Scientific report: WR 97-02

KNMI

Optimal interpolation of partitions: A data assimilation scheme for NEDWAM-4

Description and evaluation of the period November 1995 - October 1996

A.C. Voorrips

De Bilt, 1997

Scientific report = wetenschappelijk rapport; WR 97-02

De Bilt, 1997

P.O. Box 201
3730 AE De Bilt
The Netherlands
Wilhelminalaan 10
Telephone +31 30 220 69 11
Telefax +31 30 221 04 07

Author: A.C. Voorrips

UDC: 551.466.1
551.46.062.5
ISSN: 0169-1651
ISBN: 90-369-2120-1

Optimal Interpolation of Partitions: A data assimilation scheme for NEDWAM-4

Description and evaluation of the period

November 1995 - October 1996

A.C. Voorrips

Abstract

A new data assimilation scheme for NEDWAM4, optimal interpolation of partitions (OIP), is presented, including a detailed description of the software. Results are evaluated of a test over 12 months, in which two parallel wave model versions were run 4 times a day, one with and one without assimilation. First, some examples are described. Then, statistics over the whole period are given. OIP turns out to give better wave analyses in terms of integral parameters, in a comparison with wave buoy observations which were not assimilated (Eierland). In the forecast, the main improvement is seen in swell cases, and especially for the low-frequency wave height H_{10}. The improvement is systematic in most wave parameters up to at least 12 hours in the forecast. Based on these and previous results, it has been decided to incorporate the OIP scheme in the operational wave forecast suite at KNMI.

Contents

1 Introduction

The goal of this report is twofold. First, a description is given of a new data assimilation scheme for NEDWAM4, which is the new operational wave forecast model at KNMI. Second, the performance of the scheme is evaluated by comparing runs with and without the scheme for a period of 12 months. Although initially meant for those at KNMI who are going to operate the scheme or evaluate its results, it is hoped that it may be of interest to the wider range of people who are interested in operational application of wave data assimilation.

The basic idea of wave data assimilation is simple. The estimate which a wave model makes of the present sea state (the analysis) is usually determined by running the model forced by analyzed wind fields from a meteorogical model. It can be improved however, if wave observations can be sent in near-real time to the computer which runs the wave model. The process of correcting the computed sea state with these observations is called *data assimilation*. Improvement of the wave analysis by the assimilation will also lead to the (short-time) wave forecast. This is especially so in the case of swell, because the information which is extracted from the observations will propagate further into the model area, mainly undisturbed by (possibly erroneous) wind fields.

At KNMI, wave measurements from various sources (buoys, ships, satellites) are available in near-real time (typically, 1 to 3 hours after observation). However, until now, no data assimilation is done in the operational wave forecasting suite. A first operationally feasible assimilation scheme was devised a few years ago by Burgers et al (1992). Unfortunately, it turned out that the corrections which it applied to the model sea state were lost very quickly in the forecast (Burgers et al, 1992; Mastenbroek et al, 1994), and therefore, it was not added to the operational suite at KNMI. It was suggested by Mastenbroek et al that this rapid loss of information had to do with the fact that the scheme could only assimilate significant wave height and mean wave period measurements. They argued that more detail of the wave spectrum had to be extracted from the observations in order to apply realistic corrections to the wave field in the analysis.

In this report, we present a new assimilation scheme, called Optimal Interpolation of Partitions (OIP). The scheme tries to improve on the scheme of Burgers et al by assimilating more details of the spectrum, notably the mean parameters (energy, mean period, mean direction) of all separate wave systems which can be identified in observed and modeled wave spectra. Nevertheless, it is still a rather simple scheme, which does not significantly increase the amount of computer time needed to make a wave analysis or forecast.

The OIP scheme was originally designed by Hasselmann et al (1994, 1996) to assimilate inverted synthetic aperture radar (SAR) spectra from the ERS satellites. For the North Sea, a much more important set of observations is formed by seven directional wave buoys (figure 1). Voorrips et al (1996) adapted the scheme to assimilate this type of observations. They tested the scheme for the North Sea in a few experiments, in which it appeared to give an improvement compared to the scheme of Burgers et al (1992). Later on, it was compared with the more sophisticated (and more expensive) DASWAM assimilation system of Delft Hydraulics (Voorrips and de Valk, 1997). Also this experiment showed satisfactory behaviour of the OIP scheme, especially given its low computation costs.

The present report gives an evaluation over an extended test period, from November 1995 until October 1996. In this period, a wave model run with OIP assimilation and a run without assimilation were performed in parallel, four times a day. On the basis of these and previous results, it has been decided to add the OIP assimilation scheme to the operational wave forecasting suite at KNMI.

Section 2 summarizes the OIP algorithm. Section 3 gives an evaluation of the parallel runs, first by showing some examples, then by presenting the statistics. Is is followed by conclusions and recommendations in section 4. Appendix A describes the software in some detail. In appendix B, finally, the results of the parallel run without assimilation are compared with those of the (old) operational NEDWAM run in the same period.

2 The OIP data assimilation scheme

2.1 Outline

The OIP scheme is described extensively in Hasselmann et al (1996) and in Voorrips et al (1996). In appendix A, a detailed description of the software is given. Here, we will give a short summary of the most essential features of the method.

OIP is a so-called *sequential* assimilation mthod. This means that it is called during the wave model forecast at every time at which new observations are available; it combines the model state at that time (the *first guess*) with the new observations to calculate an *analyzed* model state; and this analyzed state serves as the initital condition for a new model run, until a new set of observations is processed, etc. This set-up is different from e.g. multi-time-level variational schemes, in which all observations over a certain time period (the "assimilation window") are assimilated simultaneously.

