

Data Assimilation Tests Using NISE10 Storm Surge Model

**Jane A. Williams
Chris Wilson
Kevin J. Horsburgh**

May 2008



[page deliberately blank]

CONTENTS

Executive Summary	3
Introduction	4
Method	5
Results and Discussion	6
Conclusions	26
Acknowledgements	26
References	26
Appendix 1 – Interim results on 3D var method	28

Executive summary

Data assimilation is a process that combines observed data into a numerical forecast model in a dynamically consistent way. It can correct for errors at model boundaries and account for processes not included in the model. Thus, surge model forecasts may be improved if their initial state is optimized. However, they may still deviate from observations if the subsequent (meteorological) forcing is inaccurate. The length of time for which data assimilation will improve a numerical forecast is a key practical consideration when implementing any scheme operationally.

As part of the STFS Development Plan (2005-2006), POL was commissioned to explore the usefulness of a simple boundary correction method proposed by Flather (1993). In this approach, observed residuals derived from the tide gauge at Aberdeen are used to correct the open boundary surge input to the model whose domain is truncated along a line of latitude at Aberdeen. Such direct substitution of observed values is the crudest, but simplest, method of data assimilation. There is a risk of introducing computational instability since the corrected value will no longer be consistent with values at neighbouring grid points.

We found that the introduction of a modified residual wave at Aberdeen tide gauge improves the statistical performance at that location: however, there is no systematic improvement in performance at any more southerly locations along the east coast of the UK. Indeed, in many of the diagrams presented here the output from the standard numerical model run and that with data assimilated at Aberdeen is indistinguishable. In some runs, improved performance of the model with data assimilation over the first 12 hours of the forecast is at the expense of subsequent over-prediction of the surge at many locations.

Based on the absence of any systematic improvements, and the absence of any dynamic basis for the method, it is not recommended that the boundary-correction approach be followed. A further reason to reject the method is the need to truncate model domains such that the lateral boundaries coincide with a particular tide gauge. Restricting the domain of the 3.5 km resolution model to the latitude of Aberdeen is ill-advised both on scientific grounds (since neither surges nor tides can propagate freely around Scotland) and on organizational grounds (since SEPA is member of STFS).

A more promising approach is the use of the so-called "3D Var method" (based on variational analysis). This involves the use of sea level data throughout the entire model domain and the subsequent minimisation of a cost function which compares information in the numerical model with said data. Initial work at POL on using the method in the surge model has been promising. Progress has been made both on the mathematical methods required to evaluate the cost function and on computational issues. Our results to date are summarised in Appendix 1. The next stages of this work involve implementing the scheme in a practical manner within the surge model to emulate an operational solution. This additional work will be jointly funded by POL and the Environment Agency and will resume during 2008.

Introduction

NISE10 is a two dimensional tide-surge model covering the **North and Irish Seas** and **English Channel** with a resolution of approximately 3.5km (**10** times that of the operational shelf-wide model, CS3X). The model grid is shown in Figure 1.




Figure 1: NISE10 model grid showing location of Aberdeen.

This model has a northern open boundary which is coincident with the tide gauge at Aberdeen. This facilitates the assimilation of observed data from the tide gauge into the model if desired. The model is forced at the boundaries by 26 tidal constituents, and surge elevations and currents interpolated from a lower resolution, larger scale model (in this case CS3 which has resolution of ~12km). The model is forced by 10m u and v components of wind and mean sea level pressures from the Met Office's mesoscale atmospheric model.

This report describes an experiment to investigate the effects of assimilating observed data from the tide gauge at Aberdeen on the accuracy of the surges produced by this model using two east coast surge events as case studies.

Method

For our data assimilation tests, we use the simple boundary correction method described by Flather (1989; 1993). This method involves using observed residuals derived from the tide gauge at Aberdeen (i.e. observed water level – harmonically predicted tide) to correct the open boundary surge input to the model. The error, calculated at hourly intervals from the difference between the model and the observed surge, is assumed to decay linearly eastwards along the open boundary to zero at a point approximately midway along the boundary. When the assimilation has completed (i.e. when there are no more observations), the final error is assumed to remain constant for the remainder of the run. The open boundary surge current input to the model is similarly corrected by assuming geostrophic balance.

We have chosen two east coast surge events:

Case 1: 17th November 2004 (start of run 06z 17/11/04)
Case 2: 19th January 2005 (start of run 18z 19/01/05)

For each of these cases, we have set up the model to be run in an operational style. Therefore each run comprises a 6 hour ‘hindcast’ and a 48 (+1) hour ‘forecast’. The start times of the runs have been chosen so that the peak surge occurs in the period $t+12$ to $t+18$ of the forecast. This would provide a realistic forecasting situation i.e. where the surge occurs at a point into the forecast that would give a suitable lead-time for flood warnings, and also close enough to the start of the forecast where data assimilation could possibly be effective.

Firstly, open boundary surge data (elevations and currents) for these two events were extracted from the operational re-run archive of CS3. They were then interpolated to the NISE10 grid. Then mesoscale atmospheric model data were extracted for the same periods. Initial conditions for the model runs were also required as warm starts were essential, as the spin-up for a cold start would take up to three days and hence the forecast would be unreliable. Arrays of NISE10 which had been previously run (without assimilation) were obtained to make this possible. To enable comparisons, we also extracted surge elevations for these events from these arrays for locations down the east coast of Great Britain. For validation, observed residual elevations derived from tide gauges at these locations for both events were compiled. (As part of the monthly model validation, observed residual data are extracted from each of the A-class tide gauges, see e.g. Williams and Horsburgh, 2006.)

We then used the observed surges derived from Aberdeen for assimilation into the hindcast part of the model run. The open boundary surge input was corrected as previously described and the model was run. Model residuals were calculated by subtracting the tide only run from the tide + surge run. Time series were extracted at east coast locations and were plotted with observations and the standard NISE10 residuals.

Results and Discussion

First we consider Case 1. For the first run we assimilated observed data from 06z 17/11/04 to 12z 17/11/04, i.e. 7 hours, from $t-6$ to $t+0$. Time series of surge elevations from the model run are shown in Figure 2 along with observations and output from the standard NISE10 model.




Figure 2: NISE10 (with assimilation of 7 hours of standard residuals at Aberdeen) v standard NISE10 (no assimilation) v observed residuals.




Figure 2: NISE10 (with assimilation of 7 hours of standard residuals at Aberdeen) v standard NISE10 (no assimilation) v observed residuals.

