k - h_{k,k}(\mathbf{x}_k^f) \end{pmatrix}, \tag{13}$$

which is very similar to Equations (5) and (6).

In practice, the observation error evolution is parametrized using a simple time weighting that increases the error linearly according to the time since the observation reported. The observation error variance is increased by a factor of 1.5 at the end of one time correlation scale. The time correlation scales are described for each of the observation types in Section 3.2.

Giving the background field at previously assimilated observation locations weight in the analysis has the effect of updating the background error covariance without having to explicitly change the \mathbf{B} matrix. This method does not rely on the forecast being very accurate at the locations of previously assimilated data: the error evolution at these locations is taken into account by increasing the observation error variance as a function of the time since the observation reported.

2.3. Overview of the analysis algorithm

The assimilation scheme used in FOAM can be summarized as the following sequence of steps:

1. Identify the observations that have previously been used, and set their observational increments to zero. This will result in an innovation vector such as that given in Equation (13).
2. Update the observational error covariance using a time weighting that increases the error linearly according to the time since the observation reported, in order to approximate the changes in the errors given in Equation (9).
3. Calculate a new analysis using both new and previously used observations, using the AC scheme described by Equations (2) and (3).
4. Integrate the model forward to the next analysis time, incorporating the analysis increments during that period, using incremental analysis updates, as described in Section 3.4.

3. Practical implementation

3.1. Error covariances

The analysis equations (2) and (3) require the a priori specification of the error covariances for the model forecast and observations. The observation error covariance matrix \mathbf{R} contains estimates both of the errors in the measurements and of the errors of representativity, which are due to processes not represented by the model. The representativity errors can often dominate the measurement errors, and vary depending on the model resolution. In the FOAM system, the observation error covariance matrix is currently assumed to be diagonal, i.e. observations are uncorrelated with each other. This approximation is reasonable for some observation types, though not perhaps for satellite data such as along-track altimeter data (or some *in situ* data types).

The specification of the model forecast error covariance matrix is crucial to the effectiveness of any data-assimilation scheme, as it determines how the observation

increments are spread onto the model field. The size of the \mathbf{B} matrix is of the order of $10^7 \times 10^7$ (for the $\frac{1}{3}^\circ$ North Atlantic and Arctic FOAM configuration), and is therefore too large to be specified completely. Even if matrices of this size could be stored, the information required to specify it accurately is not available. In the FOAM system, the background error covariance matrix is approximated by specifying the variances and approximating the off-diagonal elements using a particular functional form (described later in this section).

The model forecast is likely to contain errors on many different horizontal and vertical scales, and it is important to use knowledge about the dynamics of the ocean to specify the type of function used to represent these scales. The number of different correlation scales that are to be extracted from the data depends on the type of application. In short-range ocean forecasting, such as in the FOAM system, it is thought that there are two main sources of model forecast error. The first arises from errors in the forcing of the model by atmospheric fields such as the wind, heat or freshwater forcing, or in the response of the model to such large-scale forcing. Errors in the wind and heat forcing are likely to occur on horizontal scales similar to those for synoptic weather systems in midlatitudes (of the order of a few hundred kilometres), with fairly small correlation with errors in the deep ocean. Errors in the freshwater fluxes applied to the ocean model may have smaller temporal and horizontal scales, as they will mainly affect errors in the mixed layer. These types of errors are expected to be small compared with those for the large-scale wind forcing, and so are ignored here. Errors that occur on scales similar to those of atmospheric synoptic-scale systems are called 'synoptic scale' errors.

The second source of forecast error arises from errors in the internal dynamics of the model. These internal errors are likely to be associated with the baroclinic modes of the ocean, and hence consist of several horizontal and vertical scales. At present, the resolution in FOAM is such that many of the higher baroclinic modes are not resolved, and so these are included among the errors of representativity. The first baroclinic mode is likely to have horizontal errors on scales of a few tens of kilometres, with large vertical scales, and is represented in the background error covariance. These types of errors are termed 'mesoscale' errors.

The functional form for the off-diagonal elements of \mathbf{B} is specified to be the combination of two SOAR functions, each with an associated variance and length scale:

$$f(r) = V_{\text{syn}} \left(1 + \frac{r}{L_{\text{syn}}} \exp\left(-\frac{r}{L_{\text{syn}}}\right) \right) + V_{\text{mes}} \left(1 + \frac{r}{L_{\text{mes}}} \exp\left(-\frac{r}{L_{\text{mes}}}\right) \right), \quad (14)$$

where V is the variance, r is the separation distance, and L is the length scale; subscript 'syn' indicates

the synoptic-scale component and subscript 'mes' the mesoscale component.

There are various methods in the literature for estimating the error covariance matrices. The method of Hollingsworth and Lönnberg (1986) can be used to estimate both the forecast and the observation error covariances. Here, pairs of observation-minus-model-forecast values are used to form the error statistics. This method has been used to estimate the error covariances for FOAM by Martin *et al.* (2002). The forecast values were taken from a run of the system in which data were assimilated using previous estimates of the error covariances. This work provided a set of spatially inhomogeneous observation and model forecast error variances and correlation scales for each observation type. An example of the functional fit to the data is given in Figure 1(a) for sea surface height (SSH) data at a location in the Gulf Stream for the $\frac{1}{3}^\circ$ North Atlantic and Arctic FOAM configuration. The two scales can be clearly seen. Similar calculations were performed throughout the model domain. The field of spatially-varying mesoscale variances is shown in Figure 1(b). The areas of high variance correspond to areas of high mesoscale variability, such as the Gulf Stream, although there are also areas of high variance on the northwestern European shelf and on Hudson Bay, where tides become important (tides are not included in the model and are difficult to remove from the data).

The part of the background error covariance matrix that determines the vertical correlations is also required by the analysis. These correlations are also split into two components, associated with the mesoscale and synoptic-scale errors, and are estimated using the data described in the preceding paragraph. For the temperature- and salinity-profile data, vertical correlation length scales of 200 m for the mesoscale component and 100 m for the synoptic-scale component are applied everywhere.

3.2. Observation-group-specific aspects

The result from Equations (4)–(7) can be applied within an analysis step in order to split the analysis into separate observation groups. Within FOAM, the observation groups that are assimilated are sea surface temperature (SST), SSH, temperature profile, salinity profile, and sea-ice data. The iterations described in Equation (2) are performed over each observation group, but also over all observation groups, so that the entire analysis procedure is performed a number of times. In this way, the information provided from the analysis of temperature and salinity data will feed back into the analysis of SSH, and vice versa. Provided the observations are consistent with each other and are assimilated in a consistent manner, this procedure will converge over observation types. If there are inconsistencies in the data, iterating over observation types should mean that no one type is given more weight than any other.

