Estuarine, Coastal and Shelf Science (2001) **53,** 459–466
doi:10.1006/ecss.1999.0626, available online at http://www.idealibrary.com on IDEAL®

Hindcasting Coastal Sea Levels in Morecambe Bay

J. D. Annan

Proudman Oceanographic Laboratory, Bidston Observatory, Prenton, Wirral, CH43 7RA, U.K. and FRSGC, 3173-25 Showa-machi, Kanazawa-Ku, Yokohama City, 236-0001, Japan

Received 31 December 1998 and accepted in revised form 7 July 1999

A method of variational data assimilation (by means of an adjoint model) is used to improve the boundary conditions for a high resolution depth-averaged two dimensional non-linear hydrodynamic model of Morecambe Bay. The modelled region contains a large inter-tidal area and has an unusually high tidal range. The assimilation procedure is robust and effective in reducing model error. This system performs considerably better than the previous simple method used for assimilating data into the same numerical model. Height errors as low as 10 cm (root mean square) can be achieved.

Keywords: hydrodynamic numerical model; variational data assimilation

Introduction

The coastal zone of the U.K. has great economic and environmental significance, with 20% of the U.K. population living within 1 km of the coast. Maps of the intertidal zone are required to manage and monitor this region effectively, for example as part of statutory Shoreline Management Plans. There are several possible methods for mapping the intertidal zone including *in situ* surveying and remote sensing techniques. The aim of the INDUS project is to investigate the effectiveness of the ' waterline ' method (Koopmans & Yang, 1994) for mapping the intertidal zone. This method combines a satellite image (which locates the instantaneous position of the coastline) with modelled sea surface elevation (which indicates the height along that coast-line). A set of these heighted coastlines, distributed over a range of tidal phases, can then be used to produce a Digital Elevation Model (DEM) of the intertidal region. A detailed description of this process can be found in Mason *et al.* (1995). The resulting maps, some examples of which are displayed in Mason *et al.* (1998), have a variety of uses, including the monitoring of bathymetric change and evolution, and the estimation of sediment transport (Mason *et al.*, 1999).

One of the fundamental limitations on the accuracy of the resulting map is the accuracy with which the numerical model can describe the water level along the coastline. This factor is particularly critical in areas with gently sloping beaches. On steep beaches the primary source of error in the DEM is caused by the resolution of the satellite image which necessarily limits the accuracy with which the coastline can be positioned. For example, in an ERS SAR image, the pixel size is 12·5 m. For a typical steep beach slope of 1:30, the associated height error due to the pixel width is about 40 cm. For a more shallow beach (say 1:500 slope) the error associated with the pixellation is in the region of 2 cm. In these areas the accuracy (and therefore value) of the DEM depends largely on the accuracy with which the sea surface elevation is modelled. Achieving a heighting accuracy in the region of 10 cm has been identified as a key objective in meeting the requirements of the potential end-users of the DEM (Davenport *et al.*, 1996).

The purpose of the work described in this paper is to investigate the performance of a high-resolution implementation of the hydrodynamic model, and to explore to what extent its results can be improved by the assimilation of local *in situ* measurements of water level. The assimilation technique used is a variational method implemented by means of an adjoint model. Previous work in this field has demonstrated the value of assimilation for improving the performance of similar but simpler numerical models by optimizing model parameters such as friction coefficients and bathymetry (Spitz & Klink, 1998; Lardner *et al.*, 1993), as well as by varying the boundary conditions (Lardner, 1993). The work described here tests the adjoint method with a fully non-linear high resolution depth-averaged two dimensional numerical model of a region with a large tidal range and substantial wetting and drying areas. The goals are firstly to ensure that the method is robust and effective in assimilating

