RMetS
Royal Meteorological Society

Implementing a variational data assimilation system in an operational 1/4 degree global ocean model

Jennifer Waters,* Daniel J. Lea, Matthew J. Martin, Isabelle Mirouze, Anthony Weaver and James While

Met Office, Exeter, UK

*Correspondence to: J. Waters, Met Office, Fitzroy Road, Exeter, EX1 3PB, UK. E-mail: jennifer.waters@metoffice.gov.uk

This article is published with the permission of the Controller of HMSO and the Queen's Printer for Scotland.

This article describes the implementation of an incremental first guess at an appropriate time three-dimensional variational (3DVAR) data assimilation scheme, NEMOVAR, in the Met Office's operational 1/4 degree global ocean model. NEMOVAR assimilates observations of sea-surface temperature (SST), sea-surface height (SSH), *in situ* temperature and salinity profiles and sea ice concentration. The Met Office is the first centre to implement NEMOVAR at 1/4 degree and the required developments are discussed, with particular focus on the specification of the background-error covariances.

Background-error correlations in NEMOVAR are modelled using a diffusion operator. The horizontal background-error correlations for temperature, salinity and sea ice concentration are parametrized using the Rossby radius, which produces relatively short correlation length-scales at mid to high latitudes, while a flow-dependent mixed-layer depth parametrization is used to define the vertical length-scales for the 3D variables.

Results from a one-year reanalysis with NEMOVAR are presented and compared with the preceding operational data assimilation scheme at the Met Office. NEMOVAR is shown to provide significant improvements to SST, SSH and sea ice concentration fields, with the largest improvements seen in regions of high variability such as eddy shedding and frontal regions and the marginal ice zone. This improvement is associated with shorter correlation length-scales in the extratropics and an improved fit to observations in NEMOVAR. Some degradation to subsurface temperature and salinity fields where data are sparse is identified and this will be the focus of future improvements to the system.

Key Words: operational oceanography; data assimilation; 3DVAR

Received 23 September 2013; Revised 17 April 2014; Accepted 22 April 2014; Published online in Wiley Online Library 14 June 2014

1. Introduction

The Met Office's global Forecasting Ocean Assimilation Model (FOAM) has been run operationally since 1997. The model produces analyses and forecasts of ocean currents, temperature, salinity, sea-surface height (SSH) and sea ice concentration. These products are used by the Royal Navy, commercially, for research purposes and for initializing the Met Office's seasonal prediction system, the Global Seasonal forecast system (GloSea5).

The current configuration of the operational FOAM system was implemented in early 2013. The system uses the hydrodynamic model Nucleus for European Modelling of the Ocean (NEMO: Madec, 2008) and the Los Alamos coupled sea ice model (CICE: Hunke and Lipscomb, 2010). The global FOAM system is run in the ORCA025 configuration, which is a 1/4 degree tripolar grid developed at Mercator–Ocean. The grid resolution ranges from 28 km at the Equator to 6 km at high latitudes. The global configuration has 75 vertical levels, with 1 m vertical resolution at the surface and decreasing resolution with increasing depth. The FOAM system assimilates satellite and *in situ* sea-surface temperature (SST) observations, *in situ* temperature and salinity profiles, altimeter sea-level anomaly (SLA) observations and satellite sea ice concentration data using a 24 h data assimilation window. A detailed description of the operational FOAM system is provided by Blockley *et al.* (2014).

An important component of the FOAM system is a data assimilation scheme. Data assimilation allows the model fields and observations to be combined to create an improved estimate of the ocean state. Data assimilation is crucial for initializing ocean model forecasts, understanding model deficiencies and constraining drift in ocean-model simulations. There are various centres producing global operational forecasts and analyses. These centres employ a variety of data assimilation techniques. The Australian Bluelink Ocean Data Assimilation System (BODAS) uses an ensemble optimal interpolation system in a near-global system (Oke *et al.*, 2008), the French Mercator system uses a static singular evolution extended Kalman (SEEK) filter (Brasseur *et al.*, 2005; Lellouche *et al.*, 2013), the US Navy oceanographic centres

© 2014 Crown Copyright, Met Office
Quarterly Journal of the Royal Meteorological Society © 2014 Royal Meteorological Society

use a multivariate three-dimensional variational data assimilation system called NCODA 3DVAR (Cummings and Smedstad, 2013) and the Japanese Meteorological Agency also use a multivariate three-dimensional variational data assimilation scheme (Usui *et al.*, 2006).

A multivariate variational data assimilation scheme, NEMOVAR (Mogensen *et al.*, 2009, 2012), has recently been implemented in the FOAM system. This replaces the Analysis Correction (AC) data assimilation method (Martin *et al.*, 2007). NEMOVAR is an incremental 3DVAR First Guess at Appropriate Time (FGAT) data assimilation scheme for the dynamical ocean model in NEMO and is developed in collaboration with the Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique (CERFACS), European Centre for Medium-Range Weather Forecasts (ECMWF) and Institut National de Recherche en Informatique et en Automatique (INRIA)/Laboratoire Jean Kuntzmann (LJK). A variational scheme provides more efficient minimization and therefore better convergence than the preceding AC scheme. NEMOVAR has the particular advantage that it has been developed specifically for NEMO and the collaboration ensures continual development of the system. The flexible formulation of the NEMOVAR control vector makes it relatively simple to include new variables. This has been exploited in extensions to NEMOVAR, such as sea ice concentration assimilation and SST bias correction. There is also ongoing work at the Met Office to implement NEMOVAR in the Operational Sea Surface Temperature and Sea Ice Analysis (OSTIA: Stark *et al.*, 2007). A variational scheme is chosen rather than an ensemble method, partially because of the limitations on achievable ensemble size at the Met Office. This would severely limit the performance of an ensemble data assimilation system. NEMOVAR has some specific attractive features, which include multivariate balance relationships (Weaver *et al.*, 2005) and a background-error correlation model based on an implicitly formulated diffusion operator (Mirouze and Weaver, 2010). Balmaseda *et al.* (2012) describe the recent operational implementation of NEMOVAR in the ECMWF seasonal forecasting system at 1 degree resolution.

The Met Office is the first centre to implement NEMOVAR at 1/4 degree resolution. A comprehensive technical description of NEMOVAR at the Met Office is provided by Waters *et al.* (2013). This article describes the specific developments to the NEMOVAR system for operational, global implementation. These include developments to the background-error standard deviations, incorporating flow-dependent vertical length-scales and implementing a simple and efficient technique for estimating the normalization factors required by the background-error correlation model based on the diffusion operator. The performance of NEMOVAR is analyzed and comparisons are made with the preceding AC data assimilation system.

In section 2, the NEMOVAR system is briefly described. The main developments to the NEMOVAR system for the Met Office 1/4 degree implementation are discussed in section 3 and in section 4 the key differences between NEMOVAR and AC implementations are listed. A one-year re-analysis with NEMOVAR is used to assess the system and the experimental set-up is described in section 5, along with the experiment for the comparative AC re-analysis. In section 6, the results from the re-analyses are assessed and the conclusions and focus for future developments are discussed in section 7.

2. NEMOVAR overview

NEMOVAR is described fully in the work of Mogensen *et al.* (2012) and its predecessor, OPAVAR, is described by Weaver *et al.* (2003, 2005). An overview of NEMOVAR is given here to provide a background to the developments described in the next sections.

The fundamental equation in NEMOVAR is the incremental cost function. This is written as follows:

$$J(\delta\mathbf{x}) = \frac{1}{2}\delta\mathbf{x}^{\mathrm{T}}\mathbf{B}^{-1}\delta\mathbf{x} + \frac{1}{2}(\mathbf{d} - \mathbf{H}\delta\mathbf{x})^{\mathrm{T}}\mathbf{R}^{-1}(\mathbf{d} - \mathbf{H}\delta\mathbf{x}), \qquad (1)$$

where the increment $\delta\mathbf{x} = \mathbf{x} - \mathbf{x}_{\mathrm{b}}$ is the difference between the state vector $\mathbf{x}$ and its background estimate $\mathbf{x}_{\mathrm{b}}$, $\mathbf{d} = \mathbf{y} - H(\mathbf{x}_i)$ is the innovation vector, $\mathbf{y}$ is the observation vector and $\mathbf{x}_i = M_{t_0 \to t_i}(\mathbf{x}_{\mathrm{b}})$. In this notation, $M_{t_0 \to t_i}$ is the nonlinear propagation model that propagates the background state to the state at time i. The ocean state vector contains temperature (T), salinity (S), sea-surface height (η), horizontal velocities (u, v) and sea ice concentration (C). The observation-error covariance matrix is denoted by $\mathbf{R}$. The observation errors are assumed uncorrelated and therefore $\mathbf{R}$ is a diagonal matrix. The background-error covariance matrix, $\mathbf{B}$, is a large ($\mathcal{O}(2 \times 10^{17})$ elements in the 1/4 model used here), full rank matrix. This makes it impossible to calculate and store $\mathbf{B}$ or its inverse matrix $\mathbf{B}^{-1}$. In NEMOVAR, $\mathbf{B}$ is defined by a parametrized covariance model, as described in this section. The operator H is the observation operator, while the matrix $\mathbf{H}$ denotes the linearized observation operator.

The cost function is minimized iteratively using a $\mathbf{B}$-preconditioned conjugate gradient (BCG) algorithm (Algorithm 2 in Gürol *et al.*, 2014). BCG requires matrix–vector multiplications with $\mathbf{B}$, not its inverse $\mathbf{B}^{-1}$.

An important feature of NEMOVAR is the balance operator, which allows covariances between different ocean variables to be accounted for. The balance relationships are specified through the balance operator K. The inverse operator K^{-1} is used to transform the state vector of mutually correlated variables, $\mathbf{x} = (T, S, \eta, u, v, C)$, into a vector of approximately uncorrelated variables, $\mathbf{x}_{\mathrm{U}} = (T, S_{\mathrm{U}}, \eta_{\mathrm{U}}, u_{\mathrm{U}}, v_{\mathrm{U}}, C_{\mathrm{U}})$, where the subscript 'U' denotes the unbalanced component of the variable. This transforms the multivariate problem into a univariate problem and thus only the univariate error covariances for $\mathbf{x}_{\mathrm{U}}$ need to be specified. The linearized balance relationships for the increments $\delta\mathbf{x} = (\delta T, \delta S, \delta\eta, \delta u, \delta v, \delta C)$, defined by Weaver *et al.* (2005), are written below. Weaver *et al.* (2005) chose temperature as the lead variable and therefore temperature is treated in totality throughout. The sequence of balance transformation can be written as

$$\begin{aligned}
\delta T &= \delta T, \\
\delta S &= \mathrm{K}^{\mathrm{b}}_{\mathrm{ST}}\delta T + \delta S_{\mathrm{U}}, \\
\delta\eta &= \mathrm{K}_{\eta\rho}\delta\rho + \delta\eta_{\mathrm{U}}, \\
\delta u &= \mathrm{K}_{up}\delta p + \delta u_{\mathrm{U}}, \\
\delta v &= \mathrm{K}_{vp}\delta p + \delta v_{\mathrm{U}}, \\
\delta C &= \delta C_{\mathrm{U}}, \qquad (2)
\end{aligned}$$

where the density increment $\delta\rho$ and the pressure increment δp are specified as follows:

$$\begin{aligned}
\delta\rho &= \mathrm{K}^{\mathrm{b}}_{\rho T}\delta T + \mathrm{K}^{\mathrm{b}}_{\rho S}\delta S, \\
\delta p &= \mathrm{K}_{p\rho}\delta\rho + \mathrm{K}_{p\eta}\delta\eta. \qquad (3)
\end{aligned}$$

The operator K_{ij} is the linearized balance operator and denotes the transformation from variable j to i. The superscript 'b' denotes cases when the operator has been linearized about the background state at the beginning of that data assimilation window. In this implementation, the state vector has been extended to incorporate sea ice concentration. This is currently treated as a totally unbalanced variable in the linearized balance relationships.

The temperature–salinity balance, K_{ST}, is a linearized version of the water–mass property conservation scheme of Troccoli and Haines (1999), where the assumption is that the temperature–salinity relationship in the water column is

© 2014 Crown Copyright, Met Office
Quarterly Journal of the Royal Meteorological Society © 2014 Royal Meteorological Society

preserved. For SSH, the balanced component, $K_{\eta\rho}\delta\rho$, is baroclinic while the unbalanced component is barotropic. The SSH balance, $K_{\eta\rho}$, is a linearized version of the Cooper and Haines (1996) scheme, whereby changes in the SSH are associated with lifting and lowering of the water column. The velocity balance, K_{up} and K_{vp}, is specified using the geostrophic relationship. The balanced density equation is the linearized equation of state, while the balanced pressure equation is the vertically integrated hydrostatic equation. For a more detailed description of the balance relationships, refer to Weaver *et al.* (2005) and Mogensen *et al.* (2012).

In NEMOVAR, the full multivariate background-error covariance matrix **B** can be written as a combination of the univariate background-error covariance matrix $\mathbf{B}_U$ and the balance operator matrix **K**:

$$\mathbf{B} = \mathbf{K}\mathbf{B}_U\mathbf{K}^T, \tag{4}$$

where $\mathbf{B}_U$ is a block diagonal matrix with each block containing the spatial covariances of the mutually uncorrelated variables $(T, S_U, \eta_U, u_U, v_U, C_U)$. We will refer to $\mathbf{B}_U$ as the unbalanced background-error covariance matrix. In the absence of velocity observations, there is no information to correct the unbalanced components of velocity in 3DVAR, so these components can be ignored in $\mathbf{B}_U$.

The unbalanced background-error covariance matrix is rewritten as follows:

$$\mathbf{B}_U = \mathbf{D}^{1/2}\mathbf{C}\mathbf{D}^{1/2}, \tag{5}$$

where $\mathbf{D}^{1/2}$ is the diagonal matrix of standard deviations and **C** the block diagonal matrix of correlations.

In NEMOVAR, the background-error correlations are modelled using a diffusion equation. Successive applications of an implicit diffusion operator are used to simulate the matrix multiplication of an autoregressive correlation matrix. The method is easily able to account for complex boundary conditions, can be used to model various different autoregressive functions and allows for geographically varying length-scales (Mirouze and Weaver, 2010). As the number of applications of the implicit diffusion operator tends to infinity, the modelled correlation function is Gaussian. In our application, we use ten applications of the diffusion operator, which provides a good approximation to a Gaussian correlation function.