Each group also has time correlation scales associated with it. These are specified to be 5 days for the surface observations (i.e. SST, SSH and sea ice), and 10 days

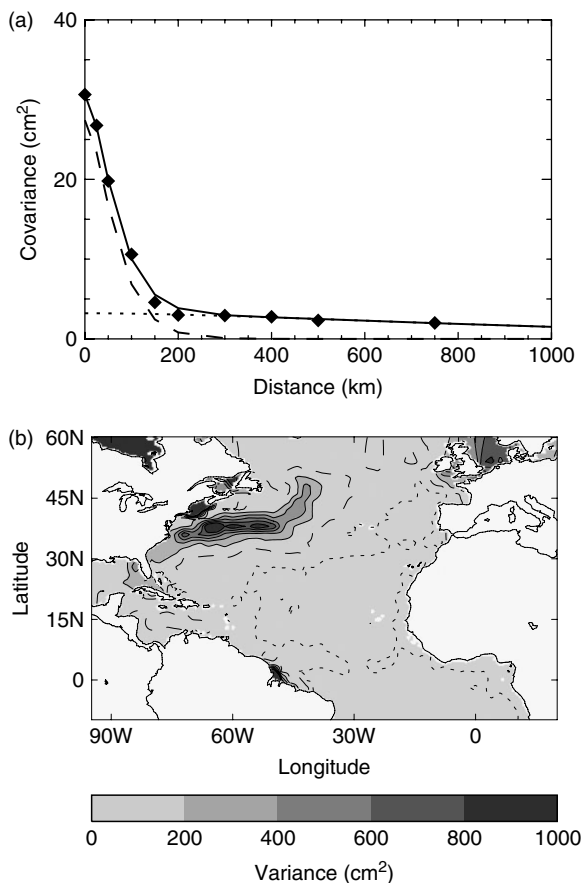


Figure 1. (a) Example of the functional fit to SSH covariance data using two SOAR functions at (30°W, 40°N). Shown are the covariance data (diamonds), the mesoscale errors (dashed line – length scale is 37.13 km), the synoptic-scale errors (dotted line – length scale is 561.63 km), and the total errors (solid line). (b) Spatially-varying SSH mesoscale forecast error variances for the $\frac{1}{3}^\circ$ North Atlantic and Arctic FOAM configuration. The solid lines have contour intervals of 200 cm², the dashed line marks the 50 cm² contour, and the dotted line marks the 10 cm² contour.

for the temperature- and salinity-profile data. These time correlations are included by increasing the errors associated with each observation in accordance with these time-scales, as described at the end of Section 2.2. Aspects of the data assimilation that are specific to each observation group are now described.

3.2.1. SST data

In situ SST measurements from moored and drifting buoys, ships and other platforms are assimilated into the FOAM system. These data are available in real time over the Global Telecommunications System (GTS). An example of the data coverage in the North Atlantic for one day is given in Figure 2. This shows that SST is fairly well observed by *in situ* measurements, although these data alone do not generally resolve mesoscale features. Satellite SST data are also assimilated in FOAM. Currently the satellite data used are from a 2.5°-resolution Advanced Very High Resolution Radiometer (AVHRR) product produced by the National Environmental Satellite, Data and Information Service. The type of satellite data used is

currently being upgraded to include high-resolution measurements produced by the GODAE high-resolution SST pilot project (GHRSSST-PP) (Donlon *et al.*, 2005). This will provide data that will resolve mesoscale ocean features and will significantly improve the SST analyses produced by FOAM.

In order to assimilate the SST data in real time, an automatic quality-control procedure is required. This includes a Bayesian procedure described by Ingleby and Lorenc (1993) for checking the compatibility of the observations with the current model forecast, given a priori error statistics for the observations and model fields. In order to achieve consistency between the quality-control and data-assimilation steps, the error covariance information described in Section 3.1 is used for the description of the probability density functions of the observations and model fields in the quality control.

The SST analysis in FOAM is calculated by comparing the temperature in the top model level (0–10 m depth) with the observations. The *in situ* observations are made at various depths, and so can include diurnal effects that are not reproduced by the model. Satellite SST data also contain significant biases compared to *in situ* data. These biases are analysed and removed using a procedure described in Section 3.5. Once this bias correction has been performed, a two-dimensional analysis of SST is produced using the method described in Section 2.

The SST analysis produces a two-dimensional field of temperature increments. These are applied throughout the mixed layer of the ocean as diagnosed by the mixed-layer model in FOAM. The depth to which these increments are applied is limited to 660 m in order to prevent increments from being applied at depth in regions of deep convection where the vertical correlations are likely to be less reliable.

3.2.2. Temperature- and salinity-profile data

All *in situ* temperature- and salinity-profile data that are available in near-real time on GTS are assimilated into FOAM. These include the large amount of Argo temperature- and salinity-profiling float data (Argo Science Team, 1998), which are now the dominant source

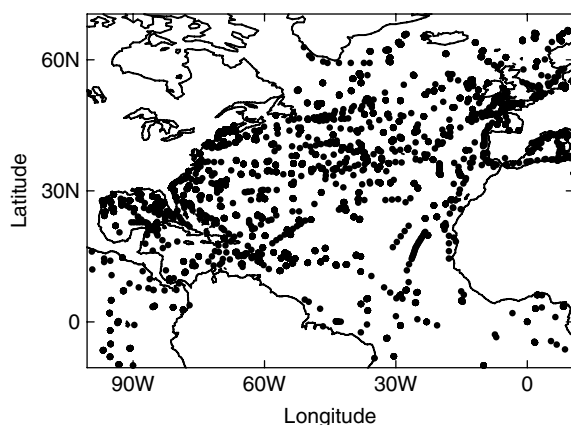


Figure 2. *In situ* SST data coverage for the North Atlantic for 20 March 2006. There are 9113 observations.

of *in situ* profile data. Other conductivity–temperature–depth measurements are also assimilated, together with the moored buoy data from the TAO/TRITON and PIRATA arrays, and temperature data from expendable bathythermographs (XBTs). Figure 3(a) shows the profile data reporting temperature to a depth of at least 1000 m for a 10-day period in March 2006. This shows the wide coverage now available thanks to the Argo array. Figure 3(b) shows the temperature profile data reported during one day, 20 March 2006; this gives an idea of the amount of new information available at each assimilation time. Similar coverage of salinity profile data reporting to at least 1000 m is also available (not shown).

