Implementation and tuning of an altimeter data assimilation scheme for high resolution FOAM models

Ocean Applications Technical Note No. 26

Adrian Hines

January 2001

Summary

The Cooper and Haines (1996) scheme for assimilating satellite altimeter sea surface height data has been implemented in the FOAM high resolution nested models. Tests of the scheme in the 1/3° North Atlantic and 1/9° Caribbean and Gulf of Mexico models highlighted deficiencies in the implementation of the scheme. Results are presented from a series of tuning experiments which has led to considerable improvements in the effectiveness of the scheme. Data usage issues, specifically quality control and the choice of mean sea surface height, are considered. The tuning has produced significant improvement in the ability of the assimilation to initialise mesoscale features.

Table of Contents

1. Introduction

Satellite measurements of the sea surface height (SSH) and sea surface temperature (SST) now provide accurate observations of ocean parameters with near-global coverage. These observations complement the relatively sparse network of measurements of subsurface ocean properties. The coverage of subsurface measurements is likely to increase dramatically in the next few years as a result of the ARGO project (ARGO Science Team, 1998), and it is the coherent use of the observations from the various sources in combination which offers the greatest potential for ocean state estimation.

The first satellite altimeter missions, which flew in the 1980's, enabled the potential of satellite SSH data to be explored, and laid the foundations for current missions. Three satellite altimeters are currently in operation, two of which provide data daily on an on-going pre-operational basis: Topex/Poseidon and ERS-2. The third, Geosat Follow-On, has recently been accepted as operational, with data soon to be available. Numerous corrections are required to produce measurements of the dynamic SSH signal from the satellite data. Corrections for atmospheric effects are frequently derived from atmospheric numerical weather prediction (NWP) models. Orbit errors, however, provide a bigger challenge. Best orbit solutions rely upon a cross-over analysis of several satellite tracks, and hence cannot be provided in near real-time. Fast delivery orbit solutions are less accurate, but recent progress has produced solutions of sufficient accuracy to allow effective use of the data. A further difficulty is caused by poor knowledge of the Earth's gravitational field at the surface, the geoid. Variations in the geoid over a wide range of spatial scales cause variations in SSH, which need to be removed to obtain the dynamic SSH. Current geoid models are not sufficiently accurate to enable the geoid to be removed, and despite dedicated future missions to map the geoid such as GRACE and GOCE, it will be several years before this is a practical option. Instead, the fact that the geoid is a time invariant field on the time-scales of interest is exploited by subtracting a multi-year mean from the along-track data. However, in doing so the time mean dynamic SSH is also removed, and hence the resulting data anomalies must be added back to a suitable mean dynamic SSH. A mean dynamic SSH derived from model output or hydrographic data is frequently used to reconstruct the full SSH. The choice of mean dynamic SSH has been shown to have a significant impact on model results (Killworth et al., 2000), with the potential to alter the location of major current systems by hundreds of kilometres.

Meanwhile, whilst greater computing power has enabled higher resolution models to be run, the upper practical limit on model resolution is still dictated by computing capacity. To overcome this restriction, limited area models at higher resolution have been developed, driven by boundary data provided by lower resolution wider area models. At the Met Office, a 1/3° North Atlantic model has been nested within the 1° global Forecasting Ocean Assimilation Model (FOAM). A further 1/9° model covering the Caribbean and Gulf of Mexico has been nested within the 1/3° model. The FOAM system and the nested models are described in more detail in Section 3. The nested model system is designed with a view to providing forecasts of the mesoscale ocean structure, and will rely upon the use of altimeter data to provide accurate initial states from which to forecast.

In parallel to the development of the high resolution modelling capability, an altimeter SSH data assimilation scheme has been implemented within FOAM. The scheme is based on the method of Cooper and Haines (1996), and was initially implemented in a tropical Pacific version of FOAM

(Forbes, 1996). Initial development work, and subsequent tuning have been required to adapt the scheme for use in the high resolution nested models.

Section 2 summarises the development of techniques for assimilating altimeter SSH data. Section 3 gives details of the FOAM model formulation and the assimilation scheme, including the implementation of altimeter SSH data assimilation. Present and future sources of altimeter data, and their suitability for observing the mesoscale ocean structure are also discussed. Results from the initial implementation of the altimeter data assimilation scheme are discussed in Section 4, which highlights the weaknesses in the initial configuration. The tuning of the scheme is described in Section 5, and the extent to which the weaknesses are adressed is considered. Section 6 presents the conclusions.

2. Development of techniques for assimilation of altimeter SSH data

Assimilation and analysis of SST and temperature profile data is well established. Well-tuned traditional statistical assimilation schemes based around Optimal Interpolation (e.g. Clancy et al, 1992; Bell et al, 2000) and more recently developed advanced assimilation techniques, such as the Ensemble Kalman Filter (e.g. Evensen and Van Leeuwen, 1996) both show skill in producing analyses of the ocean state. In contrast to satellite SST data, satellite altimeter SSH data contains information about the structure of the whole water column, since the SSH signal is related to the vertical integral of the ocean density. The inference of subsurface properties from the SSH field has posed a new challenge for data assimilators.

Early studies looked to exploit observational evidence identifying strong correlations between the depth of the 15°C isotherm and the sea surface height in regions such as the western North Atlantic (Cheney, 1982), employing models of reduced complexity, in terms of either dynamics or vertical resolution. Hurlburt (1986) investigated the impact of simulated altimeter SSH data in a 2-layer, free-surface primitive equation model formulated on a β-plane. With direct insertion of the simulated SSH field the model was able to reproduce the pycnocline depth, and reconstruct the subsurface fields effectively. Furthermore, with initial errors of 30-50% in the subsurface pressure field, the model showed forecast skill over time periods of 2 months or more. Moreover, the study demonstrated the feasibility of extracting sub-surface information from SSH data.

The use of satellite SSH data in quasi-geostrophic (q-g) layered models was considered by Haines (1991), using simulated SSH data in a 4-layer q-g model. Noting that potential vorticity anomalies in the deep ocean layers are likely to be weak, Haines (1991) used direct data insertion to update the surface layer potential vorticity field so as to obtain the observed SSH values, whilst preserving potential vorticity in lower layers. The effect of updating the surface layer fields is then propagated to lower layers by the model itself. The technique was shown to be effective, even with a 40-day interval between data insertion. This study provided a basis for further work investigating the use of altimeter SSH data in q-g models.

Blayo et al. (1994) applied the combination of a 4-layer, 1/6° horizontal resolution q-g model with a simple SSH assimilation scheme to the simulation of the North Atlantic circulation. Despite the sub-optimality of their assimilation scheme, they found the approach to be feasible: the model and assimilation successfully reproduced the gross features of the North Atlantic circulation. The potential of the system to provide effective forecasts was subsequently examined by Brasseur et al. (1996) through a series of identical twin and predictability experiments. They

conclude that forecasts of mesoscale eddies remain realistic for only a couple of weeks in advance.

A more sophisticated variational assimilation technique was applied to a q-g model of the Newfoundland Ridge and Basin by Cong et al. (1998). Their variational approach used both initial conditions and boundary conditions as control variables, and was used to assimilate real altimeter SSH data. Results showed significant correlation between the reconstructed velocity field and data from near-surface and sub-surface current meters. Although the assimilation of SSH data has shown success in q-g models, using the data in models with more general dynamics, and greater vertical resolution requires alternative techniques.

Techniques for assimilating altimeter SSH data in full primitive equation ocean models have so far mostly relied upon projecting SSH increments on to subsurface temperature and salinity increments. Various approaches have been successfully adopted, which can be broadly categorised as statistically or dynamically based schemes. Statistical schemes directly exploit statistical correlations between SSH increments and either subsurface temperature and salinity increments (Mellor and Ezer, 1991), or subsurface pressure increments (Smedstad and Fox, 1993) to project the increments downwards. Other schemes use Empirical Orthogonal Functions (e.g. Gavart and De Mey, 1997; Verron et al., 1999) to calculate the subsurface increments. These approaches have been applied to the problem of simulating the Gulf Stream, with identical twin experiments showing the schemes to be effective. Mellor and Ezer (1991) observe a reduction in nowcast error of around 50% when assimilating simulated data along satellite tracks, and note that with perfect knowledge of the SSH field errors would be reduced to about 15%. Forecasts show skill (compared to persistence) out to 10-20 days, but little skill beyond 30 days. Smedstad and Fox (1993) observed a reduction in the error in the upper layer pressure field of a two-layer model to about 5% after assimilating complete SSH fields. Both studies consider the impact of multiple satellite altimeters, and conclude that two altimeters with appropriate orbit characteristics would lead to significant improvement in analyses. Similar reductions in errors are found by Verron et al. (1999) using an Extended Kalman Filter in twin experiments in a primitive equation model of the Tropical Pacific. Using perfect data with full spatial coverage, errors are reduced to about 5%; adding artificial noise of up to 5cm rms to the data increases the residual error to about 10%. They note, however, that the effect of erroneous wind forcing is more significant, increasing residual errors to around 15%.

Whilst the statistical techniques show success in identical twin experiments, it is clear that, even with perfect knowledge of the ocean SSH, error remains, and the sub-surface ocean state is not uniquely determined by the SSH observations. Furthermore, deficiencies in surface forcing fields may also act to degrade the quality of ocean analyses. In addition, the statistical schemes use correlations derived from the model being employed, which also provides the simulated data for the identical twin experiments. Hence the correlation statistics used in these experiments are necessarily correct for the simulated data. Statistical correlations between real surface data and subsurface anomalies are likely to differ from those derived from model output (and subsurface data coverage is not currently sufficient to allow their computation). Hence the identical twin experiments are likely to significantly over-estimate the ability of the schemes to reduce errors. Assimilating real data therefore, with limited sampling in both time and space, inevitably produces more limited success. Ezer and Mellor (1994) apply their statistical technique to assimilate Geosat altimeter data in a model of the Gulf Stream. Although they observe that errors with SSH assimilation are consistently smaller than when not assimilating, they note that some mesoscale meanders and rings are not well captured by the assimilation. Nowcasts of Gulf Stream axis location show greater skill than those of temperature variations.

Meanwhile, alternative approaches have been developed which rely upon relationships between physical properties to compute subsurface increments. Cooper and Haines (1996) (hereafter CH) developed a method for projecting the SSH increments that preserves temperature and salinity on isopycnal surfaces. The scheme is based on the observation that surface height anomalies are not generally well correlated with subsurface potential vorticity anomalies on isopycnal surfaces, but tend to be better correlated with anomalies in isopycnal depth. The scheme seeks to preserve potential vorticity on subsurface isopycnal surfaces. Surface height anomalies are balanced by subsurface temperature and salinity anomalies arising from a vertical displacement of the water column through a distance calculated from the hydrostatic relation (see Section 3 for a full description of the scheme). The scheme relies upon the assumption that the SSH increments are entirely baroclinic, which is not always appropriate, particularly at high latitudes. Furthermore, the choice of vertical displacement of the model water masses is somewhat arbitrary; horizontal displacement could be equally effective, although perhaps more difficult to implement. Despite these drawbacks, CH observed some success in applying the scheme to an idealised 21-level primitive equation model of the Gulf Stream. Using a re-initialisation approach, with data assimilated every 9 days, and geostrophic balancing velocity increments applied, global temperature rms errors and errors in the deep ocean currents are reduced to around 40%. Further identical twin experiments using the scheme in a Mediterranean general circulation model (Drakopoulos et al., 1997) showed similar improvements in the temperature and velocity fields. Additionally, the impact of using daily wind forcing was assessed, and was found to lead to reductions in temperature and velocity errors of around 50% without assimilation. The addition of altimeter data assimilation using an extension of the CH scheme to include an update to bottom pressure produced a further 30% error reduction, but further experiments suggested that such an update may be unnecessary if wind stresses are correctly known.

One of the key advantages of the CH scheme lies in the fact that, since it is based on physical principles, it is model independent. Hence it is insensitive to model changes, and can be readily incorporated in an existing model. This is in contrast to the statistically based schemes, which rely upon correlations to project the surface data downward. Due to the sparse availability of sub-surface observations, such correlations are necessarily derived from model data, and hence are model-specific. Significant changes to model physics would necessitate a re-calculation of the underlying correlations. There are, however, regions where the assumptions behind the CH method are inappropriate, in particular high latitudes. At present, data in these regions are eliminated, but further development of the method may extend its practicality for these locations.

More recently the CH scheme has been implemented in the OCCAM 1/4°, 36 level global model. In the first part of a two-part paper, Fox et al. (1999a) consider a series of identical twin experiments, assimilating simulated data using a re-initialisation approach. Simulated mapped SSH data is assimilated every 10 days, with balancing geostrophic velocity increments applied. They observe a 60% reduction in error in temperature, salinity and current fields within the top 1000m of the ocean, and a 40% reduction below, with errors remaining low in a subsequent 60 day forecast run. They note problems in weakly stratified regions of the Arctic and Antarctic seas, which they avoid through applying a scale selective filter to the altimeter data which removes scales much larger than the local deformation radius. These identical twin experiments formed the basis for the 2nd part of the paper, which describes an assimilation of real altimeter data into OCCAM (Fox et al., 1999b). Prior to assimilation, mapped sea level anomalies are added to a mean SSH which is modified in the Gulf Stream and Kuroshio regions (see Section 5.5 for a detailed discussion of the choice of mean SSH). Sea level updates are weighted by sea level errors

computed at low resolution using an adaptive method. As a consequence of the choice of mean SSH, the assimilation is successful in improving separation points of western boundary currents. Altimeter data sea level anomalies are successfully propagated between assimilation times, although damping of small scale features is noted. Predictions out to 40 days show skill compared to persistence.

Using the combined statistical-dynamical assimilation method of Oschlies and Willebrand (1996), Killworth et al (2000) have shown further evidence that the choice of mean SSH for use with the data anomalies is significant. Using a level model of the North Atlantic, and various representations of the mean SSH, they observe the rapid adjustment towards reality of the mean interior flow in the top 1000m of the ocean when a more accurate mean SSH is used. They observe that the location of major current systems can be adjusted by hundreds of kilometres through the choice of mean SSH. These impacts are in addition to the improved variability produced by asismilation of the altimeter data.