Each block of the correlation matrix has the form

$$\mathbf{C} = \mathbf{\Gamma}^{1/2}\mathbf{F}\mathbf{\Gamma}^{1/2}, \tag{6}$$

where **F** is a diffusion operator and $\mathbf{\Gamma} = \mathbf{\Gamma}^{1/2}\mathbf{\Gamma}^{1/2}$ is a matrix of normalization factors to ensure that **C** has values of one along its diagonal. The 3D diffusion operator is formulated as

$$\mathbf{F} = \mathbf{F}_x^{1/2}\mathbf{F}_y^{1/2}\mathbf{F}_z^{1/2}\mathbf{W}^{-1}(\mathbf{F}_z^{1/2})^T(\mathbf{F}_y^{1/2})^T(\mathbf{F}_x^{1/2})^T, \tag{7}$$

where $\mathbf{F}_{x,y,z} = \mathbf{F}_{x,y,z}^{1/2}\mathbf{F}_{x,y,z}^{1/2}$ are 1D implicitly formulated diffusion operators applied along the component directions x, y or z and **W** is a diagonal matrix of volume elements. Ten implicit iterations are applied in each direction.

Calculating the normalization factors is a non-trivial and potentially computationally expensive task. Analytical estimates of the normalization factors can be useful when the specified length-scales (diffusion coefficients) are constant or slowly varying (Mirouze and Weaver, 2010; Yaremchuk and Carrier, 2012). The exact normalization factors can also be determined by applying $\mathcal{O}(N)$ applications of the diffusion equation on an initial field of Dirac delta functions, where N is the number of grid points. This is computationally infeasible for realistic ocean models. An alternative method for determining the normalization factors is the randomization method (Weaver and Courtier, 2001). The randomization method calculates the normalization factors by applying the diffusion operator to an ensemble of random vector fields. This method can produce more accurate results than the analytical method when length-scales vary rapidly, but requires a sufficiently large ensemble size. Further information on this aspect of the scheme is described in section 3.

The major developments made to NEMOVAR for the Met Office implementation are the extension to assimilate sea ice concentration, the incorporation of an SST bias correction scheme (Martin *et al.*, 2007; Donlon *et al.*, 2012), the inclusion of an altimeter bias correction system, which is incorporated as an extra term in the cost function (Lea *et al.*, 2008), and the specification of the error covariances.

Sea ice concentration assimilation is incorporated by adding sea ice concentration to the state vector (as described above). When the ice concentration increment is positive, it is added to the first ice category of CICE only with a specified thickness of 50 cm. For negative ice concentration increments, ice is removed from the lowest ice category and then, if necessary, from progressively higher ice categories.

In the SST bias correction scheme, mismatches between observations and reference data are used to calculate a smoothly varying bias field that is then subtracted from the observations. For our experiments, the biases are calculated using *in situ* and Advanced Along Track Scanning Radiometer (AATSR) data as the reference, with biases assumed to have a 7 degree length-scale. A separate run of NEMOVAR, external to the main minimization, is used to calculate SST bias corrections, which are then applied during the FGAT step. Note that, with AATSR no longer available, the operational system currently uses only *in situ* data as a reference dataset for bias correction. OSTIA has begun using a subset of the MetOp AVHRR data for reference, but this did not improve the results in FOAM so has not been implemented.

We focus here on the developments to the error covariances in NEMOVAR. These are discussed in the next section.

3. NEMOVAR error covariances

For the background-error covariance model (5), it is necessary to define, for each unbalanced variable, the standard deviations and the length-scales of the implied correlation functions of the diffusion operator.

For the Met Office implementation, there has been particular development of the temperature and salinity background-error covariances, in order to account for errors that are both temporally and spatially varying. In an ensemble system, the evolution of the error covariances can be estimated and updated using ensemble statistics. However, for our variational system, non-static background-error covariances are achieved by parametrizing the change in the background-error standard deviations and correlations based on knowledge of the physical processes and using the model's background field. This may not capture all the spatial (and temporal) variability in the errors; it is therefore also useful to consider climatological error covariances calculated from statistical methods. The climatological error covariances can provide us with information on the 3D spatial variation of the background errors, which can highlight regions and situations where larger errors occur.

Another important development is the inclusion of a flow-dependent vertical length-scale parametrization and an efficient method for estimating the normalization factors at each analysis step. These developments are discussed below. The sea ice concentration background-error standard deviations are also presented in this section, since the assimilation of sea ice concentration is new in this implementation of NEMOVAR.

3.1. FOAM statistical background-error standard deviations and correlation length-scales

FOAM statistical background-error standard deviations have been calculated using the National Meteorological Center (NMC)

© 2014 Crown Copyright, Met Office
Quarterly Journal of the Royal Meteorological Society © 2014 Royal Meteorological Society

method (Parrish and Derber, 1992) on one year of 24 and 48 h forecast fields. The NMC method is known to underestimate the background-error standard deviation, as it does not account for model error. Therefore the global average NMC background-error standard deviations are scaled using background-error standard deviations calculated from the Hollingsworth and Lönnberg method (Hollingsworth and Lönnberg, 1986). The Hollingsworth and Lönnberg method uses observation minus background values (the innovations) to estimate the error covariances. Making the assumption that observation and background errors are uncorrelated and that observation errors are spatially uncorrelated allows for the estimation of the background- and observation-error standard deviations. In this study, the Hollingsworth and Lönnberg error standard deviations are calculated from one year of innovation data on 1–10 degree resolution grids (depending on observation type). The global average estimates of the background- and observation-error standard deviations were used to scale the NMC results. It is necessary to use the global values for the scaling, as the observation coverage is not sufficient to produce a robust estimate of the Hollingsworth and Lönnberg errors at all grid points. The calculation of the error standard deviations was separated into four seasons to capture seasonal variability. For the AC application, a second-order autoregressive (SOAR) correlation function (see for example Gneiting, 1999) with two length-scales was fitted to the NMC statistics at all global locations. The two length-scales were fixed at 40 and 400 km. These length-scales were chosen to represent the small-scale (mesoscale) and large scale (synoptic-scale) processes. The function fitting produced fields of mescoscale and synoptic-scale NMC background-error standard deviations. Martin *et al.* (2007) provides more detail on the method used to calculate these statistical error covariances.

For this article, the term 'FOAM statistical' is used to define background-error standard deviations calculated using the combination of the NMC method and Hollingsworth and Lönnberg method described above.

In NEMOVAR, a combination of these statistical error background standard deviations and vertical parametrizations are used to define the temperature and unbalanced salinity background-error standard deviations (see sections 3.2 and 3.4). The sea ice concentration background-error standard deviations are prescribed as the FOAM statistical errors, while the unbalanced SSH errors are defined as the synoptic-scale FOAM background errors. We asume that the mesoscale errors correspond to the baroclinic errors, while a dominant proportion of the synoptic-scale errors correspond to the barotropic errors.

3.2. *Temperature background-error standard deviations*

In NEMOVAR, the temperature background-error standard deviations are parametrized based on the vertical gradients of the temperature background field (Weaver *et al.*, 2005; Daget *et al.*, 2009). The temperature background errors are assumed to be dominated by errors due to vertical displacement of the water mass. The largest background-error standard deviations therefore occur in regions where the temperature is varying rapidly in the vertical (i.e. in regions with large vertical temperature gradients).

The original parametrization was defined as follows.

(1) The temperature depth gradient field, $\mathrm{d}T/\mathrm{d}z$, is calculated from the background fields.
(2) The depth at which the maximum value of $|\mathrm{d}T/\mathrm{d}z|$ occurs is calculated. If the maximum occurs below 300 m, this depth is set at 300 m.
(3) The vertical gradient field is multiplied by a vertical scaling factor δz, such that the temperature background-error standard deviation is $\sigma_T = \delta z(\mathrm{d}T/\mathrm{d}z)$. Where σ_T exceeds the specified maximum, $\sigma_{\max}$, it is set to $\sigma_{\max}$. A lower bound σ_{ml} is set in the mixed layer and a lower bound σ_{do} is set in the deep ocean.
(4) Finally, a vertical smoothing is applied.

We have adapted this parametrization for the 1/4 degree model. The scaling factor, δz, is set equal to $20m$. This value was chosen after analytical comparisons with statistical error standard deviations of the 1/4 degree model. The maximum bound is the NEMOVAR default of $\sigma_{\max} = 1.5$ K (see Daget *et al.*, 2009). The lower bound defined above is a discontinuous function, which in the case of very low temperature stratification will cause undesirable vertical structures in the background-error standard deviations and a rapid convergence to the deep ocean value below 300 m. Since the deep ocean value is assumed to be small, this causes small background-error standard deviations at intermediate depths and under-weights the observations that have static error standard deviations. In the Met Office implementation, the lower bound has been redefined as a continuous function. This decreases smoothly from σ^T_{surf} at the surface to σ_{do} in the deep ocean and is defined as follows:

$$\sigma^T_{\min}(i,j,k) = \sigma_{\mathrm{do}} + (\sigma^T_{\mathrm{surf}}(i,j) - \sigma_{\mathrm{do}}) \exp\left[\frac{d(1) - d(k)}{L}\right], \tag{8}$$

where d is the level depth, L is a length-scale, i and j are the horizontal model coordinates and k is the vertical coordinate. The parameters in Eq. (8) are chosen to be consistent with the FOAM statistical background-error standard deviations. These statistical error standard deviations have a global deep ocean value of 0.098 K and this is used to set σ_{do} in Eq. (8). The length-scale, L, is set at 500 m and was chosen from qualitative comparisons with the statistical background-error standard deviations. The original parametrization has the advantage that it is flow-dependent, but since it is a vertical parametrization it does not account for the horizontal variability in the background-error standard deviations. In the Met Office's implementation, horizontal variability is incorporated through $\sigma^T_{\mathrm{surf}}(i,j)$ in Eq. (8), which is set equal to the FOAM statistical surface temperature background-error standard deviations.

Figure 1(a) compares the use of the original lower bound and the new continuous lower bound at 170°W, 25°N for a single day. The background-error standard deviations when using the new continuous lower bound compare more closely with the statistical error standard deviations in Figure 1(c). The new lower bound produces background-error standard deviations that are larger near the surface and converge less quickly to the deep ocean value below 600 m. The final step in the generation of the background-error standard deviations is the application of a vertical smoothing. The smoothing in this implementation is achieved with the same diffusion filter that is used to model the vertical background-error correlations. The vertical scales of the implied smoothing kernel are also the same as those used for the temperature background errors in **B** (see section 3.6). This reduces the impact of noise in the vertical temperature gradients, accounts for errors in the position of the thermocline in the background field and is internally consistent, since the smoothing depends on the assumed length-scale of the temperature background errors. Figure 1(b) shows the impact of the vertical smoothing. The resulting background-error standard deviation profile compares well with the statistical background-error standard deviations (Figure 1(c)). The peak in the parametrized values has a larger magnitude and smaller spread; this is as expected, since the parametrized error standard deviations represent the errors for one day, while the statistical error standard deviations are averaged over many days.

The parametrization has been tuned to produce similar results to the statistical error standard deviations, but has the advantage that it is flow-dependent and able to respond to changes in the depth of the thermocline.

3.3. *Unbalanced salinity background-error standard deviations*

The unbalanced salinity background-error standard deviations in NEMOVAR are vertically parametrized based on the

© 2014 Crown Copyright, Met Office
Quarterly Journal of the Royal Meteorological Society © 2014 Royal Meteorological Society

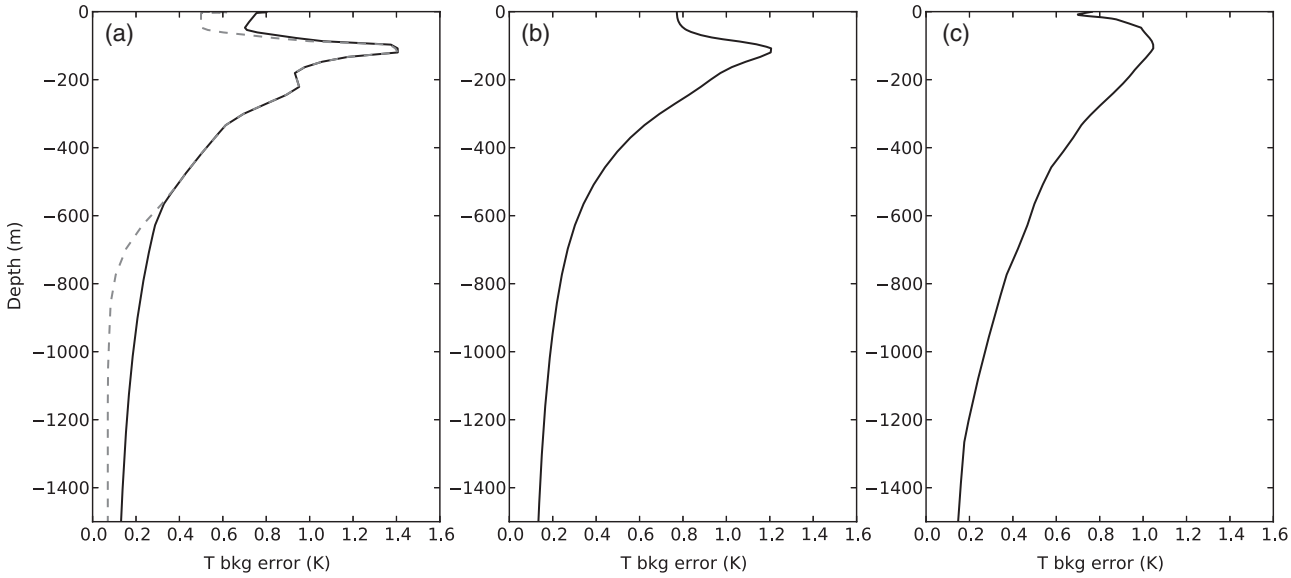

Figure 1. Temperature background-error standard deviation profiles at 170°W, 25°N on 1 January 2011. In (a), the dashed grey line shows the original NEMOVAR parametrization with the default values $\sigma_{\mathrm{ml}} = 0.5$ K and $\sigma_{\mathrm{do}} = 0.07$ K, while the black line shows the parametrization using the lower bound defined in Eq. (8). Panel (b) shows the same result as the black line in the left-hand plot, but with vertical smoothing, while (c) shows the statistical error standard deviations.

temperature–salinity gradients ($\mathrm{d}S/\mathrm{d}T$) of the background fields. The parametrization is an empirical formulation defined by Daget *et al.* (2009), which assumes that unbalanced salinity background-error standard deviations are large in the mixed layer but small in the deep ocean. The structure of the background-error standard deviations is designed to be consistent with the T–S balance. The salinity balance in NEMOVAR is a T–S balance that approximately conserves T–S water mass properties. From Ricci *et al.* (2005), the T–S structures are strong in isentropic regions such as the thermocline, making T–S property conservations important in these regions. However, in the mixed layer temperature and salinity are uncorrelated and the T–S balance does not apply. Therefore, unbalanced salinity dominates in the mixed layer, while balanced salinity dominates in the thermocline.