A comprehensive automatic quality-control system for profile data is used in FOAM, as described by Ingleby and Huddleston (2007). This includes various checks on the data, including track checks, spike checks, stability checks, duplicate checks, background checks (as with SST), and buddy checks. This system is crucial for the operational FOAM system, particularly for checking the Argo salinity data, which can be prone to sensor drift. A balance has to be struck between rejecting good data and including bad data. In an operational setting, the requirement of a robust system means that this balance is weighted so that the scheme may reject some good data in order to ensure that no bad data are allowed to be assimilated. This is important, as once a bad platform

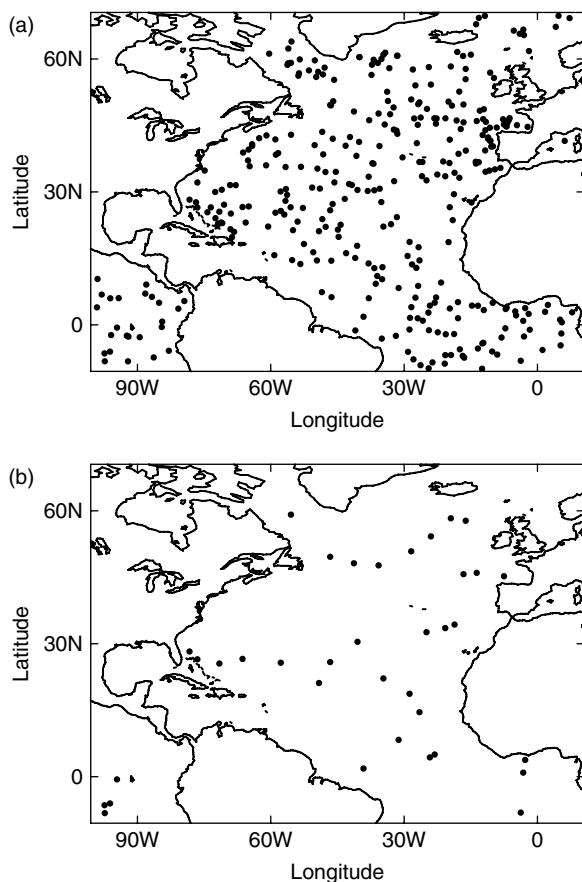


Figure 3. *In situ* temperature-profile data coverage for the North Atlantic at 1000 m depth (a) between 15 and 24 March 2006 (341 observations) and (b) on 20 March 2006 (36 observations).

has been assimilated it is then more likely to be accepted next time it reports, and can have a significant impact on the model fields.

A three-dimensional analysis is produced separately for the temperature-profile and salinity-profile data. In each case, the analysis is obtained by assuming that the solution can be separated into vertical and horizontal components, as described in Section 2.1. Within each iteration of the analysis described in Equation (2), a vertical analysis is performed at observation locations, prior to a horizontal analysis at each vertical level. The vertical correlations take the same form as the horizontal correlations described in Equation (14). Both the vertical and the horizontal analyses use the error covariance information described by Martin *et al.* (2002).

3.2.3. Altimeter SSH data

The FOAM system assimilates the along-track sea-level-anomaly (SLA) altimeter data processed in near-real time by Collecte Localisation Satellites (CLS) (<http://www.cls.fr/>). These currently consist of data from the Jason-1, Envisat and Geosat Follow-On instruments. An example of the data coverage for a 5-day period in March 2006 is given in Figure 4(a), and an example of the new data available for one assimilation period on 20 March 2006 is shown in Figure 4(b). The data are quality-controlled during the CLS processing, but an additional background check is performed before assimilation into FOAM, using a method similar to that used for the SST data.

The SLA observations are provided as anomalies relative to a time-mean field. This is necessary because

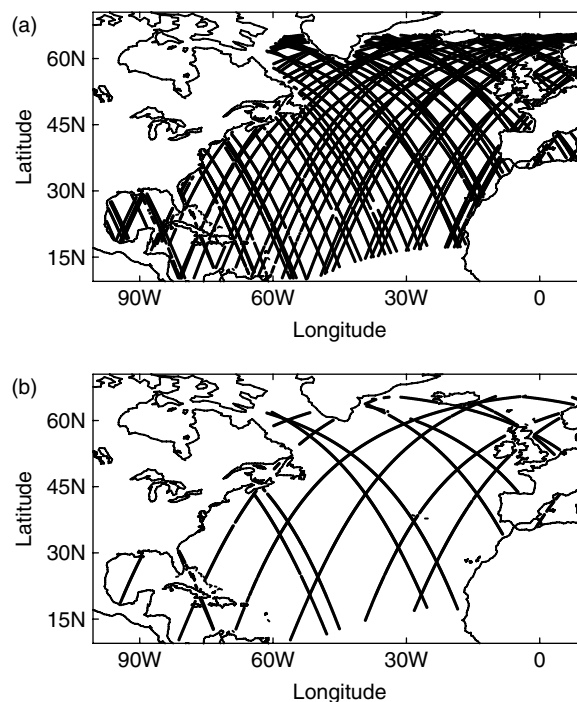


Figure 4. Altimeter SSH data coverage for the region covered by the North Atlantic $\frac{1}{2}^\circ$ -resolution FOAM configuration (a) between 18 and 22 March 2006 (15 731 observations) and (b) for 20 March 2006 (2975 observations).

the mean height of the ocean (relative to a reference ellipsoid), which includes the Geoid (a fixed gravity-equipotential surface) and the mean dynamic topography (MDT), is not known accurately enough relative to the variability in the dynamic topography (typically of the order of 10 cm), mainly because of insufficient knowledge of the Geoid. The time mean is removed from the SSH observations; this removes the Geoid, but also removes the MDT. A reference MDT therefore needs to be added to allow us to compare the observations to the model fields. Various options for a reference MDT are available. These include the model's mean SSH from a hindcast run, a mean derived from hydrographic data, a mean derived from geodetic measurements (for example, from GRACE and GOCE), information from the altimeter data themselves, or any combination of these. The MDT currently used in the North Atlantic FOAM configurations is described in detail in (Hines, 2001). It is based on the dataset of Singh and Kelly (1997) in the northeastern Atlantic, which is merged with data from a long integration of the FOAM system.

The altimeter data are assimilated using a two-stage approach. First, a horizontal analysis is performed, using the method described in Section 2. This results in a two-dimensional field of SSH increments, which is split into a part due to errors in the mesoscale ocean dynamics and a part due to larger-scale errors in the model (for instance due to steric effects not captured by the model), calculated as part of the analysis using the error variance information described in Section 3.1. These two components of the increments are used differently. The part due to mesoscale errors is projected onto the subsurface temperature and salinity using a modified version of the scheme introduced by Cooper and Haines (1996), which is described in more detail in the next paragraph. The part of the SSH increments that is due to the larger-scale errors could then be used to update the barotropic stream function (though at present this part is discarded). The global 1°-resolution FOAM configuration does not resolve mesoscale dynamics, so the altimeter data are not assimilated in this configuration. The higher-resolution regional FOAM configurations either permit or resolve the ocean eddy field, and so the altimeter data are assimilated into these models.

The scheme used to project the mesoscale part of the SSH signal onto the subsurface density field is based on that of Cooper and Haines (1996). The model's density profile is raised or lowered in order to produce the equivalent change in hydrostatic surface height while preserving the bottom pressure. Observation and model density profiles have been used by Martin *et al.* (2002) to calculate the errors in the vertical placement of the density profile from a hindcast run of the FOAM system. These show that most of the errors occur in the region of the thermocline, with little error in the deep ocean or near the surface mixed layer. For this reason, the scheme of Cooper and Haines (1996) has been modified by applying a change to the displacement as a function of depth. This function produces a displacement that is maximal

in the thermocline, decays below the thermocline depth to zero at the bottom of the model, and decays above the thermocline depth to zero at the base of the mixed layer. The surface-height increment produced using this method is the same as that produced using the Cooper and Haines (1996) method, but the vertical profile of the density increments is different.