Development of operational systems for assimilation of altimeter data has followed from the development of the basic techniques. The US Navy's Modular Ocean Data Assimilation System (MODAS), based upon a statistical technique of Carnes et al. (1990, 1994), produces daily analyses of the subsurface ocean structure on a 1/8° grid using real-time altimeter data, whilst versions of the NRL Layered Ocean Model (Metzger et al., 1998) assimilating altimeter data are being transitioned to operational use. The Harvard Ocean Prediction Scheme (HOPS) provides operational forecasts of the Gulf Stream region using a q-g model initialised using 'feature models' which are constrained by altimeter SSH data (Glenn and Robinson, 1995). A q-g model of the Azores region forms the basis of the French SOAP93 operational forecast system (Le Squere and Dombrowsky, 1994) which assimilates altimeter SSH data using the EOF assimilation method of Dombrowsky and De Mey (1992).

3. The FOAM ocean model, assimilation scheme and data sources

3.1 FOAM Ocean Model

The FOAM ocean model is a primitive equation, level model based upon the original formulation of Bryan (1969). The model forms the ocean component of the Met. Office Unified Model system, and hence is largely the same code as used for seasonal forecasting and climate prediction. Physical parametrisations incorporated in the model used in these experiments are the vertical diffusion scheme of Pacanowski and Philander (1981); a mixed layer scheme based on Kraus and Turner (1967) and Davis et al. (1981); vertical mixing of horizontal momentum based upon Large et al. (1994); lateral mixing of tracers along isopycnals using the method of Redi (1982). The FOAM system includes a thermodynamic sea-ice model based on Semtner (1976) and Hibler (1979), and a dynamic sea-ice component based upon Bryan et al. (1975). Full details of the basic model formulation used in these experiments are given in Bell et al. (2000) and the formulation of the nested models is described in Storkey (2001).

A global 1° resolution version of the FOAM system has been used to provide boundary conditions for a nested 1/3° North Atlantic model, which in turn provides boundary conditions for a 1/9° model of the Caribbean and Gulf of Mexico. The nested model domains are shown in Figure 1. All three models have 20 vertical levels which are unevenly spaced. Vertical resolution decreases with depth, from 10m at the surface, to 600 m at the lowest level which is at a depth of 5500m. Bathymetry for the global model has been interpolated from the DBDB5 bathymetry

data set (U. S. Naval Oceanographic Office, 1983), whilst that for the 1/3° and 1/9° models is derived from the Smith and Sandwell (1997) data set. Both nested models make use of a rotated grid, with the north pole at 17°N, 56°E. In the 1/3° model this enables the inclusion of the whole of the North Atlantic and the Arctic basin (Figure 1(a)). Boundary conditions at the open boundaries are applied using the Flow Relaxation Scheme, relaxing the model tracer, velocity and streamfunction fields.

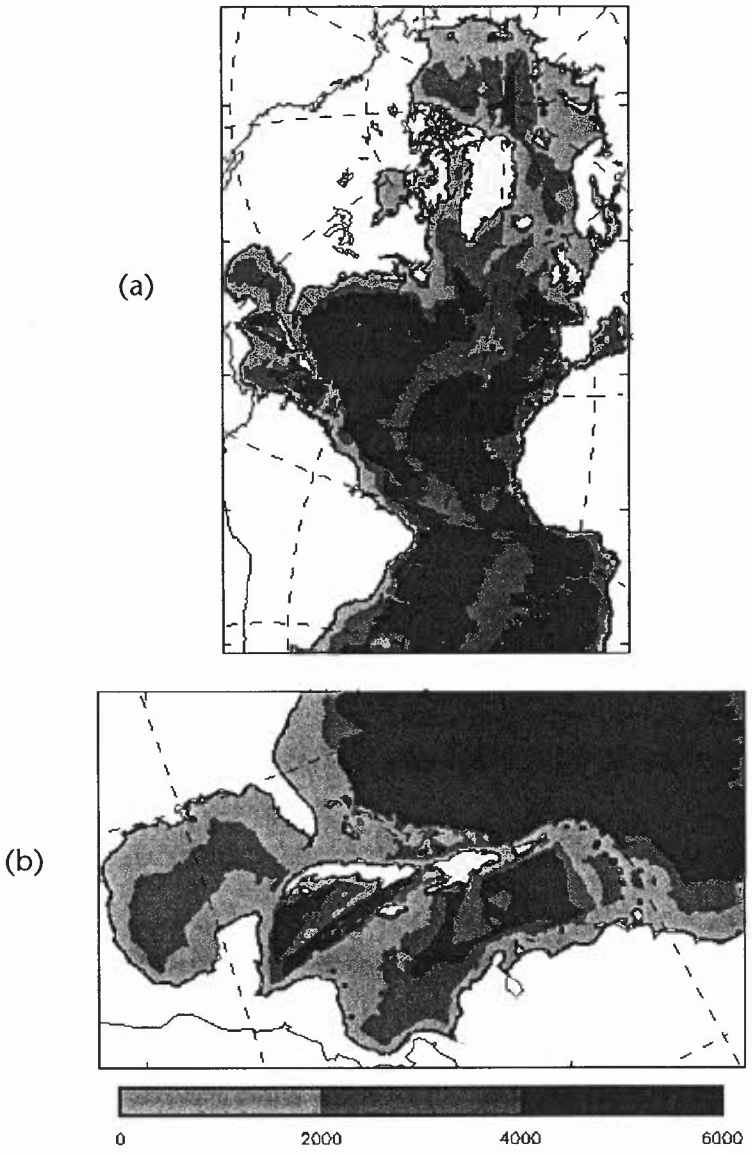

Figure 1: Nested model domains and bathymetry (depth in m). (a) 1/3° North Atlantic model; (b) 1/9° Caribbean and Gulf of Mexico model.

Surface fluxes for the three models are provided by monthly mean climatologies of heat fluxes (Oberhuber, 1988), wind stress (Hellerman and Rosenstein, 1983), evaporation (Esbensen and Kushnir, 1981) and precipitation (Jaegar, 1983). For future experiments we intend to use 6-hourly fluxes from the Met. Office NWP system, and work has been completed to set up the infrastructure required to do this.

3.2 Basic assimilation scheme

The assimilation scheme within the FOAM system is based upon the Analysis Correction method of Lorenc et al. (1991). The scheme is an iterative successive correction method, which approximates Optimal Interpolation. Full details of the implementation of the scheme in the FOAM system are given in Bell et al. (2000).

Given an estimate of the optimal analysis on the u^{th} iteration of the scheme, x[u], and a vector of observations, y_o, the analysis on the subsequent iteration is given by

$$\mathbf{x}[u+1] = \mathbf{x}[u] + \lambda \mathbf{W}\tilde{\mathbf{Q}}\left(\mathbf{y}_o - \mathbf{k}\left(\mathbf{x}[u]\right)\right), \qquad (3.1)$$

where W is a matrix containing the weight relative to the background field given to each observation at each gridpoint, $\tilde{\mathbf{Q}}$ is a diagonal matrix of normalisation factors which take into account the spatial density of observations, k is an interpolation operator from analysis gridpoints to observation locations and λ is a relaxation factor which controls the step size of the increment.

In practice, equation (3.1) is applied at the m^{th} gridpoint as

$$\mathrm{x}_m[u+1] = \mathrm{x}_m[u] + \lambda \sum_i \left(\mu_{mi}\tilde{\mathrm{Q}}_i \mathrm{R}_i \left(\mathrm{y}_i - \mathrm{k}_i(\mathbf{x})\right)\right), \qquad (3.2)$$

where i is an index over observations. Here, μ_{mi} represents the estimated forecast error covariance (normalised by the variance), modelled as a second-order auto-regressive function,

$$\mu_{mi} = \left(1 + \frac{\mathrm{h}_{mi}}{\mathrm{s}}\right)\exp\left(-\frac{\mathrm{h}_{mi}}{\mathrm{s}}\right), \qquad (3.3)$$

where h_{mi} is the distance between observation i and gridpoint m, and s is the specified correlation scale. In the experiments described here, the correlation scales are taken to be uniform and isotropic. The choice of correlation scale is discussed in Section 5. $\tilde{\mathrm{Q}}_i$ in 3.2 is the normalisation factor, which is determined from a weights function, D, as

$$\tilde{\mathrm{Q}}_i = \left\{\varepsilon_i^2\left(1 + \mathrm{k}_i(\mathbf{D})\right)\right\}^{-1}.$$

The weights function at the m^{th} gridpoint is given by

$$D_m = \sum_i |\mu_{mi}| \mathrm{R}_i / \varepsilon_i^2.$$

Here, ε_i^2 is the ratio of the observation error to the background error,

$$\varepsilon_i^2 = (\mathrm{O+F})_{ii} / \mathrm{B}_{ii}, \qquad (3.4)$$

where the the observational error matrix O+F is taken to be diagonal. R_i is a time window weighting function which takes the form of a symmetric linear ramp, with maximum value of 1

at the observation validity time, falling to zero at either end of the specified observation time window.

For efficiency, a recursive filter is used to evaluate the analysis equations. Two passes of the filter are required to approximate the second-order auto-regressive function, (3.3).

The system deals with multiple observation types in sequence, and hence effectively interleaves analysis steps for different data types. In the experiments described here, SST observations, temperature profile observations, and altimeter SSH data analysis steps are performed in turn. The system is also able to assimilate salinity profiles, but none are used here. Correlation scales, time windows, and weighting factors can be set independently for each different data type. In these experiments, for SST and temperature profile data correlation scales are set to 300 km everywhere; equal weighting is given to the model and observations; and the time windows used for SST and temperature profile data are 5 and 10 days either side of the observation's validity time respectively. The choice of parameters for the altimeter SSH data is discussed in Sections 4 and 5.

3.3 Altimeter SSH data assimilation scheme

The altimeter SSH data assimilation scheme consists of two main steps: first, an analysis of the altimeter SSH data is performed using the Analysis Correction scheme, to compute gridpoint SSH increments; these increments are then projected onto sub-surface temperature and salinity increments using the method of Cooper and Haines (1996).

The CH scheme is based upon the observation that surface anomalies are likely to be more strongly correlated with isopycnal depth anomalies, than with subsurface property anomalies on isopycnal surfaces. Hence the scheme seeks to impose the SSH increments required by the data through a re-arrangement of water parcels without changing their T, S characteristics. This is

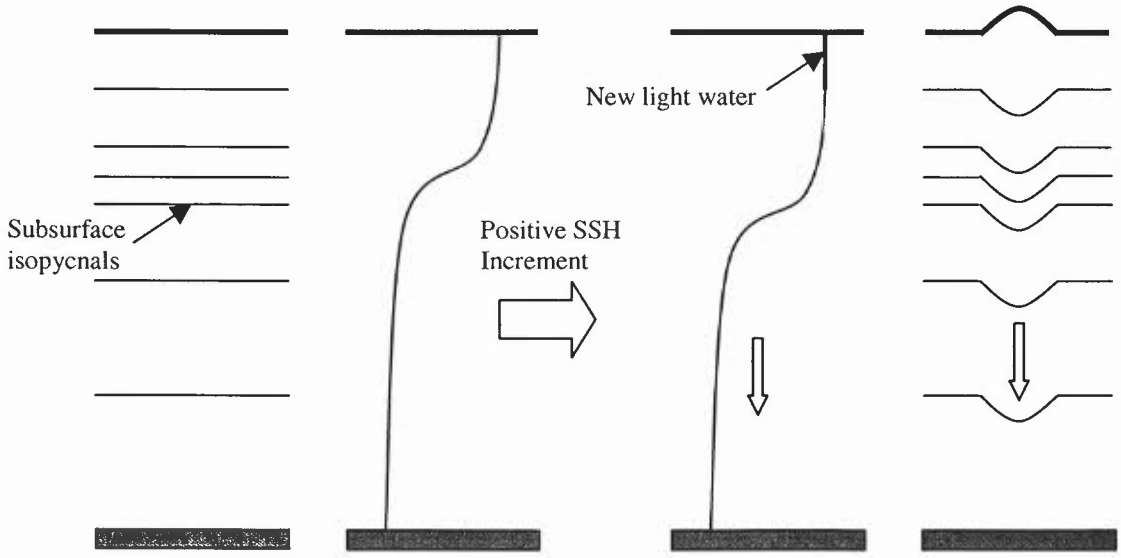

Figure 2: Schematic illustration of the assimilation method of Cooper and Haines (1996) for a positive SSH increment. (a) Water column before assimilation; (b) water column after assimilation.

achieved in the simplest manner possible, through a vertical displacement of the local water column, exploiting the hydrostatic relation between surface and subsurface pressure updates:

$$\Delta p(z) = \Delta p_s + g\int_z^0 \Delta\rho\,\mathrm{d}z,$$

where Δp, Δp_s and $\Delta\rho$ are increments to the subsurface pressure, surface pressure and subsurface density respectively.

Thus the pressure increment at any level can be related to the vertically integrated density increments in the fluid above. To close the problem, an assumption is made that the SSH (or equivalently, surface pressure) increment derives from entirely baroclinic changes to the water column. Barotropic changes can then be discounted by setting the pressure change at the ocean floor at $z = -H$ to zero, and hence

$$\Delta p_s = g\int_0^{-H} \Delta\rho\,\mathrm{d}z,$$

i.e. the change in surface height is compensated entirely by the integral of the subsurface density increments. Although there are potentially any number of possible re-arrangements of the water parcels that would produce the required SSH increment (including both local and non-local re-arrangements), CH adopt the simplest choice, a vertical displacement of the local water column. Reducing the method to a local problem renders the scheme practical to implement in sophisticated models. A schematic illustration of the scheme is shown in Figure 2. Note that water will be removed or added from the top or bottom of the water column, depending on the direction of displacement. Water that is added to the column is given the same T-S properties as the water directly above or below it in the water column.