Two major approximations are made in formulating the OIP method. The first is to

define pre-calculated (and constant) forecast and observation error covariances, which are used at every assimilation time step. This distinguishes Optimal Interpolation (OI) schemes from the far more costly Kalman filter method, in which the error covariances are explicitly propagated by the model dynamics. The second approximation is to apply a technique called *spectral partitioning*, with which the number of free parameters in a wave spectrum is effectively reduced by more than an order of magnitude (see subsection 2.2). With these two approximations, a scheme can be constructed which is very cost-effective: an assimilation step takes considerably less time than the propagation of the wave model between two consecutive assimilation times.

The scheme consists of the following subsequent steps:

1. *Partitioning* of all observed and model spectra, i.e. division of each spectrum into a few distinct segments. The physical interpretation of each segment ("partition") is that it represents a wave system, corresponding to a certain meteorological event (e.g., swell from a distant storm in the past, or wind sea which is generated by local wind). Every partition is described by three mean parameters: its total energy, mean direction and mean frequency.

2. *Cross-assignment* of partitions of different spectra. If, in two corresponding spectra (e.g., observed and first-guess at the same location), partitions occur which are sufficiently close to each other in mean spectral parameters (frequency, direction), they are supposed to represent the same wave system, and a cross-assignment is made.

3. *Optimal interpolation* of the mean parameters from observed and model partitions which are cross-assigned. Thus, an analyzed field of partition parameters is obtained.

4. *Update* of each spectrum locally, based on the first-guess spectrum and on the analyzed partition parameters.

5. *Update* of the driving wind field as well, if there is a wind sea-component in the spectrum.

These steps are explained in a little more detail below.

2.2 Spectral partitioning

The concept of *spectral partitioning* was introduced by Gerling (1992). It is a method to describe the essential features of a two-dimensional wave variance spectrum $F(f,\theta)$ (f frequency, θ direction) with only a few parameters. This is done by separating the spectrum into a small number of distinct segments, so-called partitions, which correspond to the various "peaks" in the spectrum. The partitioning is a purely formal procedure; however, the partitions can be interpreted physically as representing independent wave systems. Details of the formalism used to calculate the partitions can be found in Hasselmann et al (1994, 1996) and in Voorrips et al

(1996).

As said, the physical interpretation of a partition will be that of a wave system, which has a different meteorological origin than other partitions in the spectrum. The data assimilation scheme makes use of this interpretation. The underlying assumption is that components of the discretized spectrum which lie within a partition are fully correlated with each other, whereas components from different partitions are uncorrelated. One can then limit oneself to calculating only a few integrated parameters of every partition, and perform the assimilation on these integrated parameters rather than on the full spectrum. The assimilation scheme as designed by Hasselmann et al (1994, 1996) uses three parameters per partition: the total energy of each partition, the mean frequency and the mean direction. Typically, one will not find more than three or four partitions in a spectrum, so the number of partition parameters (around 10) is more than an order of magnitude less than the number of parameters to represent a wave spectrum (300 in the case of WAM).

Each partition is defined to be either swell, or wind sea, or mixed wind sea/swell according to a criterion involving the wind speed and the mean wave vector of the partition (cf Voorrips et al, 1996).

2.3 Partitioning of buoy spectra

The original partitioning scheme as was devised by Hasselmann et al (1994, 1996) can only be applied to a full two-dimensional wave spectrum, such as a model spectrum or an inverted SAR spectrum. Pitch-and-roll buoy data, however, contain only the one-dimensional wave spectrum $E(f)$, plus limited information about the directional distribution. To assimilate these data as well, an adapted version of the partitioning scheme was developed (Voorrips et al, 1996) which needs only $E(f)$ and the mean wave propagation direction $\overline{\theta}(f)$ as a function of frequency. Tests with synthetic buoy spectra which were extracted from full spectra showed very good agreement between the two partitioning schemes.

2.4 Cross-assignment of partitions

The next step in the assimilation procedure is to merge the model first-guess and observed partition parameters into an analyzed field of parameters. We have assumed that different partitions within a spectrum are uncorrelated, since they are created by different meteorological events. So, we want to treat these partitions separately from each other in the assimilation. On the other hand, partitions in different spectra (e.g., model and observed spectra, or two model spectra at different locations) *are* correlated if they are created by the same event. Therefore,

we have to define a cross-assignment criterium between the partitions of two different spectra, in order to decide whether a partition in one spectrum represents the same wave system as a partition in the other spectrum.

The criterium which is used is based on the distance in spectral space between the mean parameters of two partitions. The ones which are closest to each other are cross-assigned. In case the number of partitions in the observed and model spectra do not match, additional assumptions are needed. For details, we refer to Voorrips et al (1996), and to appendix A.

2.5 Optimal interpolation of partition parameters

When the cross-assignment is done, the mean parameters of the model and observed partitions can be combined to obtain an analyzed field of partition parameters. An important input for the OI procedure are the error covariances of the errors in the observed and the model parameters. The covariances were obtained by calculating long-term statistics of differences between observations and NEDWAM model forecasts. The observation errors were assumed to be spatially independent. For the model forecasts, an error correlation length of 200 km was found.

2.6 Update of wave spectra and wind field

The analyzed partition parameters from the optimal interpolation are now combined with the first-guess spectra to obtain analyzed spectra. Every first-guess partition is multiplied by a scale factor and shifted in the (f, θ) plane such that its mean parameters are equal to the parameters obtained by the optimal interpolation. Small gaps in the spectrum which arise by the different shifts for different partitions are filled by two-dimensional parabolic interpolation.

When a wind sea partition is present in the spectrum, the driving wind field is modified using a simple growth curve relation (Voorrips et al, 1996). The new winds are used until the next wind field is read in, which is half a wind time step later (90 minutes).