The parametrization of Daget *et al.* (2009) can be written as follows:

$$\sigma^{S_{\mathrm{U}}}(k) = \begin{cases} \sigma^{S_{\mathrm{U}}}_{\mathrm{max}} & d(k) \geq z_{\mathrm{max}}, \\ \sigma^{S_{\mathrm{U}}}_{\mathrm{max}}\alpha(k) & d(k) < z_{\mathrm{max}}, \end{cases} \tag{9}$$

where $\alpha(k) = (0.1 + 0.45[1 - \tanh(2\ln[d(k)/z_{\mathrm{max}}])])$ and $\sigma^{S_{\mathrm{U}}}_{\mathrm{max}} = 2.5$ psu. Daget *et al.* (2009) defined z_{max} as the depth of the maximum $\mathrm{d}S/\mathrm{d}T$. Some modifications to the parametrization have been made for our FOAM–NEMOVAR implementation. In order to make the parametrization consistent with the SSH balance, we redefined z_{max} to be equal to the mixed-layer depth. The above scheme is a flow-dependent vertical parametrization. In order to incorporate horizontal variability, $\sigma^{S_{\mathrm{U}}}_{\mathrm{max}}$ is modified so that $\sigma^{S_{\mathrm{U}}}_{\mathrm{max}} = \sigma^{S_{\mathrm{U}}}_{\mathrm{max}}(i,j) = \sigma^{S}_{\mathrm{surf}}(i,j)$, where $\sigma^{S}_{\mathrm{surf}}(i,j)$ is the FOAM statistical surface salinity background-error standard deviation.

3.4. *Sea ice concentration background-error standard deviations*

The sea ice concentration background errors used in FOAM–NEMOVAR are the FOAM statistical errors. Figure 2 shows the December–January–February sea ice concentration background-error standard deviation in the Arctic and Antarctic. In the Arctic, the largest errors are in the marginal ice zone. The errors are generally large throughout the domain in the Antarctic; this may be because these errors are for the Southern Hemisphere summer, when melt ponds are present.

3.5. *Horizontal background-error correlation length-scales*

The horizontal correlation length-scales for temperature, unbalanced salinity and sea ice concentration in FOAM–NEMOVAR are prescribed through the first baroclinic Rossby radius (as in Cummings, 2005). This is calculated from annual mean fields of temperature and salinity. The Rossby radius is assumed to correspond to the scales of the ocean mesoscale processes (for example, eddy and frontal features). The Rossby radius-dependent horizontal correlation length-scales are shown in the left-hand plot in Figure 3. The dependence of the Rossby radius on $1/f$ (where f is the Coriolis parameter) becomes a problem near the Equator, where $f \rightarrow 0$. Thus the Rossby radius has been capped at 150 km. This value was chosen from analysis of correlations computed using the NMC method at the Equator. A minimum cap of 25 km is also applied, to prevent the length-scales from being shorter than the horizontal resolution at high latitudes. The average horizontal length-scales are compared with the horizontal grid resolution in the right-hand plot in Figure 3. This illustrates that the average Rossby radius correlation length-scales exceed the grid resolution at all latitudes. For unbalanced SSH, the correlation length-scales are set at 4° to correspond with the synoptic-scale background-error standard deviations (discussed in section 3.1).

3.6. *Vertical background-error correlation length-scales*

The background vertical correlation length-scales should relate closely to the local physical ocean conditions. The temperature and salinity in the mixed layer are assumed to be highly correlated with the surface values. For SST assimilation, in particular, it is beneficial to spread the information to the base of the mixed layer, so that the wealth of SST data can lead to improvements in the modelling of mixed layer temperatures. However, the structure below the mixed layer is very different and therefore it is important that information from the surface and mixed layer is not extrapolated to the thermocline and deep ocean.

The vertical correlation length-scales developed for the FOAM–NEMOVAR implementation are based on the length-scales applied in Cummings (2005). Cummings (2005) used vertical length-scales, which are inversely dependent on the vertical density gradients. The vertical length-scales are therefore short in highly stratified regions (i.e the thermocline) and long in unstratified regions (i.e. the mixed layer and deep ocean). Some of the difficulties in applying the Cummings (2005) parametrization are in the specification of a suitable factor for scaling the inverse

© 2014 Crown Copyright, Met Office
Quarterly Journal of the Royal Meteorological Society © 2014 Royal Meteorological Society

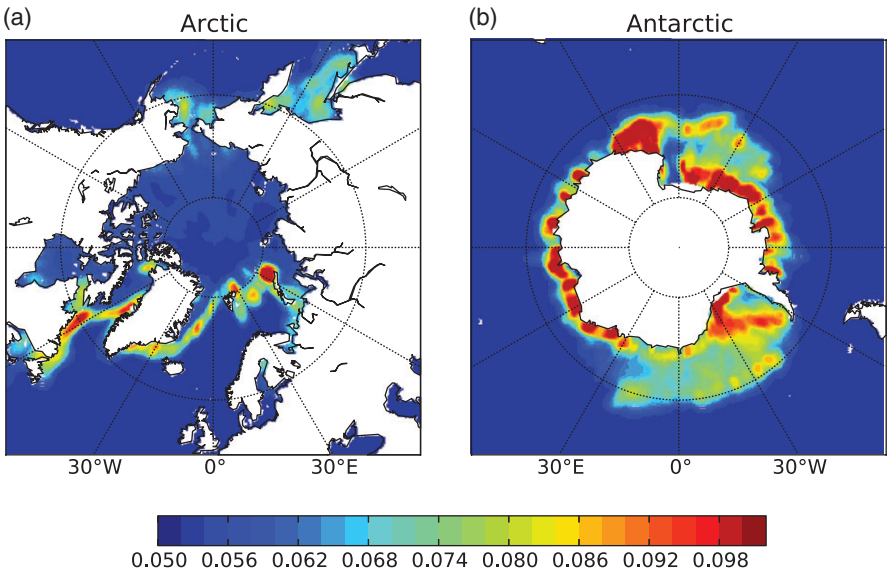

Figure 2. Sea ice concentration background-error standard deviations in fractional units. (a) Arctic and (b) Antarctic regions are shown.

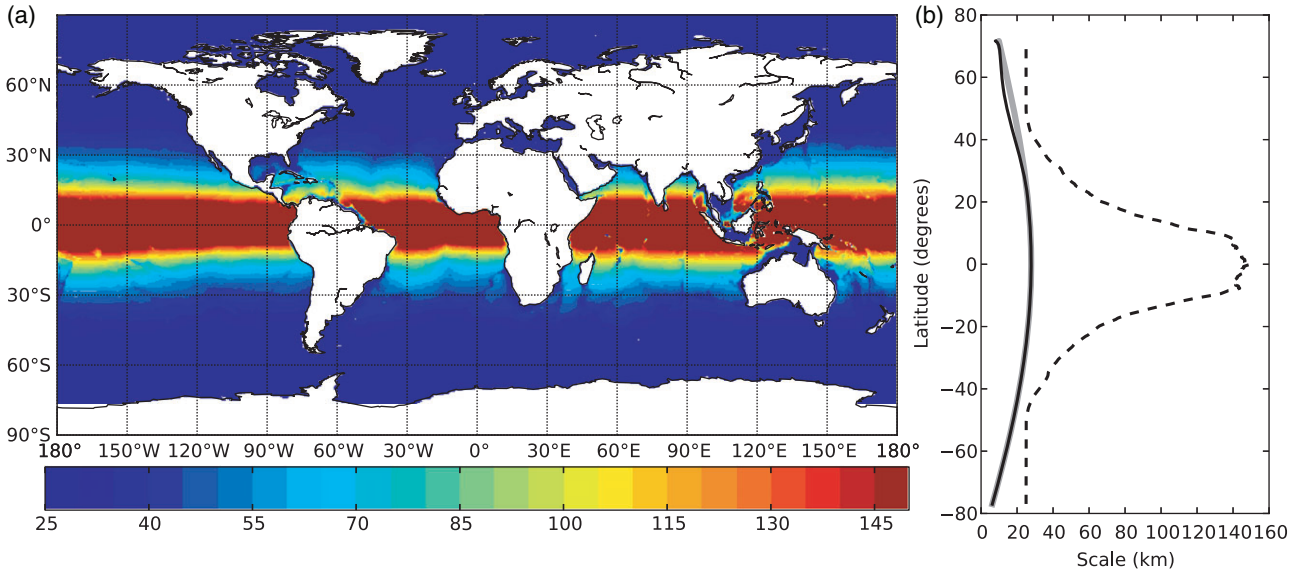

Figure 3. (a) The horizontal correlation length-scales in km. (b) The zonally averaged horizontal correlation length-scales (black dashed line) and average zonal and meridional grid resolution (solid grey and black lines, respectively), plotted against latitude.

vertical density gradients and applying a sensible cap for regions where the density gradients are small. Tests with this method found that it was difficult to find an appropriate value for capping large values and that length-scales were sensitive to very small features in the density gradients. Therefore, in the Met Office implementation the vertical length-scales have been parametrized to capture the key features of inverse vertical density gradients. At the surface, the length-scales are set as the mixed layer depth of the background state and then decrease linearly to $2\beta(k_{\mathrm{mxl}})$ at the mixed layer depth, where k_{mxl} denotes the k index associated with the mixed layer depth. The mixed layer depth is defined using the density-dependent method of Kara *et al.* (2000), which uses a density difference that is equivalent to a 0.8 K temperature change at the surface. The function $\beta(k)$ is defined as follows:

$$\beta(k) = \left[a_0 + a_1 \tanh\left(\frac{k - 15.35}{7.}\right) + a_2 \tanh\left(\frac{k - 48.03}{13.}\right)\right], \qquad (10)$$

where $a_0 = 103.95$m, $a_1 = 2.43$m and $a_2 = 100.76$m. This is the same function used to define the vertical mesh length for the 75 level version of NEMO. Below the mixed layer depth, the vertical length-scales are set equal to $2\beta(k)$. The parametrization is chosen to converge to $2\beta(k)$, as this is consistent with the vertical scales the model is able to resolve. A five-point moving average vertical filter of the correlation length-scales is then applied to provide a smooth transition of the length-scales in the thermocline. An example of the vertical length-scales when the mixed layer depth is 97 m is shown in Figure 4.

The dependence of the vertical length-scales on the background mixed layer depth ensures that the length-scales are flow-dependent. While there are clear advantages to specifying temporally varying length-scales, there are also some difficulties in its practical application. These are largely associated with the calculation of the normalization factors for the diffusion operator employed in **C**. As discussed in section 2, the analytical estimates are only suitable for constant or slowly varying length-scales. In this parametrization, the length-scales vary too quickly, especially near the base of the mixed layer. However, it would be too computationally expensive to apply the randomization method at each analysis step. The proposed solution is a normalization factor look-up table. Dobricic and Pinardi (2008) used a look-up table to generate horizontal normalization factors for a global ocean model with latitudinally varying grid resolution. They calculated the exact normalization factors for a discrete number of grid distances and then used these to interpolate to the correct

© 2014 Crown Copyright, Met Office
Quarterly Journal of the Royal Meteorological Society © 2014 Royal Meteorological Society

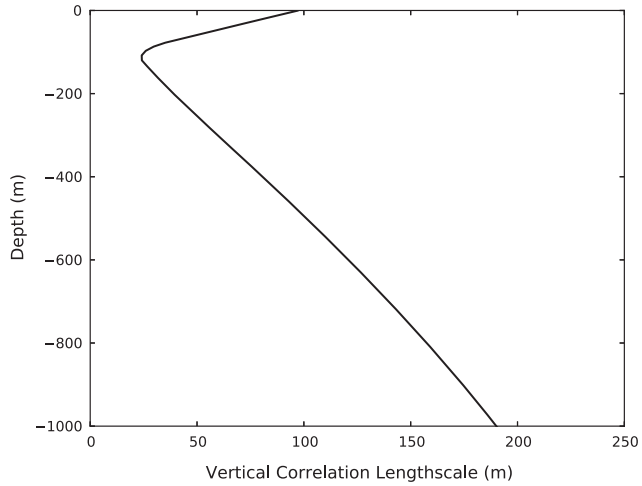

Figure 4. Vertical Correlation length-scales when the mixed layer depth is 97 m.

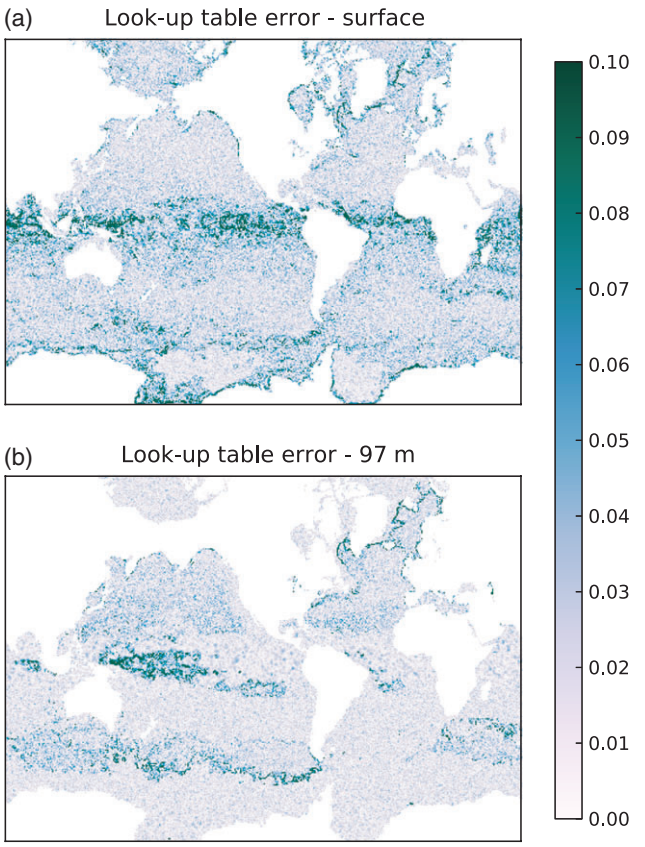

Figure 5. Relative error in the normalization factors computed using the look-up table on 1 January 2011. Panel (a) shows the surface errors, while panel (b) shows errors at 97 m.

normalization factor for each grid location. Here, a look-up table is applied to allow for variable vertical length-scales and fixed horizontal length-scales and is implemented in a different way. The first step is to limit the number of vertical length-scale parametrizations. This is done by allowing only a finite number of mixed layer depths for the vertical parametrization. In this implementation, 41 different mixed layer depths are used and these are coincident with the depths of the top 41 vertical levels of the model, with the deepest level at 565 m. The mixed layer depth is fixed at one level for all global locations and the vertical length-scales are calculated using the parametrization described above. The normalization factors for this level are then calculated using the randomization method with 3000 ensemble members; this process is then repeated for all the other discrete levels. At the end of this process there are 41 files, each containing normalization factors corresponding to 1 of 41 mixed layer depths, and this constitutes the normalization factor look-up table.