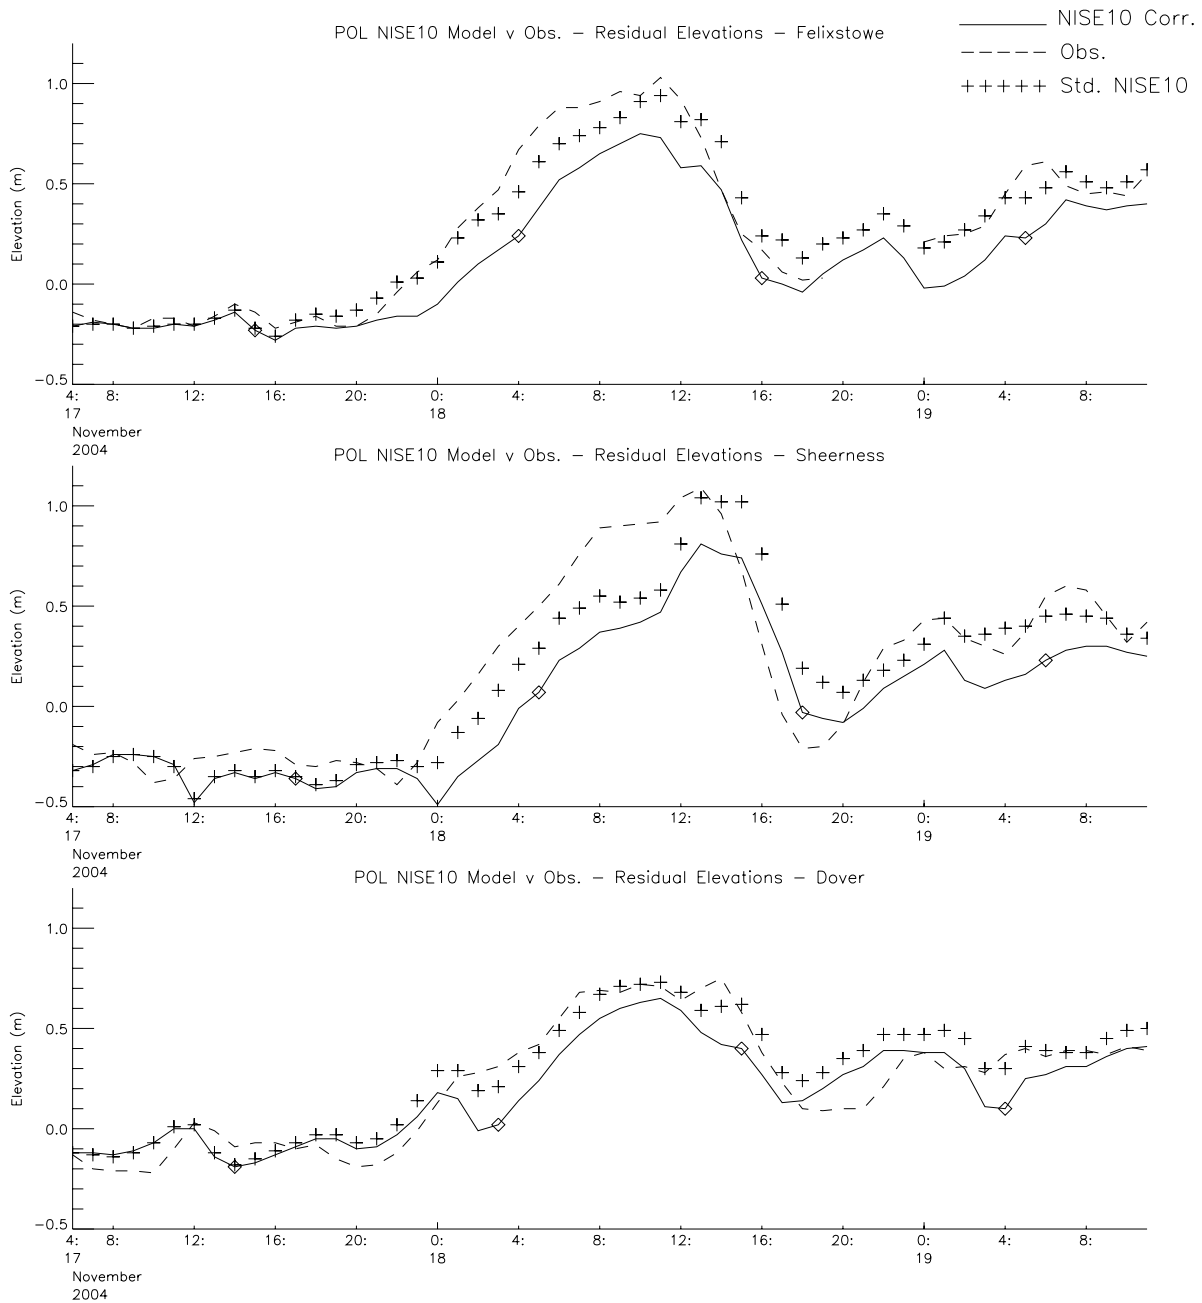


Figure 2: NISE10 (with assimilation of 7 hours of standard residuals at Aberdeen) v standard NISE10 (no assimilation) v observed residuals.

Figure 2 shows a considerable offset between NISE10 with assimilation and the standard NISE10 run and observations which is due to seasonal variations in the harmonically predicted tide. Due to prevailing weather systems, locations in Scotland have a significant seasonal signal which is consistent with the distribution and magnitude of the seasonal signals in the tides (Sa and Ssa). Previous investigations have also attempted to address this problem of seasonal variations (Flather et al., 2000).

This is manifested in a number of ways, the most significant being in the tidal predictions used to derive the observed surge. A harmonically predicted tide is based on the harmonic analysis of a number of years of observations from a tide gauge. The observations include meteorological effects as well as others (e.g. waves, density, etc.). The harmonic analysis includes long-period terms which capture these signals and they then form part of the predicted tide. In fact these signals are not really tidal, but meteorological in origin, so there is double accounting to some extent. The model tide however does not include any seasonal signals.

It is essential for NISE10 to have an observed ‘surge’ at the boundary that has been derived as closely to the model surge as possible for assimilation to be successful. Any included seasonal effects specific to e.g. Aberdeen will be treated as an error, and an inappropriate correction will be propagated through the model resulting in overcorrection to locations further south. A set of observations at Aberdeen was especially provided for us derived using a harmonically predicted tide without the long-period components. The previous model run was repeated and time series of the residuals plotted (Figure 3). It can be seen from Figure 3 that the offset has reduced. The mean error for this run and the previous run for each port are shown in Table 1.

Port	Mean Error (Run 1)	Mean Error (Run 2)
Aberdeen	-0.14	-0.04
North Shields	n/a	n/a
Whitby	-0.22	-0.14
Immingham	-0.21	-0.13
Cromer	-0.13	-0.06
Lowestoft	-0.18	-0.11
Felixstowe	-0.15	-0.09
Sheerness	-0.17	-0.11
Dover	-0.04	0

Table 1: Mean errors in surge (NISE10 – Observations) for Run 1 (assimilating standard observed surges at Aberdeen) and Run 2 (surges derived from tide without seasonal components).

Table 1 (and Figure 3) shows that the offset between model and observations is much reduced (by up to 10cm) by using observations derived from a harmonically predicted tide that does not include the long-period constituents. Therefore all further runs will be done using these observations at Aberdeen.

