PERGAMON

Continental Shelf Research 19 (1999) 1507–1520

CONTINENTAL SHELF RESEARCH

# Sea surface temperature assimilation for a three-dimensional baroclinic model of shelf seas

J.D. Annan*, J.C. Hargreaves

*Proudman Oceanographic Laboratory, Bidston Observatory, Birkenhead, Merseyside CH43 7RA, UK*

Received 23 April 1998; received in revised form 14 December 1998; accepted 13 April 1999

**Abstract**

In this paper, a novel interpolation technique based on a greatly simplified Kalman filtering approach is presented which enables satellite sea surface temperature observations to be assimilated into a three-dimensional baroclinic model of the North Sea. A series of numerical experiments of the annual cycle of seasonal stratification demonstrate a large improvement in the predictive ability of the model, and show good agreement with theoretical calculations of expected performance. 

*Keywords:* Numerical model; Data assimilation; Sea surface temperature; AVHRR; North sea

## 1. Introduction

Three-dimensional baroclinic finite-difference models of shelf seas which can calculate the evolution in time of physical features of the North Sea (such as currents and temperatures) have been developed recently (Proctor and James, 1996; Pohlmann, 1996). Runs with such models have shown that the general features of seasonal stratification can be successfully modelled, with the well-mixed southern region and the seasonally stratified northern region separated by a front.

Initialisation of the numerical model is critical for accurate operational modelling, and assimilation of observational data can greatly reduce forecast errors. Much research has been performed into assimilation techniques in meteorology, and more recently in oceanography (see Ghil and Malanotte-Rizzoli, 1991 for a review).

* Corresponding author.
*E-mail address:* jdan@pol.ac.uk. (J.D. Annan)

0278-4343/99/$ - see front matter 
PII: S0278-4343(99)00033-3

Techniques for data assimilation range from simple and computationally undemanding schemes, to complex but theoretically superior methods which are not yet computationally practical for operational modelling.

Work on assimilating conventionally made temperature measurements into deep ocean models has been performed by Derber and Rosati (1989). Their model had a horizontal resolution of 1° and 15 vertical layers. They assimilated data collected by in situ means during a 30-day window around the current model time, and used a simple horizontal correlation but no vertical correlation to smooth the data fields. In shelf seas with much shallower water (roughly 100 m depth), the greater vertical resolution attainable enables a seasonal thermocline to be more accurately resolved.

Conventionally collected surface temperature data are quite sparse in both space and time, and temperature-depth profiles are even less frequent. These data are also often not available in near real-time as would be required for operational modelling. Synoptic regular sea surface temperature (SST) measurements from satellite-borne advanced very high resolution radiometer (AVHRR) instruments provide regular and complete surface coverage (although cloud cover can restrict the data to roughly one half of the potential volume) but no subsurface information. This suggests that the transfer of information vertically through the water column might be of greater importance than horizontal smoothing for satellite SST assimilation in shelf sea models, and therefore there is a need for proper consideration to be given as to how best to assimilate these data into the models.

The purpose of this paper is to present a practical method for the assimilation of satellite SST observations into a pre-operational baroclinic model of the North Sea. In Section 2 the essential features of the finite-difference models are outlined. In Section 3 the theoretically optimal Kalman filter is described and it is shown that when some reasonable simplifications are introduced this method reduces to a simple one-dimensional interpolation scheme which is readily implemented. Tests with a one-dimensional model are described in Section 4. Tests with a three-dimensional model are described in Section 5. These demonstrate that the assimilation technique works well. The improvement in model performance is in good agreement with a theoretical calculation of the expected results. In Section 6 a brief summary of these results is given.

## 2. Numerical model

The baroclinic three-dimensional model used here is based upon the model of Proctor and James (1996), with modifications as described in Holt and James (1997). The three-dimensional incompressible Boussinesq equations are solved on a regular (in spherical coordinates) horizontal grid with a resolution of 1/3° longitude by 1/5° latitude (approximately 20 km), and covers the part of the southern North Sea bounded by 50° 41′N, 55° 41′ N, 2° 05′W and 9° 35′E. The model grid is shown in Fig. 2. The vertical discretisation is based on a $\sigma$-coordinate system with 13 levels. For convenience and computational efficiency, the horizontal equations are split into the

depth-mean components and the deviations around the mean. The numerical method uses a split time step scheme whereby the depth-mean equations are solved using a time step of 300 s, and the deviation equations have a much larger time step of 900 s. Vertical mixing is controlled by the one-dimensional solution of a variant of the Mellor-Yamada Level 2.5 Scheme which is described by Sharples and Simpson (1995), using the approach described by Annan (1999).