3 Evaluation of the test period November 1995 - October 1996

3.1 Set-up of the parallel runs

The influence of the assimilation on the wave analyses and forecasts was assessed by running the wave model with assimilation (ASS run) for a period of 12 months, from November 1995 until October 1996. Since the wave model used was not identical to the operational model

at that time (see subsection A.2 of the appendix, and below), a parallel series of runs was performed with the wave model without assimilation (NOASS runs). Four runs per day were done, with a 6-hour analysis and a 12-hour forecast period. Runs were performed with the operational HIRLAM (High Resolution Limited Area Model; Källberg, 1990) wind fields from KNMI's forecast suite, the Automatic Production Line (APL). In case the HIRLAM fields were missing, no run was done. In cases when only one or two subsequent HIRLAM fields were missing, the forecast cycle was not interrupted: then, the 6-hour or 12-hour forecast restart fields of the previous run were used to initialise the model. In case of a longer interruption, a "cold" restart had to be done, and a two-day warm-up period was necessary. These warm-up periods have been removed from the evaluation.

There are both software and resolution differences between the operational NEDWAM model and the pre-operational NEDWAM4 model which is used in this study. Another important difference is that in the operational model, an experimental correction to the HIRLAM wind speed in near-shore areas is applied (the "Kok-correction"), which has been neglected in the runs described here. In appendix B, a concise evaluation is done between the operational NEDWAM and semi-operational NEDWAM4 NOASS runs for a part of the test period. The main conclusion from that comparison is that the quality of these runs is comparable, so the effect of OIP assimilation in the operational model can be expected to be of the same order of magnitude as found in the results below.

For the assimilation, spectral WAVEC wave observations from the Dutch North Sea Measurement Network (Meetnet Noordzee) were extracted from the LDS and LFD files prepared at KNMI by Bouws (1997). Data were used for the locations North Cormorant (NOC), Auk Alpha (AUK), K13, Euro Platform (EPF), IJmuiden (IJM), and Schiermonnikoog Noord (SON); see figure 1. Measurements from Eierland (ELD) were only used for validation.

LDS files contain high-resolution frequency spectra of wave variance, mean wave direction and directional spread, with bandwidths varying from 5 mHz at low frequencies to 10 mHz at high frequencies, in a frequency range from 30 mHz up to 500 mHz. LFD files contain the same parameters but on a coarser frequency grid, with bandwidth increasing from 15 mHz at 30 mHz center frequency, until over 100 mHz at 500 mHz. For all stations except NOC and ELD, high resolution spectra from the LDS files could be used. For NOC and ELD, only the coarse LFD files were available. In cases when LDS files were missing, LFD files were used for the other stations as well.

For practical reasons (the runs had to be done on two DEC Alpha workstations which were in use during day time), all four daily runs were done at night. To prevent use of "future"

information, data assimilation was done only at -3 hour and 0 hour analysis time. In a real operational setting, often the 3 hour-forecast observations are also available, because the wave model run starts normally some 4 hours after analysis time due to the HIRLAM run which it has to wait for.

Below, first some examples of the the parallel runs are shown. In the next subsection, a statistical analysis over the whole 12-months period is given.

3.2 Examples

3.2.1 December 1995

The situations in which the OIP data assimilation method does and does not improve the wave forecasts were clearly illustrated shortly after each other in December 1995. Figures 2 and 3 show timeseries of the low-frequency wave height H_{10} for NOC, AUK and K13, for 19-25 December. Figure 2 gives analysis results, figure 3 the 12-hour forecast. The period from December 19 to December 21 is a classic example in which data assimilation can be of value. Approximately from December 19, 1200 GMT until December 20, 0600 GMT, a strong northerly wind field was passing over the southern Norwegian Sea / northern North Sea (figure 4). In the central and southern North Sea, the wind speed was relatively low, and it remained so for at least one more day afterward. So, waves which were generated in the Norwegian Sea entered the North Sea as swell. Figure 2 shows the severe underprediction by the NOASS run, and the correction by the assimilation. In figure 3, it is clear that the correction at the northerly buoys greatly improves the 12-hour forecast more to the South. Only NOC does not profit from the assimilation in the forecast, since it is the most "upstream" station.

Around 23 December, low-frequency wave energy arrives at AUK, and somewhat later, at K13. Again, the wave energy is underestimated by the first guess, and corrected by the assimilation (figure 2). In this period, however, the wind and wave directions (not shown) were mainly westward, so the information is lost very quickly in the forecast (figure 3), since there is hardly any "upstream" assimilated information to profit from.

3.2.2 June 1996

Another period in which swell was present in the southern North Sea, was from approximately 19 - 23 June, 1996. Figure 5 shows four subsequent wind fields for this period. A northerly wind field remains for several days above the northern North Sea, with the wind maximum between NOC and AUK. Figures 6 and 7 show H_{10} time series for AUK, K13 and Eierland, with analysis and 6-hour forecasts, respectively. Unfortunately, no data from NOC were available for this

period, and AUK was missing for part of the time. Still, for the period the observations at AUK were available, a clear improvement of the underestimation by WAM can be seen for the 6-hour forecasts at K13 and Eierland. Also, the analysis of Eierland, of which the observations were not assimilated, shows an improvement with respect to the first guess, showing that the wave field in the vicinity of the observations is also improved.

3.2.3 July 1996

The example of 13-16 July 1996 is one of the few events where NEDWAM4 *overestimated* the low-frequency wave height, see figure 8. As can be seen from figure 9, at 14 and 15 July an interesting meteorological situation occurs: high, mainly westerly wind speeds near Scotland, and much lower wind speeds in the central and southern North Sea. Some swell can be expected to enter the southern North Sea, because of the rather broad directional distribution of the wind sea which is generated in the northern North Sea. However, in such a situation the so-called garden sprinkler effect limits the quality of the WAM prediction. Clearly, H_{10} is overpredicted at AUK and K13. The assimilation corrects the underprediction, and especially at K13 this correction is still present in the 6-hour forecast.