(a)

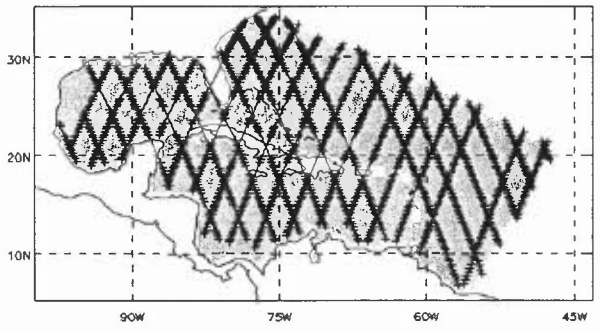

(b)

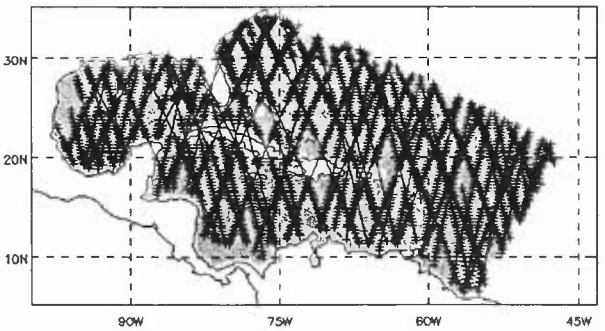

Figure 3: Satellite ground tracks over the 1/9° model area during a typical 15 day period. (a) Topex/Poseidon only; (b) Topex/Poseidon and ERS-2. Repeat periods for Topex/Poseidon and ERS-2 are 10 and 35 days respectively.

3.4 Development of altimeter SSH scheme

The Cooper and Haines (1996) scheme was initially coded within a tropical Pacific version of the Met. Office Unified Model (UM). The scheme achieved some success in increasing the variability in the model, and in improving the repesentation of equatorial Kelvin waves (Forbes, 1996). However, the choice of the tropical Pacific as a test bed for the scheme avoided many of the difficulties encountered when applying the scheme in other regions. In particular, the Tropical Pacific region is characterised by a distinct thermocline, the displacement of which produces density, and (through the hydrostatic relation) sea surface height increments which are able to balance the assimilated increments. In other regions of the ocean, the thermocline is less pronounced, or may not exist at all. Such regions cause problems when using the CH scheme, and hence the code has subsequently been developed to avoid such problems. In particular, regions where the ocean structure is largely barotropic such as high latitudes are excluded from the altimeter data assimilation to avoid generating unphysically large vertical displacements. This is achieved through the elimination of model gridpoints where the top-to-bottom temperature contrast is less than 0.05°C. Fox et al. (1999a,b) adopt an alternative approach based on the belief that barotropic sea surface height changes are predominantly large-scale, and hence they apply a filter to the altimeter data to remove large-scale features. The cut-off scale for the filter is determined by the local deformation radius.

Further changes to the original code were required to account for the presence of bottom topography, and to correct several minor coding errors.

The altimeter data assimilation scheme potentially has to compute vertical displacements at every model grid point. In the original implementation of the scheme, the appropriate vertical displacement was computed iteratively, starting with a maximum allowable displacement, then using a bisection algorithm to home in on the required displacement. For example, given the required SSH increment, dh, and two estimates of the required vertical displacement, v_1 and v_2, with resulting SSH increments dh_1 and dh_2 such that $dh_1 < dh < dh_2$, the scheme would proceed as follows:

Calculate new displacement, $v_3 = (v_1 + v_2)/2$.

Calculate resulting SSH increment, dh_3.

If $dh_3 > dh$, then use v_1 and v_3 for the next iteration.
If $dh_3 < dh$, then use v_2 and v_3 for the next iteration.

This algorithm was inefficient, and hence has been replaced by a gradient-based algorithm. Again, given initial estimates v_1 and v_2 of the required vertical displacement, with resulting SSH increments dh_1 and dh_2 the scheme would:

Calculate new displacement, $v_3 = \dfrac{\left[v_2(dh_1 - dh) - v_1(dh_2 - dh)\right]}{\left[dh_1 - dh_2\right]}$

Calculate resulting SSH increment, dh_3.

Use v_3 and whichever of v_1 and v_2 gives a SSH increment closest to dh for the next iteration. If convergence is slow, i.e. v_3 is very close to v_1 or v_2, then the interval is bisected, i.e. $v_3 = (v_1 + v_2)/2$.

This change of algorithm resulted in an overall speed-up of the altimeter data assimilation code by a factor of around 6.

Following the technical preparation and speed-up of the code, the scheme was suitable for implementation in the FOAM high-resolution models. The initial set of year-long integrations of the models therefore included altimeter assimilation runs in the 1/3° North Atlantic model, and the 1/9° Caribbean model. Analysis of the results from these runs identified a number of problems, and led to the series of 'tuning runs' to be described here. One of the key problems identified was a relative lack of impact from the addition of the altimeter data assimilation. The main emphasis of the tuning runs has therefore been to attempt to increase this impact.

Following these tuning experiments, a further set of integrations is underway with a view to providing a thorough test of the altimeter data assimilation system. These integrations will be 3 years long, and are ongoing.

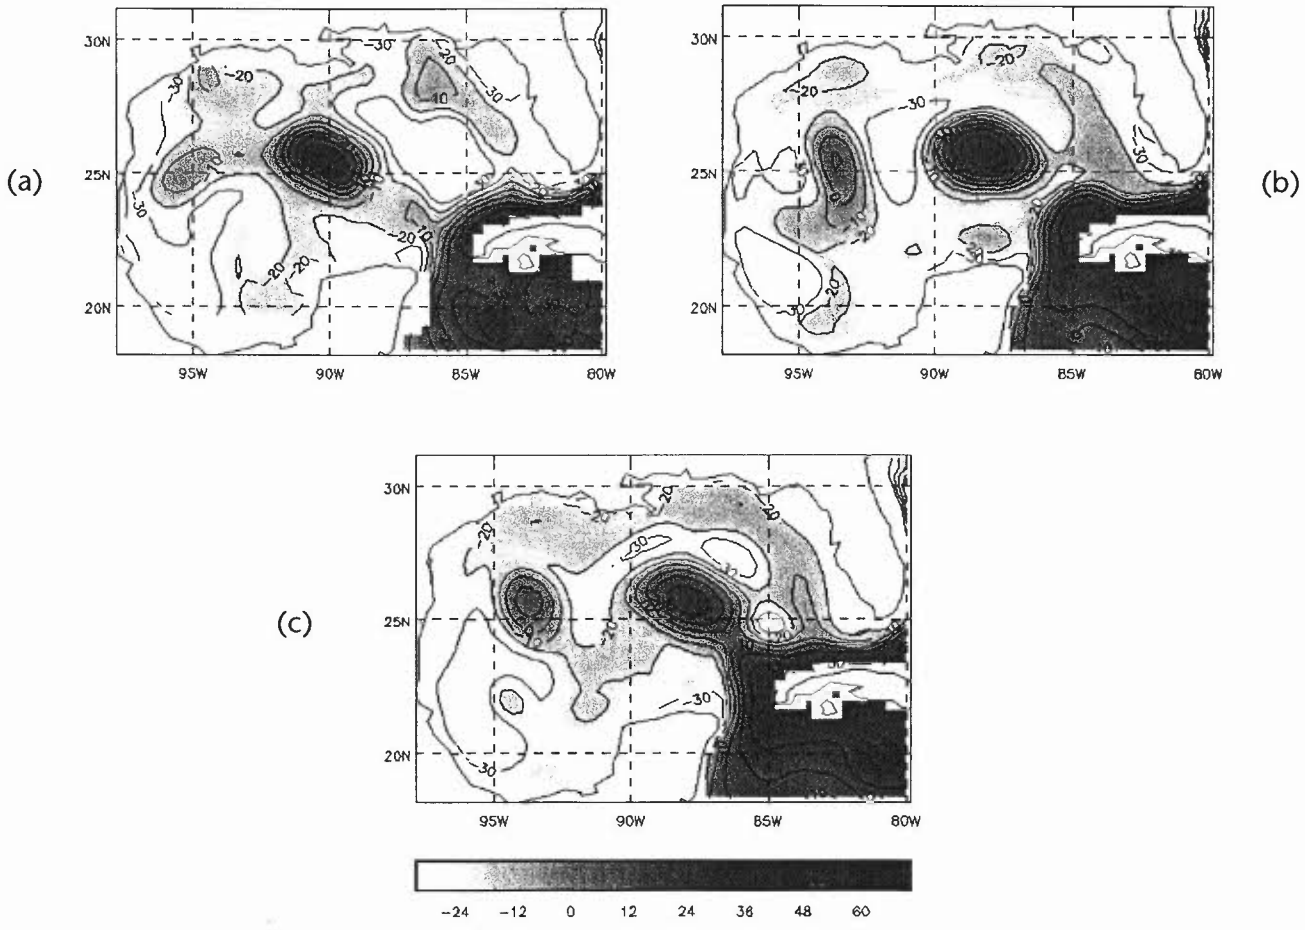

Figure 4: Sea surface height (cm) in the Gulf of Mexico from the 1/9° model after 6 months integration prior to tuning. (a) CLS mapped sea level anomaly added to 1/9° mean SSH; (b) run assimilating SST and temperature profile data; (c) run assimilating SST, temperature profile and altimeter SSH data.

3.5 Data sources

The development of satellite altimetry as a tool for ocean observation began in the 1980s with the launch of Seasat. Seasat provided the first source of global SSH height data, and although the accuracy of Seasat data was relatively poor in comparison to the current generation of altimeters, it highlighted the potential of altimeters to provide a useful insight into the structure of the ocean (e.g. Bernstein et al, 1982; Cheney et al, 1983). Geosat, launched in 1985, offered data of greater accuracy, which was used in the development of assimilation techniques for altimeter SSH data (e.g. Ezer and Mellor, 1994; Evensen and Van Leeuwen, 1996).

Following on from the Geosat mission, there are currently two satellites providing altimeter SSH data on a regular basis: the joint US - French Topex/Poseidon (T/P) launched in 1992, and the European Space Agency's ERS-2 launched in 1995. In addition, the US Navy Geosat Follow-On (GFO) mission was launched in February 1998. However, due to initial problems with the GFO platform, the satellite was not accepted as operational until November 2000, and hence regular access to GFO data is not yet available.

The T/P platform consists of two separate altimeters – the conventional Topex radar altimeter, and the more experimental solid-state Poseidon instrument. Poseidon is active for approximately one of every ten complete 10-day cycles, with Topex reporting for the remainder. Estimated error budgets for SSH measurements from T/P are around 5cm r.m.s. (Fu et al, 1994), most of which arises from the orbit correction.

The orbit characteristics of T/P and ERS-2 differ significantly. T/P has good temporal resolution with a 10-day repeat cycle, but poor spatial resolution, with cross-track spacing around 200 km at the equator. In consequence, T/P is not well suited for observing mesoscale phenomena (Jacobs et al, 1999). This can be seen in the plot of satellite tracks in the 1/9° model domain for a typical 15 day period shown in Figure 3(a). There are large unobserved areas between the satellite tracks. In contrast, ERS-2 has a longer, 35-day repeat period, and correspondingly better spatial resolution, with cross-track spacing of around 60 km at the equator. ERS-2 is therefore much better suited to observing mesoscale phenomena. Taking the two satellites together, in a typical 15 day period we get far smaller unobserved areas, as can be seen in Figure 3(b).

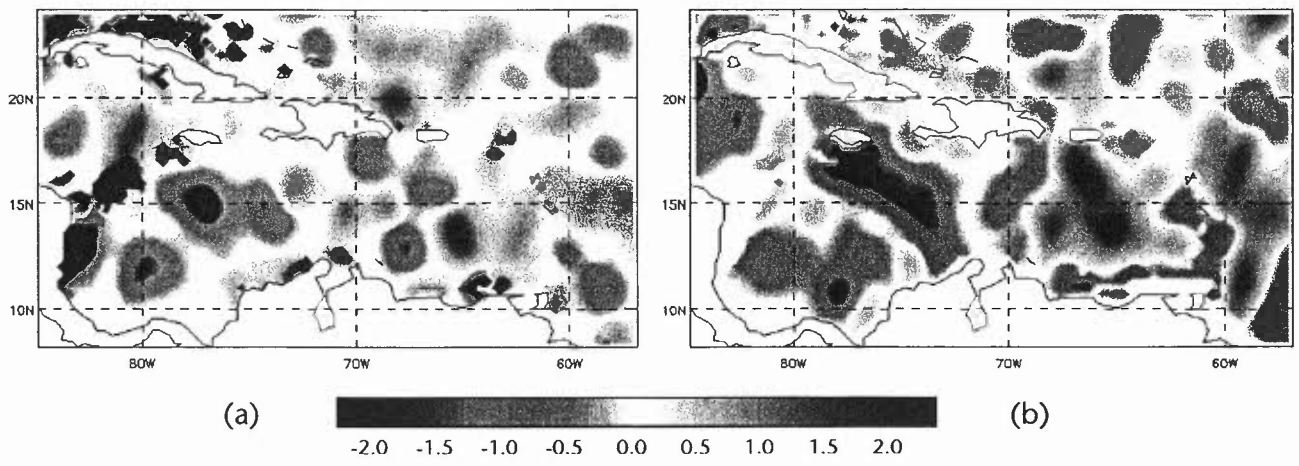

Figure 5: Vertical displacement (cm per timestep) applied by the Cooper and Haines scheme (a) with water shallower than 200m included; (b) with water shallower than 200m eliminated.

Several further altimeter missions are planned to ensure the continuity of satellite SSH data coverage. GFO data should be available routinely in the near future. Topex/Poseidon will be followed by Jason-1, due to be launched in 2001, and subsequently by Jason-2, planned for 2004. Envisat will also carry a radar altimeter (RA2) at its launch in 2001.