In applying this look-up table, the first step is to calculate the mixed layer depth at each grid point. This field is then smoothed horizontally using a Shapiro filter. The filtering is performed to ensure that the vertical length-scales vary smoothly in the horizontal. The next step is to generate a field of mixed layer levels corresponding to the smoothed mixed layer depths. These levels are currently calculated as the next level below the mixed layer depth (but could be calculated as the closest level). If the mixed layer depth for a particular location is deeper than the depth of level 41, the mixed layer level is set as 41. Once a field of mixed layer levels has been determined, the parametrized vertical length-scales can be calculated at each location. The final step is then to generate a field of normalization factors by extracting a profile of normalization factors at each global position from the look-up file with the corresponding mixed layer level. This generates a 3D field of normalization factors composed of profiles from the look-up table.

Figure 5 shows the relative error in the normalization factors when they are calculated using the look-up table. The errors are calculated as

$$\varepsilon = \left| \frac{\tilde{\Lambda} - \Lambda}{\Lambda} \right|,$$

where $\tilde{\Lambda}$ are the estimated normalization factors from the look-up table and Λ are the normalization factors calculated using the randomization method with 10 000 ensemble members. The errors in the look-up table are generally around 1% and increase to around 10% in some parts, such as the equatorial region and the Antarctic Circumpolar Current (ACC). In the equatorial region the horizontal length-scales are long, while in the ACC region the mixed layer depth varies sharply in the horizontal. At these locations the assumed separability of the horizontal and vertical normalization factors is less valid, as the vertical length-scales are varying too quickly over the horizontal length-scales. Table 1 shows the results from a correlation test. Correlations are modelled at several different locations using the estimated normalization factors from both the analytical method and the look-up table. The maximum amplitudes of the correlations are shown in the table; given the exact normalization factors, these should be equal to 1. For these sample locations, the error in the maximum correlation in the look-up table is no more than 5%. This is consistent with the results in Figure 5. The results are more variable for the analytical estimates. For locations near the surface, errors in the correlation are very large (greater than 50%). At depth, these errors reduce to a similar magnitude to the errors in the look-up table. The errors in the analytical estimates are larger at the surface, as this is a region where the vertical length-scales are varying quickly. Below the mixed layer depth, the vertical length-scales converge to 2β (see Eq. (10)) and the errors in the analytical normalization factors are reduced. The results here show that the look-up table provides a good approximation to the randomization method with 10 000 ensemble members and produces a vast improvement near the surface when compared with the analytical normalization factors. It is therefore found to be a viable and efficient alternative to applying the randomization method at each analysis step.

4. Summary of differences between NEMOVAR and AC schemes

In this section, the key differences between the NEMOVAR and AC implementations are summarized. The differences are presented in Table 2. In the AC scheme, two horizontal correlation length-scales are used: a 40 km mesoscale component, which approximately represents the first baroclinic Rossby radius, and a large scale of 400 km associated with errors in the atmospheric forcings at atmospheric synoptic scales. NEMOVAR can presently only specify one correlation length-scale for each control variable.

© 2014 Crown Copyright, Met Office
Quarterly Journal of the Royal Meteorological Society © 2014 Royal Meteorological Society

Table 1. Correlation test results.

		Maximum value of correlation	
Location	Depth (m)	Look-up table	Analytical method
170°W 0°N	0	0.98	1.63
170°W 30°N	0	1.04	1.76
170°W 30°S	0	0.98	1.64
25°W 0°N	97	0.99	1.04
25°W 30°N	97	0.96	0.95
25°W 30°S	97	0.98	1.10
14°W 52°N	0	0.95	1.78

The temperature, unbalanced salinity and sea ice concentration horizontal length-scales are based on the Rossby radius (as in Cummings 2005). From Figure 3, the NEMOVAR horizontal length-scales are much shorter than the AC length-scales in the extratropics. For unbalanced SSH, the correlation length-scales are defined as 4°, similar to the 400 km used in the AC system.

For temperature and salinity, there are two vertical correlation length-scales in the AC method: these are 200 and 100 m and correspond to the mesoscale and synoptic components, respectively (Martin *et al.*, 2007). For SST, the information is spread to the bottom of the mixed layer with a correlation of 1. In NEMOVAR, all temperature observations are assimilated in the same manner, so SST observations are simply treated as single-point temperature profile observations. The vertical length-scales are parametrized based on the mixed layer depth and were described in detail in section 3.6.

The same observation errors are used in NEMOVAR and the AC system. These are statistical errors determined in a similar way to the FOAM statistical background errors (Martin *et al.*, 2007). They are the NMC errors scaled using the Hollingsworth and Lönnberg observation errors.

In the AC system, the SSH balance is applied through the Cooper and Haines (1996) scheme, while geostrophic balance is used to produce balanced velocity increments. NEMOVAR is a fully multivariate system. Temperature, salinity, SSH and velocities are coupled at each iteration of the minimization of the cost function via the linearized balance operator and its adjoint.

Both NEMOVAR and AC are implemented within a 24 h window FGAT framework and both use the Incremental Analysis Updates (IAU) method to apply the increments to the model fields. For each analysis cycle, a 24 h background model run is performed from time T_0 to produce match-ups between the observations and model at the correct observations times (the FGAT step). These match-ups are used to calculate the innovations, which are used in the NEMOVAR step. The innovations are also used to calculate the innovation statistics presented in section 6. A final 24 h IAU model run is performed from T_0 to apply all the increments $(\delta T, \delta S, \delta\eta, \delta u, \delta v, \delta C)$ gradually to the model throughout the period. This is intended to reduce shock to the system. A 24 h analysis cycle is chosen to balance the requirements for short-range forecasting, while allowing enough time for reasonable spatial coverage of observations (particularly profile observations).

5. Experiment set-up

Two experiments are presented in this study: FOAM with the Analysis Correction Scheme (FOAM−AC) and FOAM with NEMOVAR (FOAM−NEMOVAR). The experiments use the same forcings, ocean model and ice model and assimilate the same observations. The model used here is the ocean component of NEMO (Madec, 2008) at version 3.2, coupled to the CICE sea ice model (Hunke and Lipscomb, 2010). This is forced by surface fields from the Met Office numerical weather prediction (NWP) system at three-hourly frequency, with the fluxes calculated online in the NEMO model using the CORE bulk formulae and interpolated in time to each model time step.

Both models were initialized using 3D temperature and salinity fields obtained from the operational FOAM system on 10 June 2010, interpolated from the 50 vertical levels used operationally at that time to the 75 vertical levels used here. SSH and velocity fields were initialized from zero and sea ice concentration was initialized from a previous two and a half year reanalysis (with sea ice concentration data assimilation). A model-only spin-up of

Table 2. Main features of the NEMOVAR and AC data assimilation systems. Entries that span the center between the AC and NEMOVAR columns represent common system features.

	AC	NEMOVAR
Data assimilation method	Analysis correction scheme	3DVAR
Background-error standard deviations	Statistical FOAM errors	*T* and unbalanced *S* use a combination of statistical FOAM errors and parametrization. Sea ice and SSH use statistical FOAM errors.
Horizontal background-error correlations	Two length-scales: 400 km (synoptic scale), 40 km (mesoscale).	*T*, unbalanced *S* and sea ice: single Rossby-radius-dependent scale. Unbalanced SSH: 4 degrees 400 km (synoptic scale).
Vertical background-error correlations	*T* and *S*: two length-scales (100 and 200 m). SST: apply increments to the base of the mixed layer.	*T*, *S* and SST: parametrized mixed layer dependent vertical length-scales.
Observation error	Statistical errors calculated with the NMC method (Parrish and Derber, 1992) and scaled using observation errors calculated from the Hollingsworth and Lönnberg (1986) method.	
Multivariate balance	Cooper and Haines (1996) scheme for SSH balance. Geostrophic balance for velocities. No *T*/*S* balance.	Multivariate system: linear balance applied at each iteration of the minimization of the cost function.
Initialization	Incremental analysis updates (Bloom *et al.*, 1996) with a 24 h time-scale	
Observation operator	Horizontal: bilinear Vertical: spline Temporal: FGAT	

© 2014 Crown Copyright, Met Office
Quarterly Journal of the Royal Meteorological Society © 2014 Royal Meteorological Society

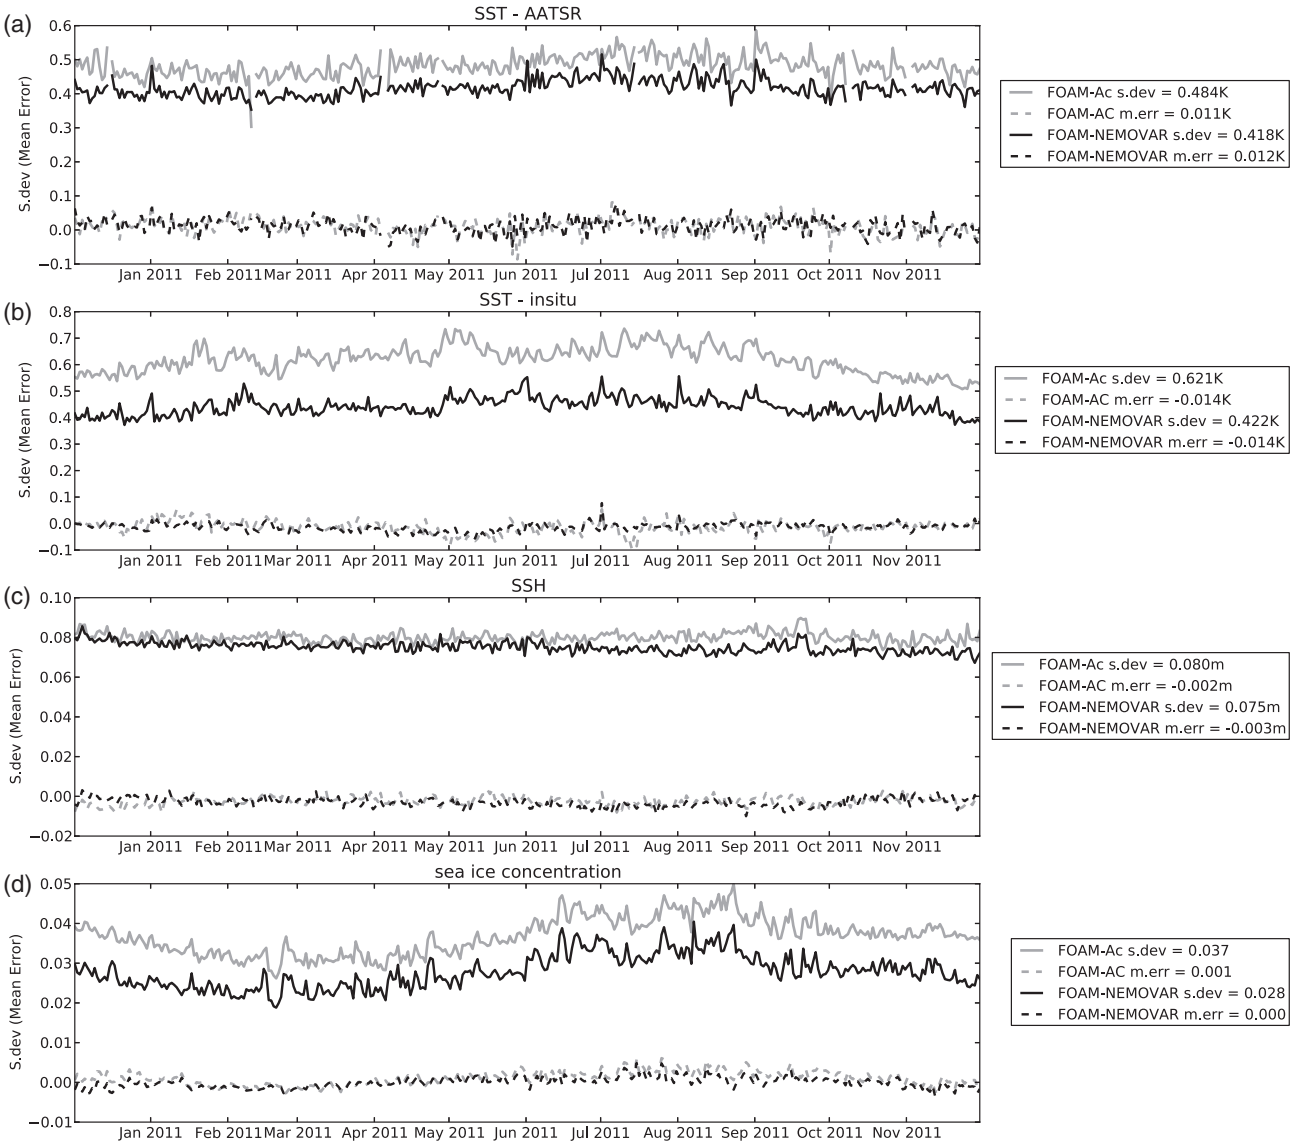

Figure 6. Global surface innovation statistics for (a) AATSR-SST, (b) insitu-SST, (c) SSH and (d) sea ice concentration. The solid line is the innovation standard deviation and the dashed line is the innovation mean. The legend shows the innovation standard deviation and innovation mean statistics for the full one-year period. The average number of observations per day for each observation type is as follows: AATSR = 95 898, *in situ* SST = 36 515, SSH = 53 003, sea ice = 883 040.

three weeks was carried out to allow the velocity and SSH fields to adjust, followed by five months with data assimilation from 1 July 2010–30 November 2010.

The main reanalysis period covers the period 1 December 2010–30 November 2011. The observations assimilated during that period are as follows.