The "slow reaction" of the scheme (at AUK, 12 and 15 h GMT of July 14, the H_{10} is not yet corrected sufficiently, and the same happens at K13 a few hours later) in this particular case is a rather subtle effect. First, the low-frequency energy at this time is only a few percent of the total wave energy, and it is very sensitive to the mean wave period. This means that the uncertainty of this parameter, both in the observation and in the model first-guess, is rather large, so the correction of this parameter by the assimilation scheme may be not optimal. Note that H_{10} itself is not a parameter which is assimilated: only the total energy and the mean frequency of every partition is corrected, and if the true spectral form of the partition is different from the model, this may have a detrimental effect on the H_{10} correction. Also, the rapid change in performance of the first-guess NEDWAM4 run can be of importance: WAM is *underpredicting* the mean wave period and wave height a few hours earlier, and at 12 GMT WAM starts to *overpredict* the mean wave period and wave height. So, at first, the assimilation scheme has increased the wave energy in the surroundings of AUK. This still causes an extra overprediction later on at 12 GMT, which has to be reduced by new measurements. In general, cases where the first-guess error varies rapidly in time and space (for instance, in the vicinity of a front), the assumption of fixed error covariances underlying the optimal interpolation approach is violated, so the scheme will not work in an optimal way.

3.2.4 October 1996

As the last example, we take the period of 28-31 October 1996. Wind fields are given in figure 10. From October 29, 0000 GMT, strong northerly winds occurred in the northern North Sea; 12 hours later, the northerly flow was all over the North Sea, and at October 30, 0000 GMT, the wind had already turned to west in the southern North Sea. This is not a typical swell situation, since the wind speed is high in the southern North Sea as well as in the north; nevertheless, wave energy is propagating from north to south, so some impact of the assimilation can be expected. Figure 11 shows that this is indeed the case: the corrections at analysis time lead to improved H_{10} forecasts at AUK and K13. After October 30, however, the second peak at K13 is missed, because this extra wave energy comes aleady too much from the west/northwest to be captured by AUK.

3.3 Statistics

3.3.1 Availibility of observations

On average, the timely availability of the wave observations was very good, see figure 12: only North Cormorant was missing for a few extended periods. North Cormorant and AUK had, on the other hand, most of the time no wind velocity data available.

3.3.2 Monthly averages

Figures 13, 14, and 15 show monthly statistics of H_s for NOC, AUK, and K13, respectively, validated against the buoy observations.

The first thing which strikes the eye is, naturally, the large seasonal variability of wave height. Also, the relatively large negative bias of the NOASS runs is apparent, and it is clearly correlated with the absolute wave height. The quality of the NOASS runs is discussed in appendix B.

The large reduction of both systematic (bias) and random (standard deviation) errors by the assimilation at analysis time is to be expected, since the observations have been used in the assimilation. In the 6-hour and 12-hour forecast, most of the reduction has disappeared. Still, some improvement remains, and it is systematic over the whole period. Only at NOC, the improvement is almost completely lost after 6 hours, which can be explained by the fact that it cannot easily profit from assimilation of measurements from the other buoys: most of the errors in the wave height will be caused by either local wind, or by wave energy propagating from the Norwegian Sea or the Atlantic.

The improvement of the wave forecast by the data assimilation is on average small. This is quite understandable, since one can only profit from data assimilation if (a) propagation of wave energy is more important than local wave generation (swell cases), (b) the mean wave direction is such that some "upstream" observations are available, and (c) the accuracy of the first guess is comparable to or worse than the accuracy of the observation. This situation occurs not very often, but when it happens, it can be of great importance (swell prediction in the southern North Sea is one of the main purposes of NEDWAM). In the next paragraph, we will therefore not only give overall statistics, but also statistics for the swell conditions only.

3.3.3 Averages over the whole period

In this section, we present results of the performance averaged over the whole 12-month period. Averages are calculated over all data, but also over those situations which are considered to be swell-dominated. The criterion for a swell situation is

$$1.3\frac{U_{10}}{c_m}\cos(\Delta\theta) < 1 \tag{1}$$

where U_{10} is wind speed, c_m is the mean phase velocity, and $\Delta\theta$ is the angle between the wind direction and the mean wave direction. Although this does not restrict the selection only to the cases (a), (b) and (c) of the last paragraph, it is still a selection of cases which are more favourable for data assimilation.

Figures 16 - 23 show results for significant wave height H_s, low-frequency wave height H_{10}, mean wave period T_m, and mean wave direction θ_w, for the locations NOC, AUK, K13, EPF, and ELD. Data from the latter location have not been assimilated, and are therefore useful for the validation of the analysis; results for the other buoys are mainly for validating the forecasts. For T_m and θ_w, only those cases were selected where the observed significant wave height was more than 1 m. This was done because for very low wave energies, the observed mean wave period and direction become very unreliable.

For all four parameters, bias, standard deviation of error and root mean square (RMS) error are given. For H_s, H_{10}, and T_m, also the "normalized RMS error" (NRMS) is given, defined as the RMS error divided by the RMS of the observation (in %). NRMS is a more significant parameter than the scatter index (RMS error divided by the mean of the observations) for parameters of which the standard deviation is very large compared to the mean, which is especially the case for H_{10} in the southern North Sea; this parameter is almost zero most of the time.

For H_s and H_{10}, we see that some impact of the assimilation remains until the 12-hour

forecast. With the exception of NOC, selecting only swell cases generally enlarges the influence of the assimilation. E.g. for K13, the NRMS of the 12-hour H_{10} forecast reduces from 0.45 to 0.36, an improvement of 20 %.

The improvement in mean wave direction is lost much faster in the forecast than the other parameters. One explanation can be that the directional resolution of the WAM spectrum (30°) is much coarser than the average correction by the assimilation (around 5°); the rather subtle directional correction is usually lost quickly in the propagation due to numerical diffusion and the "garden sprinkler"-effect.