4. Results from initial integrations

The initial set of year-long integrations of the nested models included runs assimilating temperature profile and SST data, and runs additionally assimilating altimeter SSH data. The temperature profile and SST assimilation runs have therefore been used as a baseline for assessing the impact of the altimeter assimilation. The model's temperature and salinity fields were initialised to the Levitus (1994) climatology, and the models were then spun-up for 6 months from rest assimilating temperature profile and SST data.

In the initial altimeter data assimilation runs, along-track data from the Topex/Poseidon (T/P) satellite, provided in near-real-time by CLS, France, were used. At this stage, given the preliminary nature of the experiments, ERS-2 data were not assimilated. Due to lack of detailed knowledge of the geoid, the data are reported as sea level anomalies relative to a three-year mean from the period 1993-1995. In order to assimilate the data we need to reconstruct the full SSH by adding the data anomalies to a suitable mean SSH. For the initial integrations, the annual mean SSH from the year-long run assimilating temperature profile and SST data has been used. The choice of a one year mean from an integration through 1997 was not necessarily the most appropriate choice. The choice of mean state is discussed further in Section 5. Note, however, that the consistency of this choice could be improved by re-aligning the mean of the data anomalies to the period chosen for the mean state. The data anomalies are reported as

$$\eta^{anom}_{93-95} = \eta^{total} - \bar{\eta}_{93-95}$$

where η^{anom}_{93-95} are the reported data anomalies realtive to the 1993-95 mean, $\bar{\eta}_{93-95}$, and η^{total} is the total measured SSH.We want to obtain

$$\eta^{anom}_{97} = \eta^{total} - \bar{\eta}_{97}$$

But

$$\overline{\left(\eta^{anom}_{93-95}\right)}_{97} = \overline{\eta^{total}}_{97} - \bar{\eta}_{93-95}$$
$$= \bar{\eta}_{97} - \bar{\eta}_{93-95}$$

So

$$\eta^{anom}_{93-95} - \overline{\left(\eta^{anom}_{93-95}\right)}_{97} = \eta^{total} - \bar{\eta}_{93-95} - \bar{\eta}_{97} + \bar{\eta}_{93-95}$$
$$= \eta^{total} - \bar{\eta}_{97}$$

Hence the data anomalies can be adjusted to give anomalies relative to the mean for 1997 (or any other period), for consistency with the choice of mean SSH used, simply by subtracting the mean of the reported anomalies for the period concerned from the data values. For future experiments the data will be adjusted to account for any mismatches between the period of the

mean state used and the reference period of the altimeter data anomalies in this way. The impact on results will be investigated in the course of current and future work.

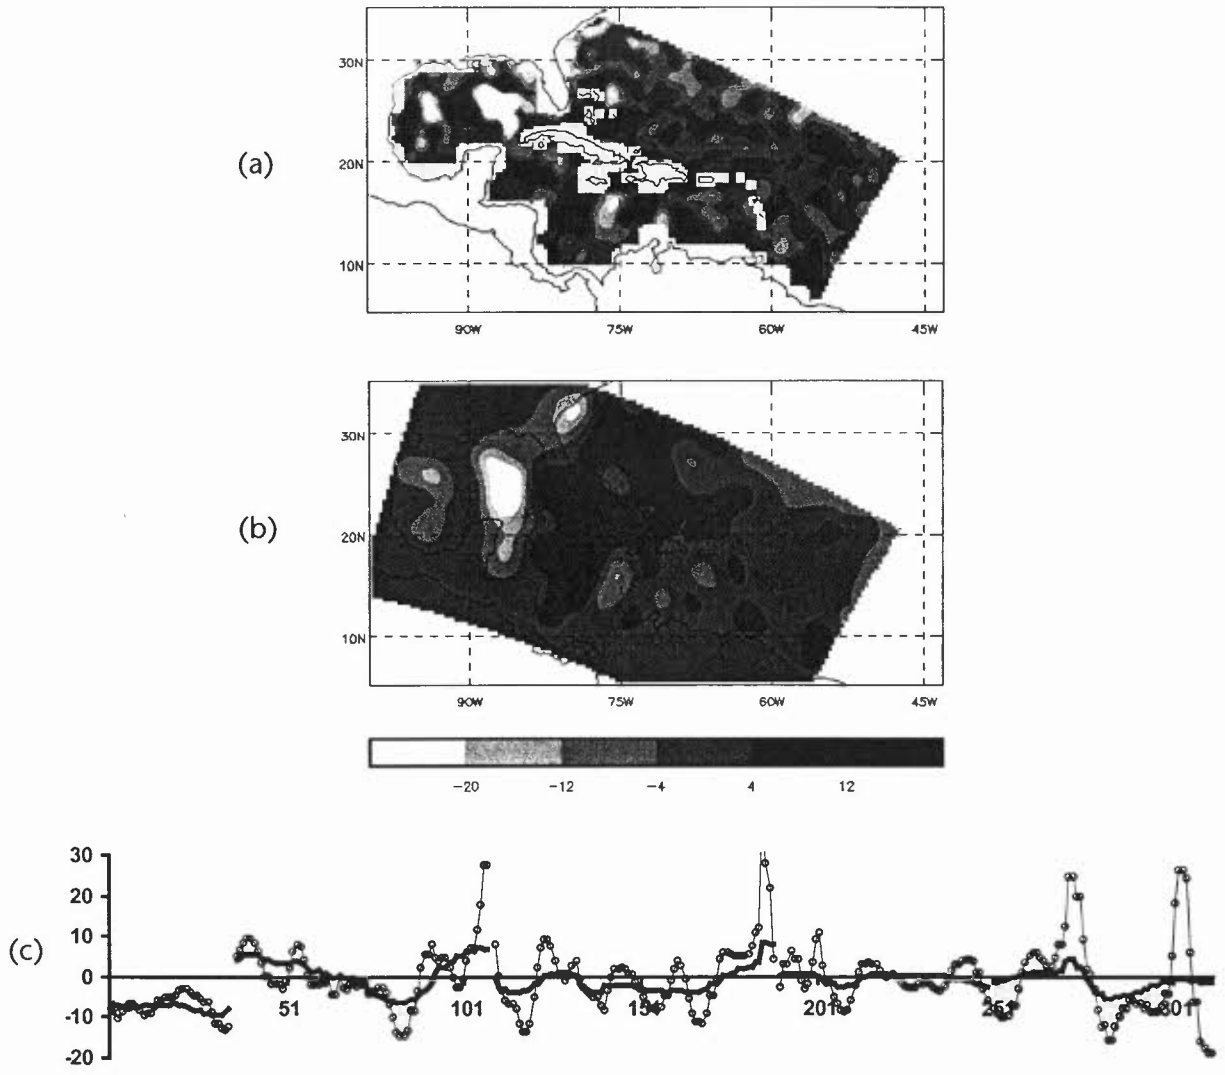

Figure 6: (a) Difference (m) between 1/9° model instantaneous SSH and CLS mapped sea level anomaly added to 1/9° mean SSH; (b) typical SSH increments applied by the AC scheme using 100km correlation scales; (c) data along five Topex/Poseidon tracks. The circles are the model SSH minus the observed anomaly added to the 1/9° mean; the solid line is the increment applied by the AC scheme using 100km correlation scales.

In order to compute gridpoint SSH increments using the AC scheme, several parameters controlling the scheme need to be specified. In the initial runs, these parameters are chosen based upon experience of assimilating other data types, and on our beliefs about the underlying ocean structures that we hope to capture through the assimilation. These 'first-guess' parameters are a necessary first step, as they provide a basis from which we can hope to refine their specification. Indeed, given model results, we can compute estimates of the parameters. However, in order to get useful results, we first need to ensure that our 'first-guess' parameters are sensible. It is, therefore, an important function of the test runs described here to refine our

initial estimates to the extent that we can get results that are useful for calculating proper estimates of the assimilation parameters.

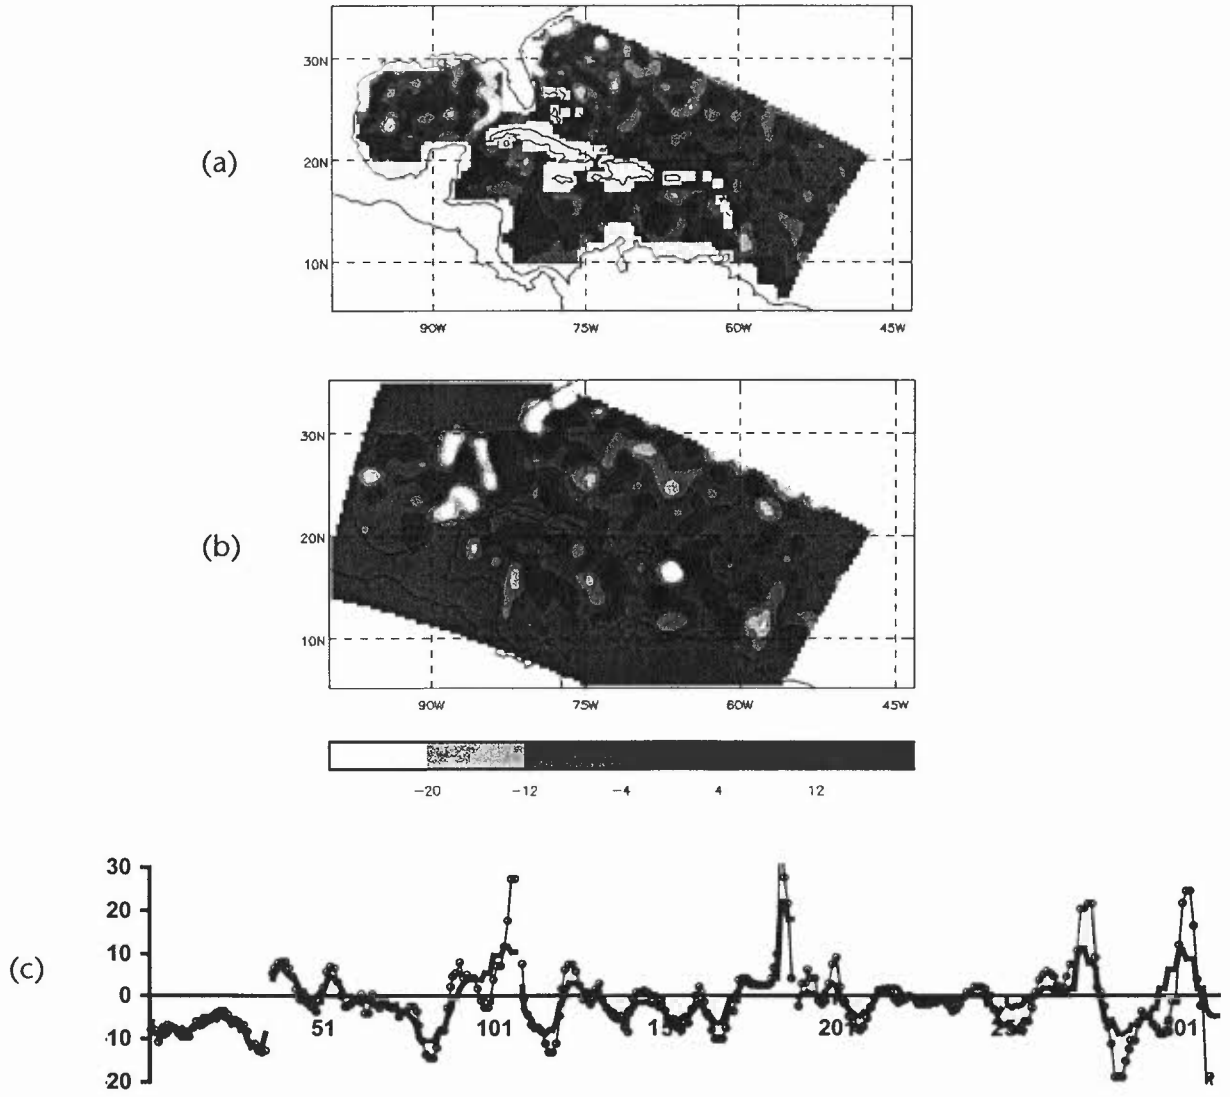

Figure 7: As Figure 6, but for AC scheme using 40km correlation scales.

For the initial integrations, our estimates of the assimilation parameters for the altimeter data were as follows: firstly, we chose to give equal weight to the model and the observations; secondly, we chose to use correlation scales of 100km throughout the model domain; thirdly, we chose to use observations over a time window of 10 days either side of their validity time (with a linear 'ramp' weighting function).

To assess the output from the initial integrations we use the mapped sea level anomaly product from CLS. The mapped anomalies are added to the mean SSH used with the along-track data. Note that the CLS mapped data is a statistical mapping of 21 days altimeter data (with more recent data having greatest weight), hence care is needed when comparing directly with instantaneous model output. Many of the results presented here focus on the Gulf of Mexico.

This is a good area for assessing the impact of the altimeter assimilation for several reasons. Firstly, the loop current that flows through the Gulf of Mexico periodically sheds eddies, that then flow westward through the Gulf. The area therefore frequently has large isolated eddies, the initialisation of which provides a good test of the assimilation. Whilst other areas are also characterised by transient eddies, in many regions these eddies are obscured somewhat by other mesoscale activity. Secondly, the Gulf is common to all three model domains, so is a good location for comparison of different resolutions. Finally, the spatial scales of eddy features in the Gulf of Mexico are comparatively large due to low latitude location hence it provides a suitable test bed for an initial implementation of the assimilation scheme.