E-mail: jdan@pol.ac.uk, jdannan@jaustec.go.jp

0272–7714/01/100459+08 $35.00/0

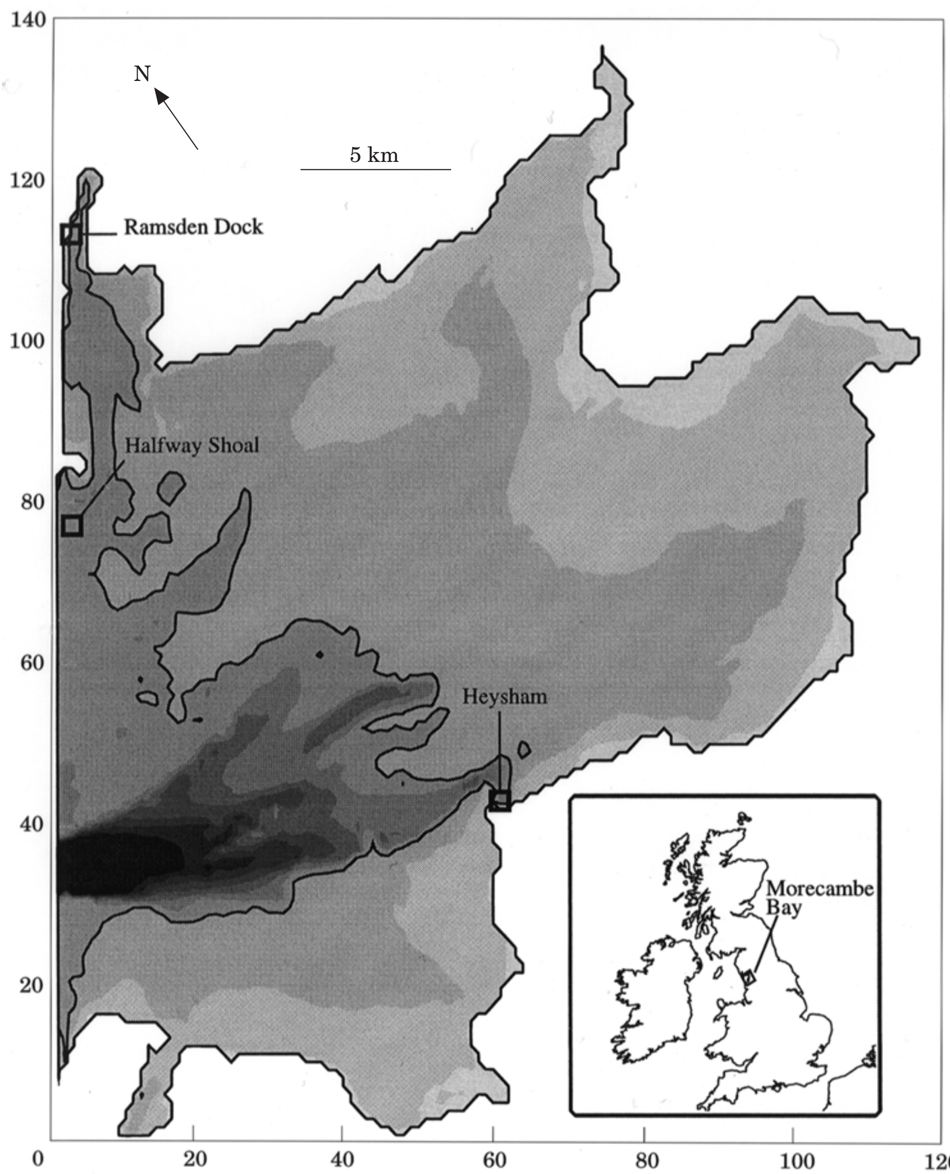

FIGURE 1. Model domain and bathymetry. Grid size is 250 m, depth contours are at 4 m intervals. Solid lines indicate approximate positions of high and low water levels.

in situ data into this numerical model, and secondly to evaluate the resulting accuracy of the numerical model in describing sea surface elevation at the coast.

Morecambe Bay

The site for this work is the Morecambe Bay area in the U.K. This region, located on the eastern side of the Irish Sea, contains an intertidal area of approximately 340 000 ha which represents roughly 12% of the U.K.'s entire inter-tidal area. The bay has an unusually large spring tidal range of around 8·2 m, with a substantial tidal asymmetry. The mean spring high and low water levels at Morecambe are 4·6 m and −3·6 m respectively (all elevations are given relative to Ordnance Datum Newlyn) and this asymmetry increases further within the bay. The model domain and bathymetry are shown in Figure 1. The gridded bathymetry was generated from a high resolution (50 m) DEM of the area constructed from a variety of sources including an earlier application of the waterline method. This grid generation is described by Mason *et al.* (1999). The accuracy of the depth values in the inter-tidal zone was estimated to be in the region of 40 cm (RMS) but the strong temporal decorrelation indicated by the data implies that any map will rapidly lose accuracy as the region evolves over time. Subtidal depth data, collected from a variety of sometimes inconsistent sources, is generally of lower quality. Discrepancies between the model bathymetry and the geographical location of the tidal gauges required that the grid points used for comparison with the tide gauge data are not the nearest to the correct geographical location, but have been selected to lie in locations with suitable channels and water depth (as far as this is possible). The locations used for the tide gauges are indicated on the figure. Since the model domain does not include the tide gauge at Halfway Shoal, the nearest internal grid point is used for comparison with the data. The errors arising from these lateral displacements will be only a few centimetres since the sea surface slope is generally small. For example, the typical speed of propagation for a shallow water wave (tide), given by $\sqrt{gD}$, is around 8–9 ms^{-1} for water of 7 m depth and so a displacement by a single grid box leads to a phase error of only 30 s (and around 1 cm RMS height error over the tidal cycle). Some manual adjustment to the bathymetry was made to ensure continuity of deep channels which are barely resolved by the model grid, and the channel between Walney Island and the mainland was also closed. This channel is in any case so shallow that it is effectively closed for most of the tidal cycle and is also extremely narrow, so the error introduced by this procedure will be small.