The heat flux at the surface is specified using the algorithm of Goldsmith and Bunker (1979) and there is no heat flux through the sea bed. Meteorological forcing is provided by the UK Meteorological Office model.

Although models of this type are effective in providing a qualitative description of the thermal structure of the sea (Holt and James, 1997), they can be seen to drift away from the true sea state over time and so would be of limited use in an operational scenario without a method for constraining the model by reinitialising it with the use of observational data.

## 3. Assimilation

The Kalman filtering method of data assimilation (Kalman, 1960) requires knowledge of the error statistics of both the model and observations. Calculation of the error covariance matrices is required, but this is a formidable task which is only computationally feasible for the very simplest of models. If, however, enough simplifying assumptions are made, the problem can be reduced to a manageable level. The Kalman filter is outlined below, and then some reasonable simplifications are suggested which enable implementation with minimal computational effort.

### 3.1. The Kalman filter

Assume that the model is linear, without forcing. This can be written as

$$x^f_{t+1} = Ax^a_t, \tag{1}$$

where $x^f_{t+1}$ is the state vector which represents all model variables model in the forecast for time $t+1$, $x^a_t$ is the state vector of the analysed (i.e. after assimilation) model output for time $t$ and $A$ is the matrix which describes the propagation of the model forwards through a single time step. For a linear model, the particular forced solution can be added to the homogeneous solution at a later stage. For a non-linear model, a linearised set of model equations can be derived and used for the assimilation scheme and this will normally give good results since the non-linear terms are usually of limited significance over the short term (Ghil and Malanotte-Rizzoli, 1991, p. 164).

If there is a vector of observations at time $t+1$, denoted by $x^o_{t+1}$, then the analysis step is performed by

$$x^a_{t+1} = x^f_{t+1} + K_{t+1}(x^o_{t+1} - Hx^f_{t+1}), \tag{2}$$

where $K_{t+1}$ is the Kalman gain matrix and $H$ is a matrix which converts model variables to observed variables (perhaps by interpolating between grid points, or by calculating a linear function of model variables). The bracketed expression is sometimes referred to as the innovation vector.

Note that (2) is essentially the same analysis step as is used for simple statistical interpolation schemes (e.g. direct insertion or "optimal interpolation"). The difference between the Kalman filter and the simpler schemes lies in the way in which the gain matrix is calculated. In simple interpolation schemes, the gain matrix is typically assumed to have a very simple form and does not vary in time. The full Kalman filter is implemented by the following set of equations. Firstly, the model forecast covariance matrix $P^f$ is calculated from the previous analysis covariance matrix $P^a$ and the system (or model) error covariance matrix $Q$ by

$$P^f_{t+1} = AP^a_t A^{\mathrm{T}} + Q_t. \tag{3}$$

Then the Kalman gain matrix is calculated via

$$K_{t+1} = P^f_{t+1} H^{\mathrm{T}} [HP^f_{t+1} H^{\mathrm{T}} + R_{t+1}]^{-1} \tag{4}$$

where $R$ is the observation error covariance matrix.

The analysis step (2) is performed and then the analysed error covariance matrix is calculated by

$$P^a_{t+1} = (I - K_{t+1} H) P^f_{t+1}. \tag{5}$$

For a state vector (model size) of $n$ variables, $A$ is a $n \times n$ matrix, but the evolution of the model variables only depends on their near neighbours and so the model propagation only requires $O(n)$ operations. However, the calculation of the error covariance matrices would appear to require $O(n^2)$ calculations, and in any case the system and observational error covariance matrices $Q$ and $R$ are not known in practice. For a workable application, therefore, it is clear that simplifications have to be made. Fortunately, a careful consideration of the three-dimensional numerical model reveals that some significant simplifications can be made without creating large errors.