The scheme just described (and the unmodified scheme of Cooper and Haines (1996)) can produce unrealistic displacements to the density profile when the stratification of the model's density profile is weak. In the absence of knowledge about how the SSH error is correlated with errors in the subsurface fields in these regions, a weighting factor is applied to the displacements: this weighting is equal to one where the top-to-bottom temperature difference is greater than 10 °C and decreases linearly to zero when this temperature difference is less than 5 °C. There are also problems with assimilating the altimeter SSH data in shallow water, mainly due to the contamination of the observed signal by tides, which are not represented by the model. For this reason, another weighting factor is applied to the vertical displacements: this weighting is equal to one for depths greater than 1000 m and decreases linearly to zero when the depth is less than 100 m.

3.2.4. Sea-ice data

The sea-ice analysis currently used in FOAM is a separate scheme from that used for the other observation types. This scheme nudges the FOAM sea-ice concentration fields towards a daily sea-ice analysis produced by the Canadian Meteorological Centre, with a time-scale of one day. We plan to upgrade the scheme to assimilate the satellite sea-ice concentration data produced by the Special Sensor Microwave/Imager (SSM/I) instrument, using the same method as for the other observation types. Details of the new scheme for assimilating sea-ice data are described by Stark *et al.* (2005).

3.3. Multivariate aspects

The assimilation procedure described above (except for the sea-ice analysis) produces three-dimensional increments to the temperature and salinity fields at the analysis time. In order to produce changes to the other prognostic model variables, multivariate balance relationships are used. This is done because the ocean is generally in geostrophic balance (apart from the Ekman component) and the model should not be pushed out of balance by the assimilation as this may cause undesirable or unphysical responses in the model. Increments to the baroclinic velocities are obtained using the geostrophic relationship, except within a few degrees of the Equator, where this relationship does not apply. The baroclinic velocity increments are ramped to zero within 5° of the Equator. No balancing increments are made to the barotropic stream function, and there is no special treatment of the surface wind-driven velocities.

3.4. Application of the increments

The result of the assimilation is a set of three-dimensional updates to the model prognostic fields. Direct insertion of these analysis increments can cause shocks to the model, and may generate within it spurious and undesirable gravity waves. To reduce the shock to the model, the assimilation increments are applied slowly during the next day of the forecast, using a method known as ‘incremental analysis updates’ (Bloom *et al.*, 1996). This scheme applies the increments as a constant forcing to the model equations between analysis times. The temperature, salinity and baroclinic-velocity increments are applied in this manner.

3.5. Bias-correction issues

The assimilation scheme as described in Section 2 assumes that there are no systematic errors in either the model or the observations. In practice, this assumption is not valid, and the systematic errors in both the model and observations can cause significant problems for the assimilation scheme. Dee and Da Silva (1998) discuss the general theory for dealing with systematic errors in the data-assimilation framework.

Observational biases can be dealt with if they are known a priori. For instance, the low-resolution AVHRR satellite SST data described in Section 3.2 are known to contain systematic errors. In this case a large-scale two-dimensional analysis is performed on the satellite data and on the *in situ* data separately. The difference between these two analyses is then taken to be the bias (the *in situ* data are assumed to be unbiased), and the satellite data are adjusted accordingly. More sophisticated techniques are being developed for estimating the biases in the high-resolution SST data provided by GHRSSST-PP, so that data from different satellite sensors (which measure the temperature at different depths) can be assimilated in a consistent manner.

Other observation types also contain biases. For instance, the salinity data from some Argo floats are known to have drifts and offsets, which are corrected in delayed-mode processing. Some XBTs are also thought to have small warm biases. These errors are not currently estimated or corrected as part of the real-time operational FOAM system, although data with significant errors will be rejected by the quality control.

The raw altimeter SLA data also contain biases, although most of these are removed during the processing by CLS. However, the mean dynamic topography used to enable the assimilation of these data can cause systematic errors in the observations (SLA and MDT). A method that calculates and removes these biases online has been implemented in the development version of the FOAM assimilation system by Drecourt *et al.* (2006). This method assumes the model sea level to be unbiased and estimates the bias in the MDT using a procedure that behaves similarly to a running mean over the sea-level residuals. This procedure is efficient at improving the sea-level forecast.

Systematic errors in the model can cause significant problems for the assimilation scheme. An example of this is in the tropical oceans, where the main large-scale dynamical balance is between the surface wind stress and the subsurface pressure gradients. When assimilating temperature, salinity and altimeter data, the subsurface pressure gradients are altered to be closer to the true state, but systematic errors in the wind forcing and the way in which it is input into the model can disrupt this balance (Huddleston *et al.*, 2004). This causes spurious circulation cells to be set off by the assimilation; these tend to oppose the increments being put in by the assimilation. The assimilation of data in the presence of systematic model error therefore has a detrimental impact on the model dynamics. Bell *et al.* (2004) describe how FOAM attempts to overcome these problems by including a correction to the pressure-gradient term, calculated using information from the model and observation differences. An alternative, and complementary, method to alleviate these problems has been proposed by Burgers *et al.* (2002): in this method, an extended ‘geostrophic’ relationship is used to make balancing east–west velocity increments. This is not currently implemented in FOAM.

In the extratropics there are other systematic errors in the model. For instance, many models are unable to reproduce the true separation point of the Gulf Stream from the east coast of North America, or the northward turn of the Gulf Stream extension. The assimilation of data can correct for these errors during an analysis, but the model will soon drift when producing a forecast. The semi-prognostic method introduced by Greatbatch *et al.* (2004), which uses ideas similar to the pressure-correction method of Bell *et al.* (2004), has been shown to correct such errors, and could be used to deal with problems of this nature.

4. Demonstration of the impact of data assimilation

The performance of the data-assimilation scheme is assessed using two experimental frameworks. The first set of experiments is based on an identical-twin setting, in which the impact of the altimeter data assimilation scheme is assessed. The second set of experiments assimilates real data over a 5-year period, and the results are shown for the $\frac{1}{9}^\circ$ configuration.

4.1. Identical-twin experiments

In order to show the impact of the assimilation of altimeter data in a setting in which the ‘truth’ is known, identical-twin experiments have been performed. These involve running a two-year spin-up of the FOAM system (1° global model, $\frac{1}{3}^\circ$ North Atlantic and Arctic model and $\frac{1}{9}^\circ$ North Atlantic model) starting from rest, with climatological temperature and salinity fields taken from (Levitus *et al.*, 1998). A one-way nesting is used, in which the temperature, salinity, baroclinic velocity, stream function and sea-ice fields of the smaller domain are relaxed

towards those of the outer domain near the boundary, using the flow relaxation scheme (Martinsen and Engerdahl, 1987). The models are forced by 6-hourly surface fields from the Met Office Numerical Weather Prediction (NWP) system between May 2002 and June 2004. At the end of this spin-up period, a year-long integration is run, which assimilates the real temperature and salinity data into each of the model configurations, as described in Section 3.2. The fields from this run are taken to be the 'true' state of the ocean. The conditions at the end of this year-long run are then used as the initial state for another run of the same period in which the temperature and salinity data are assimilated. The fields from this run are taken to be the 'control'. No SST observations are assimilated in any of these integrations. An assimilative model is used as a control in these experiments in order to observe the impact of the altimeter assimilation in the presence of temperature and salinity assimilation, rather than in isolation.