Model details

Hydrodynamic model

The numerical model is a simplified version of the model described in detail by Flather (1993). The model integrates the equations for conservation of depth-averaged mass and momentum on a regular Cartesian grid.

The sea surface elevation is calculated from the continuity equation which expresses the conservation of volume:

$$\frac{\partial \zeta}{\partial t}+\frac{\partial (Du)}{\partial x}+\frac{\partial (Dv)}{\partial y}=0 \tag{1}$$

In this equation, ζ denotes the sea surface elevation, D is the total water depth and the velocities in the x and y directions are u and v respectively.

Velocities are calculated at interior grid points from the conservation of momentum equations:

$$\frac{\partial u}{\partial t}+u\frac{\partial u}{\partial x}+v\frac{\partial u}{\partial y}+g\frac{\partial \zeta}{\partial x}-fv+\frac{Ku(u^2+v^2)^{1/2}}{D}-A_h\nabla^2 u=0 \tag{2}$$

$$\frac{\partial v}{\partial t}+u\frac{\partial v}{\partial x}+v\frac{\partial v}{\partial y}+g\frac{\partial \zeta}{\partial y}+fu+\frac{Kv(u^2+v^2)^{1/2}}{D}-A_h\nabla^2 v=0 \tag{3}$$

Here f is the Coriolis parameter, g is the acceleration due to gravity, K is the friction coefficient for the quadratic bottom friction term and A_h is the horizontal eddy viscosity. At closed boundary points, the velocity normal to the boundary is zero (wetting and drying is discussed below), and a long the open boundary, the normal velocity is calculated from a radiation condition

$$u=\hat{u}+\sqrt{\frac{g}{D}}(\hat{\zeta}-\zeta) \tag{4}$$

where $\hat{u}$ and z are the prescribed value for velocity and elevation at the boundary. Thus any excessive elevation over the prescribed value is treated as a gravity wave propagating out across the open boundary.

There are no surface forcing terms in this model. The implicit assumption is that the surface forcing has a very minor effect over the small model domain, and that the model variables within the model domain are determined by the open boundary conditions alone.

The model grid is arranged as a regular Arakawa C grid in Cartesian coordinates, with 120 × 140 height points at a spacing of 250 m. The axes are rotated through an angle of 30° anticlockwise from grid North. The u velocities are calculated at the mid-points between pairs of adjacent ζ points in the x direction, and similarly for the v velocities. The numerical discretization follows the explicit method of Flather and Heaps (1975), with the trivial simplifications necessary to adapt their regular latitude-longitude grid to a Cartesian coordinate system. The time step, limited by the CFL criterion, is chosen to be 5 s. The changing water level results in many grid points flooding and drying during the model integration. In the simple method used by Flather and Heaps (1975) to determine which velocity points are wet and which are dry, the velocity points are treated as wet if either the relevant adjacent elevation points have non-zero depth, or if one of these neighbouring elevation points has a non-zero water depth and the sea surface slope across the velocity point exceeds a threshold value (in the direction required to generate flow from the wet to dry elevation point). A minimum water depth is also imposed in the denominator of the friction term in the momentum equation in order to avoid the singularity that could otherwise occur. This simple wetting and drying algorithm has been found to generate a large amount of numerical noise on this model grid which would completely swamp the model output were it not for the substantial horizontal diffusion coefficient of 10^2 m^2 s^{-1}. Various modifications have been suggested which improve the performance of the wetting and drying algorithm, including the use of sub grid-scale ‘ area factors ’ (Flather & Hubbert, 1990) or by the use of ‘ sloping facets ’ (George & Stripling, 1997). It should be noted, however, that the horizontal diffusion term in Equations (2) and (3) is in fact generally lower than the numerical diffusion that would be generated by the use of an implicit integration scheme with no specific additional diffusion term, and the diffusion coefficient is also well below the limit for numerical stability in this numerical scheme. This amount of diffusion does not significantly affect the tidal propagation. For these reasons, the original scheme of Flather and Heaps (1975) described above for wetting and drying is considered to be adequate for this work.