### 3.2. Simplifications

The residual horizontal currents are, in general, small (rarely exceeding 0.05 $\mathrm{ms}^{-1}$), so the time scale of error propagation across the grid is large compared with that of vertical mixing. The first simplification, therefore, is to assume that the horizontal correlations are small enough to be ignored. This means that, within the assimilation routine, the model can be treated as an array of one-dimensional vertical mixing models, and any relationship between horizontal temperature gradients and density-driven flows can be ignored. It should also be noted that the processing of satellite SST data tends to introduce horizontal correlations in the observations, through spatial interpolation, smoothing, and error correction. This reduces the need to include a horizontal correlation in the gain matrix. If significant large-scale gravity

waves (driven by horizontal density differences) were to occur in practice, then horizontal smoothing could be added, but this has not proved to be necessary during this work.

The second simplification is based upon the fact that the turbulent kinetic energy is usually close to a quasi-equilibrium level — indeed, many simpler three-dimensional baroclinic models include this assumption explicitly (Proctor and James, 1996; Poulmann, 1996) — and will rapidly readjust to any imposed vertical temperature gradient. It follows that the temperature can be adjusted independently of the turbulent kinetic energy without interfering with the model dynamics in the vertical. No attempt will be made to alter the position of the thermocline during the analysis step, although it remains free to evolve according to the model dynamics. This decision is virtually forced since the vertical resolution of the model (of the order of 5 m) does not enable small adjustments to be made to the position of the thermocline, and SST observations cannot in any case provide much useful information on its depth. This simplification reduces the problem still further to one of the assimilating SST information in a vertical water column with a fixed structure of mixed and stratified regions.

When a thermocline exists, mixing is relatively rapid within the upper and lower layers (on the time scale of hours or days) but there is very little mixing a cross the thermocline itself. For example, a typical diffusion coefficient of $10^{-5}\,m^2s^{-1}$ within the thermocline implies a time constant of almost 350 h for mixing across a 5 m width, whereas a 20 m wide mixed layer with a diffusion coefficient of $10^{-2}\,m^2s^{-1}$ has a time constant of under 6 h. The minimum inter-observational time for successive satellite observations is 12 h. The third simplification is, therefore, to assume that the areas above and below the thermocline are each well mixed, but that there is negligible mixing across the thermocline. In the absence of a thermocline, the entire water column is well-mixed and if multiple thermoclines were to occur, then these would divide the water column into separate well-mixed regions.

All of these simplifications are not actually imposed on the model itself, which is free to evolve according to its underlying dynamical equations. The simplifications are merely used in the calculation of interpolation coefficients of the Kalman gain matrix. The next section shows how as a result this calculation is reduced to an easily manageable level.

### 3.3. Assimilation in the simplified system

The first simplification described above reduces the problem to one of assimilating SST information into a one-dimensional vertical mixing model. The quasi-stationarity of the turbulent energy means that only the heat diffusion equation needs to be considered in the one-dimensional model. The error in the modelled temperature fields is due primarily to errors in the modelled heat flux through the sea surface (due to inaccuracies in both the meteorological forcing and the heat flux sub-model). The upper mixed layer is rapidly mixed, so this heating error is uniformly distributed across the upper grid points. As the temperature observation is limited to the top

grid point, only the first column of the Kalman gain matrix is needed. Some simple algebra reveals that this is given by

$$\mathbf{k} = \begin{pmatrix} \frac{\alpha}{\alpha + r} \\ \vdots \\ \frac{\alpha}{\alpha + r} \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \tag{6}$$

where $\alpha$ is the forecast model error variance at all grid points within the upper mixed layer immediately prior to assimilation and $r$ is the observational error variance for the sea surface temperature measurement. The analysed error variance is given by $\alpha r/(\alpha + r)$ at each grid point in the upper layer. If the error statistics are in a quasi-equilibrium state and that the model error variance at each grid point in the upper layer increases by an amount $d$ between analysis steps, then the analysed error variance will also be given by $\alpha - d$. Therefore, $\alpha$ can be estimated from

$$\alpha - d = \frac{\alpha r}{\alpha + r} \tag{7}$$

which can be rearranged to give

$$\alpha = \frac{d + \sqrt{d^2 + 4rd}}{2} \tag{8}$$

and the gain coefficients in (6) can readily be calculated. The computational effort required by this analysis routine is insignificant by comparison with the model integration. The increase in model error variance $d$ is a priori unknown. Any error in the modelled sea surface temperature is due primarily to innacuracy in the heat flux model (and meteorological forcing data) and the resulting heat surplus or deficit is spread evenly throughout the upper mixed layer. Therefore, $d$ will have the form