Figure 4 shows the SSH in the Gulf of Mexico from the 1/9° model in July 1998 after 6 months of integration. The mapped observations (Fig 4(a)) show a single isolated eddy located in the centre of the Gulf at 90°W, 26°N, which has been shed by the loop current. Whilst both the temperature profile and SST assimilation (Fig 4(b)), and the run additionally assimilating altimeter data (Fig 4(c)) also exhibit a large isolated eddy, in both cases the eddy is located too far to the east. Furthermore, both assimilations show a second, weaker eddy in the west of the Gulf (94°W,26°N), which is not present in the data. Whilst the altimeter data assimilation has adjusted the shape and strength of both eddies, it has been ineffective at initialising the features contained in the data. Some of the lower amplitude features, such as the arc of low SSH in the NE of the Gulf are initialised better by the altimeter assimilation runs. However, the overall impact of the altimeter assimilation has been poor. This poor impact persists throughout the integrations, and is also a feature of the 1/3° model results.

5. Tuning the altimeter data assimilation scheme

The lack of impact due to the altimeter assimilation in the original year-long integrations motivated a series of tuning experiments to attempt to increase the impact. The tuning covered three main areas: basic code changes to improve the robustness of the assimilation method; tuning of the assimilation parameters; consideration of data sources and usage.

5.1 Basic code amendments

5.1.1 Corrections for shallow water

The initial integrations identified particular problems with the Cooper and Haines scheme in regions of shallow water. Figure 5(a) shows the mean vertical displacement applied by the scheme off the north coast of South America for a 5 day period during the integration. Large vertical displacements can be seen as black patches. These patches correspond closely to the regions of shallow water. In shallow water, the top to bottom temperature contrast is typically relatively small, and hence large vertical displacements are needed to achieve the required SSH increments. The Cooper and Haines (1996) scheme is based upon the assumption that the thermocline can be displaced to achieve the SSH increments, but in shallow water regions the water column is frequently well mixed, and hence the scheme is inappropriate. In addition, tides produce large surface displacements in shallow water, which may not be accurately removed by the tidal corrections applied to the data. Hence the data in altimeter SSH data in shallow water tend to be unreliable. We therefore adopt a similar approach to Fox et al (1999a,b), and do not apply the CH scheme in water shallower than 200m depth. The resulting vertical displacement field is shown in Figure 5(b), which shows that problems with excessive displacement have been eliminated. Whilst this approach is effective, it is somewhat ad hoc, and the cut off at 200m

depth is somewhat arbitrary. A more robust approach would be to penalise the vertical displacement through the assimilation cost function, but this is beyond the scope of the work presented here.

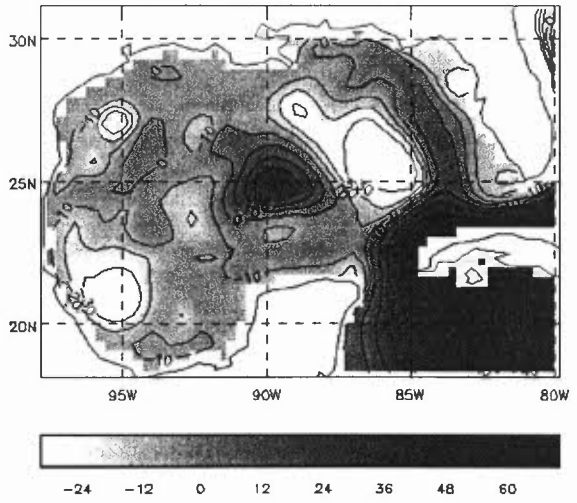

Figure 8: Sea surface height (cm) in the Gulf of Mexico for 1/9° model after 6 months integration with tuned assimilation parameters, and Topex/Poseidon and ERS-2 data.

5.1.2 Corrections to rigid lid pressure computation

A further problem was identified relating to an inconsistency in the computation of the rigid lid pressure from timestep to timestep. This inconsistency results in the model losing information about the rigid lid pressure increments made at a timestep, and hence increments at subsequent timesteps can be considerably over- or under-estimated. Furthermore, this problem prevents the model from representing seasonal signals contained in the data.

The model SSH (or equivalently, rigid lid pressure) is computed from the streamfunction after the model timestep. However, the SSH is only calculated to within an arbitrary constant, so an extra condition is needed to complete the computation. Originally, the condition chosen was that of zero spatial mean SSH over the whole domain i.e. $\langle h \rangle = 0$, where $\langle\ \rangle$ denotes the mean over all model grid points. It is this condition that causes problems at subsequent timesteps.

The assimilation scheme computes a set of grid point sea surface height increments, $\delta h_{i,j}^{assim}$ from the observation increments. Subsurface temperature and salinity increments are then computed for a vertical displacment of the water column by the amount needed to provide the required SSH increment.

At the next timestep, the SSH is again computed from the streamfunction, which will have been modified by both the temperature and salinity increments, and the model dynamics. The SSH will therefore have been changed by the assimilation increment from the previous analysis step, and also by the model dynamics. Hence the new SSH will be given by

$$h_{i,j}^{t+1} = h_{i,j}^{t} + \delta h_{i,j}^{assim} + \delta h_{i,j}^{dyn}$$

where t denotes the time step and $\delta h_{i,j}^{dyn}$ is the change in SSH due to the model dynamics. However, the SSH is again normalised to have zero mean, whereas

$$\left\langle h^{t+1} \right\rangle = \left\langle h^{t} \right\rangle + \left\langle \delta h^{assim} \right\rangle + \left\langle \delta h^{dyn} \right\rangle = \left\langle \delta h^{assim} \right\rangle + \left\langle \delta h^{dyn} \right\rangle .$$

Normalising the SSH field to have zero spatial mean is equivalent to assuming that the mean of the increments applied at the previous timestep is zero, and that the SSH is unaffected by the model evolution between timesteps. These assumptions are clearly incorrect. In particular, seasonal effects will produce increments with non-zero means. Although it may be difficult to adjust the mean SSH to account for variations introduced by the model dynamics, it is clearly straightforward to account for changes due to assimilated increments, and to adjust the spatial mean SSH accordingly.

This problem has been resolved by explicitly retaining the information about the mean SSH increment. This information is stored internally at one timestep, and passed to the SSH calculation routine on the next timestep. The mean SSH is then updated by the mean increment to ensure temporal consistency in the SSH field i.e. we set

$$\left\langle h^{t+1} \right\rangle = \left\langle h^{t} \right\rangle + \left\langle \delta h^{assim} \right\rangle$$

The mean rigid lid pressure is also written to the model dump to enable the model to retain the information through restarts.

5.2 Assimilation parameters

Having resolved the basic coding issues, the natural next step was to investigate the sensitivity of the results to the parameters specified for the assimilation scheme. Note that although the parameters will be considered separately below, it is the combination of the different parameters that is most significant. For example, increasing the weight given to the observations in isolation will not necessarily lead to a better analysis, particularly if the correlation scales are such that the SSH analysis is over-smoothed. However, the combination of increased weight to the observations with well chosen correlation scales would be expected to lead to significant improvements in the quality of the analysis.

5.2.1 Correlation scales

The correlation scales chosen to represent the model background error covariances should be representative of the scales of the model errors. Figure 6(a) shows the difference between the 1/9° model instantaneous SSH and the CLS mapped field part way through the initial year long altimeter data assimilation. The actual SSH increments computed by the AC scheme, using the 100km correlation scales, are shown in Figure 6(b). Clearly, the increments being applied are far larger scale than the model errors. Note also that the typical scale of the applied increments is somewhat longer than the 100km scale specified. The assimilation is over-smoothing the data increments, and the choice of 100km scales therefore appears inappropriate. This is confirmed by

the plot of along-track data for 5 satellite tracks shown in Figure 6(c). Comparison with along-track rather than mapped data removes any uncertainties introduced by the mapping of the data, and hence should be more reliable. The along-track data increments show considerable small scale structure, all of which is smoothed out by the assimilation. The fact that the increments being applied by the assimilation are of somewhat larger scale than the specified correlation scale is a characteristic of the assimilation scheme, or more specifically of the current implementation of the scheme, that may warrant further investigation. The AC scheme will converge towards the Optimal Interpolation analysis after several iterations (Lorenc et al., 1991). In the current implementation, however, only a single iteration is applied on each timestep, and hence convergence to the OI solution is not achieved. It is probable that this feature of the current implementation may explain the discrepancy between the scales specified, and the effective scales that the scheme applies.

Ideally, the correlation scales should be estimated from calculations based upon model output from an assimilation. However, the results of such calculations depend upon the scales used in the assimilation, hence in order to get meaningful results from such calculations we need to have a reasonable initial estimate of the scales. A series of runs were therefore carried out to investigate the sensitivity to the scales used, which suggested that scales of around 40km were most appropriate. Figure 7(a) shows the model differences from the mapped data when 40km scales are used; the spatial scales of the errors are similar to those for the run using 100km scales. The SSH increments applied (Figure 7(b)) are, however, much smaller scale, and the scales of the increments agree well with those of the model error. This is reflected in the along track data (Figure 7(c)), which shows that the assimilation retains much more of the small scale information in the data, whilst avoiding over-fitting the higher amplitude features. In subsequent experiments, we have therefore chosen to use scales of 40km everywhere. Note that we would expect significant spatial variation in the true scales, but at present we do not have enough information to enable us to model this effectively. This will be pursued in detail at a later date.

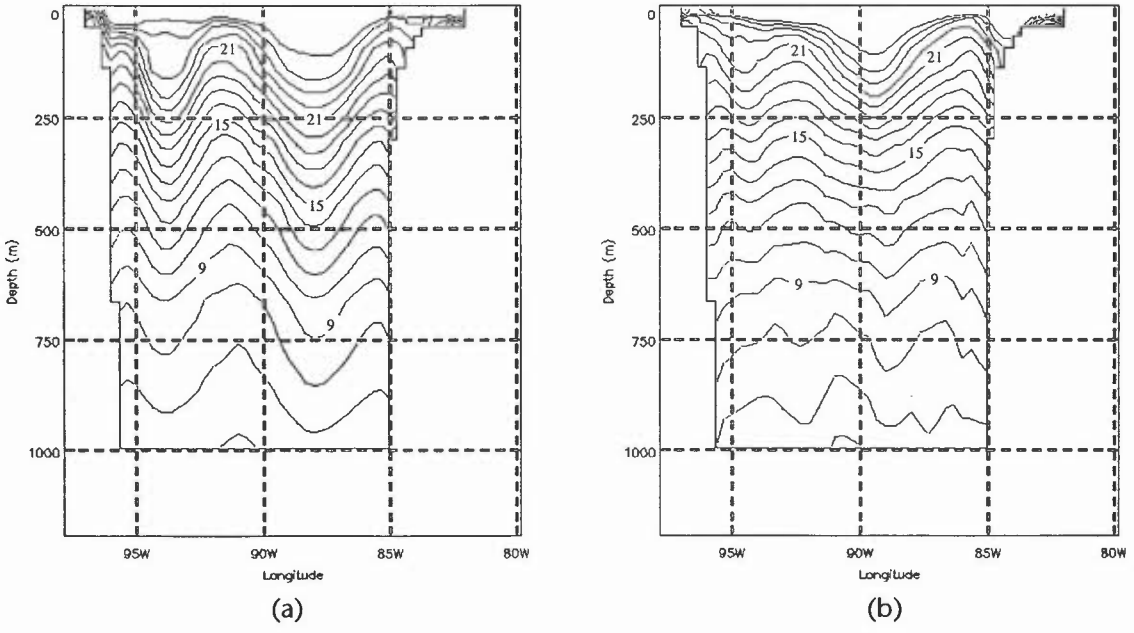

Figure 9: Sub-surface temperature sections (°C) through the Gulf of Mexico along 26°N after 6 months of integration in the 1/9° model (a) SST and temperature profile assimilation; (b) SST, temperature profile and altimeter SSH data assimilation. Contour interval is 1.2°C.

5.2.2 Relative weighting of model and observations

Given a better estimate of the correlation scales, it was then possible to assess the relative weight to give to the data. Again, the relative weightings should ideally be estimated from calculations based on model output, and should reflect the relative size of the errors in the model and the data. However, at this stage we are not able to compute such estimates, and instead we have carried out sensitivity studies to investigate the ideal weighting. Not surprisingly, these studies showed improvement in the fit to the data through giving the data twice the weight given to the model, and further improvement by giving the data 4 times the weight given to the model. Increasing the weight further led only to very marginal improvements, and hence for subsequent runs the observations are given 4 times the weight given to the model.

5.2.3 Time window

Of the assimilation parameters, the model is least sensitive to the time window used. Topex/Poseidon has a repeat period of 10 days, so with the original time window of 10 days either side of validity time, 2 complete repeat cycles would be assimilated at every analysis step. There is therefore some redundancy in the data being assimilated. Reducing the time window to 5 days either side of validity time did not lead to any detriment to the model results, so for subsequent experiments the shorter time window was used. Note that for the experiments described here, the full time series of data is available, so observations ahead of the model validity time can be assimilated. In an operational configuration, this would not be possible, and a longer time window may therefore prove advantageous by increasing data coverage.

5.3 Baseline results

The basic tuning of the assimilation parameters has been used to provide a solid basis for examining in greater detail the effect of the altimeter data assimilation. In order to do this, a new baseline for comparison was required, and hence the initial integration was repeated using the tuned assimilation parameters. At this stage, ERS-2 data were also included, and hence the impact on the results obtained is not only due to the improved specification of assimilation parameters, but also due to the greatly improved data coverage. These baseline results have also been used to undertake a preliminary assessment of the validity of the choices of assimilation parameters.

The impact of assimilating both the Topex/Poseidon and ERS-2 data using the tuned assimilation parameters is illustrated by Figure 8, which shows the SSH in the Gulf of Mexico after 6 months assimilation in the 1/9° model. The mapped CLS data, and the SSH from the original altimeter assimilation run are shown in Figures 3(b) and (c) respectively. In the new altimeter assimilation, the single large eddy in the centre of the Gulf has been successfully shifted to the west, in closer agreement with the observations. Note, however, that the amplitude of the SSH signature of the eddy has been reduced. The second eddy to the far west seen in the original altimeter assimilation has been removed, again in closer agreement with the observations. The new altimeter data assimilation therefore has been considerably more successful than the original assimilation in terms of correcting the model SSH field.