Adjoint model

The adjoint model was derived directly from the discretized model equations via Lagrangian multipliers. This process is entirely mechanical and the resulting equations are lengthy (due to the discretization method on the staggered grid), so they are not repeated here. A fuller explanation of the technique and some examples of the resulting sets of equations are given by Lardner (1993) and Das and Lardner (1991). An alternative approach would have been to take the adjoint of the continuous model equations, and then these could have been discretised in an arbitrary manner. This method could have resulted in much simpler and faster code, but the results of Chertok and Lardner (1996) suggest that this approach can degrade the results. Another computationally significant issue is that, in principle, the complete set of model variables at each grid point and each time step from the forward model run are required in order to correctly integrate the adjoint model backwards in time. In practice, this is not computationally feasible (writing to and reading from a hard disk is too slow, and the data cannot be stored in a reasonable amount of RAM). One approach to this problem would be to adopt the checkpoint method as outlined by Griewank (1992), but this method has the drawback that the forward model run must be performed at least twice per iteration. A faster and commonly used alternative is simply to store the model fields at regular intervals during the forward model run and then interpolate between these values in order to generate the intermediate data values. This introduces a discretization error into the adjoint model which depends on the spacing between the storage of the data points, but this is unlikely to be too harmful if the interval is small relative to the physical scales of evolution of the model. In similar circumstances, Lardner *et al*. (1993) found that an interval of 24 min (4 of their model time steps) between data fields did not degrade the performance of their adjoint significantly. An interval of 15 min (180 time steps) was used in this work, and again this did not appear to degrade either the speed of the optimization or the ultimate result.

The optimization routine used is the conjugate gradient descent method implemented as E04DGF in the NAG library routines. Comparisons between different optimization methods (Navon & Legler, 1987; Das & Lardner, 1991) have suggested that the choice of optimization algorithm can be a significant factor in the efficiency of the assimilation, but only more rarely affects the results and the optimal choice is in any case problem-specific. The number of iterations required for each of the model runs described below was of the order of 25, and on average over 90% of the improvement came within the first five iterations, so there is little opportunity for a superior algorithm to improvement performance significantly. For a 12·5 h model run (9000 time steps) on the 120×140 grid, the forward model took approximately 6 min to run on a 195 MHz SG Indigo2 R10000, and the adjoint took roughly three times as long. The code was written with clarity and simplicity in mind rather than speed and numerical precision, and these figures could probably be improved if required.

The correctness of the adjoint model was checked by comparing the calculated gradient for each control variable with a finite difference approximation to the gradient. This procedure (contained within the routine E04DGF) demonstrated the correctness of the code when run for a small number of time steps. Over longer integrations, the rounding error due to the short time step and single precision arithmetic resulted in some inaccuracy, but this was never sufficient to hinder convergence.

Experiments

The basic experimental procedure was to spin up the model from rest, and then to undertake a set of experiments to compare the effects of assimilating different data sets. The control variables used included both the open boundary conditions over the model run, and the initial conditions (water depth and currents) for the interior of the model domain. The model equations were used as a strong constraint. The cost function which was to be minimized contained a model error term (which penalized the difference between model output and a particular data set) and terms which penalized deviation of the control variables from their initial guesses. These latter terms are required in order for the problem to be well-posed. The weights attached to the various terms should be chosen so as to reflect their relative uncertainty, but tests showed that the results were rather insensitive to the values used. Tide gauge data has an error of at most a few centimetres, and the uncertainty of the boundary conditions and initial model state are at least an order of magnitude larger, so these penalty terms were very small and barely affected the model results.