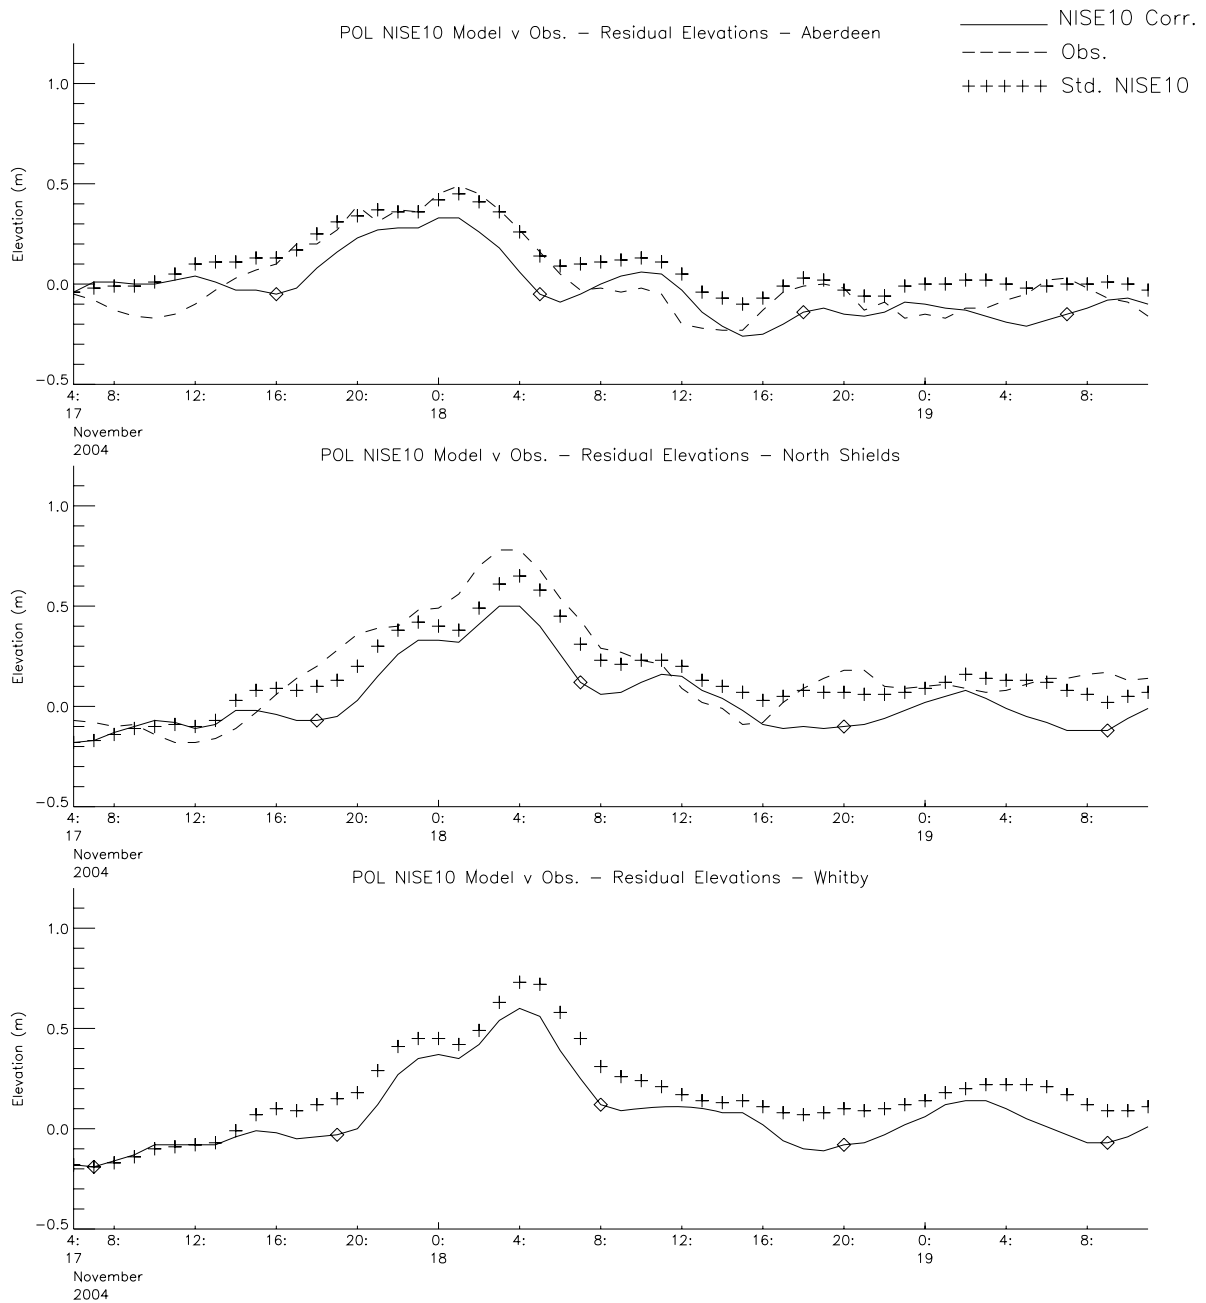


Figure 3: NISE10 (with assimilation of 7 hours of residuals at Aberdeen with no long-period tidal constituents) v standard NISE10 (no assimilation) v observed residuals.

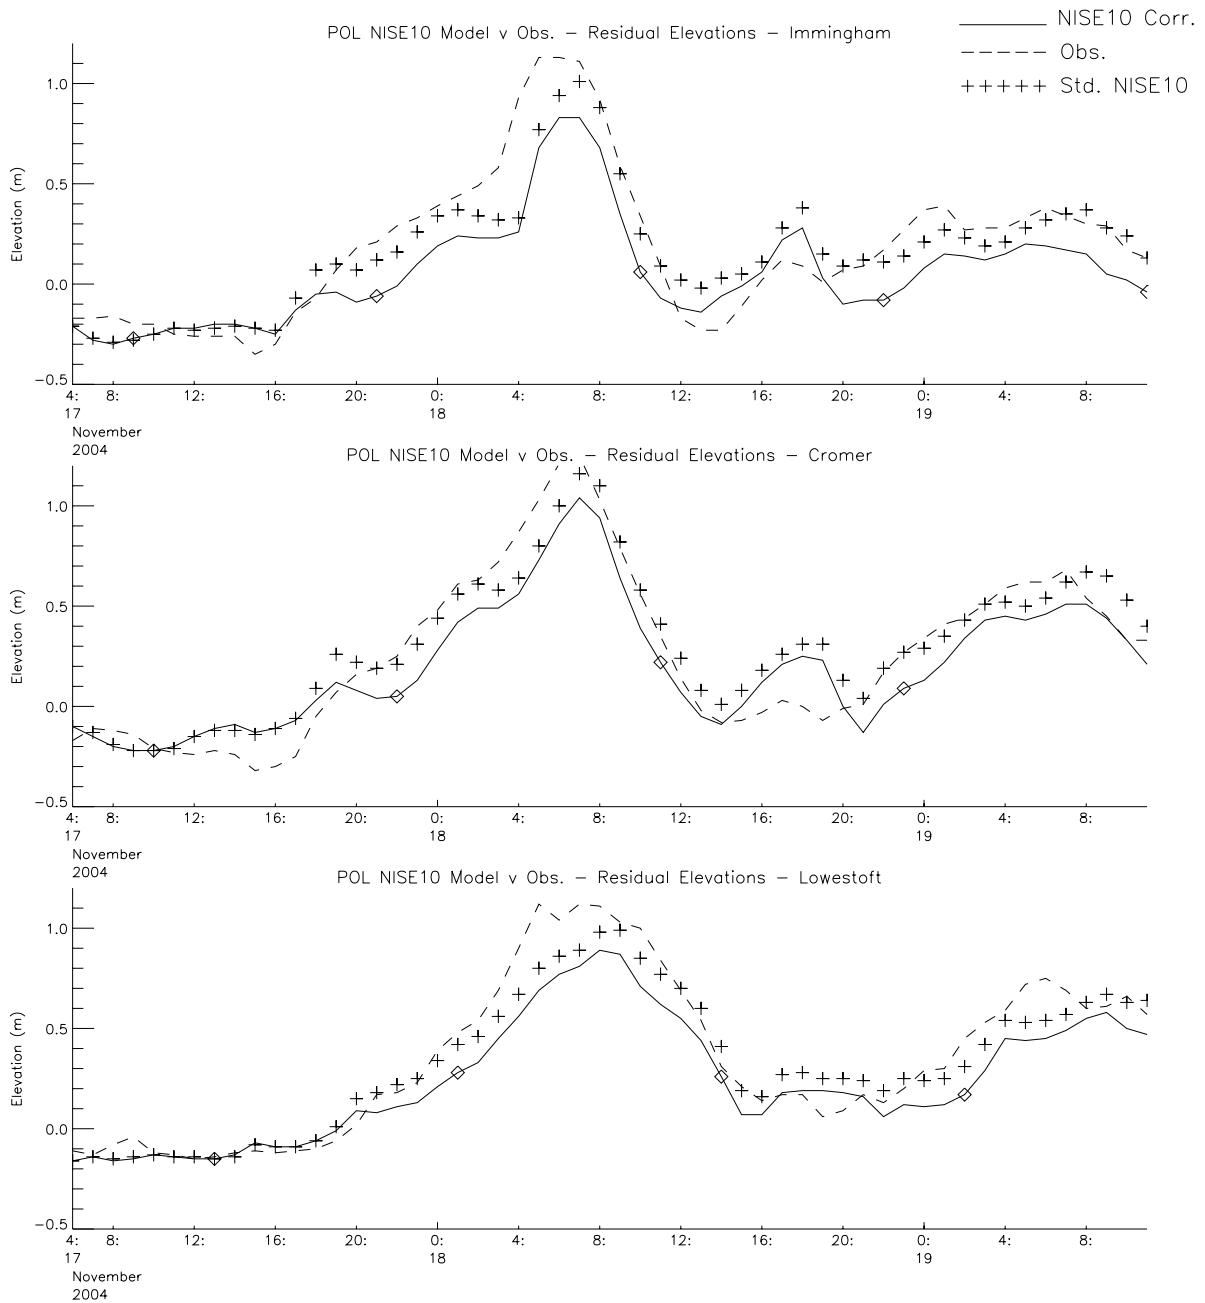


Figure 3: NISE10 (with assimilation of 7 hours of residuals at Aberdeen with no long-period tidal constituents) v standard NISE10 (no assimilation) v observed residuals.