Finally, also the quality of the wind speed corrections made by the assimilation scheme has been assessed. Table 1 shows the RMS error of the HIRLAM analysis (NOASS) and of the corrected winds (ASS analysis) with respect to platform wind speed measurements (note that the wind speed observations are not used in the assimilation). It turns out that the wind speed corrections applied by the assimilation scheme *increase* the RMS error, instead of improving the winds! There are a few possible explanation for this behaviour. First, the wind speed correction which is estimated from the wind sea correction is calculated using a simple growth curve relation (Voorrips et al, 1996), which is really only valid in case of stationary and homogeneous wind fields. Second, the wind sea correction is always influenced by previous wind sea corrections at the same measurement site (the "memory" of the wave correction is longer than the 3 hours between consecutive assimilation times). The wind speed correction algorithm does not take this memory into account. The algorithm is expected to work much better in the open ocean, and with measurement locations which vary with time, like satellite measurements. Indeed, in such cases better results have been obtained (P.A.E.M. Janssen, personal communication) with a comparable algorithm (Janssen et al, 1989).

The erroneous wind speed corrections can have only a limited influence on the quality of the forecast, since they are only used until a new wind field is read in, which is already 90 minutes after assimilation. Nevertheless, it may be advisable to remove or improve the wind speed correction in the scheme.

	NOC	AUK	K13	EPF	IJM	ELD	SON
N	577	1213	2392	2525	2504	2194	2294
NOASS	2.52	1.43	1.67	2.07	1.76	1.81	2.17
ASS	2.52	1.70	1.96	2.05	2.02	2.11	2.31

Table 1: RMS errors of the HIRLAM (NOASS analysis) and corrected (ASS analysis) wind speed (in m/s), over the period November 1995 - October 1996.

4 Conclusions and recommendations

We have presented and evaluated a new assimilation scheme for NEDWAM4. Evaluation of individual swell situations, as well as a statistical analysis over a one-year period, shows clear improvement of the analysis and the 6-hour and 12-hour forecasts by the assimilation of pitch-and-roll buoy observations. The impact on the forecast is largest in swell cases, and for significant wave height, low-frequency wave height and mean wave period. The correction of the mean wave direction is quickly lost in the forecast. The update of the wind field by the assimilation, on the other hand, decreases the quality of the wind.

Because of the negligible computation cost of the scheme and the positive impact on analysis and short-term forecast (shown in this and other studies), it has been decided to implement the system in the operational forecasting suite at KNMI.

In the present implementation, only measurements from the Meetnet Noordzee WAVEC buoys are assimilated routinely. Addition of other measurement types should be relativily easy. Experiments with inverted SAR spectra from the ERS satellites (Voorrips and de Valk, 1997) have shown that adding these measurements can have a (small) positive impact in the northern part of the model region. To make this operational, a SAR inversion scheme should be implemented operationally as well. Another possible extension would be the assimilation of non-spectral observations, like the ERS altimeter wave height measurements. This requires some development of the present scheme, but should be rather straightforward. Experience with assimilation of these observations in the past (Burgers et al, 1992; Mastenbroek et al, 1994) was not too promising, however.

Due to the simplicity of the scheme, several ad-hoc decisions have been made which may be improved on. Additional experiments with different wind correction algorithms should point out whether it is useful to retain the wind correction in the scheme. The current treatment of spectra with different numbers of partitions is robust but might be optimized, in order to reduce

the number of observations which is rejected in case of non-matching partitions. The quality control system is very simple and should be refined, especially if less reliable observations will be used in the future. A good candidate to build on is the quality control system of Etala and Burgers (1994). Finally, it is recommended to continue research on more advanced assimilation schemes. They may become feasible for operational purposes within some years, because of the continuous increase of available computer power.

Acknowledgments

The author would like the thank Evert Bouws for the supply of the wave observations, Frits Koek for the operational NEDWAM results, and Jeanette Onvlee and Gerbrand Komen for stimulating discussions. He acknowledges support from the Technology Foundation (STW). He is partly affiliated to Delft University of Technology, Department of Technical Mathematics and Informatics.

References

[1] Bouws, E. (1997), "CICwaves: Wave data from CIC in March", WM Data Handbook, http://bcorda.knmi.nl:8000/WM/gegevensatlas/handboekdata/overige/cicw.

[2] Burgers, G. (1990), "A guide to the Nedwam wave model", *KNMI Scientific Report, WR-90-04*, 81 p.

[3] Burgers, G., V.K. Makin, G. Quanduo and M. de las Heras (1992), "Wave data assimilation for operational wave forecasting at the North Sea", *3rd International Workshop on Wave Hindcasting and Forecasting, May 19-22, 1992, Montreal, Canada*, 202-209.

[4] Etala, M.P., and G. Burgers (1994), "Real-time quality control of wave observations in the North Sea", *J. Atm. Oc. Techn. 11 (6)*, 1611-1624.

[5] Gerling, T.W. (1992), "Partitioning sequences and arrays of directional ocean wave spectra into component wave systems", *J. Atm. Oc. Techn. 9*, 444-458.

[6] Günther, H., S. Hasselmann and P.A.E.M. Janssen (1992), "Wamodel Cycle 4", *DKRZ Tech. Report 4*, 102 p.

[7] Hasselmann, S., C. Brüning and P. Lionello (1994), "Towards a generalized optimal interpolation method for the assimilation of ERS-1 SAR retrieved wave spectra in a wave

model", *Proc. Second ERS-1 Symp., Oct. 11-14, 1993, Hamburg, Germany, ESA SP-361*, 21-25.

[8] Hasselmann, S., P. Lionello and K. Hasselmann (1996), "A wind- and wave- data assimilation scheme", to appear in *J. Geophys. Res.*

[9] Janssen, P.A.E.M., P. Lionello, M. Reistad and A. Hollingsworth (1989), "Hindcasts and data assimilation studies with the WAM model during the Seasat period", *J. Geoph. Res. C94*, 973-993.

[10] Källberg, P. (editor) (1990), "The HIRLAM forecast model, level 1", documentation manual, SMHI, Norrköping, Sweden.

[11] Komen, G.J., L. Cavaleri, M. Donelan, K. Hasselmann, S. Hasselmann and P.A.E.M. Janssen (1994), "Dynamics and modelling of ocean waves", Cambridge University Press, 532 p.