The open boundary of the model is less than 20 km wide, and initially a uniform boundary condition was applied across it. The tide gauge at Halfway Shoal is very close to the open boundary, and the data from this tide gauge is used for the elevation boundary condition. In order to generate a normal velocity for the radiation boundary condition, a preliminary assimilation was performed, in which the elevation data from Halfway Shoal were assimilated into the model (by comparison with the internal point indicated in Figure 1) using only the normal velocity at the boundary as a control variable. An initial guess for this velocity was required, but using a negligible weight in the cost function ensured that this did not materially affect the final result. The optimal set of boundary conditions from this assimilation run was then used as the initial guess for subsequent runs. In the final experiment, a more complex boundary condition (described below) was used.

The usual method of generating boundary conditions, which would be required without the fortuitous location of the Halfway Shoals tide gauge, is to use a series of nested models to generate the open boundary conditions from (ultimately) a global ocean model. This technique was used for previous modelling work in the same area and it is interesting to compare the accuracy of those results with the work presented here. As part of this previous work a simple assimilation technique was also applied, and these results are also useful for comparison.

Results and discussion

The model is spun up from rest for a period of 22 h starting from 23:15h on the 22 August 1997. The open boundary elevation is close to 0 (and rising) at the start time, and the model run ends at low water. Extending this spin up time does not affect the results.

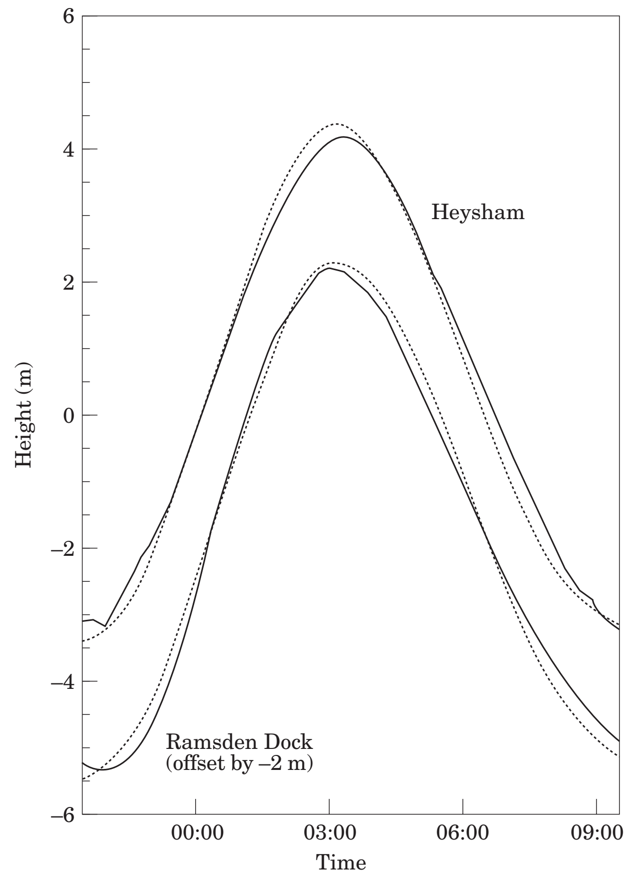

FIGURE 2. Model results with no assimilation. Dotted lines are tide gauge data, solid lines are model results.

TABLE 1. RMS errors for various model runs

Experiment details	RMS error at Heysham (cm)	RMS error at Ramsden Dock (cm)
No assimilation	25	24
Ramsden Dock only	37	4
Ramsden Dock and Heysham	18	17
Linear boundary condition	10	10

The assimilation experiments were performed on the subsequent tidal cycle, for the 12·5 h period starting at 21:15h on 23 August and ending at 09:45h on the 24th. The results of the control run, when no assimilation was performed (other than to generate the boundary conditions as described above), are shown in Figure 2. RMS errors for all the experiments performed are listed in Table 1. The RMS errors at Heysham and Ramsden Dock are 25 and 24 cm respectively, and the mean biases are 5 cm and 3 cm.