Figure 3: NISE10 (with assimilation of 7 hours of residuals at Aberdeen with no long-period tidal constituents) v standard NISE10 (no assimilation) v observed residuals.

Next we looked at how sensitive the assimilation method was to the number of hours observations assimilated and how the last error value would influence the forecast, as this value is assumed constant for the remainder of the forecast part of the run. We re-ran Case 1 this time assimilating data from $t-0$ to $t+6$ i.e. 13 hours of observations. This is unrealistic in an operational situation but will demonstrate how assimilating observations closer to the event may change the forecast results. The output time series are shown in Figure 4.

The model time series in Figure 3 show smaller surge elevations for the run with assimilation compared to the standard NISE10 run. When comparing the results, we disregard Sheerness and Dover because there are known problems with the propagation of the surge in the model which will be the subject of further investigation. Comparing Figure 4 and Figure 3, it is interesting to see the difference in the forecast residuals due to including more observations in the assimilation process. This shows that the model is very sensitive to the amount of data assimilated (i.e. the final observed value assimilated) and can give very different results depending on how many hours of observations are available for assimilation. Obviously the closer to the event those observations are available; the more effective the assimilation will be within the model. Figure 4 gives surges much closer to, and in some cases larger than the observations.

Considering this we perform two final experiments. Within the operational schedule of the Met Office's supercomputer, the surge model is one of the last processes to be run. So, in reality the 12z run will actually be run at about 1530, which means that there will be tide gauge observations available up to 1500 i.e. $t+2$. In a typical model run with 6 hours hindcast this will make the total number of observations to be assimilated up to 9 hours. The next run (Run 5) will assimilate 9 hours observations in the standard way. Run 6 will attempt to override the sensitivity of the assimilation method to the final error by applying the mean error over the 9 hours at the point of the final data assimilation i.e. $t+3$ onwards. The corresponding time series from Runs 5 and 6 can be seen in Figures 5 and 6 respectively. It can be seen in figures 5 and 6 that the surge elevations produced by these runs are much closer to the standard NISE10 values.

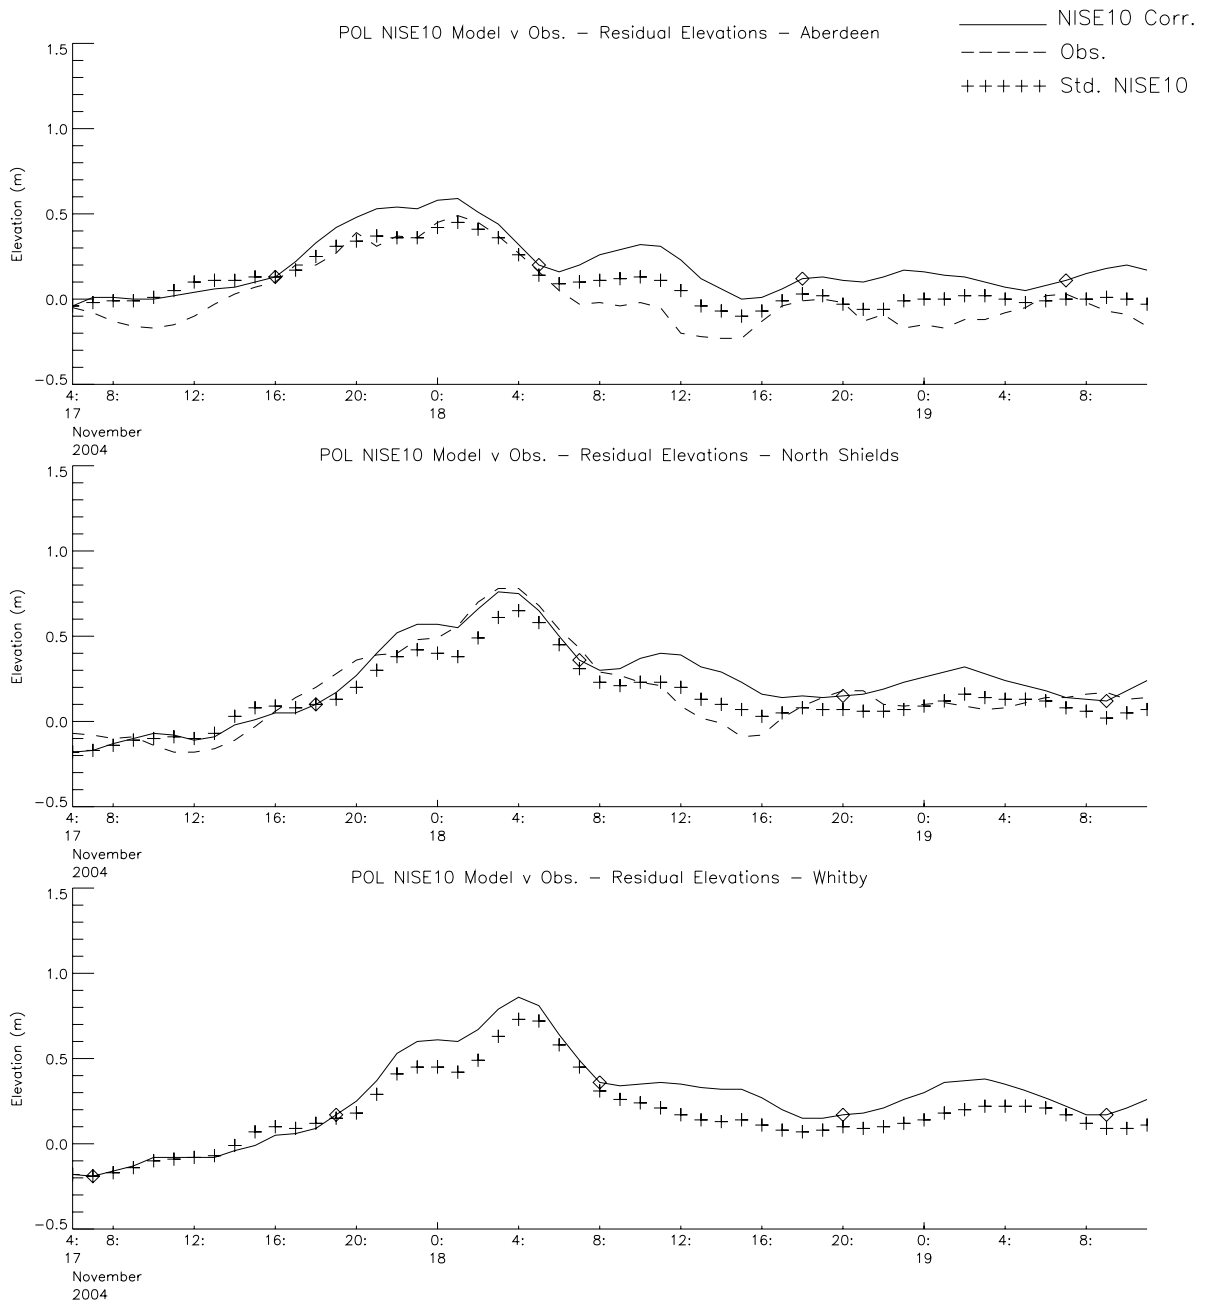


Figure 4: NISE10 (with assimilation of 13 hours of residuals at Aberdeen) v standard NISE10 (no assimilation) v observed residuals.