$$d = \frac{k\Delta t}{\Delta z}, \tag{9}$$

where $\Delta t$ is the inter-assimilation time interval (not the model time step), and $\Delta z$ is the depth of the upper mixed layer (not the vertical grid spacing in the model). The base of the upper mixed layer was taken to be the highest grid point with a diffusion coefficient of $10^{-4}\,\mathrm{m^2 s^{-1}}$ or lower, but the results showed limited sensitivity to this choice of value. There remains the unknown constant $k$, and a range of plausible values for this constant was tested in the following experiments.

## 4. One-dimensional model experiments

### 4.1. Model and data

The original three-dimensional problem has been reduced to an array of one-dimensional problems and so it is logical to make the first tests of the technique outlined above using a vertical one-dimensional model. The model used is located at the point 55°30′N, 0 55′E in the North Sea. This model, which is thoroughly described in Warrach (1998) and Annan (1999), has the advantage of running very much faster than a full three-dimensional model, but the vertical turbulent mixing and thermal diffusion is modelled in essentially the same manner. Moreover, a time-series of temperature profiles was collected at this point for much of 1989 as part of a major experimental study in 1988–1989 (the North Sea Community Project; see Charnock et al., 1994). This data set enables the assimilation system to be independently validated.

During the North Sea Project, temperature-depth profiles were routinely monitored by a thermistor chain at station CS (55°30′N, 0°55′E) during late 1988 and much of 1989. The record for 1989 is almost continuous from early April to the end of September, with the exception of most of June, and so it contains the onset and erosion of thermal stratification in the seasonal cycle. The highest point at which measurements were routinely made was at 10 m depth. The thermistors were recalibrated at approximately monthly intervals. Some drift could be seen to occur between calibrations. The accuracy of these measurements is in the region of 0.1°C RMS error.

Satellite-derived sea surface temperature observations from the NOAA/NASA Pathfinder Program (Smith et al., 1996) are freely available to researchers in a variety of products. Full details can be found from the web page http://podaac-www.jpl.nasa.gov/ where the archive is located. The data are derived from 5 channel AVHRR on board the NOAA-7, -9, -11 and -14 polar orbiting satellites, and calibrated globally using buoy data. The data used here are provided on a nominal 9 km grid twice daily at approximately 2 a.m. and 2 p.m. local time (cloud cover permitting). We use the "best SST" product which has been processed to eliminate the worst measurement errors. In a fully operational setting, where a lower quality product might have to be used, a simple quality control check could be implemented by for instance rejecting data where the satellite measurement differs from the model prediction by more than a specified margin. Even with the "best SST" product, there are still significant errors. The accuracy of the satellite data (which is only quoted to a resolution of 0.15°C) depends in part on the accuracy with which the difference between the bulk sea surface temperature and the skin surface temperature can be estimated. It has been found that this difference is more variable during the day than the night, due to the occasional presence of a hot surface skin layer (Schluessel et al. 1990; Donlon and Robinson, 1997). For this reason the night time and day time satellite data sets have been considered both separately and together. The NOAA/NASA Pathfinder SST Validation web page http://www.ccpo.odu.edu/ ~ lizsmith/ contains information on the large-scale spatial and temporal errors in the Pathfinder SST data set, and shows that persistent regional biases of up to 1°C may occur. A direct comparison of the

satellite data and coincident in situ measurements at station CS during 1989 reveals a bias of $-0.41 \pm 0.06$°C for the night time observations, and $-0.34 \pm 0.09$°C for the day time observations (errors, where quoted, are the standard deviation). In both cases the satellite observations are too cold by comparison with the in situ measurements. These biases are highly significant (but the difference between the day and night figures is not) and since they exceed the known accuracy of the thermistor chain by a large margin, this bias is attributed entirely to the satellite data. For use in the model tests, the bias was subtracted from the observational data since the theory of the assimilation assumes unbiased data. The residual RMS errors were 0.39°C for the night time observations and 0.57°C for the day time observations. These figures are broadly consistent with the work of Schluessel et al. (1990) and Donlon and Robinson (1997) who found that satellite SST observations made during the night are normally more reliable than those made during the day.