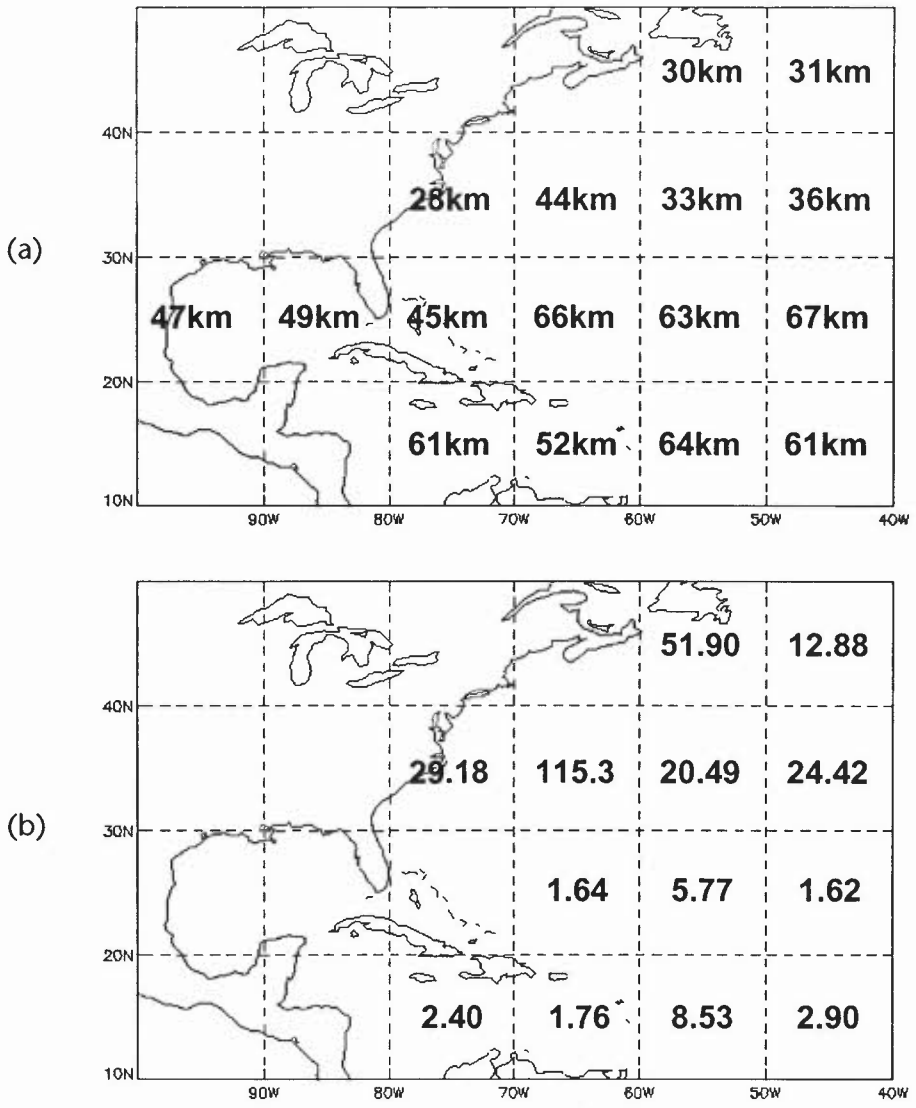

Figure 10: Estimates of (a) correlation scales and (b) error ratio (background error / observation error) for the Gulf of Mexico and Gulf Stream region, computed from 1/3° model output.

Figure 9 shows the subsurface impact of the new altimeter data assimilation. Zonal temperature sections through the Gulf of Mexico on 26°N are shown for the same time as the fields in Figures 4 and 8 for the temperature profile and SST assimilation (Figure 9(a)) and the new altimeter data assimilation (Figure 9(b)). In the temperature and SST assimilation, the eddy located too far to the east has a strong subsurface signature down to the ocean floor. The second, smaller eddy to the west is also clearly visible throughout the full depth. The results for the altimeter assimilation run show clearly that the core of the eddy to the east has been successfully shifted to the west, although the strength of the temperature signal of the eddy has also been reduced throughout the ocean depth. The second eddy to west has been almost eliminated, although a small subsurface temperature signal does remain. The new altimeter assimilation has been quite successful in initialising the eddy features seen in the Gulf of Mexico at this time. Similar results have been seen in the 1/3° model runs.

5.3.1 Estimates of correlation scales and error variances

Results from these integrations have been used to estimate correlation scales by fitting a functional form to the background error covariance calculated from the data. The background error covariance is calculated as a function of separation distance for each grid box on a 10°x10° grid, and a second-order auto-regressive (SOAR) function is fit to the error covariance in each gridbox. The scale parameter of the best fit SOAR function represents the correlation scale as specified in the assimilation scheme. The consistency of the a-priori specified correlation scale can therefore be checked against the scale derived from the model output. Figure 10(a) shows the scales computed from the 1/3° model output for gridboxes in the Gulf Stream region. The scales computed from the output show considerable spatial variation, most noticeably decreasing with latitude. Shorter scales are also seen in regions of the most intense currents. The use of spatially varying correlation scales will be covered in detail in future work. As already noted, the effective scales applied by the assimilation scheme are longer than the specified correlation scale. When taking this into account, the scales calculated from the output seem broadly consistent with the 40km scales specified, and hence the 40km scales should provide a solid basis for future work including further investigation of the correlation scales.

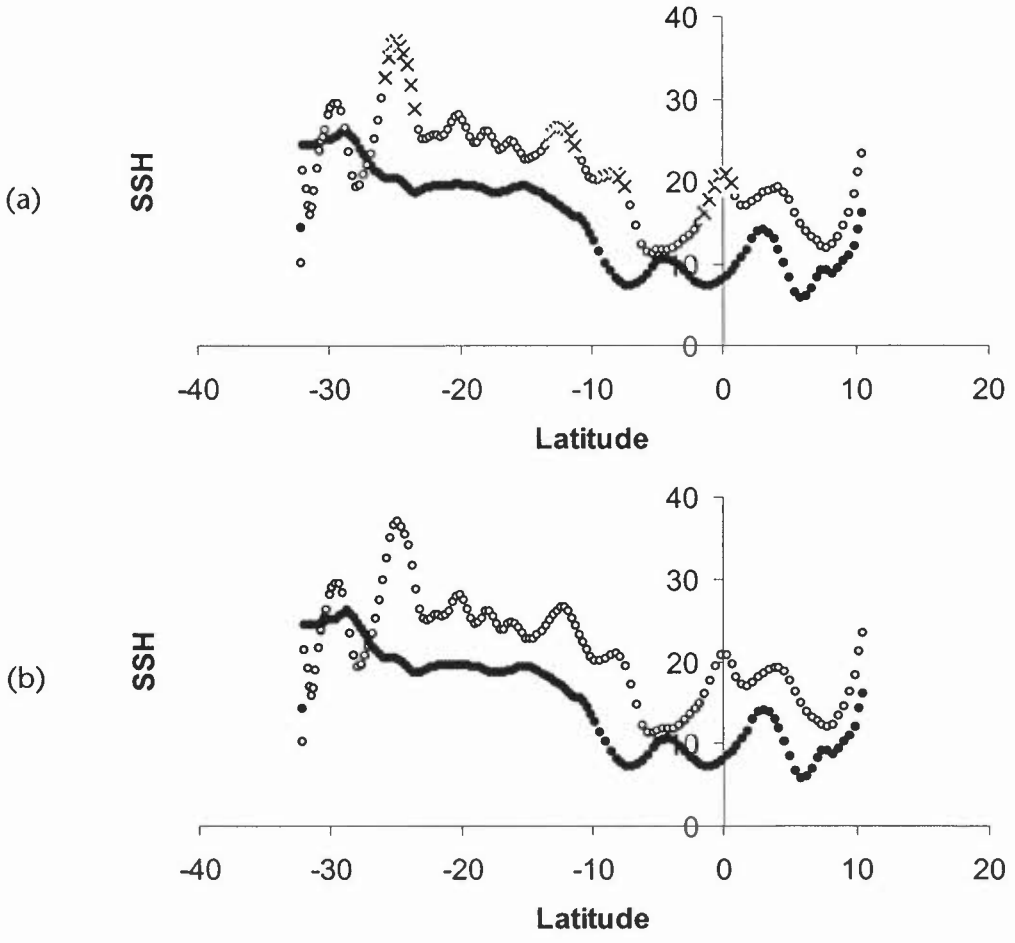

Figure 11: Impact of quality control background check on a typical satellite track. Solid circles are the background field, open circles are observations that would pass the background check, and crosses are observations that would fail the background check (a) using $\alpha = 2$ (b) using $\alpha = 3$. The offset between the background field and observed values is discussed in the text.

Estimates of the ratio of error variances (ε_i^2 in (3.4)) have been computed using the model error variance computed from the output, and the SOAR functions fitted to the error covariance data. Extrapolation of the SOAR functions to zero separation gives an estimate of the background error variance, whilst the computed model error variances give the sum of the background and observation error variances. It should be noted, however, that calculation of observation error variance by this method is recognised to be unreliable, since results are dependent upon the functional form fitted to the covariance data (Versal et al, 1988). Despite this, such results do give an indication of the magnitude of the ratio of the error variances, and hence enable the suitability of the specified relative weighting of the model and observations to be assessed. The ratio of error variances calculated using this method are shown in Figure 10(b). The computed values are clearly very variable, which in part reflects the unreliability of the method. Furthermore, the clear distinction between the values north and south of 30°N may in part be explained by a problem in the calculations which resulted in short separations being significantly under-sampled at low latitudes. However, the variation in values suggests that the ability to specify spatially varying error ratios in the assimilation scheme would be advantageous. This will be addressed in future work. The range of values also suggests that for much of the domain the choice of an error ratio of 4 is unlikely to be over-fitting the data. Increasing the error ratio beyond 4 was seen to have little further impact on the analysis, since giving the observations 4 times the weight given to the model already constrains the analysis to be a reasonably close fit to the observations. The calculation of error variances will be repeated with output from future runs, with the sampling issues highlight above addressed.

5.4 QC of data

Investigation of the reduction in the amplitude of the SSH signal when assimilating data from Gulf of Mexico eddies highlighted a problem with the quality control checks done during the altimeter data processing. For many satellite tracks passing over mesoscale features, observations near the peaks of high amplitude features were being eliminated by the quality control. Hence the centres of eddy features were frequently effectively unobserved, resulting in the assimilation failing to introduce features with sufficient amplitude. The data concerned was being flagged by the background check, which eliminates data that differs from a background field by more than a given number of standard deviations. More specifically, an observation is rejected if:

$$\left|\hat{\eta}_{\text{obs}} + \overline{\eta} - \eta_{\text{model}}\right| > \alpha\sigma\left(\hat{\eta}_{\text{T/P obs}}\right),$$

where $\hat{\eta}_{\text{obs}}$ is the observed SSH anomaly, $\overline{\eta}$ is the mean SSH used in the data processing, η_{model} is the background SSH, α is the quality control ratio, and $\sigma\left(\hat{\eta}_{\text{T/P obs}}\right)$ is the standard deviation of T/P observations (note that the model error is not taken in to account). For the runs reported above, η_{model} was taken from a run of the model assimilating temperature profile and SST data only, and $\sigma\left(\hat{\eta}_{\text{T/P obs}}\right)$ was computed from a 5 year time series of Topex/Poseidon observations. For the initial integrations, α was set to 2, and the result of the background check for a typical satellite track is shown in Figure 11(a). The observations have much more small scale variability than the background SSH, and the background SSH appears to be offset below the observed values along most of the track. This leads to the rejection of the data around the peaks of several of the small amplitude features.

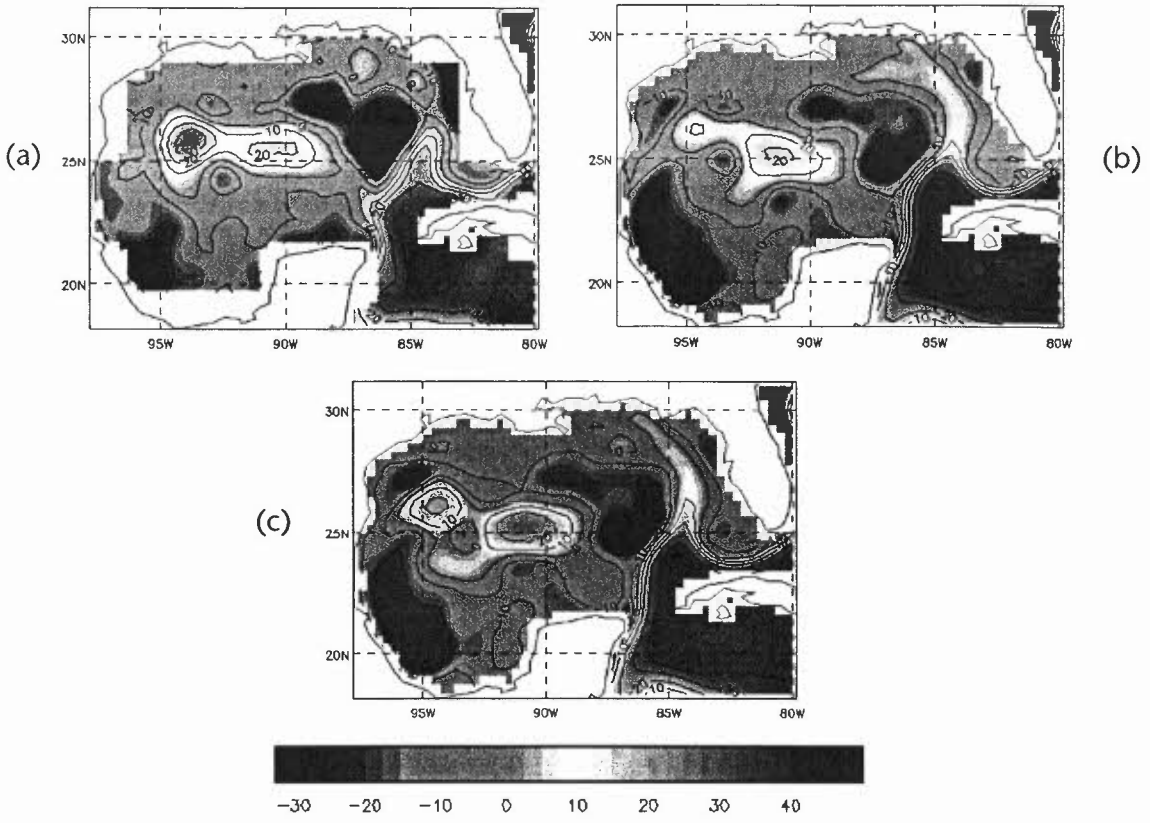

Figure 12: Sea surface height (cm) in the Gulf of Mexico from the 1/3° model assimilating SST, temperature profile and altimeter SSH data after 6 months integration. (a) CLS mapped sea level anomaly added to 1/3° mean SSH; (b) with altimeter data background check using $\alpha = 2$; (c) with altimeter data background check using $\alpha = 3$.