- Satellite SST data from the Advanced Very High Resolution Radiometer (AVHRR) onboard the National Oceanic and Atmospheric Administration (NOAA) and MetOp satellites, Advanced Along-Track Scanning Radiometer (AATSR) and Advanced Microwave Scanning Radiometer–Earth Observing System (AMSRE) instruments. These are obtained through the Group for High-Resolution Sea-Surface Temperature (GHRSST) project (http://www.ghrsst.org). These are L2P data (swath data), which have been spatially averaged prior to assimilation using a 13 km radius. The foundation SST is assimilated in FOAM, so SST observations with diurnal warming are excluded by using only observations made at night or in high winds.
- Surface temperature *in situ* data from ships, moored and drifting buoys, available over the Global Telecommunications System (GTS).
- Along-track altimeter sea-level anomaly (SLA) data from the *Jason-1*, *Jason-2* and *Envisat* satellites, provided by Aviso/CLS (http://www.aviso.oceanobs.com).
- *In situ* temperature and salinity profile data from the EN3 dataset (Ingleby and Huddleston, 2007).
- Sea ice concentration data from Special Sensor Microwave Imager (SSM/I) satellites provided by EUMETSAT OSI-SAF (http://osisaf.met.no).

6. Assessment

The assessment of the FOAM–NEMOVAR system is now presented. We compared FOAM–NEMOVAR with the preceding FOAM–AC system, as this provides a benchmark for our implementation. Intercomparison studies have shown that FOAM–AC is competitive with other operational ocean forecasting systems (Oke *et al.*, 2012, for example). We consider the innovation statistics, as well as more detailed comparisons with specific observations.

6.1. *Global innovation statistics*

In this section, the global innovation statistics are presented for the FOAM–NEMOVAR and FOAM–AC trials. Innovation statistics compare the model background with the observations prior to assimilation. The statistics considered within this study are the innovation standard deviation and the innovation mean. Figure 6 shows global time series of innovation statistics for

© 2014 Crown Copyright, Met Office
Quarterly Journal of the Royal Meteorological Society © 2014 Royal Meteorological Society

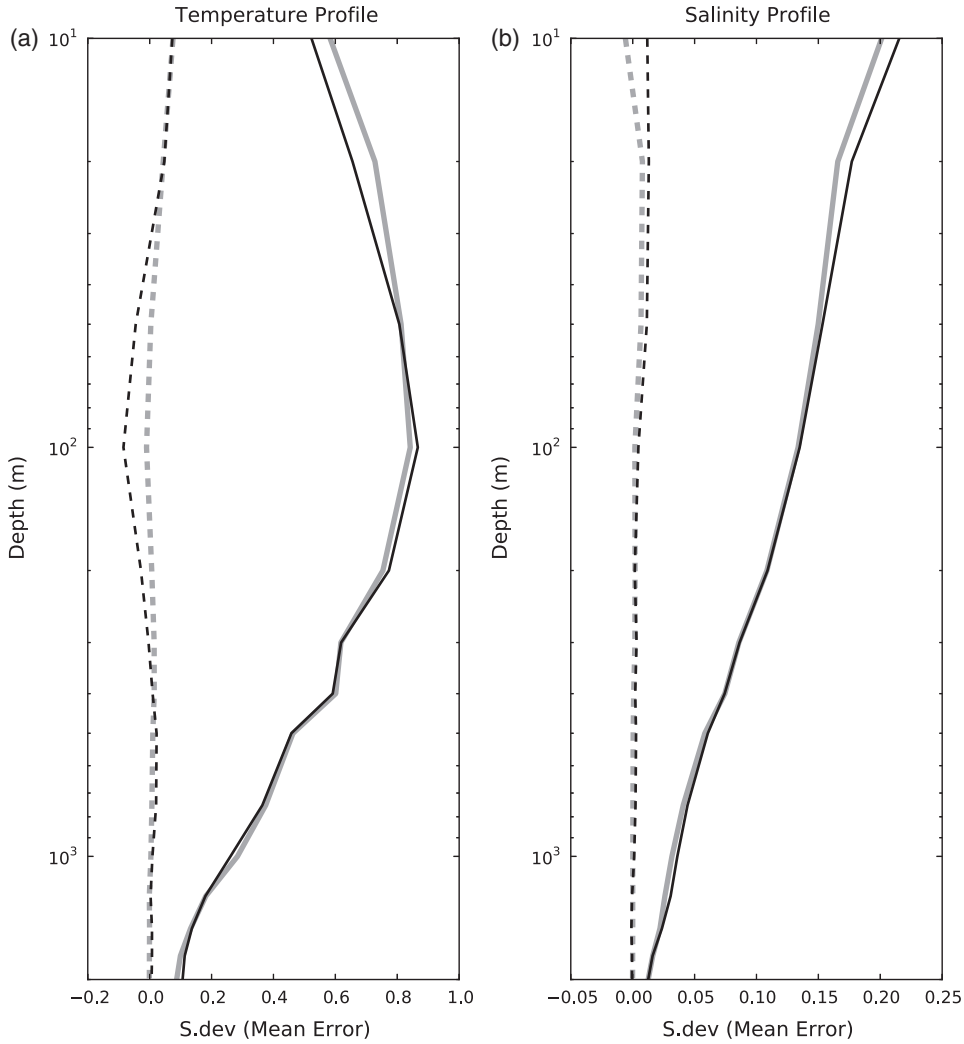

Figure 7. Global profile innovation statistics. The grey line is FOAM–AC and the black line is FOAM–NEMOVAR. The solid line is innovation standard deviation and the dashed line is the innovation mean. Panel (a) shows temperature profile statistics while panel (b) shows salinity profile statistics. For temperature, FOAM–NEMOVAR innovation standard deviation (mean) = 0.599 (0.002) and FOAM–AC innovation standard deviation (mean) = 0.610 (0.016). For salinity, FOAM–NEMOVAR innovation standard deviation (mean) = 0.121 (0.005) and FOAM–AC innovation standard deviation (mean) = 0.116 (0.001).

SST (AATSR and *in situ*), SSH and sea ice concentration for FOAM–NEMOVAR and FOAM–AC. FOAM–NEMOVAR produces a very good improvement to the innovation standard deviation for SST, with a reduction of 14% for AATSR–SST innovation standard deviation and 32% for *in situ* SST innovation standard deviation when compared with FOAM–AC. In Donlon *et al.* (2012), a target uncertainity of 0.5 K is defined for the OSTIA innovation root-mean-square error (RMS). The FOAM–NEMOVAR innovation statistics meet this target comfortably, which is a very positive result. There is also a good improvement in sea ice concentration innovation standard deviation in FOAM–NEMOVAR relative to FOAM–AC, with the innovation standard deviation reduced by 24% throughout the period. In FOAM–NEMOVAR, there is a small improvement in the SSH innovation standard deviation and this benefit increases throughout the period. This is probably related to the altimeter bias correction (Lea *et al.*, 2008). The altimeter bias correction takes several months to spin up and it seems that, while the bias correction in FOAM–AC has stabilized by the beginning of the trial period, the bias correction in FOAM–NEMOVAR is continuing to improve. The reduction of the innovation standard deviations of the surface variables in FOAM–NEMOVAR compared with FOAM–AC are thought to be related to the shorter correlation length-scales in the extratropics in FOAM–NEMOVAR and the improved fit to observations provided by the efficient minimization in NEMOVAR. The innovation means for surface variables are small and very similar for both FOAM–NEMOVAR and FOAM–AC.

Figure 7 shows the global innovation statistics over the full validation period for profile temperature and salinity binned over depth. A bootstrap method (Efron, 1979) has been applied to test the significance of the differences in these profile statistics. The differences in both the temperature mean error and innovation standard deviation from the two experiments were found to be significant at the 99% level for all depths. For salinity, the differences in mean error and innovation standard deviation are significant at the 98% level, except at 2000 m depth, where the differences for innovation standard deviation are less significant. For temperature profiles, the innovation standard deviations are generally similar for FOAM–NEMOVAR and FOAM–AC, but FOAM–NEMOVAR produces better results at the surface and slightly degraded results in the thermocline (at 100 m). However, the temperature innovation mean shows a clear degradation at 50–300 m depth for FOAM–NEMOVAR compared with FOAM–AC. This result is discussed in more detail in the next section, where it is shown that this is a model bias that FOAM–NEMOVAR is not correcting (rather than something introduced by the assimilation).

For the salinity profiles, FOAM–NEMOVAR produces a larger innovation standard deviation and mean innovation difference at the surface. This is probably associated with the shorter horizontal correlation length-scales in FOAM–NEMOVAR. While the shorter length-scales produce improvements for all other surface fields, surface salinity observations are very sparse and therefore FOAM–NEMOVAR has difficulty in constraining salinity. Future work will focus on improving the salinity statistics through developments to the background-error correlation length-scales.

© 2014 Crown Copyright, Met Office
Quarterly Journal of the Royal Meteorological Society © 2014 Royal Meteorological Society

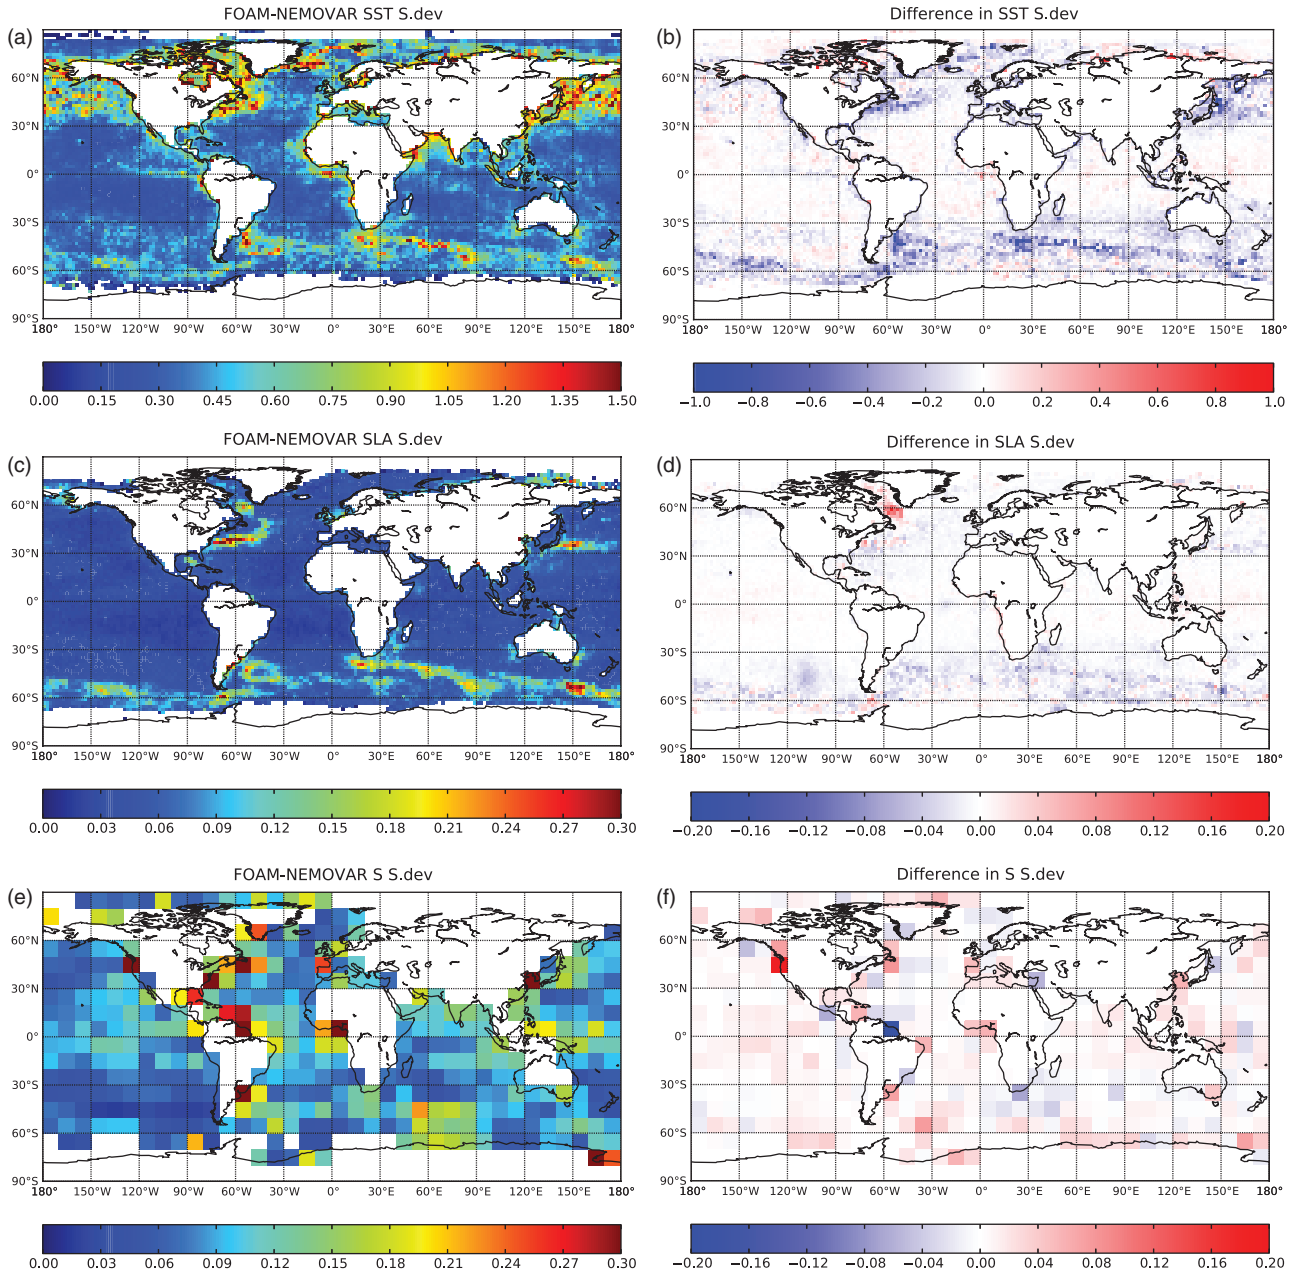

Figure 8. Binned innovation standard deviation fields for June–July–August 2011. Plots of FOAM-NEMOVAR innovation standard deviation are shown for (a) SST, (c) SLA, (e) salinity profile, (g) temperature profile and (i) sea ice concentration. Plots of innovation standard deviations from FOAM-NEMOVAR minus innovation standard deviations for FOAM-AC are also shown for (b) SST, (d) SLA, (f) salinity profile, (h) temperature profile and (j) sea ice concentration. For SST and SSH, 2° bins are used, for profile temperature and salinity 10° bins are used and data are binned over all depths and for sea ice concentration 5° bins are used.