### 4.2. Results of assimilation

It is straightforward to run the one-dimensional model using suitable meteorological data (UKMO) and assimilating the satellite observations. The error of the model output can then be calculated by comparison with the coincident in situ measurements to provide a measure of the success of the scheme. The accuracy of the satellite observations has been calculated by comparison with the in situ measurements. There were 137 night time and 135 day time satellite observations, which were not uniformly distributed in time, but appeared in clumps. The model variance ($d$ in (8)) is a priori unknown, and is calculated for each assimilation step via Eq. (9), using an assumed value for $k$ (constant during a model run) and the appropriate mixed layer depth and time interval. The assimilation scheme was tested with a wide range of values for $k$, and the best results were obtained when $k$ was set at a value equivalent to a model variance of $(1/4°C)^2$ per day for a thermocline at a depth of 20 m. The results are robust with respect to this figure, however, and there was only a very slight degradation in the overall results when values of $(1/8°C)^2$ and $(1/2°C)^2$ per day were used.

The assimilation was performed with either day or night observations, and also with both data sets. The results are shown in Table 1 which indicates that the model forecast error is greatly reduced when assimilation is performed. The more accurate night time data does have a larger effect than the day time data, but only slightly, and combining both data sets produces the most accurate results.

The results can be validated by using (8) to estimate the errors that should have been produced by these model runs. For instance, if the day time and night time satellite observations were available at each point at regular 3 day intervals (this being

Table 1
RMS Errors from one-dimensional assimilation (°C)

| No assimilation | Night time data | Day time data | All data |
|---|---|---|---|
| 1.1 | 0.43 | 0.45 | 0.35 |

the approximate average frequency at which information was obtained), then the error of the model forecasts when assimilating day time data alone, using the previously stated variances for the model and observations, should be 0.51°C, and for night time data the corresponding figure is 0.41°C. When both sets of data are used, the error should be 0.38°C. These figures are close to those obtained in the experiments above.

## 5. Three-dimensional experiments

In order to investigate whether the assimilation method causes problems with horizontal density-driven currents, it must be tested in a three-dimensional model. The model used here is the 20 km resolution baroclinic model of the southern North Sea which has already been outlined in Section 3.

### 5.1. Data

During the North Sea Project, the Royal Research Ship *Challenger* undertook a series of 15 regular surveying cruises in the model area at approximately monthly intervals from August 1988 until October 1989. The cruise track covered roughly 2000 nautical miles and lasted for 12 days. Many types of data were collected by the ship and also by permanent moorings. The one-dimensional model experiments reported above used data collected from one of these fixed sites. The sea surface temperature was measured continuously at 2 m depth by an autoranging TSG103 thermosalinograph on board the ship. This instrument was calibrated regularly against CTD-casts. The results were archived at 30 s intervals throughout the cruises, along with the ship's position. There were roughly one thousand CTD-casts taken during this period, and the data from these were also archived. Although the SST data are less accurate than the CTD-casts, they are still sufficiently accurate, when compared to the model and satellite observations, to be considered as 'truth', and these hourly figures have the advantage of providing an abundance of data points coincident with the satellite images. Furthermore, these data also provide interesting information concerning the sub-grid-scale variability of sea surface temperature. The CTD-casts can also be used to investigate the subsurface effects of the assimilation.

Initially, to assess the accuracy of the satellite observations, a direct comparison was performed between the satellite data and the ship's SST measurements whenever observations coincided. This was deemed to have occurred when the ship had measured the water temperature inside a pixel which was observed by the satellite during the same hour. For the day time data, the bias was $-0.34 \pm 0.06$°C with a residual root mean square difference of 0.47°C, and the night time data was fractionally more accurate with bias and residual RMSD of $-0.31 \pm 0.06$°C and 0.45°C, respectively. For the model runs, however, the satellite pixels (nominally 9 km square, but in fact close to 6 km in latitude) were averaged to the 20 km model grid squares. A comparison between these averaged satellite data and the ship measurements gave larger bias and residual RMSD measurements of $-0.45 \pm 0.06$°C and

0.54°C during the day, and $-0.44 \pm 0.07$°C and 0.63°C during the night, so these figures are used for the calculations of the gain required in the analysis routine.