A solution to this problem with the current system is to increase the value of α to make the background check less stringent. The result of the background check using $\alpha = 3$ for the same satellite track is shown in Figure 11(b). All data values along the track are now retained. However, it is evident from Figure 11 that the background SSH field used in the quality control is poor, and it is this that leads to the rejection of useful observations in the original case, rather than the observations themselves being questionable. The ideal solution, which is better suited to operational configurations, is to use the SSH field from the latest model analysis as the background field for the quality control. This field should already represent far more of the small scale features in the data, and hence should provide a much more robust background check. When using the SSH from the on-going model run, it should be possible to set α back to 2 without detriment to the solution.

This increase in the value of α leads to further increases in the quality of the SSH analysis, as shown in plots of the SSH in the Gulf of Mexico from the 1/3° model after 2 months integration, shown in Figure 12. Figure 12(a) shows the mapped sea level anomaly data from CLS, added to the 1/3° model mean SSH. The mapped data show two adjacent eddies in the Gulf. The results from the run using the original quality control, and the new quality control are shown in Figures 12(b) and (c) respectively. With the original quality control, although both eddies are

reproduced by the assimilation, both are considerably weaker than in the data, with the western eddy particularly poorly represented. Changing the quality control leads to data in the centre of the two eddies being retained, and hence both eddies are strengthened, in better agreement with the data.

5.5 Mean SSH

SSH measurements as reported by satellite altimeters contain not only the dynamic SSH, but also variations in SSH due to the Earth's gravitational field, the geoid. At the present time, the geoid is not accurately known, and although satellite missions such as GOCE and GRACE are planned to measure the geoid, it is likely to be several years before geoid estimates mapping a wide range of spatial scales with sufficient accuracy are available to enable the separation out of the dynamic SSH. In the meantime, given that the geoid is (virtually) time invariant, it can be removed by removing the mean SSH from the data. This, however, also removes the mean dynamic SSH, and the resulting sea level anomalies must be added to a suitable mean dynamic SSH to reconstruct SSH measurements for use in the assimilation.

In the initial integrations, the annual mean SSH from the integration assimilating temperature profile and SST data has been added to the sea level anomaly data. However, the model has systematic errors, particularly in the location of major current systems such as the Gulf Stream. Killworth et al (2000) have shown that the locations of major current systems can be adjusted by hundreds of km through use of a more accurate mean SSH, and hence systematic model errors may be reduced. Note that due to the presence of seasonal signals in the data, it is important that an annual mean SSH is used (see Section 5.1.2).

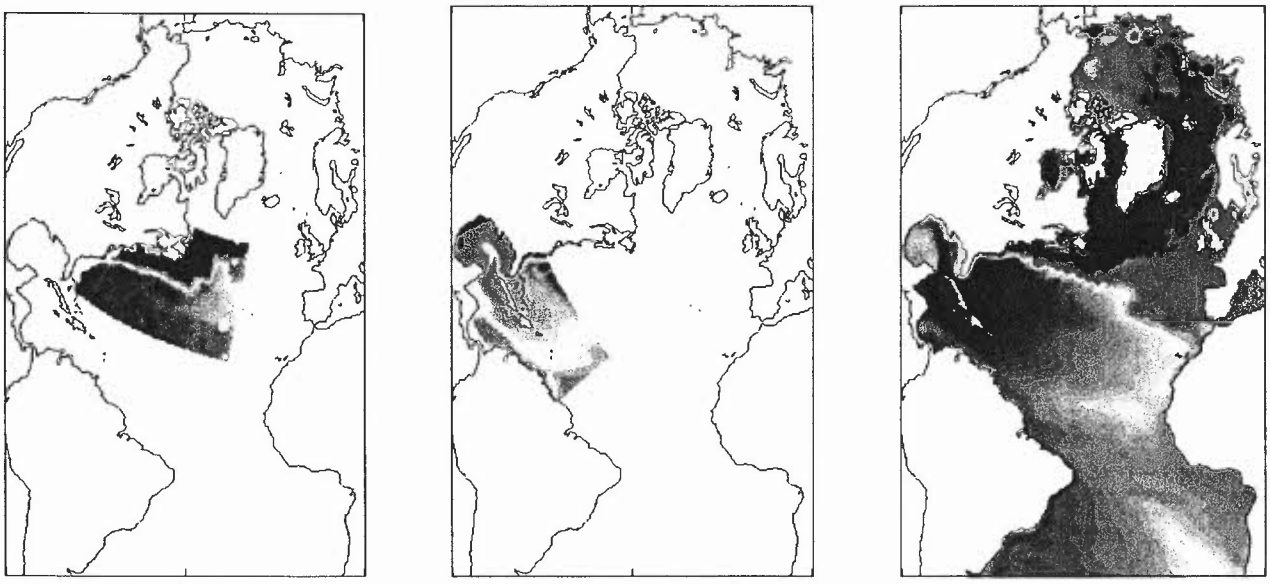

Figure 13: Mean sea surface height fields in the 1/3° model domain (a) Singh and Kelly (1997) Gulf Stream region data set; (b) annual mean from 1/9° model assimilating temperature profile and SST data; (c) annual mean from 1/3° model assimilating temperature and SST data.

A mean SSH data set for the Gulf Stream region on a 1° grid has been produced by Singh and Kelly (1997) (hereafter SK), and is shown in Figure 13(a). This mean SSH was calculated from a combination of in situ and satellite data, and is considered to be more accurate than mean SSH fields computed from in situ data alone (Killworth et al, 2000). In addition, the mean SSH from the 1/9° model (Figure 13(b)) might be anticipated to better represent the main current systems than the mean from the 1/3° model (Figure 13(c)). Hence a new mean SSH for use in the 1/3° model can be formed from a combination of the SK data set, the 1/9° annual mean SSH, and the original 1/3° annual mean SSH. The SK data set is probably the most reliable of the three, and overlaps slightly with the 1/9° model area. In this overlap region, the SK data set has therefore been used.

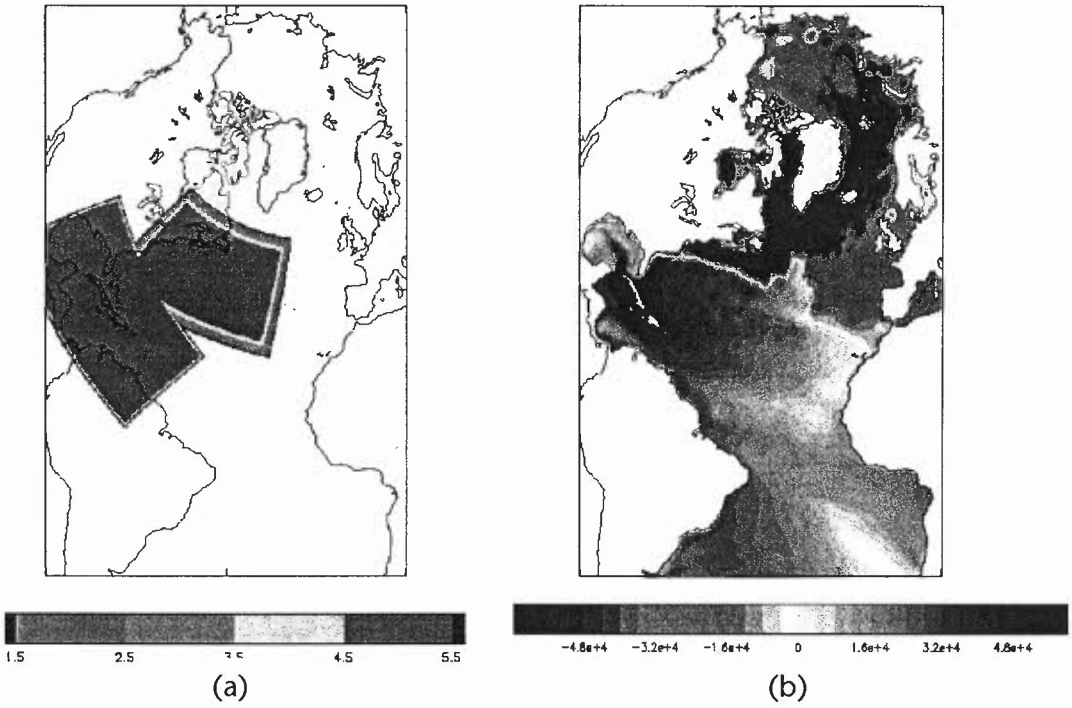

Figure 14: (a) Weights used for blending the Singh and Kelly, 1/9° annual mean, and 1/3° annual mean sea surface heights; (b) the new 1/3° model domain sea surface height.

The strategy for producing the new mean state is as follows:

1) Form 'weights' for use in blending the datasets on the SK grid and 1/9° grid. For SK, the weights are a field with value 1 at the outermost grid points of the SK domain, value 2 one row in from the outer edge etc., and value 5 at all points away from the edge of the domain. For the 1/9°, the weights are formed in the same way, but increase over the outermost 10 rows of points, to have a value of 10 in the inner parts of the domain.

2) Interpolate the SK and 1/9° means and weights to the 1/3° model grid. The interpolated weights are shown in Figure 14(a), which also indicates the areas of the 1/3° model domain covered by the SK and 1/9° means.

3) Blend the 1/3° and 1/9° means:

a) Compute the difference between the mean SSH in the 1/3° and 1/9° models in the blending region (i.e. wherever 0 < weights < 10 for the 1/9° model).

b) Offset the 1/9° SSH field by the difference in means so that the two fields have the same mean over the blending region.

c) Blend the two fields using a linear weighting over the blending region. On the outer edge, the field is set equal to the 1/3° mean SSH; at the inner points, the field is equal to the 1/9° mean SSH, with a linear transition between the two.

4) Blend the intermediate SSH with the SK field:

Repeats steps a)-c) above, but with a 5-point wide blending region.

5) Compute the spatial mean of the new SSH, and subtract from the new mean field.

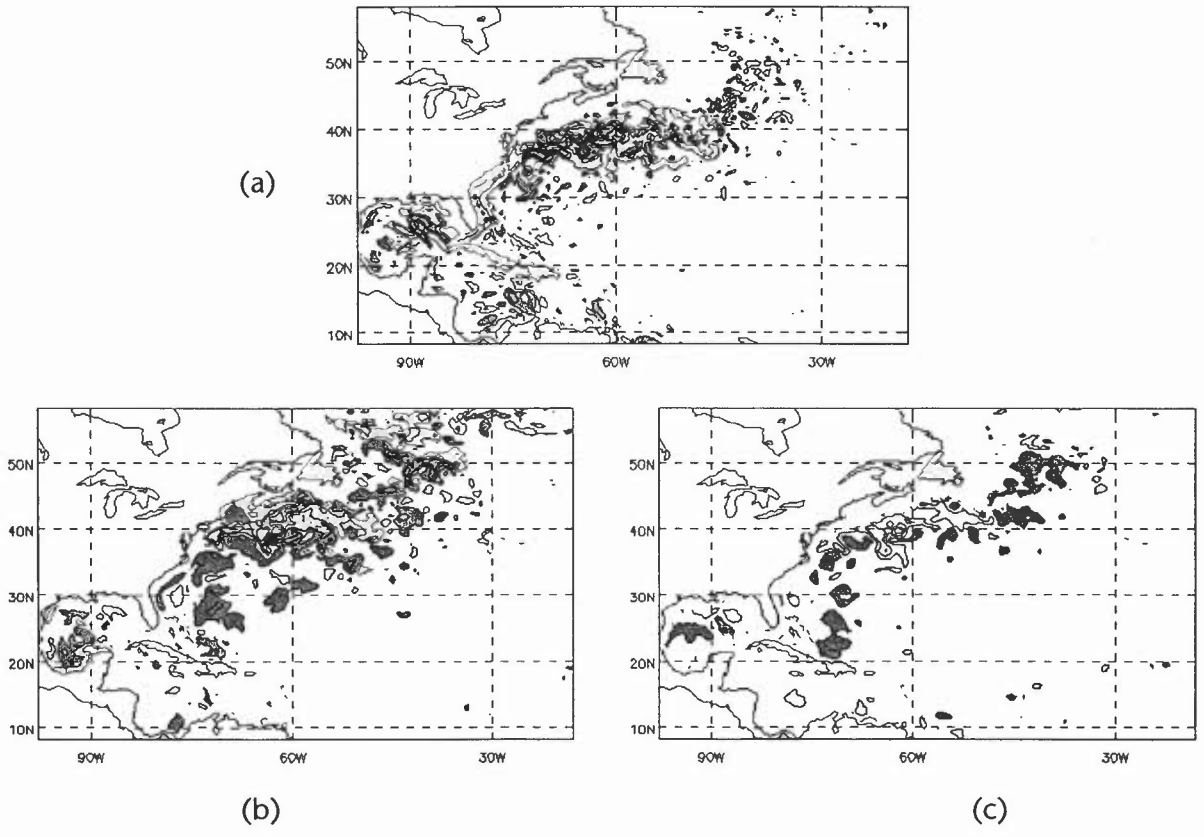

Figure 15: Differences between 1/3° model run using new mean SSH and 1/3° model run using the original mean SSH (new mean SSH run – original mean SSH run) in the Gulf Stream region. (a) surface current speed; positive (negative) differences in excess of 15 cm s-1 are shaded dark (light). Contours are plotted at ±15, ±45 and ±75 cm s-1. (b) SST (°C); (c) potential temperature at 243m depth (°C); positive (negative) differences in excess of 1.5 °C are shaded dark (light). Contours are plotted at ±1.5, ±4.5 and ±7.5 °C.