Simulated observations of SSH are created by subsampling the fields from the 'true' run according to the locations of the real along-track altimeter data available at each time. Noise is added to these observations, having a normal distribution with a standard deviation of 4 cm. An 'assimilation' experiment is then run, using the $\frac{1}{9}$ North Atlantic model in which the simulated SSH observations are assimilated using the scheme described in Section 3.2.

Figure 5(a) shows root mean square (RMS) error statistics, calculated from the difference between the model

field and the 'true' field at every model grid point, of daily mean SSH during the year of the identical-twin experiments. 'Daily mean' here refers to the 24-hour period over which the assimilation increments are nudged into the model, starting from the analysis time. The dashed line shows the errors in the 'control' run compared to the 'true' run; this shows a steady reduction in the errors in SSH when assimilating only the temperature- and salinity-profile data. However, the RMS errors at the end of the 'control' integration are still about 9 cm, indicating that the system is very sensitive to the initial conditions. When assimilating the altimeter SSH data, the RMS errors are significantly reduced, to about 6 cm averaged over the whole North Atlantic, as shown by the solid line. The same errors calculated at the along-track observation locations and averaged over the year of the integration are approximately 12 cm for the 'control' integration and 7 cm for the 'assimilation' integration. These assimilation errors are higher than the observation-error standard deviations of 4 cm, indicating that the error variances specified in the assimilation scheme are not optimal.

The evolution of errors during forecasts has been estimated by running one 30-day forecast for each month of the year (July 2004 to June 2005). Figure 5(b) shows the average errors (daily mean differences from the truth) for these forecasts for the 'control' and 'assimilation' integrations, together with the errors if the analysis is persisted into the forecast period. The small number of forecasts means that these errors should be treated with




Figure 5. RMS errors in SSH compared to the 'true' run at all model grid points, for the 'control' integration (dashed lines) and SSH-assimilation integration (solid lines), (a) for the daily mean fields, and (b) as a function of forecast day, with persistence (dotted line) also shown.

caution, but they do give an idea of the forecast errors in the system. The forecast runs from the ‘assimilation’ integration have smaller errors than those from the ‘control’ integration throughout the 30 days, showing that the analysis of SSH is providing a good initialization of the model. The errors from the ‘assimilation’ forecasts are very similar to persistence for the first 7 days of the forecast, but the model does provide an improved forecast in the subsequent 23 days.

The impact of the assimilation of SSH data on the subsurface temperature and salinity fields is illustrated in Figure 6(a) and (b) respectively. These show the RMS error statistics, calculated from a model daily average after the analysis compared to the ‘truth’ at every model grid point. The SSH assimilation significantly reduces the errors in temperature in the top 1000 m, and the errors in salinity are reduced in the top few hundred metres. Below these depths, the errors are increased by the assimilation of SSH data.

An example of the impact of the assimilation system on non-observed variables is given in Figure 7, which shows the annual mean errors in the surface velocity magnitude for the ‘control’ and ‘assimilation’ integrations. It is clear from this figure that the errors in surface velocities are reduced by the assimilation of altimeter data, although significant errors remain. The annual mean spatially-averaged absolute-surface-velocity error is 4.2 cm s^{-1} for the ‘control’ integration and 3.0 cm s^{-1} for the ‘assimilation’ integration.

The identical-twin experiments just described are useful for assessing the potential of the scheme. The results show that the altimeter data coverage has the most significant ability to resolve the mesoscale structure of the ocean, and brings complementary information about the temperature and salinity structure within the top few hundred metres. There are, however, many errors in the

real system that are not present in this idealized setting. These include random and systematic errors in the MDT used for the altimeter assimilation, inconsistency between observation types, errors in the model, and many others. We next describe some experiments in a real-world setting in which all of these errors are likely to be present.

4.2. Hindcast integrations

To show the impact of the assimilation of data in a realistic setting, a set of integrations has been performed for the period January 2001 to July 2005. The 1° global model, the $\frac{1}{3}^\circ$ North Atlantic and Arctic model and the $\frac{1}{9}^\circ$ North Atlantic model have all been run for this period, starting from operational FOAM analysis fields and forced by 6-hourly surface fluxes from the Met Office NWP system. These integrations have been run with no data assimilation (CONTROL), with assimilation of all *in situ* temperature and salinity data (ALLTS), and as a repeat of the ALLTS run but with the Argo data withheld (NOARGO).

The impact of the Argo data on the FOAM system is illustrated in Figure 8, which shows RMS errors of temperature and salinity analyses from the $\frac{1}{9}^\circ$ North Atlantic model compared to observations that have not yet been assimilated but are valid within 24 hours of the analysis. That is, the analysis fields at T_0 are compared to data profiles that are valid between T_0 and $T_0 + 24$. The ALLTS run (solid line) has much less error in temperature over all of the top 1000 m than the NOARGO integration (dotted line) and the CONTROL integration (dashed line). The salinity errors are most significantly reduced by the Argo data assimilation in the top 600 m. The salinity errors with the NOARGO integration are marginally worse than with no data assimilated in the top 400 m. This




Figure 6. RMS errors in (a) temperature and (b) salinity, compared to the ‘true’ run, for the ‘control’ integration (dashed lines) and the SSH-assimilation integration (solid lines). The errors are shown as a function of depth, and have been averaged over the year of the integrations.




Figure 7. Annual mean errors in surface-velocity magnitude compared to the 'true' run, for (a) the 'control' integration, and (b) the SSH-assimilation integration. Only regions where the magnitude exceeds 15 cm s^{-1} are shaded.




Figure 8. RMS errors in (a) temperature and (b) salinity analyses, compared to *in situ* profile observations before they are assimilated, averaged between January 2001 and July 2005 and over the model domain. Results are shown for the ALLTS run (solid lines), the NOARGO run (dotted lines) and the CONTROL run (dashed lines).

is probably because assimilating only temperature data can disrupt the density structure in the model, so that, although temperature is improved, the salinity fields are slightly worsened, with no independent salinity data to constrain them. This result indicates the importance of methods such as that of (Troccoli and Haines, 1999), in which balance relationships are used to update the salinity fields when assimilating only temperature data. Overall, these statistics show both the importance of the Argo data and the beneficial impact of the data assimilation on the model fields.

The differences between the average temperature at 1000 m depth from each of the model runs and the (Levitus *et al.*, 1998) climatology are shown in Figure 9. When no data are assimilated, the model fields have significant biases at this depth. The data assimilation corrects for most of the drift in the temperature, as shown in Figure 9(c), although some differences with the climatology remain. An area of significant bias is the Gulf of Mexico; this is probably because of the lack

of Argo data in this region. There are also significant differences in the Gulf Stream region, although these are likely to be due to the climatology not being able to resolve the temperature gradients in this region. Some of these differences with the climatology could also be due to a real interannual signal that is contained in the observations but not in the climatology. Results from the NOARGO run show that the Argo data are crucial for obtaining accurate temperature analyses at this depth.