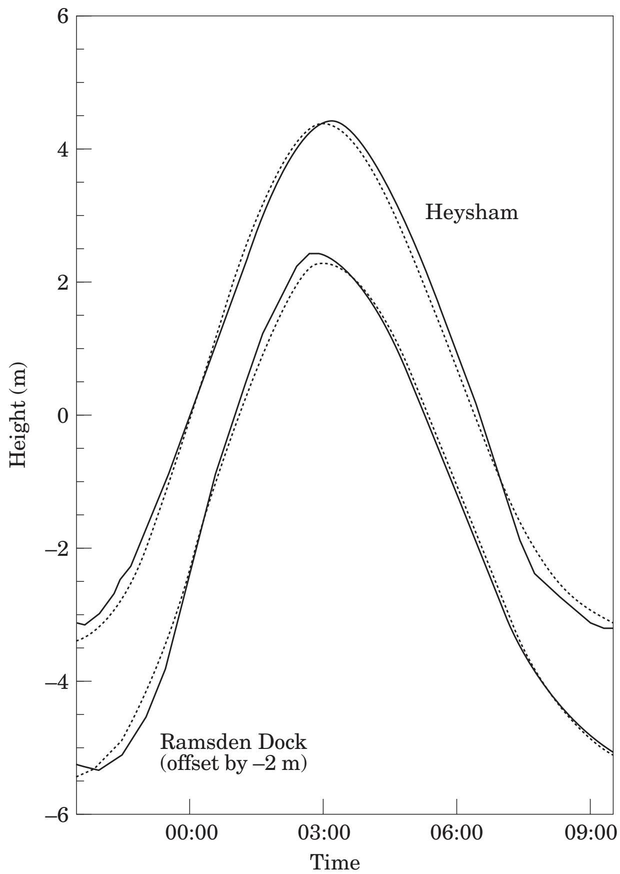

FIGURE 3. Model results with assimilation at Heysham and Ramsden Dock. Dotted lines are tide gauge data, solid lines are model results.

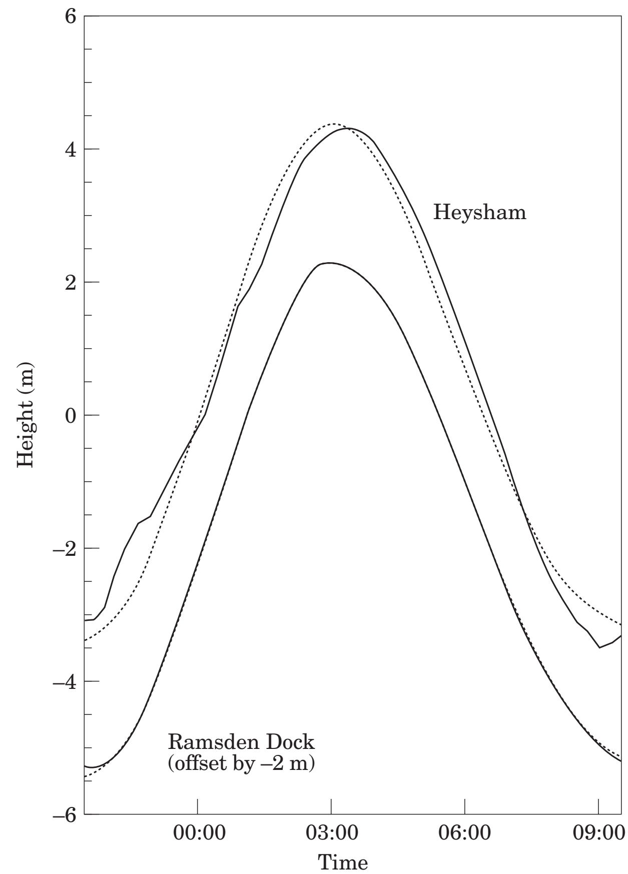

FIGURE 4. Model results with assimilation at Ramsden Dock only. Dotted lines are tide gauge data, solid lines are model results.

A small amount of residual noise from the wetting and drying routine is evident for the time series at Heysham, but a more substantial source of error is due to the different tidal profiles and the phase error at high tide. The tide initially rises too fast at Heysham and too slowly at Ramsden Dock, but at high tide this phase error is reversed.

When data are assimilated from a single tide gauge within the model, using uniform boundary conditions, the optimization generates a model run which matches that particular data set extremely closely. Unfortunately, the accuracy at the other tide gauge decreases. Figure 3 demonstrates the results when data from Ramsden Dock is assimilated. It can be seen that the phase error at Heysham has increased slightly when compared to Figure 2, and the RMS error has increased to 37 cm. The results when assimilation is performed at Heysham only are similar, with a phase shift in the opposite direction at Ramsden Dock and an increase in model error at that location (not shown). These results are not very surprising since it was apparent from the control run that the errors are generally in opposing directions.