There are several reasons for the degradation of the satellite data when averaged over the larger bins. The most significant reason is that some of the ship measurements were made in locations which were classified as land by the satellite (but not by the model). In this very shallow coastal water, the temperature can be much warmer than a few kilometers off shore where the nearest wet satellite pixel is centred. Averaging the satellite pixels to the model grid in this case results in the deeper water pixel being compared to the shallow coastal water temperature, resulting in a large bias. Secondly, there is evidence from the satellite data that pixels surrounding cloudy regions are themselves contaminated, but not enough to be identified by the cloud detection algorithm used in the processing of the satellite data. Analysis of the year's satellite data for the model region shows that, on average, pixels adjacent to cloudy regions are 0.13°C colder than pixels which are far away from cloud. This bias in the satellite measurements decreases to 0.07°C and then 0.03°C as the distance to the nearest cloud increases by single pixels. When the pixel containing the ship is cloudy, but there are clear pixels within the same model grid box, then averaging over the grid box will increase the bias since the clear pixels have in fact been contaminated by the nearby cloud. Finally, there is spatial variability in the temperature field on a smaller scale than the model grid boxes allow. For example, consecutive hourly temperature measurements from the ship's thermosalinograph show a RMS difference in excess of 0.3°C. The distance travelled during this time is only approximately 10 km and so spatial averaging of the satellite measurements within a model grid box will generate what is effectively a discretisation error. This will increase the residual RMS error but not alter the bias.

In the comparison between the ship's temperature measurement and the model results, the modelled temperature could either be taken to be that of the nearest model grid point, or a bilinearly interpolated value from the surrounding points. A more careful look at the time series of ship measurements showed that the temperature gradients encountered in successive hours were not correlated. This indicates that the sea surface temperature field is not smooth over the scale of the grid spacing. For this reason, bilinear interpolation does not improve the accuracy of the model, and the reported results are from comparison with the nearest grid point.

### 5.2. Results of assimilation

The model was run for the period 1st November 1988 – 31st October 1989. The last two months of 1988 were used as a spin-up period, and then model output was compared to the ship's data.

Model runs were performed both with and without assimilation. In all cases, the initial water temperature was taken to be a uniform 11.24°C, this being the mean of all temperature measurements collected during the month of November 1988 as part of the North Sea Project (Holt and James, 1997). Sensitivity tests demonstrated that varying this initial temperature by 1°C had a negligible effect on the results.

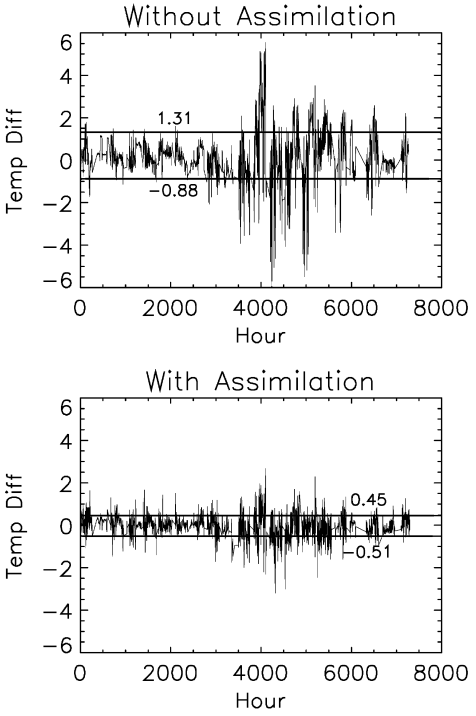


Fig. 1. Time series of temperature difference between three-dimensional model runs and ship SST measurements (degrees Celsius). Hours are counted from the start of 1989. Horizontal lines indicate one standard deviation from mean error.