The average analysed temperature increments are shown in Figure 10 for the NOARGO and ALLTS integrations. The magnitude of these increments is largest for the ALLTS run; this highlights the importance of the Argo data for capturing the errors in the model. The field shown in Figure 10(b) gives an indication of the average model drift over each day in the observed regions. It indicates a significant warm bias in the model to the west of the North Atlantic (the increments are negative, indicating that the assimilation is acting to cool the model), and a cool bias to the east. The biases in both of these

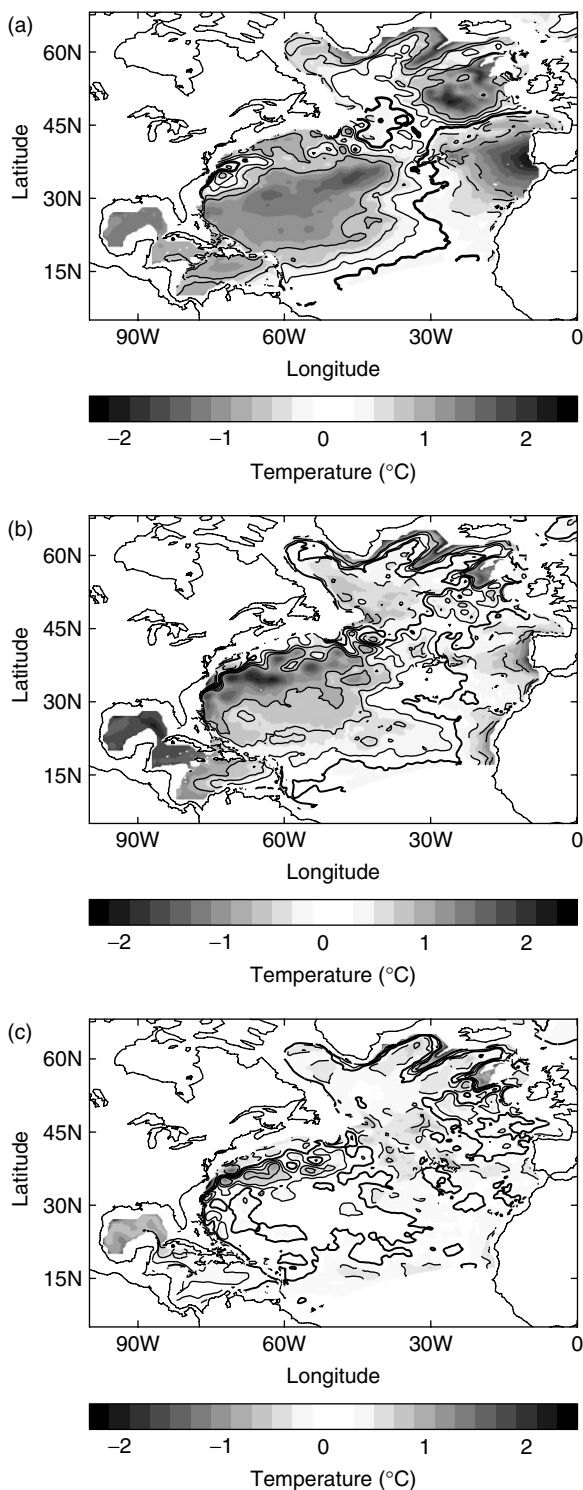


Figure 9. Difference between the (Levitus *et al.*, 1998) climatology and the temperature at 1000 m depth, averaged between January 2001 and July 2005, from (a) the CONTROL run, (b) the NOARGO run, and (c) the ALLTS run. The contour interval is 0.3 °C; negative contours are dashed, and the thick line is the zero contour.

regions are thought to be due to deficiencies in the vertical advection scheme used in these integrations. The warm bias to the west is due to the thermocline being diffused by the vertical advection scheme, and the cool bias to the east is due to the vertical gradients caused by outflow water from the Mediterranean Sea being diffused.

The vertical advection scheme will be upgraded in future versions of FOAM. An increase in the number of vertical levels would also help to alleviate some of these problems.

5. Summary and future plans

The current status of the FOAM data assimilation system has been described. The theoretical foundation of the scheme is an updated version of an optimal-interpolation scheme in which the time evolution of the observational information is taken into account. Details of the way in which this scheme is implemented have been given for each of the observation types assimilated: SST, SSH, temperature and salinity profiles, and sea ice. The basic scheme assumes that there are no systematic errors in the modelling and observational systems. This is not true in practice, so the way in which these types of errors are accounted for has also been described.

The experiments that have been performed demonstrate the impact of the data-assimilation scheme on the accuracy of the FOAM analyses. The assimilation of

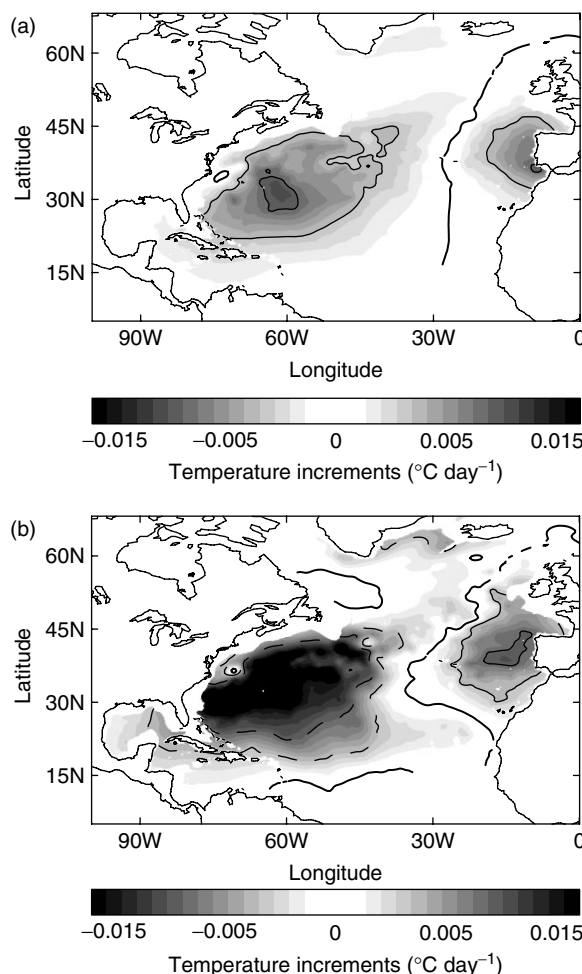


Figure 10. Average temperature increments at 1000 m depth from (a) the NOARGO run and (b) the ALLTS run. The contour interval is 0.005 °C day⁻¹; negative contours are dashed, and the thick line is the zero contour.

altimeter data is shown to improve both the surface height and the subsurface temperature and salinity fields in an identical-twin setting. Real-world hindcast experiments have shown the impact of Argo data when it is assimilated into FOAM.