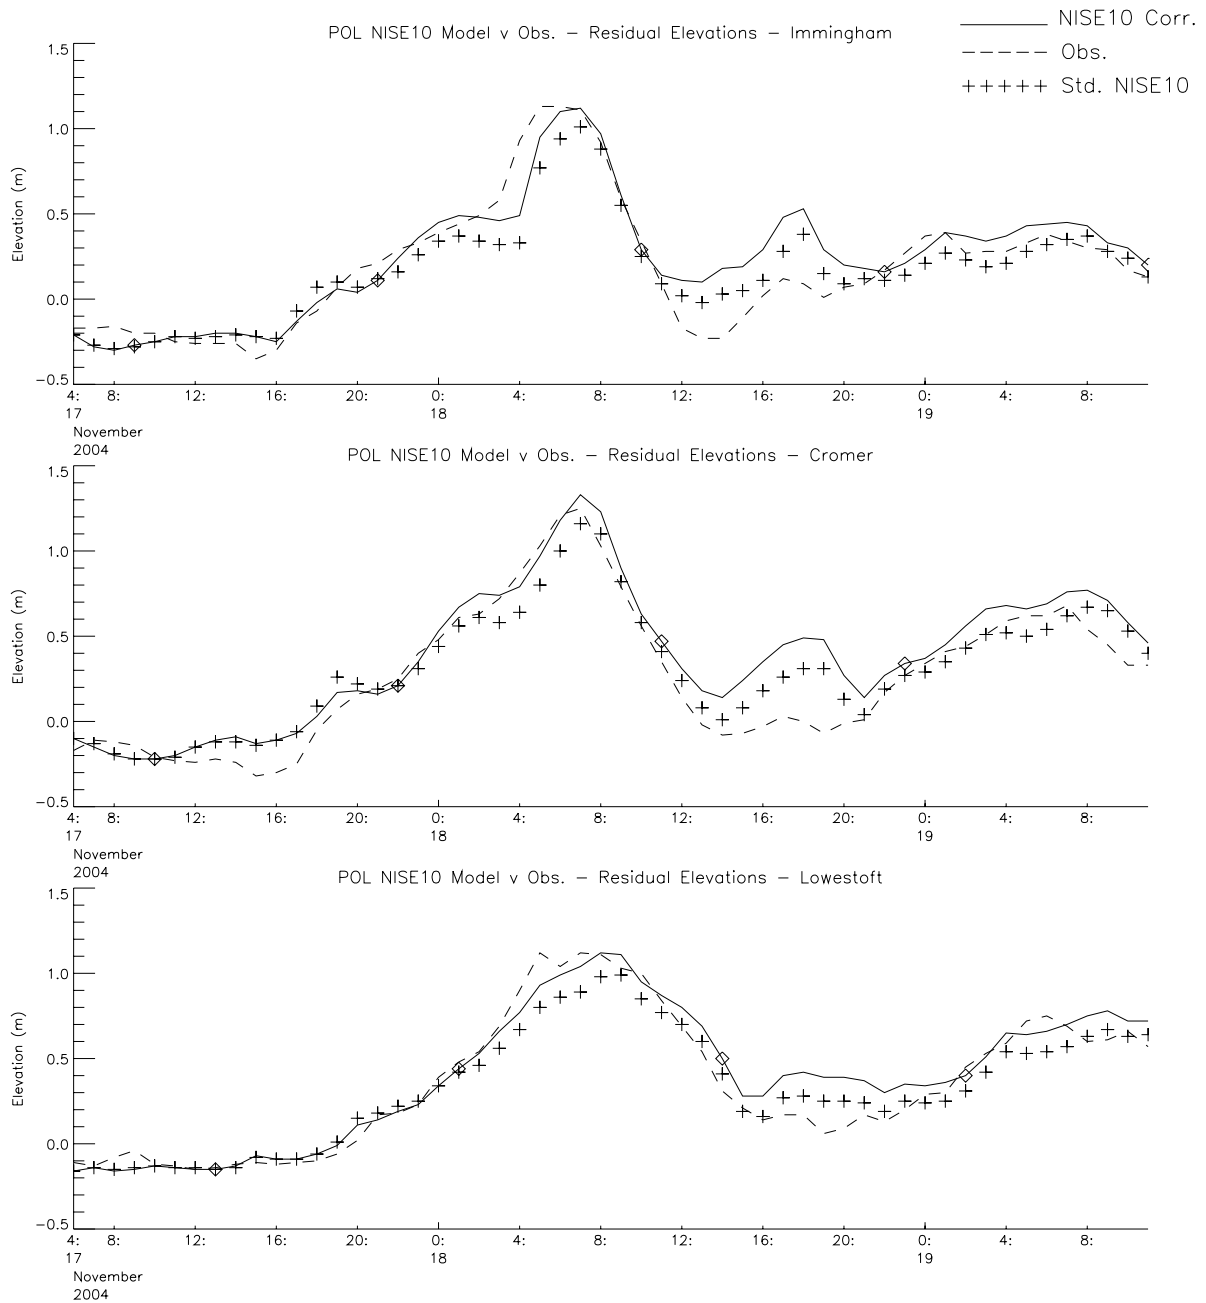


Figure 4: NISE10 (with assimilation of 13 hours of residuals at Aberdeen) v standard NISE10 (no assimilation) v observed residuals.




Figure 4: NISE10 (with assimilation of 13 hours of residuals at Aberdeen) v standard NISE10 (no assimilation) v observed residuals.

Apart from general offsets which may be seen in time series, there are two values of particular interest in surge events. The first is the maximum observed surge (peak) and the second is the surge height at high water (the most likely scenario for coastal flooding). In this case we have used the hour closest to model HW. Table 2 summarises the model runs at the hour of maximum surge for Case 1 and Table 3 the surge elevation at the hour closest to model high water. (There were no observations available from the tide gauge at Whitby for this period.)

	Obs	NISE10	Run 2	Run 3	Run 4	Run 5
Aberdeen	0.49	0.45	0.33	0.59	0.45	0.42
North Shields	0.78	0.65	0.50	0.75	0.62	0.59
Immingham	1.13	0.94	0.83	1.10	0.95	0.92
Cromer	1.25	1.16	1.04	1.33	1.17	1.14
Lowestoft	1.12	0.89	0.81	1.04	0.92	0.89
Felixstowe	1.03	0.94	0.82	1.04	0.94	0.90
Sheerness	1.09	1.04	0.92	1.20	1.05	1.01
Dover	0.72	0.61	0.69	0.84	0.76	0.74

Table 2: Surge elevations (given in metres) at closest hour to peak observed surge from experimental model runs for Case 1, observations and standard NISE10 results.

	Obs	NISE10	Run 2	Run 3	Run 4	Run 5
Aberdeen	0.16	0.14	-0.05	0.20	0.07	0.04
North Shields	0.43	0.31	0.12	0.36	0.24	0.21
Immingham	0.34	0.25	0.06	0.29	0.17	0.14
Cromer	0.35	0.41	0.22	0.47	0.33	0.30
Lowestoft	0.31	0.41	0.26	0.50	0.37	0.34
Felixstowe	0.17	0.24	0.10	0.30	0.20	0.17
Sheerness	-0.21	0.19	0.05	0.27	0.14	0.11
Dover	0.58	0.62	0.47	0.66	0.56	0.54

Table 3: Surge elevations (given in metres) at closest hour to model HW from experimental model runs for Case 1, observations and standard NISE10 results.

From Tables 2 it can be seen that Run 3 re-produces the best peak surge. Run 3 also produced the closest surge to observations at model HW at North Shields and Immingham, however it over-predicted elsewhere. Run 3 demonstrates the importance of having a many observations as possible for the assimilation. If we disregard Run 3 as having an unrealistic number of observations then Run 4 is generally better at recreating the surge peak, however at HW there is no single best solution. Standard NISE10 is better at North Shields and Immingham, Run 4 is better at Cromer, and Run 5 better at Lowestoft and Felixstowe.