When data from both tide gauges are assimilated simultaneously (Figure 4), there is a reasonable improvement in model performance, with residual RMS errors of around 18 cm at both Heysham and Ramsden Dock and very low bias. The same qualitative errors remain, but the magnitudes are reduced. It appears from this that a uniform open boundary condition cannot adequately represent the correct tidal circulation. In reality, there is an anticlockwise gyre in the bay.

In an attempt to allow the model to better represent reality, a more complex boundary condition was introduced. Conditions were not held constant across the open region, but allowed to vary linearly, and the boundary conditions at each end of the open boundary were used as control variables. As an initial guess, both endpoints for the boundary conditions were taken to be identical to the uniform condition previously calculated for the earlier experiments.

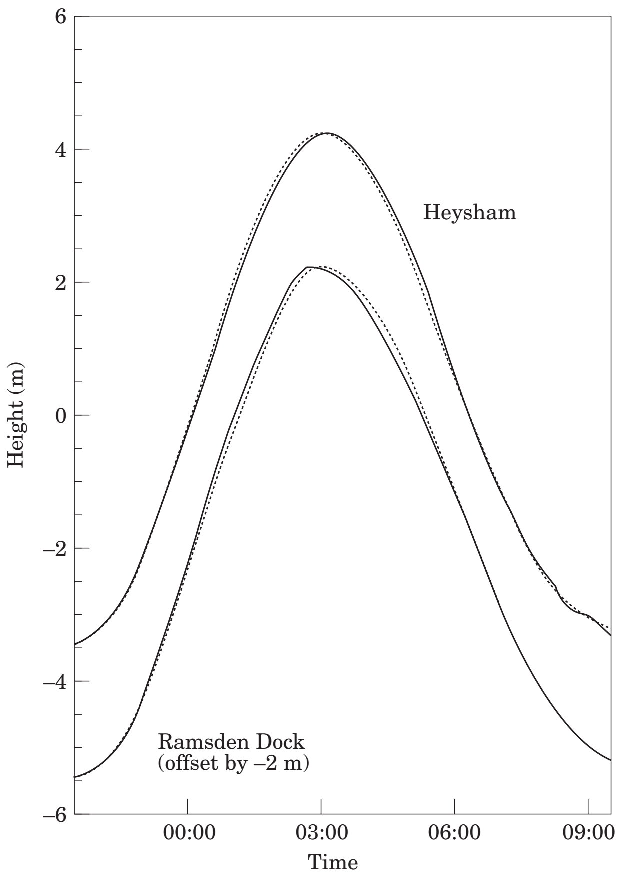

FIGURE 5. Model results with spatially varying boundary condition and assimilation at Heysham and Ramsden Dock. Dotted lines are tide gauge data, solid lines are model results.

The results from this model run (plotted in Figure 5) show an encouraging reduction in RMS model error to 10 cm at both Heysham and Ramsden Dock. The plotted results show that the model does still not quite match the timing of high tide correctly, but it is much closer than for previous model runs with a uniform boundary condition. With only two internal data sets, there is no justification for a more complex boundary condition, since the problem would tend to become underspecified (Das & Lardner, 1992), provide some examples of this). This would not necessarily be the case if the model was driven by a larger scale numerical model which might provide some useful information on the spatial structure at the boundary and therefore could permit an even greater improvement.

These results compares favourably with the results obtained by Mason *et al*. (1998*a*) when a series of nested models was used to generate the boundary conditions for a similar numerical model of this area. In that case, the RMS error was in the region of 40 cm, including a significant bias of 10–30 cm, when no data were assimilated into the model. When boundary conditions were adjusted using a simple non-optimal assimilation method, the total RMS errors were reduced to around 20 cm, consisting of a bias of around 11 cm and a residual RMS of 16 cm.

For the results when assimilation was not performed, the errors at the tide gauges can be expected to be representative of the whole model domain. When these data have been used for assimilation, however, the remaining errors will generally only provide a lower bound for the typical model error, since the model has been specifically fitted to these points. Ideally, some independent data from the interior of the bay should be used to validate the results, and further investigations are planned along these lines. For such a small model grid, however, the instantaneous sea-level variation across the entire model domain is limited. For example, the RMS height difference between the tide gauges at Heysham and Ramsden Dock is only 23 cm, and they are on opposite sides of the bay. It is therefore reasonable to hope that the average height error throughout the bay is closer to the lower limit of 10 cm than to the upper bound of 25 cm. The known inaccuracy of the model bathymetry (in the region of 40 cm height errors) is likely to be one of the limiting factors of model performance in shallow regions.