A particular area of interest is the use of two correlation length-scales (as in FOAM–AC). The application of a longer length-scale may allow locations with sparse observations to be better constrained. Results from Blockley *et al.* (2014) also suggest that changing the surface fluxes has a larger impact on the salinity statistics than changes to the data assimilation. Therefore this degradation is relatively small compared with the impact of changes to other model components.

Oke *et al.* (2012) showed in their intercomparison study that the FOAM–AC system produces the best results for temperature and salinity profiles. FOAM already performs well for profile results and we have shown here that FOAM–NEMOVAR produces improved surface innovation statistics compared with FOAM–AC. Future developments will seek to improve the small degradations in the profile statistics for FOAM–NEMOVAR, but, as previously stated, FOAM–AC produces particularly good quality profile results.

Spatial plots of the binned innovation standard deviation for FOAM–NEMOVAR, along with the innovation standard deviation difference for FOAM–AC and FOAM–NEMOVAR, are shown in Figure 8 for June–July–August. In the difference plots, regions where FOAM–NEMOVAR has improved the innovation standard deviation compared with FOAM–AC are shown in blue. These plots show the global distribution of the innovation standard deviation for the two experiments and illustrate regions where FOAM–NEMOVAR provides improvement. For SST, there is a general reduction in the innovation standard deviation in FOAM–NEMOVAR compared with FOAM–AC in regions of high variability such as the Gulf Stream, Kuroshio Current and Antartic Circumpolar Current (ACC). A similar pattern is seen in SSH, with an improvement in FOAM–NEMOVAR in regions of high variability. There is a region of degradation in FOAM–NEMOVAR in the Labrador Sea. The results are more mixed for salinity and temperature profiles. For salinity, there is an overall degradation in FOAM–NEMOVAR; this is consistent with the results in Figure 7. FOAM–NEMOVAR does produce reduced innovation standard deviation in the Amazon outflow region, Aghulus Current and Gulf Stream region, but most other regions show a small degradation in the statistics. This suggests that the degradation of salinity is a global

© 2014 Crown Copyright, Met Office
Quarterly Journal of the Royal Meteorological Society © 2014 Royal Meteorological Society

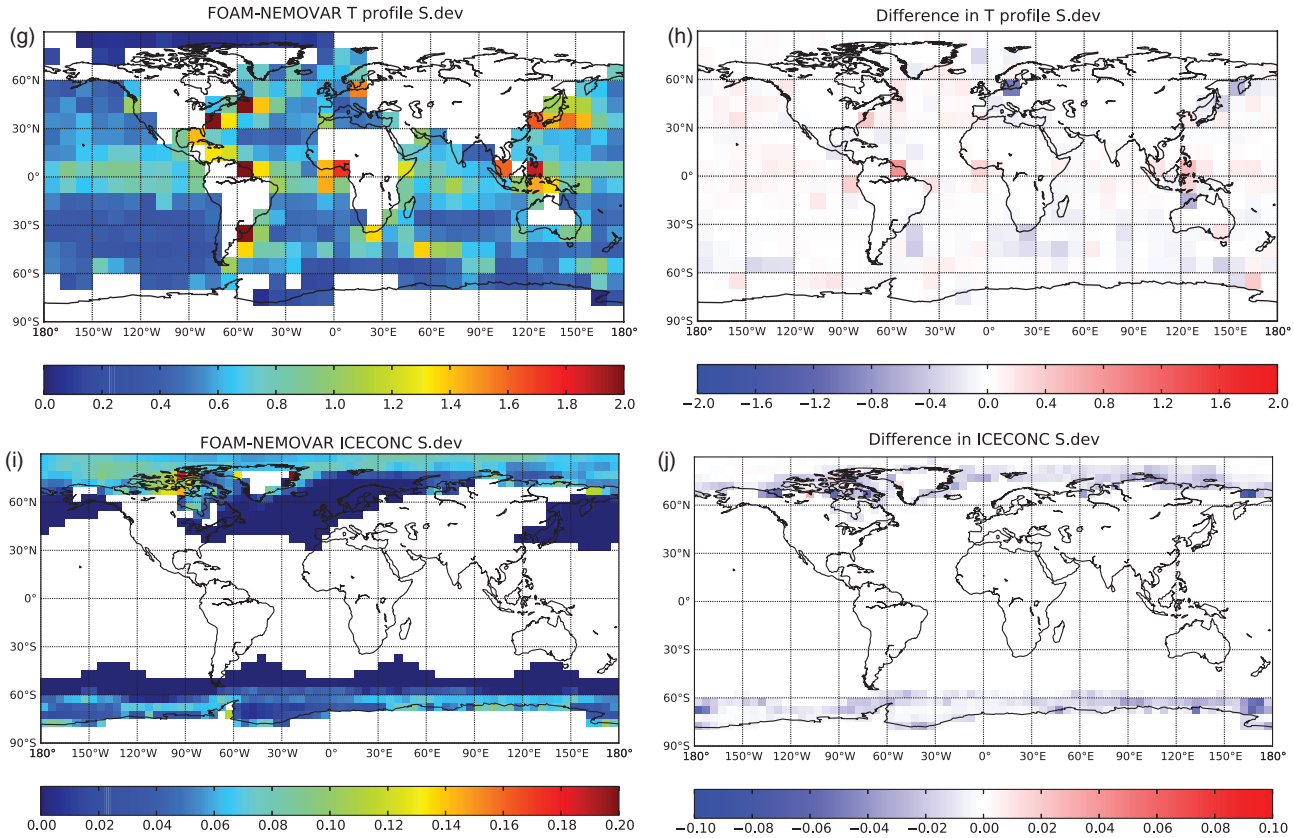

Figure 8. Continued

problem. The temperature innovation standard deviations are slightly reduced globally during the June–July–August period, from 0.63 K in FOAM–AC to 0.62 K in FOAM–NEMOVAR. Most improvements to the temperature statistics are seen across the Southern Ocean region. For sea ice concentration, there is a significant improvement in FOAM–NEMOVAR in the marginal ice zone, where sea ice is evolving most quickly.

The binned observations show that the main improvements in the surface fields in FOAM–NEMOVAR occur in regions of high variability. For SST and SSH, these are regions of frontal and eddy shedding, while for sea ice concentration these are the sea ice edge. For temperature and salinity profiles, the results are more mixed, although there is a region of improvement in FOAM–NEMOVAR in the Southern Ocean.

6.2. *Subsurface mean errors*

In the global innovation statistics, it was noted that there is a pronounced negative temperature innovation mean in FOAM–NEMOVAR at 50–300 m depth, which does not appear in FOAM–AC. This negative bias corresponds to a warm model bias relative to the observations at these depths. Investigations have shown that this bias occurs only in the extratropics. Figure 9 shows Hovmöller plots of temperature innovation mean for the Tropical Atlantic and South Atlantic regions. The Tropical Atlantic region is defined from (88.75°W, 19.75°S) to (13.75°E, 19.75°N) while the South Atlantic region is defined from (69.75°W, 59.75°S) to (29.75°E, 20.25°S). A strong negative temperature innovation bias appears in the South Atlantic plot for FOAM–NEMOVAR, but is not present in the Tropical Atlantic region.

Figure 10 presents a cross-section along 40°S of the mean temperature increments during March–April–May 2011 for FOAM–NEMOVAR and FOAM–AC. The FOAM–AC increments are primarily negative at 100 m depth, this feature is not seen in FOAM–NEMOVAR. Figure 10 also shows the March–April–May mean temperature increment fields in the South Atlantic region at 108 m. The FOAM–AC increments are predominantly negative at 108 m in the south of the domain, with an overall negative bias of −0.011 K throughout the domain. The fact that the innovation bias in FOAM–AC is significantly reduced in Figure 9 suggests that the negative increments in Figure 10 act to counter a positive model bias. In FOAM–NEMOVAR, the mean increments in the South Atlantic domain at 108 m are 0.001 K. The absence of a negative bias in the FOAM–NEMOVAR increments means that the FOAM–NEMOVAR system is not correcting this positive model bias. A key difference between the FOAM–NEMOVAR and FOAM–AC systems is the specification of the correlation length-scales. The FOAM–AC system uses two correlation length-scales and the horizontal length-scales are generally longer (particularly at high latitudes). Our investigations with FOAM–NEMOVAR suggest that increasing the horizontal length-scales while decreasing the vertical length-scales at the base of the mixed layer leads to a good reduction in the bias. A fairly extreme case was tested, where horizontal length-scales of 150 km were applied along with vertical length-scales that were set as $0.1\beta(k)$ at the mixed layer depth and in the deep ocean. These changes were implemented in the first month of the data assimilation spin-up period (July 2010). These changes reduced the minimum global mean error in the top 1000 m in FOAM–NEMOVAR from −0.114 to 0.0 K during this one-month period. However, the changes also increased the South Atlantic innovation standard deviation in FOAM–NEMOVAR to 0.785 K from 0.658 K and similar results are seen in other regions. Clearly, it is not practical to apply these changes directly in their current form, due to their negative impact on the innovation standard deviations. Shorter vertical length-scales at the base of the mixed layer are beneficial, as they prevent the spreading of information from the mixed layer into the thermocline. The vertical length-scales may need to be retuned to allow for a short length-scale at the base of the mixed layer while preserving the current length-scales near the surface and in the deep ocean. A suitable value for the length-scales at the base of the mixed layer will need further investigation; the use of of $0.1\beta(k)$ (as tested above) seems counter-intuitive, as it uses length-scales shorter than the grid

© 2014 Crown Copyright, Met Office
Quarterly Journal of the Royal Meteorological Society © 2014 Royal Meteorological Society

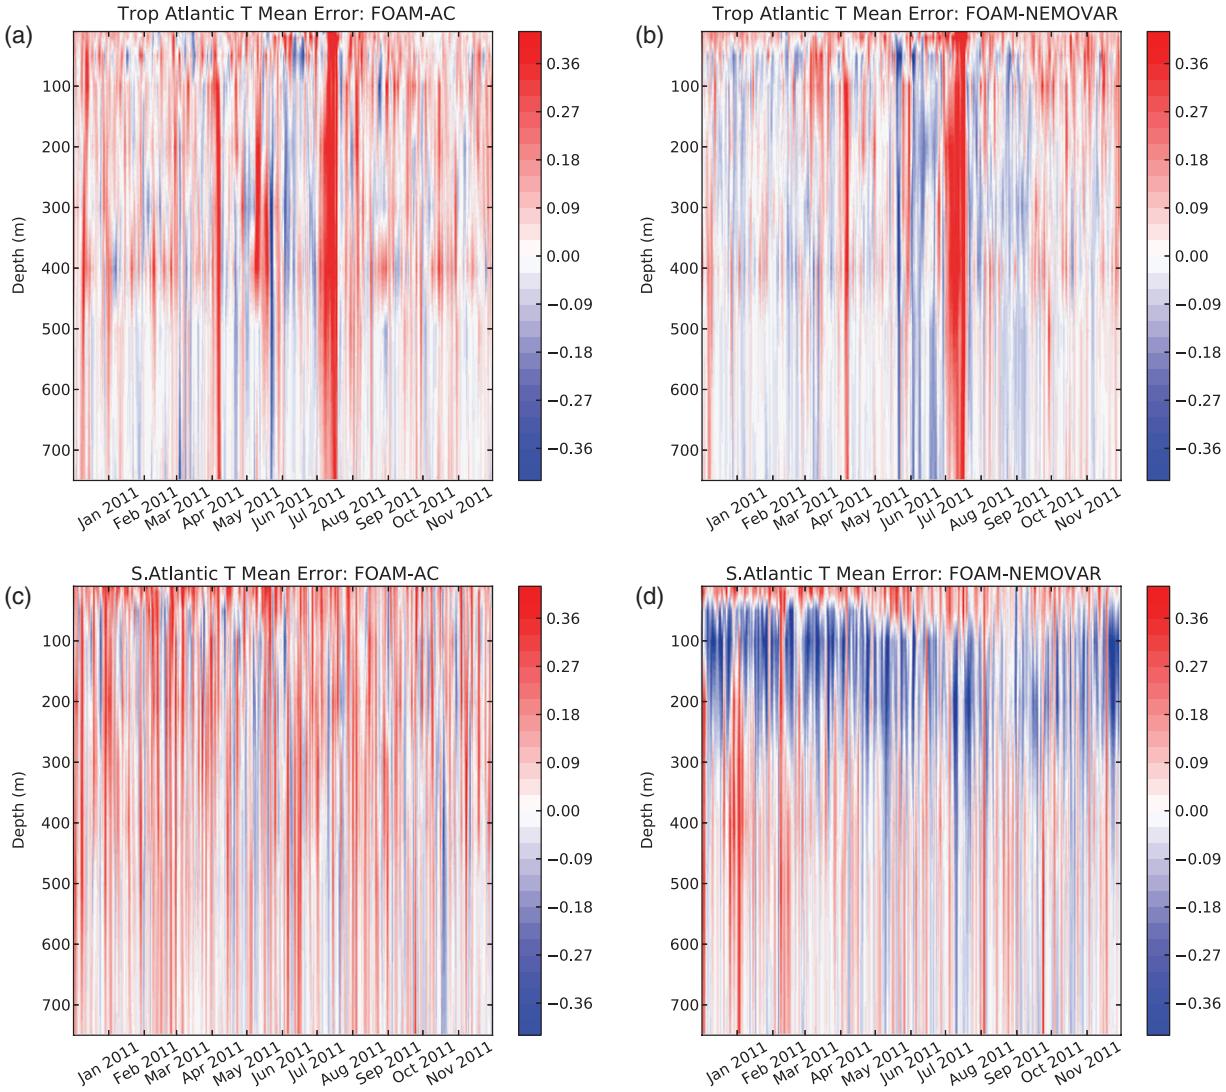

Figure 9. Temperature-profile observation minus background mean error. The top plots (a, b) show the mean error in the Tropical Atlantic, while the bottom plots (c, d) show the mean error in the South Atlantic. The right-hand plots (b, d) are from FOAM−NEMOVAR and the left-hand plots (a, c) are from FOAM−AC.