[12] Mastenbroek, C., V.K. Makin, A.C. Voorrips and G.J. Komen (1994), "Validation of ERS-1 altimeter wave height measurements and assimilation in a North Sea wave model", *The Glob. Atm. Oc. System 2*, 143-161.

[13] Press, W.H., S.A. Teukolsky, W.T. Vetterling, B.P. Flannery (1992), "Numerical Recipes in FORTRAN", second edition, Cambridge University Press, 963 pp.

[14] Voorrips, A.C., V.K. Makin and S. Hasselmann (1996), "Assimilation of wave spectra from pitch-and-roll buoys in a North Sea wave model", *KNMI Memorandum OO-96-02.* To appear in *J. Geophys. Res.*

[15] Voorrips, A.C., and C. de Valk (1997), "A comparison of two operational wave assimilation methods", *KNMI Preprint 97-06.* Submitted to *The Glob. Atm. Oc. System.*

[16] WAMDI group: S. Hasselmann, K. Hasselmann, E. Bauer, P.A.E.M. Janssen, G.J. Komen, L. Bertotti, P. Lionello, A. Guillaume, V.C. Cardone, J.A. Greenwood, M. Reistad, L. Zambresky and J.A. Ewing (1988), "The WAM model - a third generation ocean wave prediction model", *J. Phys. Oceanogr. 15*, 566-592.

A Software description

A.1 Introduction

The assimilation software is organized as a subroutine called ISARAS (Interface to SAR ASsimilation, named after the original purpose of the method developed by Hasselmann et al (1994)). ISARAS is called at observation times (every three hours) by the WAM wave model. First, we describe the implementation of the wave model before we go on to ISARAS.

A.2 The wave model

The ISARAS data assimilation package is implemented in FORTRAN as a subroutine of the NEDWAM4 program. NEDWAM4 is an implementation for the North Sea of the WAM wave model, Cycle 4 (WAMDI, 1988; Günther et al, 1992; Komen et al, 1994). At the start of the test period (November 1995), the operational wave model at KNMI was still NEDWAM (Burgers, 1990), which is based on an older version of WAM. Since the operalization of WAM Cycle 4 was already planned (it was introduced in KNMI's Automatic Production Line in October 1996), we have only worked with Cycle 4.

The model is implemented on a $1/3°$ latitude $\times$ $1/2°$ longitude grid (approximately 32 km grid spacing), ranging from 50.7° to 70° North, and from 7.5° West to 16.5° East (figure 1). This includes the North Sea and part of the Norwegian Sea. The Norwegian Sea is included mainly in order to capture wave systems which are generated in this area, and which propagate as swell into the North Sea afterwards. At each grid point, the wave variance density spectrum is discretized in 25 frequencies ranging from 0.04 to 0.4 Hz, and in 12 directions.

In the version of the model used in this study, the open boundary of the model is not forced by externally generated wave fields. Wind fields are provided by the operational HIRLAM APL runs.

A.3 Description of ISARAS

The NEDWAM-4 model calls the assimilation routine ISARAS at every assimilation time. ISARAS itself is only an interface routine between WAM and the main assimilation routine, SARAS.

Diagram 1 shows a tree of the subroutines of SARAS. SARAS starts with a dimension check (CKW2A). Next, subroutine READSPEC is called; in this routine, the observed spectra are read. Presently, these are only the spectra from the Wavec buoys of the Meetnet Noordzee, from the LDS and LFD files prepared by Bouws (1997). The high resolution spectra from LDS

files are read by READLDS and the coarser LFD spectra by READLFD (see section 3.1). LFD spectra are interpolated to the LDS frequency grid by INTERPOL_CONSERV. READSPEC proceeds by reading the corresponding model first-guess spectra from a WAM common block. QC1 does a first rough quality check (limits on wave height, etc.). It calls MANUAL_CHECK; in this routine, the user can reject data which he does not want to assimilate, or of which he knows in advance that the data will not be correct. NEW_INDEX recalculates the indices of the observations after rejection of data by QC1.

Subsequently, SWELLSEP is called to partition the WAM spectra at the observation locations, and BSEP partitions the buoy spectra. The structure of these routines is discussed below. CROSSAS cross-assigns the partitions of the WAM spectra with the buoy spectra, using the subroutine CHECKCORRELPAR to check consistency between some parameters.

As discussed in Voorrips et al (1996), observed spectra with unassigned partitions are difficult (and dangerous) to assimilate. Therefore, the following intermezzo is added at this stage of SARAS. First, RM_LOOSE_OBS tries to remove the unassigned partitions by relaxing the conditions for the combination of partitions as defined in SWELLSEP and BSEP. Next, CROSSAS is called again. If there are still unassigned partitions, the observation is rejected by routine QC2. QC2 furthermore removes observed spectra, if more than one observed spectrum at the same position has been read in: typically, the coarse spectrum from a LFD file is rejected, if the corresponding LDS spectrum exists.

After this last quality control, DEV calculates the innovations, i.e., the differences between the observations and the WAM first-guess spectra.

Now, an array of WAM spectra to be analyzed is read in by GETWSPEC (its subroutine WAMIN copies the spectra from a WAM common block). Not all WAM spectra are added to this array: only those spectra which are within a certain maximum distance to at least one observation location (GCPATH2 is used to calculate great-circle distances between grid points). In the current setting, this maximum distance is 1000 km (variable XMAX in GETWSPEC). This feature does not make a lot of difference in a North Sea setting, but for a global application of the scheme, it greatly reduces the cost of the scheme by ignoring all WAM spectra which are very far from all observations. MEANST calculates integrated parameters (significant wave height, etc.) for the selected spectra. MNINTW finds the wind sea partition in the spectra, if there is one.