Various upgrades to the data-assimilation system are envisaged, including updating the error covariance estimates. A scheme such as the NMC method (Parrish and Derber, 1992), or ensemble methods (Kucukkaraca and Fisher, 2006), will be used to provide new estimates of these, although they will be compared to the estimates currently used to ensure consistency with these 'real-world' estimates. It would be possible to update the filter to include some anisotropy in the error covariances, which should produce more realistic analyses.

The assimilation of new data types is also being investigated. These include the high-resolution SST data that are available through GHRSSST-PP, and SSM/I sea-ice-concentration data. Other future possibilities for assimilation include sea surface salinity from missions such as SMOS and Aquarius, and the velocity information available from drifters. Investigations are already underway into the assimilation of ocean-colour satellite data into a coupled physical–bio-geochemical model, using the FOAM system together with the HadOCC model (Palmer and Totterdell, 2001); this will be implemented in the operational system.

It has been shown by Haines *et al.* (2006) that a change of variable can enable the use of longer correlation scales, so that an individual observation can influence a much larger area. It would therefore be possible to obtain more information from the observations by changing the variables that are used in the assimilation. For instance, a scheme in which the control variables are density and 'spiciness' is currently under development. A scheme that uses the water mass characteristics of the model background field to update the salinity based on the temperature updates, as described by Troccoli and Haines (1999), is also currently under investigation in the development version of the assimilation scheme.

The model used in the FOAM system is currently being transitioned to use the NEMO model (<http://www.lodyc.jussieu.fr/NEMO>). The current FOAM data-assimilation system is being tested within this new modelling framework. This framework will also provide the opportunity to investigate the use of more sophisticated data-assimilation schemes that take into account the evolution of the background error covariance matrix. Such a scheme may be particularly useful when assimilating data in shallower waters, where the temporal and spatial scales of the water are much smaller than in the deep ocean.

Acknowledgements

This work is a contribution to the MERSEA project. Partial support of the European Commission under Contract SIP3-CT-2003-502885 is gratefully acknowledged.

A Appendix: Subsets of observations

The last term in Equation (6) will be manipulated using the formula:

$$\begin{aligned} \mathbf{y}_{2,k} - h_2(\mathbf{x}_{1,k}^a) &= \mathbf{y}_{2,k} - h_2(\mathbf{x}_k^f + \mathbf{x}_{1,k}^a - \mathbf{x}_k^f) \\ &= \mathbf{y}_{2,k} - h_2(\mathbf{x}_k^f) - \\ &\quad - \mathbf{H}_2 \mathbf{K}_1 (\mathbf{y}_{1,k} - h_1(\mathbf{x}_k^f)), \quad (\text{A1}) \end{aligned}$$

where

$$\mathbf{K}_1 = \mathbf{B} \mathbf{H}_1^T (\mathbf{H}_1 \mathbf{B} \mathbf{H}_1^T + \mathbf{R}_1)^{-1},$$

Equation (5) has been used to derive the second equality, and the observation operator has been assumed to be linear. Adding Equations (5) and (6) and using Equation (A1), we obtain:

$$\begin{aligned} \mathbf{x}_{2,k}^a - \mathbf{x}_k^f &= \mathbf{K} \begin{pmatrix} 0 \\ \mathbf{y}_{2,k} - h_2(\mathbf{x}_k^f) \end{pmatrix} + \\ &\quad + \left(\mathbf{K}_1 - \mathbf{K} \begin{pmatrix} 0 \\ \mathbf{H}_2 \end{pmatrix} \mathbf{K}_1 \right) (\mathbf{y}_{1,k} - h_1(\mathbf{x}_k^f)), \quad (\text{A2}) \end{aligned}$$

where

$$\mathbf{K} = \mathbf{B} \mathbf{H}^T (\mathbf{H} \mathbf{B} \mathbf{H}^T + \mathbf{R})^{-1}.$$

In order to combine terms, we reorganize the first element in the second term on the right-hand side of Equation (A2) as follows:

$$\begin{aligned} &\mathbf{K}_1 (\mathbf{y}_{1,k} - h_1(\mathbf{x}_k^f)) \\ &= \mathbf{B} \mathbf{H}_1^T (\mathbf{H}_1 \mathbf{B} \mathbf{H}_1^T + \mathbf{R}_1)^{-1} (\mathbf{y}_{1,k} - h_1(\mathbf{x}_k^f)) \\ &= \mathbf{B} \begin{pmatrix} \mathbf{H}_1^T & \mathbf{H}_2^T \end{pmatrix} \begin{pmatrix} (\mathbf{H}_1 \mathbf{B} \mathbf{H}_1^T + \mathbf{R}_1)^{-1} & 0 \\ 0 & 0 \end{pmatrix} \\ &\quad \cdot \begin{pmatrix} \mathbf{y}_{1,k} - h_1(\mathbf{x}_k^f) \\ 0 \end{pmatrix} \\ &= \mathbf{B} \mathbf{H}^T (\mathbf{H} \mathbf{B} \mathbf{H}^T + \mathbf{R})^{-1} (\mathbf{H} \mathbf{B} \mathbf{H}^T + \mathbf{R}) \\ &\quad \cdot \begin{pmatrix} (\mathbf{H}_1 \mathbf{B} \mathbf{H}_1^T + \mathbf{R}_1)^{-1} & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} \mathbf{y}_{1,k} - h_1(\mathbf{x}_k^f) \\ 0 \end{pmatrix} \\ &= \mathbf{B} \mathbf{H}^T (\mathbf{H} \mathbf{B} \mathbf{H}^T + \mathbf{R})^{-1} \\ &\quad \cdot \begin{pmatrix} I \\ \mathbf{H}_2 \mathbf{B} \mathbf{H}_1^T (\mathbf{H}_1 \mathbf{B} \mathbf{H}_1^T + \mathbf{R}_1)^{-1} \end{pmatrix} (\mathbf{y}_{1,k} - h_1(\mathbf{x}_k^f)) \\ &= \mathbf{K} \begin{pmatrix} I \\ \mathbf{H}_2 \mathbf{K}_1 \end{pmatrix} (\mathbf{y}_{1,k} - h_1(\mathbf{x}_k^f)). \quad (\text{A3}) \end{aligned}$$

Note that the penultimate step in the above manipulations uses the fact that the errors in each observation batch are assumed to be uncorrelated with each other. Using Equation (A3), we can simplify Equation (A2) to:

$$\begin{aligned} \mathbf{x}_{2,k}^a &= \mathbf{x}_k^f + \mathbf{K} \begin{pmatrix} 0 \\ \mathbf{y}_{2,k} - h_2(\mathbf{x}_k^f) \end{pmatrix} + \\ &\quad + \mathbf{K} \begin{pmatrix} \mathbf{y}_{1,k} - h_1(\mathbf{x}_k^f) \\ 0 \end{pmatrix} \\ &= \mathbf{x}_k^f + \mathbf{K} (\mathbf{y}_k - h(\mathbf{x}_k^f)). \quad (\text{A4}) \end{aligned}$$

Comparing Equation (A4) with Equation (1), one sees that $\mathbf{x}_{2,k}^a \equiv \mathbf{x}_k^a$. By induction, this result can be applied to an arbitrary number of expanding subsets for any set of observations.