Figure 5: NISE10 (with assimilation of 9 hours of residuals at Aberdeen) v standard NISE10 (no assimilation) v observed residuals




Figure 5: NISE10 (with assimilation of 9 hours of residuals at Aberdeen) v standard NISE10 (no assimilation) v observed residuals

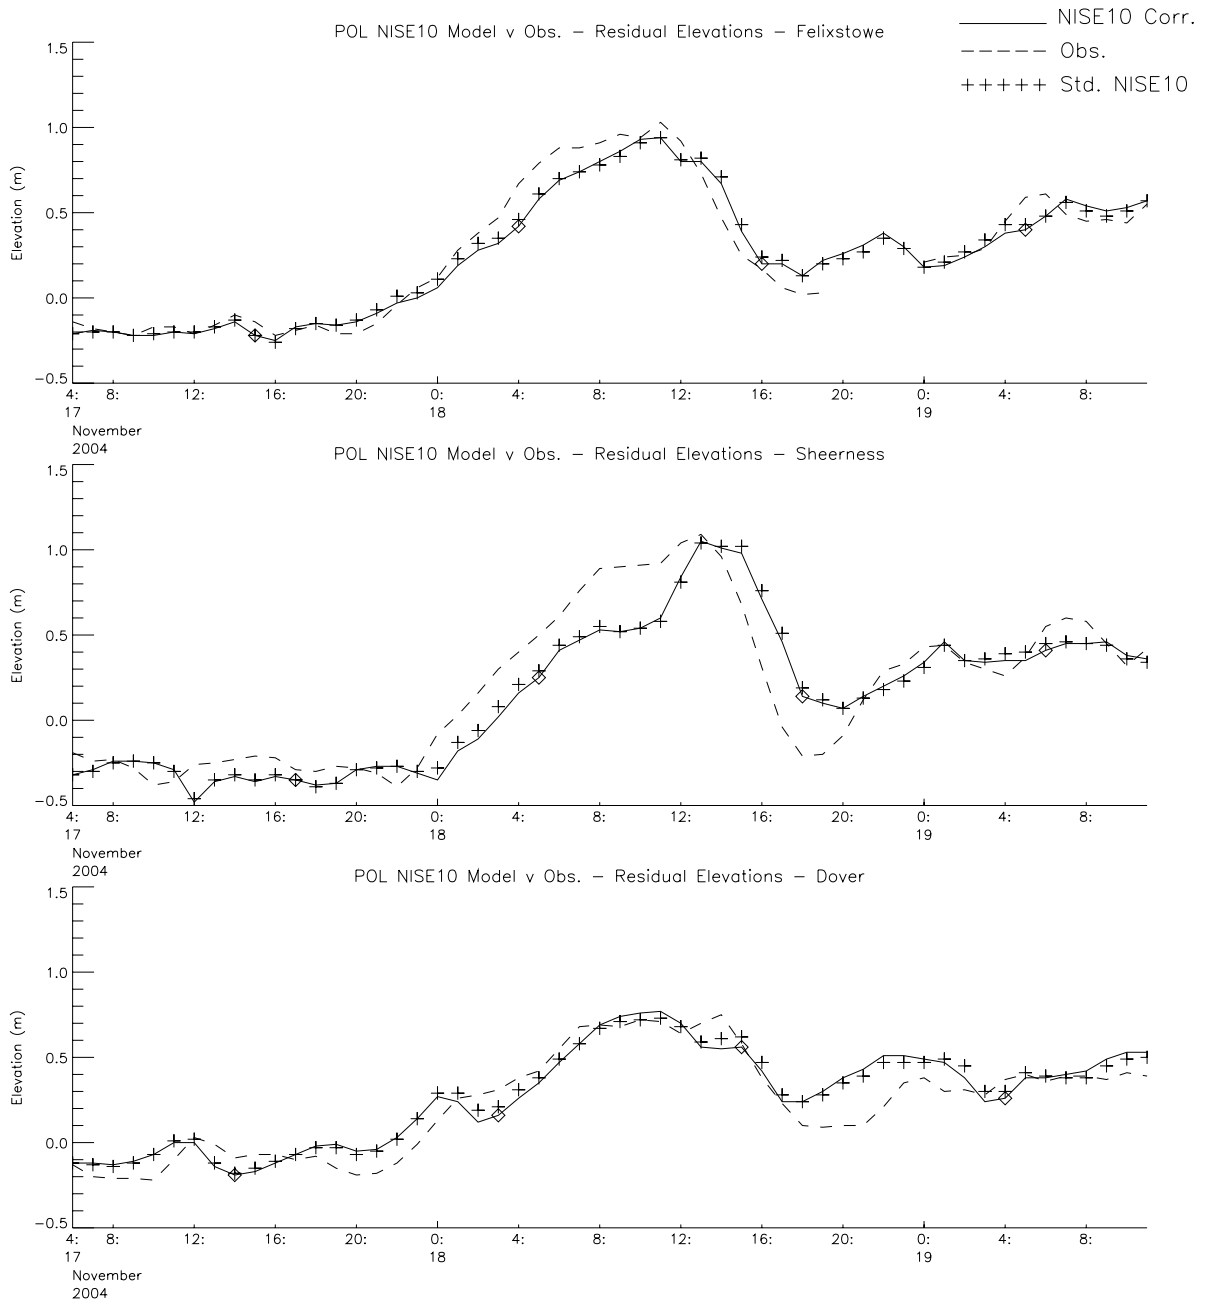


Figure 5: NISE10 (with assimilation of 9 hours of residuals at Aberdeen) v standard NISE10 (no assimilation) v observed residuals