Figs. 1 and 2 show the model performance, analysed as a function of time and space. Both figures show the difference between the ship SST measurements and the modelled SST for the nearest grid point. Fig. 1 shows this difference as a time series. The time is given as the number of hours from the start of the year, and the horizontal lines indicate one standard deviation from the mean temperature difference. The errors are greater during the summer, due to the presence of thermal stratification which means that any error in the total heat flux through the ocean surface is distributed across a smaller depth of water. Fig. 2 shows the same data, but collected by model grid box and then averaged over the whole year. A dot in the centre of a grid box indicates that it was never visited by the ship and therefore no error can be calculated for it. This spatial analysis indicates that the largest errors occur both in the northern (seasonally stratified) region and also in some grid boxes adjacent to the coast. These latter errors are probably due to a combination of the shallow water and the inability of the model to resolve the variations in bathymetry and currents close to the coast.

A comparison of the model and the CTD data confirms the improvement due to the assimilation procedure. The model error at the second grid point from the surface

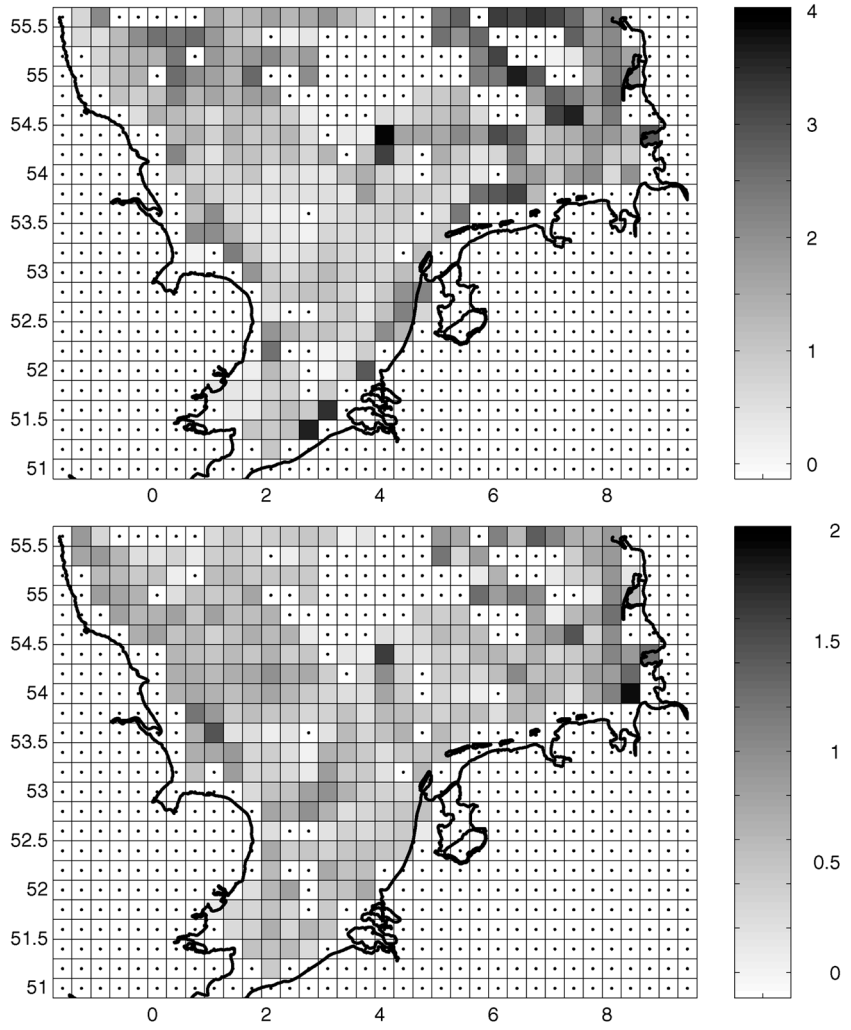


Fig. 2. RMS difference between model temperature and ship SST measurements. Top plot shows results without assimilation, bottom plot shows results with assimilation. Scale is degrees Celsius, and horizontal axes are degrees of latitude and longitude.