SWELLSEP calculates the partitions of the spectra in the analysis array. CROSSW cross-assigns the partitions of the WAM spectra in the analysis array with the WAM spectra at the observation locations, analogously to the cross-assignment in CROSSAS.

Subsequently, OPTINT is called, which performs the optimal interpolation of the mean parameters of the partitions. First, it defines P and R, which are the forecast error covariance matrix, and the observation error covariance matrix, respectively. For the current definition of P and R, see Voorrips et al (1996). OPTINT needs routine MATINV to invert the matrix $HPH^T + R$, with H the projection of the model state on the observations. MATINV calls the two Numerical Recipes routines CHOLDC and CHOLSL (Press et al, 1992) for the inversion. After calculating the corrections, analyzed values of the partition parameters are calculated.

TUSTRE calculates the analyzed WAM spectra, based on the first-guess spectra and the analyzed mean parameters of the partitions. Its structure is described below.

Based on the correction of the wind sea partition in the analysis, WINDCOR calculates a correction to the first-guess wind. It needs routine U10COR. This routine calculates a wind speed from the energy and mean period of a partition, using a scaling formula (see Voorrips et al, 1996).

Finally, routine WAMOUT is called to write the analyzed wave spectra and the analyzed wind speeds to the corresponding WAM arrays.

A.4 Description of some subroutines

A.4.1 SWELLSEP

SWELLSEP is the routine which performs the partitioning of a two-dimensional frequency / direction spectrum. Its tree structure is shown in diagram 2.

After initialisation by routines CHECKSWELLSEPPARA and INITSWELLSEP, the partitioning is done by routine PARTITIONING. As explained in Voorrips et al (1996), some recombination of small and not clearly separated partitions is needed, and this happens in COMBINEPEAKS. Its subroutines ONESTEPAPART, WINDSEA, LASTTWOBINS, HALFSPREAD, THRESHHOLD, and LOWENERGY check the following reasons for combination of partitions, respectively: peak frequency/direction of two partitions are only one grid point apart in the 2D spectrum; there are multiple windsea peaks; a partition is only present at the two highest frequency bins; the mean frequencies/directions of two partitions are separated in spectral space less then half their mean spectral width; the minimum between two peaks is more than a certain threshhold value; and the energy of one partition is less than a certain minimum, or less than a certain fraction of the total energy. The actual combination is carried out by DOCOMBINE (except in ONESTEPAPART). SUMENERGY calculates the energy of the combined partitions, and MEANS calculates other mean parameters of the partitions, like the mean direction and the mean frequency.