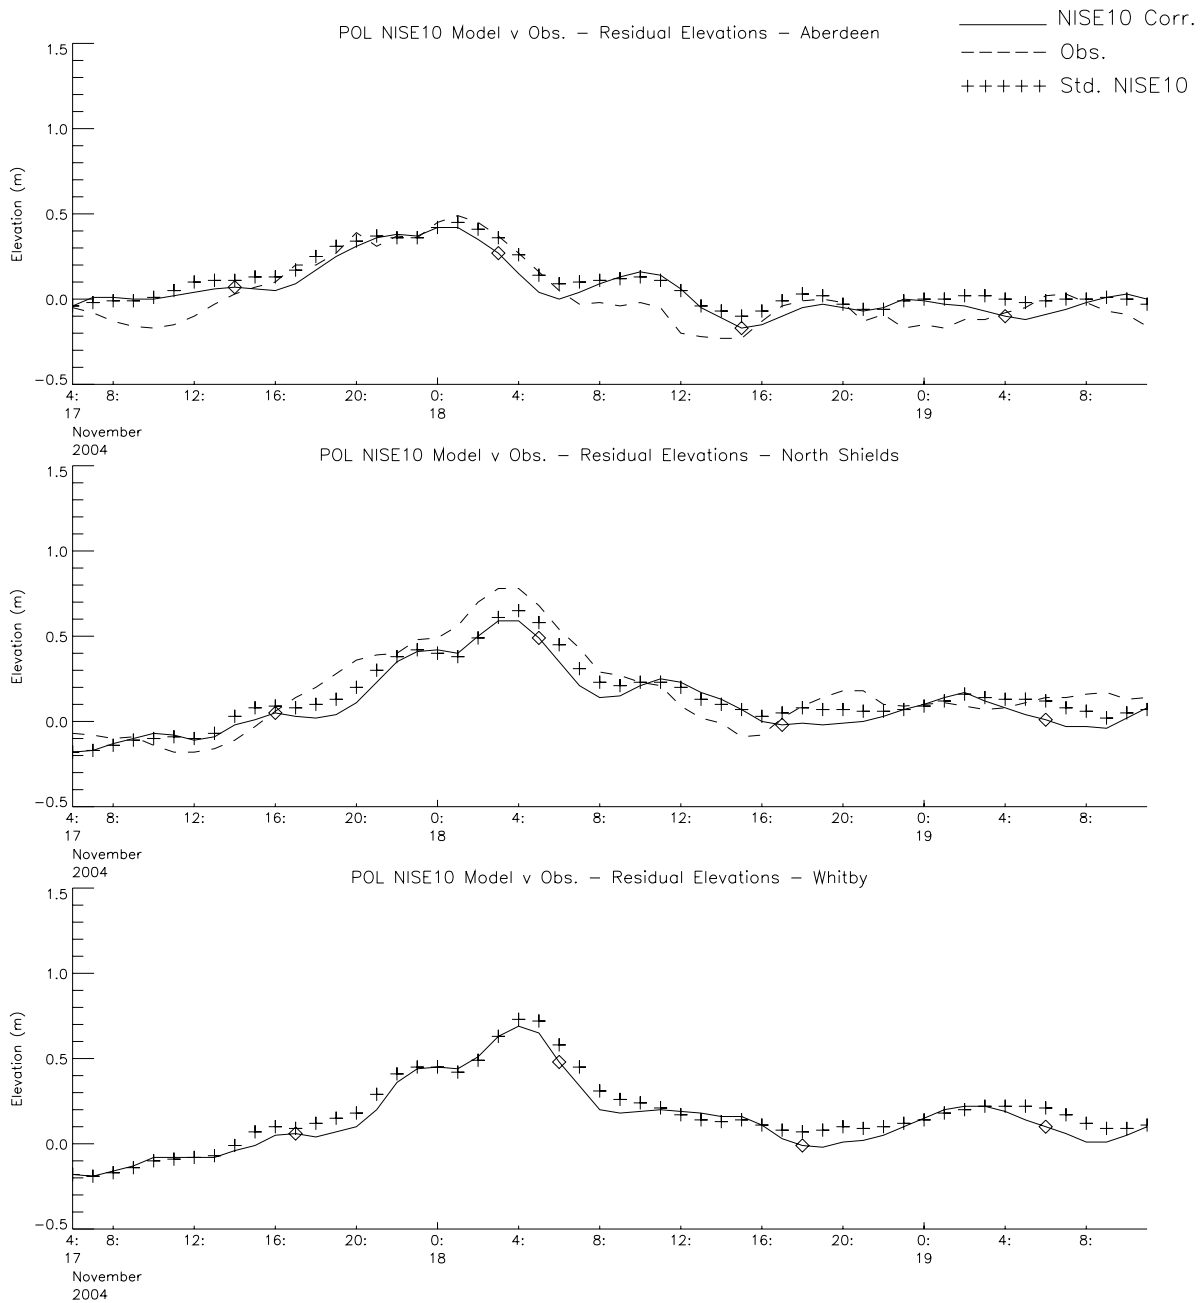


Figure 6: NISE10 (with assimilation of 9 hours of residuals at Aberdeen and mean correction applied at t+2) v standard NISE10 (no assimilation) v observed residuals

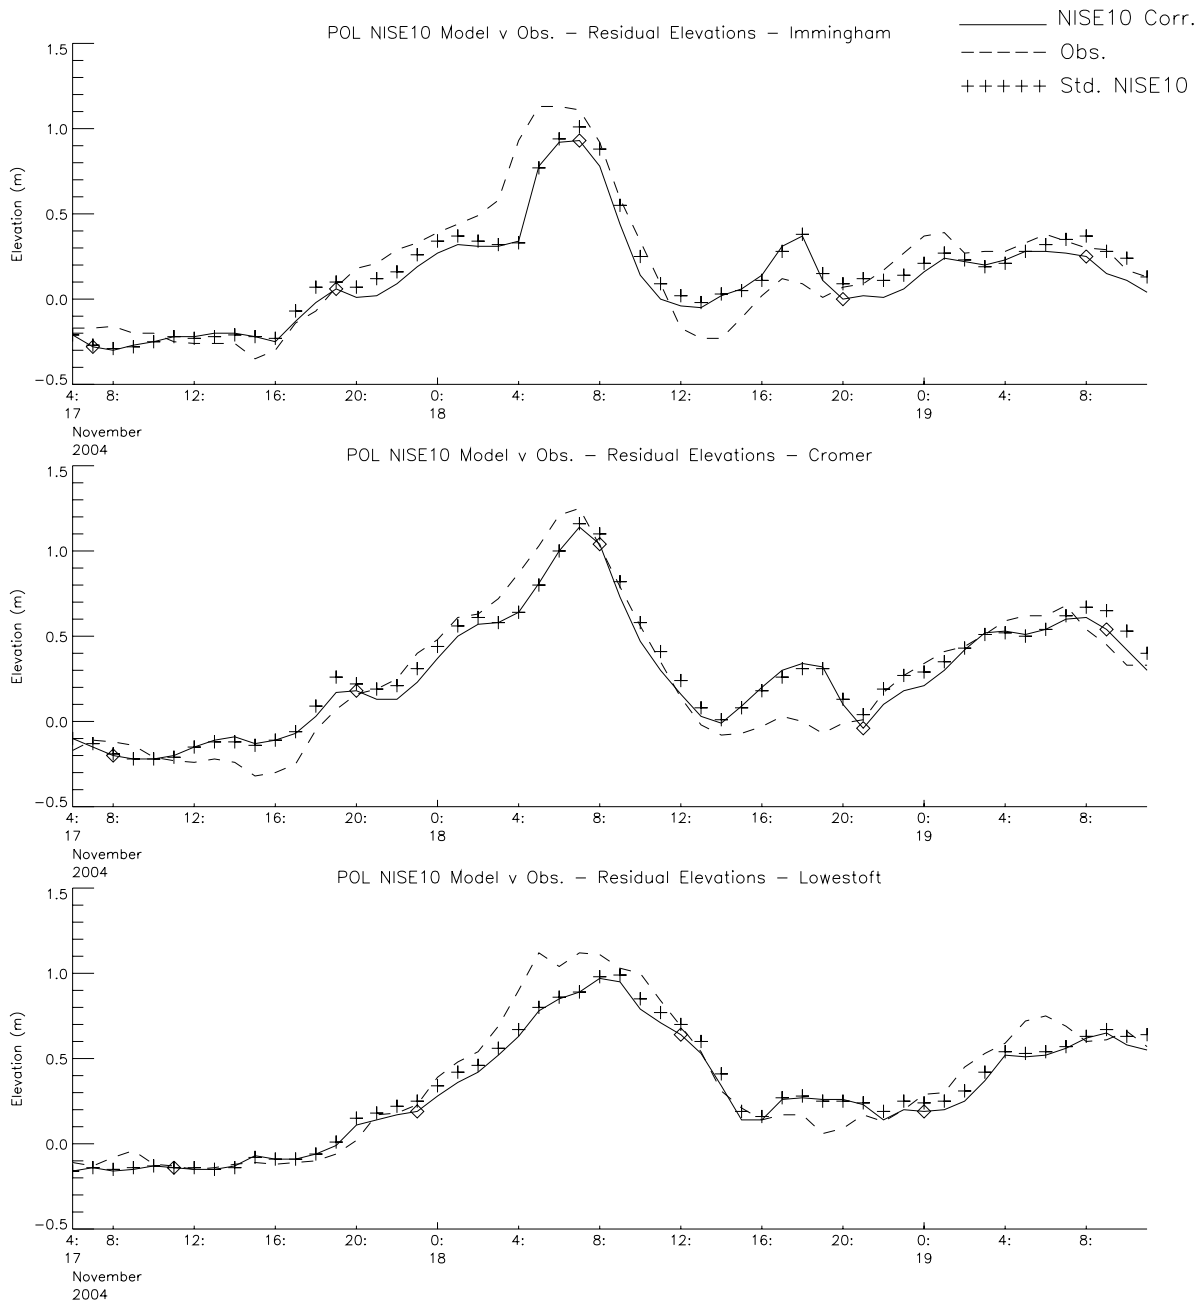


Figure 6: NISE10 (with assimilation of 9 hours of residuals at Aberdeen and mean correction applied at t+2) v standard NISE10 (no assimilation) v observed residuals.




Figure 6: NISE10 (with assimilation of 9 hours of residuals at Aberdeen and mean correction applied at t+2) v standard NISE10 (no assimilation) v observed residuals.