layer, when compared to the appropriate CTD measurements, was reduced from 1.4 to 0.5°C. At the second-to-bottom grid point, the improvement is very much reduced, only decreasing the error from 1.5 to 1.1°C when assimilation is used. This is easily explained by classifying the CTD data into two sets, depending on whether the water column is stratified (defined here by a surface-to-bed temperature difference of 0.5°C or more) or not. In both cases, the performance near the sea surface is very similar, with a very large error reduction to around 0.5°C. At the sea bed, assimilation has

essentially no effect in stratified water. This is in accordance with both the physical reality and the numerical model, since in the presence of strong stratification, heat cannot pass between the upper and lower mixed layer. In marginally stable cases (when stratification is in the process of becoming established or eroded), a small adjustment to the temperature of the upper mixed layer could lead to or prevent full mixing of the water column, but in these cases the surface to bed temperature difference is in any case small. In the unstratified case, the near-bed temperature error is substantially reduced, from 1.5 to 0.9°C, when assimilation is performed. In some cases, the model is stratified when the data shows that the water column is well-mixed in that particular grid box, and assimilation does not help on these occasions. These errors may be partly due to the limited resolution of the model both vertically and horizontally, and it would be interesting to see to what extent a higher resolution model could improve in this area.

The effect of the assimilation on currents (both barotropic and baroclinic) and sea surface elevations was extremely small. The temperature fields evolve slowly and smoothly enough that large dynamical imbalances are not generated by the assimilation procedure.

Due to the intermittent and clumped nature of cloud cover, the mean forecast error of 0.48°C (for the run with assimilation) is an average over a wide range of time intervals, with a mean interval of 1.5 days. This error is substantially lower than that of the satellite data and has very low bias. This encouraging result suggests that this system will be valuable for operational forecasting.

## 6. Conclusions

A technique based on a greatly simplified Kalman filter has been developed which enables satellite sea surface temperature observations to be assimilated into a baroclinic three-dimensional model of the southern North Sea. The simplifying assumptions reduce the analysis routine to a simple interpolation algorithm requiring minimal computational effort. Numerical experiments have been performed with a one-dimensional model, and also with a low resolution three-dimensional model of the southern North Sea. The assimilation system is robust with respect to varying parameter choices. The results show that the system produces a large reduction in model forecast error, and the performance is consistent with theoretical calculations.

## References

Annan, J.D., 1999. Numerical methods for the solution of the turbulence energy equations in shelf seas. International Journal for Numerical Methods in Fluids 29, 193–206.

Charnock, H., Dyer, K.R., Huthnance, J.M., Liss, P.S., Simpson, J.H., Tett, P.B., 1994. In: Understanding the North Sea System. Royal Society, London.

Derber, J., Rosati, A., 1989. A global oceanic data assimilation system. Journal of Physical Oceanography 19, 1333–1347.

Donlon, C.J., Robinson, I.S., 1997. Observations of the oceanic thermal skin in the Atlantic Ocean. Journal of Geophysical Research 102 (C8), 18585–18606.

Ghil, M., Malanotte-Rizzoli, P., 1991. Data assimilation in meteorology and oceanography. Advances in Geophysics 33, 141–266.

Goldsmith, R.A., Bunker, A.F., 1979. Woods Hole Oceanographic Institution collection of climatology and air–sea interaction (CASI) data, Technical Report WHOI-79-70, Woods Hole Oceanographic Institution.

Holt, J., James, I.D. (1997). A simulation of the southern north sea in comparison with the north sea project. *Continental Shelf Research*, submitted.

Kalman, R.E., 1960. A new approach to linear filtering and prediction problems. Journal of Basic Engineering 82D, 33–45.

Pohlmann, T., 1996. Predicting the thermocline in a circulation model of the North Sea—Part I: model description, calibration and verification. Continental Shelf Research 16 (2), 168–194.

Proctor, R., James, I.D., 1996. A fine resolution 3D model of the Southern North Sea. Journal of Marine Systems 8, 285–295.

Schluessel, P., Emery, W.J., Grassi, H., Mammen, T., 1990. On the bulk-skin temperature difference and its impact on satellite remote sensing of sea surface temperatures. Journal of Geophysical Research 95 (C5), 13341–13356.

Sharples, J., Simpson, J.H., 1995. Semi-diurnal and longer period stability cycles in the Liverpool Bay region of freshwater influence. Continental Shelf Research 15 (2–3), 295–313.

Smith, E., Vazquez, J., Van Tran, A., Sumagaysay, R., 1996. Satellite-derived sea surface temperature data available from the NOAA/NASA Pathfinder Program, Technical report, American Geophysical Union, http://www.agu.org/eos_elec/95274e.html.

Warrach, K.I., 1998. Modelling the thermal stratification in the North Sea. Journal of Marine Systems 14 (1–2), 151–166.