The resulting new mean SSH for the 1/3° model is shown in Figure 14(b).

The altimeter data have been reprocessed using the new mean SSH, and the 1/3° model has been integrated for several months using these data. Differences in surface current speed, SST and temperature at 250m depth between the run with the new mean SSH and the run with the original mean SSH after 2 months integration are shown in the Gulf Stream region in Figures 15(a), (b) and (c) respectively. Use of the new mean SSH has generated large temperature and current speed differences along the course of the Gulf Stream. The new mean SSH has shifted the path of the current to the south around the region of separation from the coast, as is particularly evident from the dipole structure seen in the current speed differences. In this region, the southerly shift of the current has led to a significant cooling of the SSTs, with temperature differences as large as 8-9°C. These large temperature difference persist to a certain extent throughout the ocean depth, with differences as large as 8°C seen at 250m depth.

6. Concluding Summary

The Cooper and Haines (1996) scheme for assimilating altimeter SSH data has been implemented for use in the FOAM high-resolution nested models. Following initial integrations in the 1/3° North Atlantic and 1/9° Carribean and Gulf of Mexico models in which the scheme was relatively ineffective, the assimilation parameters have been tuned, and questions concerning data usage have been addressed.

The basic tuning of the assimilation parameters has led to considerable improvements in the ability of the scheme to successfully initialise mesoscale features. The tuned parameters appear to be consistent with estimates computed from model output, and provide a suitable basis for a systematic calculation of estimates of the optimal parameter values. This will be carried out in the course of future work.

Considerations of the data processing applied have led to further improvements in the effectiveness of the assimilation. An appropriate choice of quality control ratio, together with a suitable background SSH for data processing have been seen to be extremely important in order to prevent data from high amplitude features being eliminated. Furthermore, the use of a more accurate mean SSH field with the sea level anomaly data has been shown to be capable of adjusting the path of the Gulf Stream in the 1/3° model, and offers the potential to reduce model systematic errors.

A set of 3-year long integrations pulling together all the improvements described here is currently underway. The main consequences for the use of altimeter SSH data in these longer experiments arising from the results presented here are:

- Use of correlation scales of 40km
- Observations given 4 times the weight given to the model
- Time window of 5 days either side of validity time used
- QC ratio of 3 used
- Blended Singh and Kelly, 1/9° and 1/3° mean SSH used

Results from these experiments will be used for a thorough assessment of the altimeter data assimilation scheme, and to examine the impact the altimeter data has on the forecast skill of the nested models. Results will also be used to investigate in detail the structure of correlation scales and error variances, with a view to implementation of spatial variations in both. Finally, the impact of the altimeter data assimilation in the global 1° model will also be assessed.

References

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, 32pp

Bell, M. B., R. M. Forbes, and A. Hines, 2000: Assessment of the FOAM global data assimilation system for real-time operational ocean forecasting. *J. Mar. Sys.*, **25**, 1, 1-22.

Bernstein, R.L., G. H. Born, and R. H. Whritner, 1982: SEASAT altimeter determination of ocean current variability. *J. Geophys. Res.*, **87**, 3261-3268.

Blayo, E., J. Verron, and J. M. Molines, 1994: Assimilation of TOPEX/POSEIDON altimeter data into a circulation model of the North Atlantic. *J. Geophys. Res.*, **99**, 24691-24705.

Brasseur, P., E. Blayo, and J. Verron, 1996: Predictability experiments in the North Atlantic Ocean: Outcome of a quasi-geostrophic model with assimilation of TOPEX/POSEIDON altimeter data. *J. Geophys. Res.*, **101**, 14161-14173.

Bryan, K., 1969: A numerical method for the study of the World Ocean. *J. Comput. Phys.*, **4**, 347-376.

Bryan, K., S. Manabe, and R. C. Pacanowski, 1975: A global ocean-atmosphere climate model. II. The oceanic circulation. *J. Phys. Oceanogr.*, **5**, 30-46.

Carnes, M. R., J. L. Mitchell, and P. W. Dewitt, 1990: Synthetic temperature profiles derived from Geosat altimetry – comparison with air-dropped expendable bathythermograph profiles. *J. Geophys. Res.*, **95**, 17979-17992.

Carnes, M. R., W. J. Teague, and J. L. Mitchell, 1994: Inference of subsurface thermohaline structure from fields measurable by satellite. *J. Atmos. Oceanogr. Tech.*, **11**, 551-556.

Cheney, R. E., 1982: Comparison data for Seasat altimetry in the Western North Atlantic. *J. Geophys. Res.*, **87**, 3247-3253.

Cheney, R. E., J. G. Marsh, and B. D. Beckley, 1983: Global mesoscale variability from collinear tracks of SEASAT altimeter data. *J. Geophys. Res.*, **86**, 10931-10937.

Clancy, R. M., J. M. Harding, K. D. Pollak, and P. May, 1992: Quantification of improvements in an operational global-scale ocean thermal analysis system. *J. Atmos. Oceanogr. Tech.*, **9**, 1, 55-66.

Cong, L. Z., M. Ikeda, and R. M. Hendry, 1998: Variational assimilation of Geosat altimeter data into a two-layer quasi-geostrophic model over the Newfoundland ridge and basin. *J. Geophys. Res.*, **103**, 7719-7734.

Cooper, M., and Haines, K., 1996: Data assimilation with water property conservation. *J. Phys. Oceanogr.*, **101**, 1059-1077.

Davis, R. E., R. de Szoeke, and P. Niiler 1981: Variability in the upper ocean during MILE. Part II Modelling the mixed layer response. *Deep-Sea Res.*, **28A**, 1453-1475.

Dombrowsky E., and P. de Mey, 1992: Continuous assimilation in an open domain of the northeast Atlantic. I. Methodology and application to Athena88. *J Geophys. Res.*, **97**, 9719-9731.

Drakopoulos, P. G., K. Haines and P. Wu, 1997: Altimetric assimilation in a Mediterranean general circulation model. *J. Geophys. Res.*, **102**, 10509-10523.

Esbensen, S. K., and Y. Kushnir, 1981: The heat budget if the global ocean: an atlas based on estimates from surface marine observations. Climate Research Institute, Oregon State Univ, Corvallis, Report No. 69.

Evensen, G., and Van Leeuwen, P. J., 1996: Assimilation of Geosat altimeter data for the Agulhas Current using the ensemble Kalman filter with a quasigeostrophic model. *Month. Weath. Rev.*, **124**, 85-96.

Ezer, T., and Mellor, G. L., 1994: Continuous assimilation of Geosat altimeter data into a 3-dimensional primitive equation Gulf-Stream model. *J. Phys. Oceanogr.*, **24**, 832-847.

Fox, A. D., K. Haines, B. A. de Cuevas, and D. J. Webb, 1999a: Altimeter assimilation in the OCCAM global model, part I: A twin experiment. Submitted to *J. Mar. Systems*.

Fox, A. D., K. Haines, B. A. de Cuevas, and D. J. Webb, 1999b: Altimeter assimilation in the OCCAM global model, part II: TOPEX/POSEIDON and ERS-1 assimilation. Submitted to *J. Mar. Systems*.

Forbes, R. M., 1996: Initial results from experiments assimilating satellite altimeter sea surface height data into a tropical Pacific ocean model. Ocean Applications Technical Note No. 12. Available from the U.K. Meteorological Office.

Fu, L.-L., E. J. Christensen, C. A. Yamarone Jr., M. Lefebvre, Y. Menard, M. Dorrer, and P. Escudier, 1994: TOPEX/POSEIDON mission overview. *J. Geophys. Res.*, **99**, 24369, 24381.

Gavart, M., and P. de Mey, 1997: Isopycnal EOFs in the Azores current region: A statistical tool for dynamical analysis and data assimilation. *J. Phys. Oceanogr.*, **27**, 2146-2157.

Glenn, S. M., and A. R. Robinson, 1995: Verification of an operational Gulf Stream forecasting model. In *Quantitative skill assessment for coastal ocean models*, Eds. D. R. Lynch, and A. M. Davies. American Geophysical Union.

Haines, K., 1991: A direct method for assimilating sea surface height data into ocean models with adjustments to the deep circulation. *J. Phys. Oceanogr.*, **21**, 843-868.

Hellerman, S., and M. Rosenstein, 1983: Normal monthly mean wind stress over the world ocean with error estimates. *J. Phys. Oceanogr.*, **13**, 1093-1104.

Hibler, W. D., 1979: A dynamic-thermodynamic sea ice model. *J. Phys. Oceanogr.*, **9**, 815-846.

Hurlburt, H. E., 1986: Dynamic transfer of simulated altimeter data into subsurface information by a numerical ocean model. *J. Geophys. Res.*, **91**, 2372-2400.

Jacobs, G. A., C. N. Barron, M. R. Carnes, D. N. Fox, H. E. Hurlburt, P. Pistek, R. C. Rhodes, W. J. Teague, J. P. Blaha, R. Crout, O. M. Smedstad, and K. R. Whitmer, 1999: Navy altimeter data requirements. NRL report, NRL/FR/7320—99-9696, Naval Research Laboratory, Stennis Space Center, MS.

Jaeger, L., 1983: Montly mean patterns of global precipitation. pp129-140 in Streer-Perrott, A., M. Beran, R. Ratcliffe (eds.) *Variations in the global water budget*. Dordrecht, Reidel.

Killworth, P. D., C. Dietrich, C. Le Provost, A. Oschlies, and J. Willebrand, 2000: Assimilation of altimetric data and mean sea surface height into an eddy-permitting model of the North Atlantic. Submitted.

Kraus, E. B., and J. S. Turner, 1967: A one dimensional model of the seasonal thermocline. II. The general theory and its consequences. *Tellus*, **19**, 98-106.

Large, W. G., J. C. McWilliams, and S. C. Doney 1994: Oceanic vertical mixing: a review and a model with a nonlocal boundary layer parametrization. *Rev. Geophys.*, **31**, 363-403.

Le Squere, B., and E. Dombrowsky, 1994: Description of the ocean prediction system SOAP 93. Oceans94. Oceans engineering for today's technology and tomorrow's preservation, Proceedings (1994,Volume III,Brest).

Levitus, S., 1994: World Ocean Atlas 1994, CD-ROM Data Set Doc., Volume 3: Salinity and Volume 4: Temperature. Natl. Oceanic and Atmos. Admin., Washington D. C.

Lorenc, A. C., R. S. Bell, B. MacPherson, 1991: The Meteorological Office analysis correction data assimilaiton scheme. *Quart. J. Roy. Meteor. Soc.*, **117**, 59-89.

Mellor, G. L., and T. Ezer, 1991: A Gulf Stream model and an altimetry assimilation scheme. *J. Geophys. Res.*, **96**, 8779-8975.

Metzger, E. J., H. E. Hurlburt, J. C. Kindle, R. C. Rhodes, G. A. Jacobs, J. F. Shriver, and O. M. Smedstad, 1998: The 1997 El Nino in the NRL Layered Ocean Model. 1998 NRL Review, NRL, Washington, DC, pp 63-71.

Oberhuber, J. M., 1988: An atlas based on the "COADS" data set: The budgets of heat buoyancy and turbulent kinetic energy at the surface of the global ocean. Report No. 15, Max Planck Institut fur Meteorologie, 20 pp.

Oschlies, A., and Willebrand, J., 1996: Assimilation of Geosat altimeter data into an eddy-resolving primitive equation model of the North Atlantic Ocean. *J. Geophys. Res.*, **101**, 14175-14190.

Pacanowksi, R. C., and S. G. H. Philander, 1981: Parameterization of vertical mixing in numerical models of tropical oceans. *J. Phys. Oceanogr.*, **11**, 1443-1451.

Redi, M. H., 1982: Oceanic isopycnal mixing by coordinate transformation. *J. Phys. Oceanogr.*, **12**, 1154-1158.

Singh, S., and Kelly, K. A., 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, 50pp.

Smedstad, O. M., and D. N. Fox, 1994: Assimilation of altimeter data in a two-layer primitive equation model of the Gulf Stream. *J. Phys. Oceanogr.*, **24**, 305-325.

Smith, W. H. F., and D. T. Sandwell, 1997: Global sea floor topography from satellite altimetry and ship depth soundings. *Science*, **277**, 1956-1962.

Storkey, D., 2001: Initial tuning of FOAM high resolution ocean models. Ocean Applications Technical Note No. 27, Met Office.

U. S. Naval Oceanographic Office and the U. S. Naval Ocean Research and Development Activity 1983 DBDB5 (Digital Bathymetric Data Base - 5 minute grid). U.S.N.O.O., Bay St. Louis.

Verron, J., L. Gourdeau, D. T. Pham, R. Murtugudde, and A. J. Busalacchi, 1999: An extended Kalman filter to assimilate satellite altimeter data into a nonlinear numerical model of the tropical Pacific Ocean: Method and validation. *J. Geophys. Res.*, **104**, 5441-5458.

Versal, P., H. J. Thiebaux, and L. L. Morone, 1988: Observation error variance estimates. Impact of their use on observed spatial correlation structure. Eighth conference on numerical weather prediction, February 22-26, 1988, Baltimore, Maryland.