Conclusions

An adjoint model for a nonlinear depth-averaged two dimensional numerical model of Morecambe Bay, which includes a large wetting and drying area, has been developed and tested. For this particular model grid, the boundary conditions can be partially obtained from a nearby tide gauge and this results in model performance which compares favourably with the alternative system of nested models.

Variational assimilation using the adjoint model greatly improves the accuracy of the model, and substantially outperforms the suboptimal method for adjusting the boundary condition which had been previously used. Errors can be reduced to 10 cm at the location of tide gauges within the model domain. This represents a large improvement over the most accurate modelling previously achieved in this area. An overall height accuracy in the region of 10 cm would appear to be a realistic goal for numerical modelling using this technique.

Acknowledgements

This work was funded under the BNSC Earth Observation LINK Programme. I would like to thank Roger Flather and David Mason for helpful discussions, and also two anonymous referees for their useful comments.

References

Chertok, D. L. & Lardner, R. W. 1996 Variational data assimilation for a nonlinear hydraulic model. *Applied Mathematical Modelling* **20,** 675–682.

Das, S. K. & Lardner, R. W. 1991 On the estimation of parameters of hydraulic models by assimilation of periodic tidal data. *Journal of Geophysical Research* **96,** 15187–15196.

Davenport, I. J., Mason, D. C., Flather, R. A. & Gurney, C. 1996 Foreshore study through shoreline delineation. In *Proceedings of the EUROPTO Conference on Microwave Sensing and Synthetic Aperture Radar, Taormina, Italy* 164–171.

Flather, R. A. 1993 A storm surge prediction model for the Northern Bay of Bengal with application to the cyclone disaster in April 1991. *Journal of Physical Oceanography* **24,** 172–190.

Flather, R. A. & Heaps, N. S. 1975 Tidal computations for Morecambe Bay. *Geophysical Journal of the Royal Astronomical Society* **42,** 482.

Flather, R. A. & Hubbert, K. P. 1990 Tide and surge models for shallow water—Morecambe Bay revisited. In *Modelling Marine Systems* (Davies, A. M. ed.) vol. I, CRC Press Inc, Boca Raton FL pp. 135–166.

George, K. J. & Stripling, S. 1995 Improving the simulation of drying and wetting in a two-dimensional tidal numerical model. *Applied Mathematical Modelling* **19,** 2–6.

Griewank, A. 1992 Achieving logarithmic growth of temporal and spatial complexity in reverse automatic differentiation. *Optimization Methods & Software* **1,** 35–54.

Koopmans, B. N. & Yang, Y. 1994 Satellite radar data for topographic mapping of the tidal flats in the Wadden Sea, The Netherlands. In *Proceedings, Second Thematic Conference on Remote Sensing for Marine and Coastal Environments, New Orleans*.

Lardner, R. W. 1993 Optimal control of open boundary conditions for a numerical tidal model. *Computer Methods in Applied Mechanics and Engineering* **102,** 367–387.

Lardner, R. W., Al-Rabeh, A. H. & Gunay, N. 1993 Optimal estimation of parameters for a two-dimensional hydrodynamical model of the arabian gulf. *Journal of Geophysical Research* **98,** 18229–18242.

Mason, D. C., Davenport, I. J., Robinson, G. J., Flather, R. A. & McCartney, B. S. 1995 Construction of an inter-tidal digital elevation model by the ' water-line ' method. *Geophysical Research Letters* **22,** 3187–3190.

Mason, D. C., Davenport, I. J., Flather, R. A. & Gurney, C. 1998 A digital elevation model of the inter-tidal areas of the Wash, England, produced by the waterline method. *International Journal of Remote Sensing* **19,** 1455–1460.

Mason, D. C., Amin, M., Davenport, I. J., Flather, R. A., Robinson, G. J. & Smith, J. A. 1999 Measurement of recent intertidal sediment transport in Morecambe Bay using the waterline method. *Estuarine, Coastal and Shelf Science* **49,** 427–456.

Navon, I. M. & Legler, D. M. 1987 Conjugate-gradient methods for large-scale minimization in meteorology. *Monthly Weather Review* **115,** 1479–1502.

Spitz, Y. H. & Klink, J. M. 1998 Estimate of bottom and surface stress during a spring-neap tidal cycle by dynamical assimilation of tide gauge observations in the Chesapeake Bay. *Journal of Geophysical Research* **103,** 12761–12782.