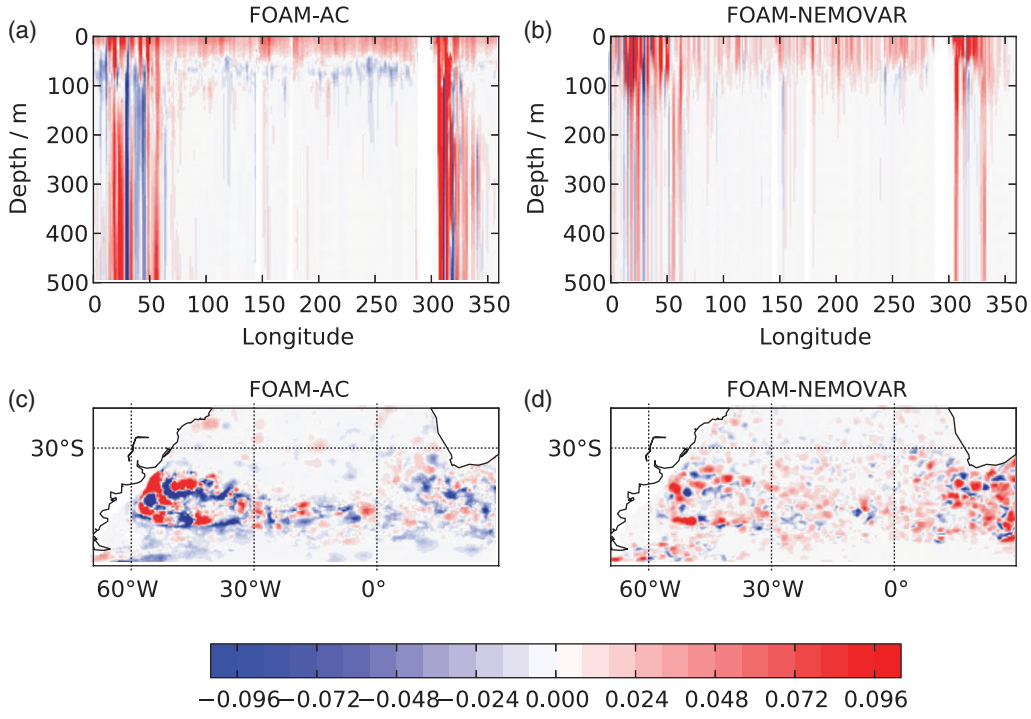

Figure 10. Mean temperature increments for March−April−May. The top plots are latitude cross-section plots at 40°S for (a) FOAM-AC, (b) FOAM-NEMOVAR. The bottom plots show the mean temperature increments at 108m in the South Atlantic domain (the same regional domain used in Figure 9) for (c) FOAM-AC, (d) FOAM-NEMOVAR.

© 2014 Crown Copyright, Met Office
Quarterly Journal of the Royal Meteorological Society © 2014 Royal Meteorological Society

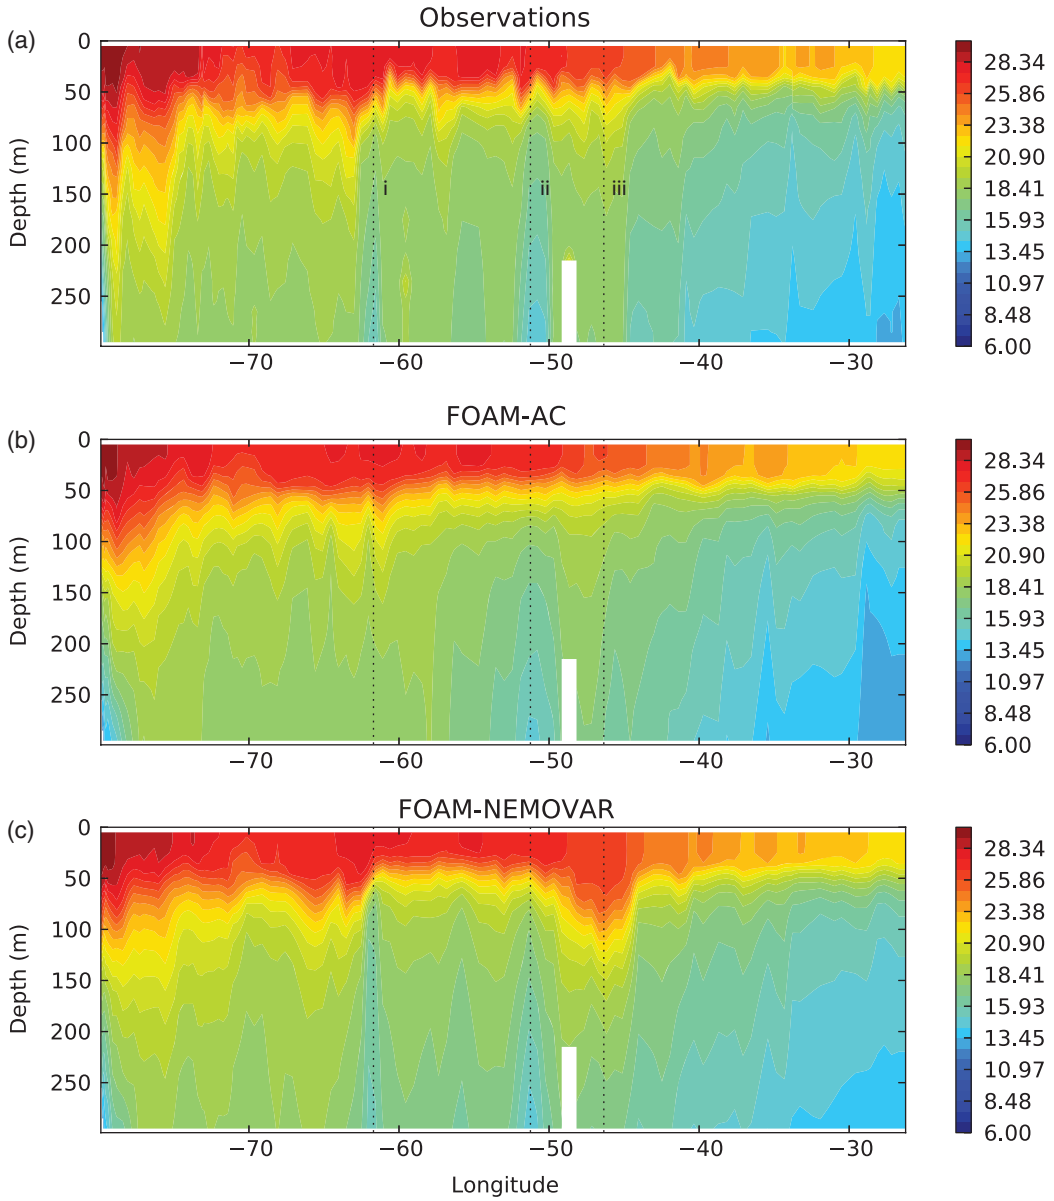

Figure 11. (a) Temperature section from an XBT track, together with the corresponding temperature sections from (b) FOAM–AC and (c) FOAM–NEMOVAR. The longitude of three specific events are marked with black dotted lines and labelled 'i', 'ii' and 'iii'. The location of the track is shown in Figure 12.

resolution. The requirement for longer horizontal correlation length-scales suggests that subsurface profile observations are currently too sparse to correct the model bias using the Rossby radius length-scales. Future developments to NEMOVAR at the Met Office include the implementation of two horizontal length-scales. This will allow for the application of a correlation function with a fatter tail. This approach may be beneficial in reducing the mean error without compromising the innovation standard deviation. Changes to allow the application of two correlation length-scales are currently being developed at the Met Office. Another potential solution is to implement a T–S bias correction scheme in NEMOVAR (Balmaseda *et al.*, 2008, 2012). This would allow for the model bias to be corrected without impacting on the T and S profile assimilation. These developments will be investigated for future implementations of FOAM–NEMOVAR at the Met Office.

6.3. *Regional case studies*

In this section, some case study comparisons with observations are made. In Figure 11, a temperature section from a cruise with multiple XBT deployments is compared with forecast temperature sections from FOAM–NEMOVAR and FOAM–AC. Note that the XBT observations are assimilated into both systems, but that the model fields presented are from the FGAT step (prior to the assimilation of the XBT observations). The location of the XBT track is shown in Figure 12, plotted on top of the SSH field. Three events have been labelled in Figures 11 and 12. A generally similar temperature structure appears in all three cross-sections. There is a distinct difference in FOAM–NEMOVAR compared with the observations at event (iii), where a warm eddy is present in FOAM–NEMOVAR that does not appear in the observations. Consideration of the SSH fields in Figure 12 shows that there is a warm eddy displaced towards the east in FOAM–NEMOVAR relative to FOAM–AC. The cold-core eddy at event (ii) in the XBT is replicated by FOAM–NEMOVAR and FOAM–AC, while the cold-core eddy at event (i) is slightly better represented in FOAM–NEMOVAR. From Figure 12, there is a visible cold-core eddy in the FOAM–AC SSH field, but it falls slightly south of the XBT track. In general, FOAM and XBT temperature sections are well matched, with the key structure in the water column and depth of the thermocline well represented in FOAM.

Figure 12 also shows the position of drifters for a seven-day period. The drifter positions suggest that both FOAM–NEMOVAR and FOAM–AC capture the position of the front well between 70°W and 80°W. In some regions FOAM–NEMOVAR produces better eddy placement, for example at (68°W, 37°N) and (45°W, 34°N), but there are a few regions where FOAM–AC is better, for example (50°W,

© 2014 Crown Copyright, Met Office
Quarterly Journal of the Royal Meteorological Society © 2014 Royal Meteorological Society

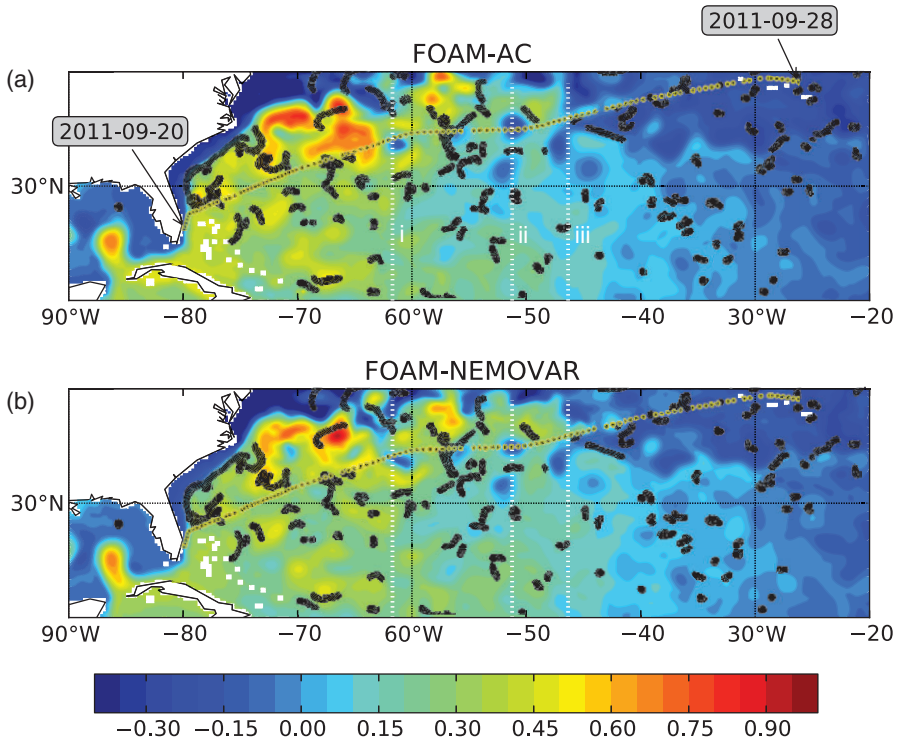

Figure 12. Model SSH fields from (a) FOAM–AC and (b) FOAM–NEMOVAR on 24 September 2011 in the Gulf Stream region. On top of the SSH fields, the XBT track from Figure 11 is plotted as the black dots with yellow outline and the location of drifters from 21 September 2011–27 September 2011 are marked by black circles. The longitudes of the specific events marked in Figure 11 are also marked on this plot by the white dotted and labelled lines.

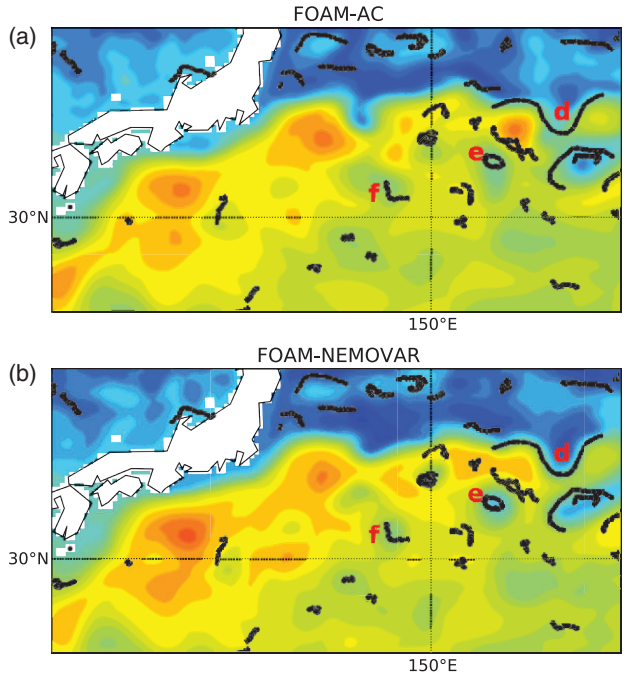

Figure 13. SSH field in the Kuroshio region on 22 November 2011 for (a) FOAM-AC and (b) FOAM-NEMOVAR. Drifter positions for 19 November 2011–25 November 2011 are plotted over the SSH fields in black. Three locations have been labelled in red as 'd', 'e' and 'f'.

37°N). Figure 13 shows a similar plot in the Kuroshio Current region. The eddy placement is better in FOAM–NEMOVAR and some regions of particular improvement are marked on the plot. Improved position of the front in FOAM–NEMOVAR is indicated by the label (d), while improved eddy positioning is indicated by the labels (e) and (f).

Overall, FOAM–NEMOVAR and FOAM–AC perform well in capturing the mesoscale features in these high variability regions and FOAM–NEMOVAR provides improved eddy and front positioning in the Kuroshio Current region. This is in agreement with the SSH and SST results in Figure 8.

6.4. *Atlantic meridional overturning circulation*

Correctly modelling processes such as the Atlantic Meridional Overturning Circulation (AMOC) is important for seasonal forecasting and climate monitoring. Since FOAM–NEMOVAR is now used to initialize the Met Office's seasonal prediction system, GLOSEA5, it is useful to consider the impact NEMOVAR has on the AMOC. An in-depth assessment of the AMOC in FOAM–NEMOVAR is presented in Roberts *et al.* (2013); here, we highlight the impact of the change in data assimilation scheme on the AMOC when all other aspects of the system are the same.