Next we consider Case 2. We have performed three runs of NISE10 based on the results from Case 1. Run 1 assimilates observed data at Aberdeen for the period of the hindcast (c.f. Case 1, Run 2). This is a 7-hour period from 18z 19/01/05 to 00z 20/01/05 i.e. t-6 to t+0. Run 2 assimilates 9 hours of observations from 18z 19/01/05 to 02z 20/01/05 i.e. t-6 to t+2 (c.f. Case1, Run 4) and finally Run 3 assimilates the same 9 hours of observations but in this case applying the mean error throughout the forecast (c.f. Case 1, Run 5). Tables have been compiled for comparison. Table 4 shows the model and observed values at east coast ports at the hour closest to maximum surge. Table 5 shows the model and observed values for each run at the hour closest to model HW.

PORT	Obs	NISE10	Run 1	Run 2	Run 3
Aberdeen	0.77	0.82	0.75	0.82	0.76
North Shields	1.05	1.02	0.96	1.03	0.98
Whitby	1.20	1.11	1.06	1.13	1.08
Immingham	1.30	1.14	1.07	1.16	1.09
Cromer	1.32	1.44	1.38	1.48	1.40
Lowestoft	1.29	1.29	1.23	1.31	1.26
Felixstowe	1.22	1.11	1.05	1.14	1.08
Sheerness	1.16	0.91	0.84	0.95	0.88
Dover	1.07	1.07	1.05	1.09	1.05

Table 4: Comparison of surge elevations at hour closest to time of maximum surge for Case 2.

PORT	Obs	NISE10	Run 1	Run 2	Run 3
Aberdeen	0.36	0.43	0.34	0.41	0.36
North Shields	0.54	0.37	0.28	0.34	0.30
Whitby	0.61	0.48	0.39	0.45	0.41
Immingham	0.58	0.51	0.42	0.48	0.44
Cromer	0.65	0.65	0.57	0.63	0.58
Lowestoft	0.70	0.75	0.68	0.74	0.69
Felixtowe	0.44	0.63	0.56	0.62	0.58
Sheerness	0.22	0.54	0.46	0.52	0.48
Dover	0.68	0.76	0.69	0.75	0.71

Table 5: Comparisons of model and observed surge elevations at hour closest to model high water.

Table 4 shows that Run 2 is generally the best for reproducing the peak surge at most of the locations. However, this is not the case at model HW. In this case standard NISE10 is closest but only by a few centimeters in each case. The difference between Run 1 and Run 2 is only the number of hours of observed data to assimilate and the differences in the results are significant. The results from the

final run in each of the cases where an alternative assumption is made after all the observations have been assimilated do not demonstrate an improvement on the standard assumption.

Conclusions

We have run some data assimilation experiments using two east coast surge case studies. We have used observations from the tide gauge at Aberdeen to correct the open boundary surge input into the model during the hindcast part of an operational-style run. We have also investigated the sensitivity of the model to the number of observations used for the assimilation and also tried an alternative assumption regarding the final error value applied to the forecast.

- For the events studied here only small enhancements were seen by assimilating data.
- The standard assumption regarding application of the final error to the forecast in the assimilation technique works better than the alternative considered here.
- Model results generally improved as more data was available for assimilation.
- As more observed data were used, the model curve became closer to the observations and the surge peak was better predicted; however in these examples overestimates were then seen in the surge at HW.

Acknowledgements

The authors would like to thank our NCOF partners at the Met Office for their continuing support, which ensures smooth running of the system and facilitates development of the operational surge model.

This work was funded by EA under the STFS Development Contract.

References

Flather, R.A. 1989. The development of a new model for Thames Barrier operations. Proudman Oceanographic Laboratory Internal Document No. 6. (Unpublished Document)

Flather, R.A., 1993. An updated forecast model for Thames Barrier operations. *POL Internal Document No. 46*. 15pp + Figs & Tables.

Flather, R. A., J. A. Williams and D. L. Blackman 2000. Causes of seasonal sea level variations and implications for surge prediction. Proudman Oceanographic Laboratory, Internal Document, No 136, 15pp, figs., tables.

Williams J.A. & Horsburgh, K.J. 2006. The Operational Storm Surge Model: Development, Performance and Maintenance: April 2005 – May 2006. *Proudman Oceanographic Laboratory, Internal Document, No. 178*, 49 p.

Appendix 1 – Interim results on POL research to implement 3D var data assimilation within the surge model

Implementing the 3DVAR data assimilation technique in the CS3X surge model involves fitting sea level to both the current model state and the observed sea level from tide gauges. This is done by minimising a 'cost function' in multi dimensions (corresponding to each of the 20,000 or so model grid points). This procedure is akin to a simple least squares linear regression, except there is extra information in the cost function to account for spatial variation in model and observational errors and how they co-vary within the model domain. For example, at Avonmouth, due to the large tidal range one would expect quite a large model error in absolute terms, compared to Cromer, where the tidal range is much smaller. Similarly, one might expect spatial variation in the error associated with tide gauge measurements - however, this error is practically uniform and assumed to be 0.02 m.

The 3DVAR cost function contains error covariance matrices which contain information about error structure in the model and the observations. This cost function is minimised with pre-conditioned conjugate gradient methods to obtain a best fit to model and observations. The solution is a pattern of sea level which is subsequently used to correct the model state as the model is integrated forward in time. In this way, the model is constrained by both observations and by its current dynamical state.

The main challenges of implementing 3DVAR are:

- i) The need to accurately calculate or construct the inverse of a very large matrix (more than 400 million elements), the background covariance matrix in the cost function.
- ii) The cost function has to be minimised in multi-dimensions and the minimisation algorithm has to be tuned to certain stopping criteria
- iii) The solution or 'analysed state' has to be injected into the model and the model has to be integrated forward in time whilst remaining numerically stable.
- iv) 3DVAR implementation does not typically transfer directly from other applications such as meteorology (where it is a mature technique), due to vastly different numbers of observations and model dimensions. Neither are meteorological observations as accurate and constraining as those at tide gauges.

Progress:

i) In collaboration with a colleague working with Gauss-Markov matrices, a precise, analytical inverse of the background error covariance matrix has been found. This constructed inverse has been successfully tested in both a small domain case of 3DVAR and in the full domain tested with a version of 3DVAR relaxed to limited constant observations.

ii) Three versions of the cost function minimiser have been used, the Powell Direct Set, NAG E04DGF (recommended by the Reading DARC group) and finally M1QN3, as used by the European Centre for Medium-range Weather Forecasting (ECMWF). The Powell method was used on the initial small domain test, but did not efficiently scale to the full domain. The latter two methods are based on pre-conditioned conjugate gradient, a multi-dimension version of the Newton-Raphson solution of polynomial equations. Results showed poor convergence with NAG

E04DGF, reflecting the prior experience of ECMWF. The M1QN3 algorithm contains subtle improvements which lead to better convergence and therefore accuracy and model performance. This algorithm is open source code, obtained under academic cost-free licence from the French authors and is used by ECMWF, Meteo France and the Canadian Meteorological Centre (<http://www-rocq.inria.fr/estime/modulopt/optimization-routines/m1qn3/m1qn3.html>).

There are several parameters to tune to stopping convergence. This tuning strikes a balance between accuracy and speed. Once the 3DVAR is close to a full working model, additional time would need to be spent on tuning.

iii) So far, the analysed state has been injected into the model without stopping and restarting the model, but rather relaxing the model sea level towards the analysed state using a smoothing function. This method is more computationally efficient and should make dynamical instability less likely to develop in cases when the analysis is not in dynamical balance. However, skill testing of the model with 3DVAR versus a model without assimilation has not identified any systematic benefits.

The next phase of this work would be to implement the analysed solution in a different way - by stopping the model and restarting from the analysed state of 'best fit' sea level. This would reduce the number of parameters to be tuned and provide a better overall skill, although possibly at the expense of computational efficiency. Some subsequent optimization may regain computational efficiency.