References

- Anderson BDO, Moore JB. 1979. *Optimal Filtering*. Prentice-Hall.
- Argo Science Team. 1998. 'On the design and implementation of Argo: An initial plan for a global array of profiling floats'. ICPO Report No. 21, GODAE Report No. 5. GODAE International Project Office: Melbourne, Australia.
- Bell MJ, Forbes RM, Hines A. 2000. Assessment of the FOAM global data assimilation system for real-time operational ocean forecasting. *J. Marine Syst.* **25**: 1–22.
- Bell MJ, Martin MJ, Nichols NK. 2004. Assimilation of data into an ocean model with systematic errors near the equator. *Q. J. R. Meteorol. Soc.* **130**: 873–893.
- Bloom SC, Takacs LL, Da Silva AM, Ledvina D. 1996. Data assimilation using incremental analysis updates. *Mon. Weather Rev.* **124**: 1256–1271.
- Bratseth AM. 1986. Statistical interpolation by means of successive corrections. *Tellus* **38A**(:): 439–447.
- Burgers G, Balmaseda M, Vossepoel F, van Oldenborgh GJ, van Leeuwen PJ. 2002. Balanced ocean-data assimilation near the equator. *J. Phys. Oceanogr.* **32**: 2509–2519.
- Cooper M, Haines K. 1996. Altimetric assimilation with water property conservation. *J. Geophys. Res.* **101**(C1): 1059–1077.
- Daley R. 1991. *Atmospheric Data Analysis*. Cambridge University Press: Cambridge, UK.
- Dee DP, Da Silva AM. 1998. Data assimilation in the presence of forecast bias. *Q. J. R. Meteorol. Soc.* **117**: 269–295.
- Donlon C, GHRSSST-PP International Science Team. 2005. 'Report of the sixth GODAE high resolution pilot science team meeting'. International GHRSSST-PP Project Office, Met Office: Exeter, UK.
- Drecourt J-P, Haines K, Martin MJ. 2006. Influence of systematic error correction on the temporal behavior of an ocean model. *J. Geophys. Res.* **111**: C11020, doi:10.1029/2006JC003513.
- Evensen G. 1994. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.* **99**: 10143–10162.
- Gordon C, Cooper C, Senior CA, Banks H, Gregory JM, Johns TC, Mitchell JFB, Wood RA. 2000. The simulation of SST, sea ice extents and ocean heat transports in a version of the Hadley Centre coupled model without flux adjustments. *Clim. Dyn.* **16**: 147–168.
- Greatbatch RJ, Sheng J, Eden C, Tang L, Zhai X, Zhao J. 2004. The semi-prognostic method. *Continental Shelf Res.* **24**: 2149–2165.
- Haines K, Blower JD, Drecourt J-P, Liu C. 2006. Salinity assimilation using $S(T)$: covariance relationships. *Mon. Weather Rev.* **134**: 759–771.
- Hines A. 2001. 'Implementation and tuning of an altimeter data assimilation scheme for high resolution FOAM models'. Ocean Applications Technical Note No. 26, NCOF, Met Office: Exeter, UK.
- Hines A, Bell MJ, Acreman D, Barciela R, Martin MJ, Sellar A, Stark J, Storkey D. 2007. 'The Forecasting Ocean Assimilation Model (FOAM) system'. In: *Proceedings of the Fourth International Conference on EuroGOOS* (to appear).
- Hollingsworth A, Lönnberg P. 1986. The statistical structure of short-range forecast errors as determined from radiosonde data. Part I: The wind field. *Tellus* **38A**(:): 111–136.
- Houtekamer PL, Mitchell HL. 2001. A sequential ensemble Kalman Filter for atmospheric data assimilation. *Mon. Weather Rev.* **129**: 123–137.
- Huddleston MR, Bell MJ, Martin MJ, Nichols NK. 2004. Assessment of wind-stress errors using bias corrected ocean data assimilation. *Q. J. R. Meteorol. Soc.* **130**: 853–871.
- Ide K, Courtier P, Ghil M, Lorenc AC. 1997. Unified notation for data assimilation: Operational, sequential and variational. WMO Int. Symp. on Data Assimilation in Meteorology and Oceanography, Tokyo, 13–17 March 1995. *J. Meteorol. Soc. Jpn* **75**(1B): 181–189.
- Ingleby NB, Huddleston M. 2007. Quality control of ocean temperature and salinity profiles – historical and real-time data. *J. Marine Syst.* **65**: 158–175.
- Ingleby NB, Lorenc A. 1993. Bayesian quality control using multivariate normal distributions. *Q. J. R. Meteorol. Soc.* **119**: 1195–1225.
- Kucukkaraca E, Fisher M. 2006. Use of analysis ensembles in estimating flow-dependent background error variances. ECMWF Tech. Memo. 492 (available from <http://www.ecmwf.int/publications>).
- Levitus S, Boyer TP, Conkright ME, O'Brien T, Antonov J, Stephens C, Stathopoulos L, Johnson D, Gelfeld R. 1998. *World Ocean Database 1998. Volume 1: Introduction*. NOAA Atlas NESDIS 18.
- Lorenc AC. 1986. Analysis methods for numerical weather prediction. *Q. J. R. Meteorol. Soc.* **109**: 701–721.
- Lorenc AC. 1992. Iterative analysis using covariance functions and filters. *Q. J. R. Meteorol. Soc.* **118**: 569–591.
- Lorenc AC, Bell RS, MacPherson B. 1991. The Met Office analysis correction data assimilation scheme. *Q. J. R. Meteorol. Soc.* **117**: 59–89.
- Martin MJ, Bell MJ, Hines A. 2002. 'Estimation of three-dimensional error covariance statistics for an ocean assimilation system'. Ocean Applications Technical Note No. 30, NCOF, Met Office: Exeter, UK.
- Martinsen EA, Engerdahl H. 1987. Implementation and testing of a lateral boundary scheme as an open boundary condition in a barotropic ocean model. *Coast. Eng.* **11**: 603–627.
- Palmer JR, Totterdell IJ. 2001. Production and export in a global ocean ecosystem model. *Deep-Sea Res. Pt I* **48**: 1169–1198.
- Parrish DF, Derber JC. 1992. The National Meteorological Center's spectral statistical interpolation analysis system. *Mon. Weather Rev.* **120**: 1747–1763.
- Singh S, Kelly KA. 1997. 'Monthly maps of sea surface height in the North Atlantic and zonal indices for the Gulf Stream using TOPEX/Poseidon altimeter data'. Woods Hole Oceanographic Institution Technical Report, WHOI-97-06.
- Stark JD, Ridley J, Hilton F. 2005. 'Synergistic use of remote sensing data in coupled ocean–ice model data assimilation'. Final Report for ESA contract number 17334/03/NL/FF, available from Met Office: Exeter, UK.
- Troccoli A, Haines K. 1999. Use of the temperature–salinity relation in a data assimilation context. *J. Atmos. Oceanic Technol.* **16**: 2011–2025.