Figure 14 shows the AMOC in FOAM–NEMOVAR and FOAM–AC compared with the RAPID observations (http://www.rapid.ac.uk/rapidmoc/). The daily model values were calculated by averaging the values calculated at every time step online in FOAM, while the observations are 12 hourly RAPID observations averaged to produce daily values. The comparison period is for 1 December 2010–30 April 2011; this is the only period of coincident model and RAPID data at the time of writing. From Figure 14, there is a clear improvement in the prediction of the AMOC in FOAM–NEMOVAR compared with FOAM–AC. FOAM–NEMOVAR captures the magnitude of the step change better during January 2011 and follows the observations more closely to the end of the period. The correlation coefficient for FOAM–NEMOVAR compared with RAPID is 0.5, compared with 0.38 for FOAM–AC compared with RAPID. The mean and standard deviation of the FOAM–NEMOVAR values (14.58/5.72 Sv) are also closer to those of the data (14.21/4.48 Sv) than the FOAM–AC run (17.05/6.74 Sv).

7. Summary and discussion

The operational FOAM system at the Met Office has been updated to use the incremental 3DVAR system, NEMOVAR, as its data assimilation component. A detailed description of the main developments to NEMOVAR error covariances for the global 1/4 degree resolution model implementation is provided and discussed. These include the specification of the error covariance and a new efficient look-up table method for determining

© 2014 Crown Copyright, Met Office
Quarterly Journal of the Royal Meteorological Society © 2014 Royal Meteorological Society

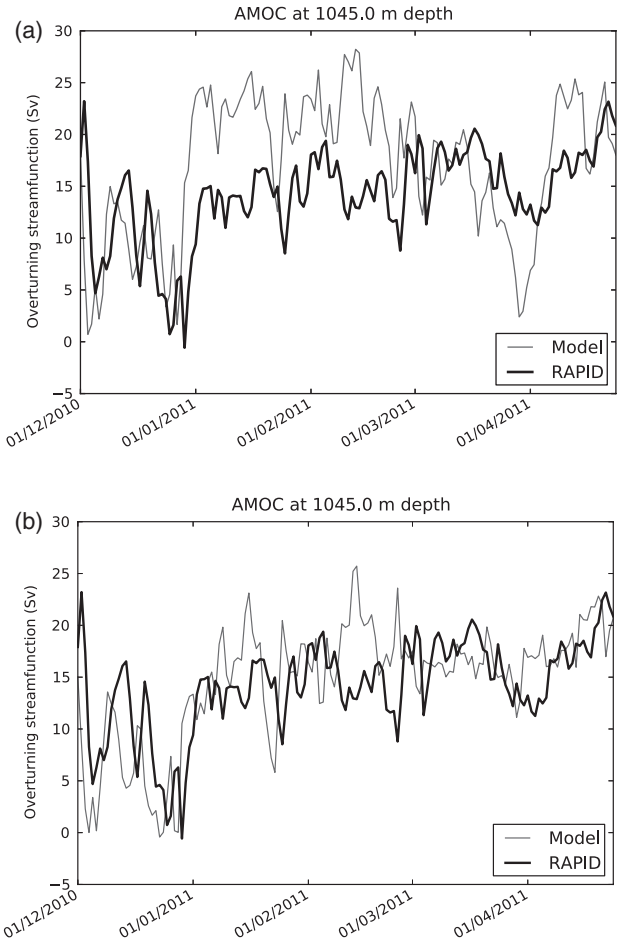

Figure 14. The Atlantic meridional overturning circulation at 1045 m: (a) FOAM–AC and (b) FOAM–NEMOVAR.

the normalization factors required by the background-error covariance operator at each analysis step when flow-dependent vertical length-scales are used.

The performance of NEMOVAR is compared with the preceding analysis correction data assimilation method within the same FOAM set-up. The two systems were compared over a one-year hindcast period from December 2010–November 2011. The assessment of the system has focused on the innovation statistics and some regional case studies, which look at comparisons with XBT and drifter observations. The innovation statistics show that the FOAM–NEMOVAR system provides significant improvement to SST, SSH and sea ice concentration. The most significant benefits to SST and SSH are seen in the frontal and eddy-shedding regions. This result is supported by comparisons of SSH fields and drifter positions in the Gulf Stream and Kuroshio current region. Meanwhile, the largest improvements to sea ice concentration are seen in the marginal sea ice regions. The temperature-profile innovation standard deviation results are similar in the two systems, with NEMOVAR showing an improvement near the surface but a small degradation in the thermocline. The surface improvements are largely attributed to shorter correlation length-scales in the extratropics in NEMOVAR and an improved global data assimilation solution in the 3DVAR framework, which allows for a closer fit to observed mesoscale features.

While there are generally good improvements to the innovation statistics, there are still some aspects requiring development. The mean error for temperature profiles is degraded at around 100 m in FOAM–NEMOVAR compared with FOAM–AC. Investigations have shown that this bias is actually a model bias that FOAM–NEMOVAR is not correcting. The salinity results are also degraded in FOAM–NEMOVAR, with an increase in innovation standard deviation in the top 300 m. Both the temperature bias and salinity innovation standard deviation results are associated with the shorter horizontal correlation length-scales in FOAM–NEMOVAR and the sparsity of profile observations. While the shorter horizontal correlation length-scales have been shown to improve the surface fields with better modelling of mesoscale features, the relative sparsity of salinity observations and subsurface temperature make it difficult for NEMOVAR to constrain these fields fully. Suggested solutions are the application of two horizontal correlation length-scales in FOAM–NEMOVAR, assimilation of satellite salinity products and the implementation of a T–S bias correction scheme. Future work will focus on developments in these areas.

The future for FOAM at the Met Office lies within a coupled ocean–atmosphere model; the improvements to the surface variables are particularly beneficial in a coupled framework and therefore the results from the surface fields are given the greater weight in this analysis. From this perspective, FOAM–NEMOVAR produces significant improvements over the FOAM–AC system. Finally, the MOC results presented in Figure 14 show a good improvement in the modelling of the AMOC in FOAM–NEMOVAR, which is particularly important in the context of seasonal prediction.

Acknowledgements

The authors acknowledge Kristian Mogensen for his significant contribution to developing the NEMOVAR system and his ongoing assistance during the FOAM–NEMOVAR implementation.

Data from the RAPID-WATCH MOC monitoring project are funded by the Natural Environment Research Council and are freely available from http://www.rapid.ac.uk/rapidmoc/.

References

Balmaseda MA, Vidard A, Anderson DLT, ECMWF. 2008. Ocean analysis system: ORA-S3. *Mon. Weather Rev.* **136**: 3018–3034.

Balmaseda MA, Mogensen K, Weaver AT. 2012. Evaluation of the ECMWF ocean reanalysis system ORAS4. *Q. J. R. Meteorol. Soc.* **139**: 1132–1161, doi: 10.1002/qj.2063.

Blockley EW, Martin MJ, McLaren AJ, Ryan AG, Waters J, Guiavarc'h C, Lea DJ, Mirouze I, Peterson KA, Sellar A, Storkey D, While J. 2014. Recent development of the Met Office operational ocean forecasting system: An overview and assessment of the new Global FOAM forecasts. *Geosci. Model Dev. Discuss.* **6**: 6219–6278, doi: 10.5194/gmdd-6-6219-2013.

Bloom SC, Takacs LL, da Silva AM, Ledvina D. 1996. Data assimilation using incremental analysis updates. *Mon. Weather Rev.* **124**: 1256–1271.

Brasseur P, Bahurel P, Bertino L, Birol F, Brankart JM, Ferry N, Losa S, Remy E, Schröter J, Skachko S, Testut CE, Tranchant B, Leeuwen PJ, Verron J. 2005. Data assimilation for marine monitoring and prediction: The MERCATOR operational assimilation system and the MERSEA development. *Q. J. R. Meteorol. Soc.* **131**: 3561–3582.

Cooper MC, Haines K. 1996. Altimetric assimilation with water property conservation. *J. Geophys. Res.* **101**: 1059–1077, doi: 10.1029/95JC02902.

Cummings JA. 2005. Operational multivariate ocean data assimilation. *Q. J. R. Meteorol. Soc.* **133**: 3583–3604.

Cummings JA, Smedstad O. 2013. Variational data assimilation for the global ocean. In *Data Assimilation for Atmospheric, Oceanic and Hydrologic Applications*, Park SK, Xu L. (eds.) II: 303–343. Springer: Berlin.

Daget N, Weaver AT, Balmaseda MA. 2009. Ensemble estimation of background-error variances in a three-dimensional variational data assimilation system for the global ocean. *Q. J. R. Meteorol. Soc.* **135**: 1071–1094.

Dobricic S, Pinardi N. 2008. An oceanographic three-dimensional variational data assimilation scheme. *Ocean Modell.* **22**: 89–105.

Donlon CJ, Martin M, Stark JD, Roberts-Jones J, Fiedler E. 2012. The operational sea surface temperature and sea ice analysis (OSTIA) system. *Remote Sens. Environ.* **116**: 140–158, doi: 10.1016/j.rse.2010.10.017.

Efron B. 1979. Bootstrap methods: Another look at the jackknife. *Ann. Stat.* **7**: 1–26.

Gneiting T. 1999. Correlation functions for atmospheric data analysis. *Q. J. R. Meteorol. Soc.* **125**: 2449–2464.

Gürol S, Weaver AT, Moore AM, Piacentini A, Arango HG, Gratton S. 2014. b-preconditioned minimization algorithms for variational data assimilation. *Q. J. R. Meteorol. Soc.* **140**: 539–556.

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.

Hunke EC, Lipscomb WH. 2010. *'CICE: The Los Alamos Sea Ice model documentation and software user's manual version 4.1'*, Technical report

© 2014 Crown Copyright, Met Office
Quarterly Journal of the Royal Meteorological Society © 2014 Royal Meteorological Society

LA-CC- 06-012. Fluid Dynamics Group, Los Alamos National Laboratory, Los Alamos, NM.

Ingleby B, Huddleston M. 2007. Quality control of ocean temperature and salinity profiles –historical and real-time data. *J. Mar. Syst.* **65**: 158–175.

Kara AB, Rochford PA, Hurlburt HE. 2000. An optimal definition for ocean mixed layer depth. *J. Geophys. Res.* **105**: 16803–16821, doi: 10.1029/2000JC900072.

Lea DJ, Drecourt JP, Haines AK, Martin MJ. 2008. Ocean altimeter assimilation with observational and model-bias correction. *Q. J. R. Meteorol. Soc.* **134**: 1761–1774.

Lellouche JM, Le Galloudec O, Drévillon M, Régnier C, Greiner E, Garric G, Ferry N, Desportes C, Testut CE, Bricaud C, Bourdallé-Badie R, Bourdall B, Tranchant M, Benkiran Y, Drillet A, Daudin A, De Nicola C. 2013. Evaluation of global monitoring and forecasting systems at Mercator Océan. *Ocean Sci.* **9**: 57–81.

Madec G. 2008. NEMO ocean engine. In *Note du Pôle de modélisation*, Vol. 27. Institut Pierre-Simon Laplace (IPSL): Paris, ISSN 1288-1619.

Martin MJ, Hines A, Bell MJ. 2007. Data assimilation in the FOAM operational short-range ocean forecasting system: A description of the scheme and its impact. *Q. J. R. Meteorol. Soc.* **133**: 981–995.

Mirouze I, Weaver AT. 2010. Representation of correlation functions in variational assimilation using an implicit diffusion operator. *Q. J. R. Meteorol. Soc.* **136**: 1421–1443.

Mogensen KS, Balmaseda MA, Weaver A, Martin MJ, Vidard A. 2009. NEMOVAR: A variational data assimilation system for the NEMO ocean model. *ECMWF Newsl.* **120**: 17–21.

Mogensen KS, Balmaseda MA, Weaver A. 2012. *The NEMOVAR ocean data assimilation system as implemented in the ECMWF ocean analysis for System 4*', Technical report 668. ECMWF, Reading, UK.

Oke R, Brassington GB, Griffin DA, Schiller A. 2008. The Bluelink ocean data assimilation system (BODAS). *Ocean Modell.* **21**: 46–70.

Oke R, Brassington GB, Cummings J, Martin M, Hernandez F. 2012. GODAE inter-comparisons in the Tasman and Coral Seas. *J. Oper. Oceanogr.* **5**: 11–24.

Parrish D, Derber JC. 1992. The National Meteorological Center's spectral statistical interpolation analysis system. *Mon. Weather Rev.* **120**: 1747–1763.

Ricci S, Weaver AT, Vialard J, Rogel R. 2005. Incorporating state-dependent temperature–salinity constraints in the background-error covariance of variational ocean data assimilation. *Mon. Weather Rev.* **133**: 317–338.

Roberts CD, Waters J, Peterson KA, Palmer MD, McCarthy GD, Frajka-Williams E, Haines K, Lea DJ, Martin MJ, Storkey D, Blockley EW, Zuo H. 2013. Atmosphere drives recent interannual variability of the Atlantic meridional overturning circulation at 26.5°N. *Geophys. Res. Lett.* **40**: 5164–5170, doi: 10.1002/grl.50930.

Stark JD, Donlon CJ, Martin M, McCulloch M. 2007. '*OSTIA: An operational, high resolution, real time, global sea surface temperature analysis system*'. In OCEANS 2007 –Europe, Aberdeen, Scotland.

Troccoli A, Haines K. 1999. Use of temperature–salinity relation in a data assimilation context. *J. Atmos. Oceanic Technol.* **16**: 2011–2025.

Usui N, Ishizaki S, Fujii Y, Tsujino H, Yasuda T, Kamachi M. 2006. Meterological Research Institute multivariate ocean variational estimate (MOVE) system: Some early results. *Adv. Space Res.* **37**: 806–822.

Waters J, Lea D, Martin M, Storkey D, While J. 2013. '*Describing the development of the new FOAM–NEMOVAR system in the global 1/4 degree configuration*', Technical report 578, Met Office, Exeter, UK.

Weaver AT, Courtier P. 2001. Correlation modelling on the sphere using a generalized diffusion equation. *Q. J. R. Meteorol. Soc.* **127**: 1815–1846.

Weaver AT, Vialard J, Anderson DLT. 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.

Weaver AT, Deltel C, Machu E, Ricci S, Daget N. 2005. A multivariate balance operator for variational ocean data assimilation. *Q. J. R. Meteorol. Soc.* **131**: 3605–3625.

Yaremchuk M, Carrier M. 2012. On the renormalization of the covariance operators. *Mon. Weather Rev.* **140**: 637–649.

© 2014 Crown Copyright, Met Office
Quarterly Journal of the Royal Meteorological Society © 2014 Royal Meteorological Society