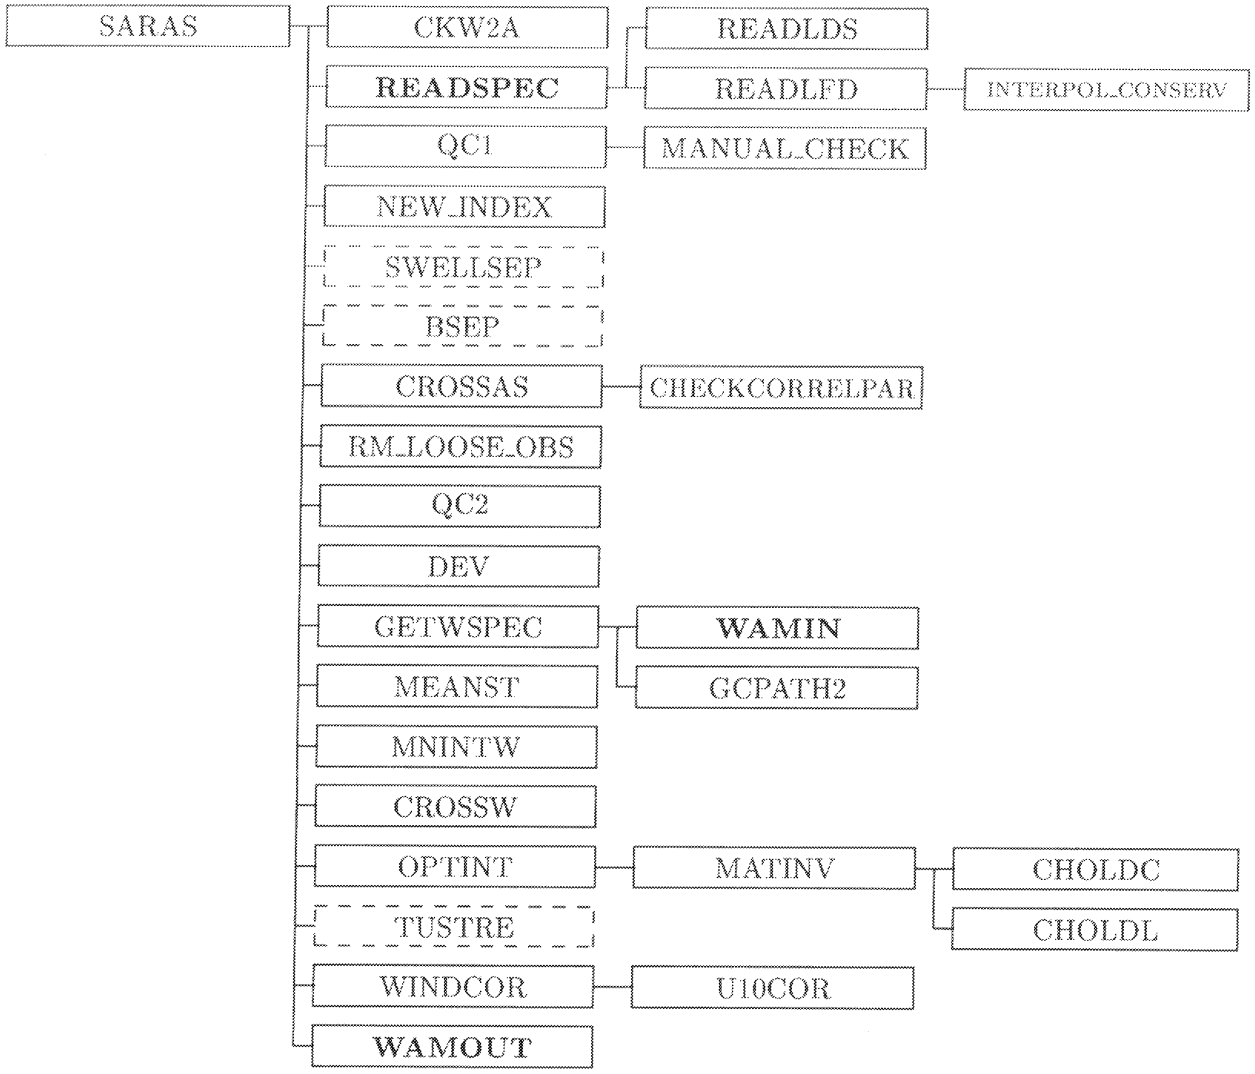

Diagram 1: Tree of the subroutines of SARAS. The bold-face routines correspond with the WAM main program via common blocks; WAMOUT is the only routine that writes to WAM. The trees of the routines with dashed boxes (SWELLSEP, BSEP and TUSTRE) are given in diagrams 2, 3 and 4, respectively.

A.4.2 BSEP

In BSEP, the partitioning is done of the buoy spectra, which consist only of the 1-dimensional wave variance and mean-direction spectra. The partitioning procedure is necessarily different from that in SWELLSEP, but results are comparable, as has been described in Voorrips et al (1996). The tree structure is shown in diagram 3.

First, partitions are calculated based on the 1-D variance spectrum alone (SEP_1D). COMB_LOW_PEAK, COMB_CLOSE_PEAKS and COMB_BAD_CONTRAST recombine partitions for the following reasons, respectively: the energy of a partition is less than a certain minimum, or less than a certain fraction of the total energy; the mean frequencies and direc-

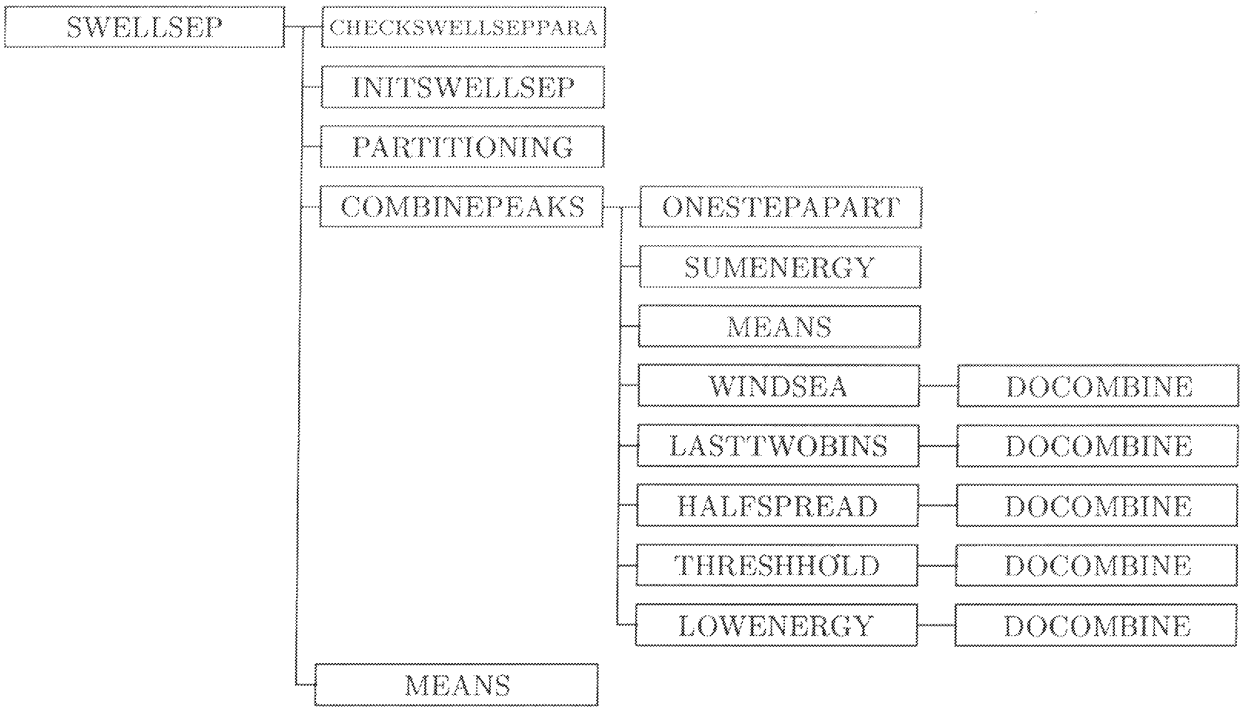

Diagram 2: Tree of the subroutines of SWELLSEP.

tions of two partitions are too close to each other; or the minimum energy density between two peaks is higher than a certain fraction of either of the peak densities. After every partitioning or recombination step, MEAN_PARAMETERS recalculates the mean parameters of every partition.

New partitions are created by SEP_DIR, if the mean direction within a partition changes too fast with frequency. SWELLWIND calculates which partitions have to be considered as swell and which as wind sea systems, and finally COMB_WINDSEA recombines all partitions which are found to be wind sea.

A.4.3 TUSTRE

TUSTRE is the routine to create new ("analyzed") spectra from first-guess spectra and analyzed mean partition parameters, by rotating and shifting each partition in the spectrum to change its mean parameters from the first-guess to the analyzed values. Diagram 4 shows its tree structure.

Some initialisation is done by CHECKTUSTREPARA and INITTUSTRE. TRANSFORM rotates and shifts the partitions in the spectrum to fit the new mean parameters. But because every partition is shifted and rotated with a different amount, some overlap and some gaps will occur in the spectrum. Overlap is accounted for by just summing up the contributions of every

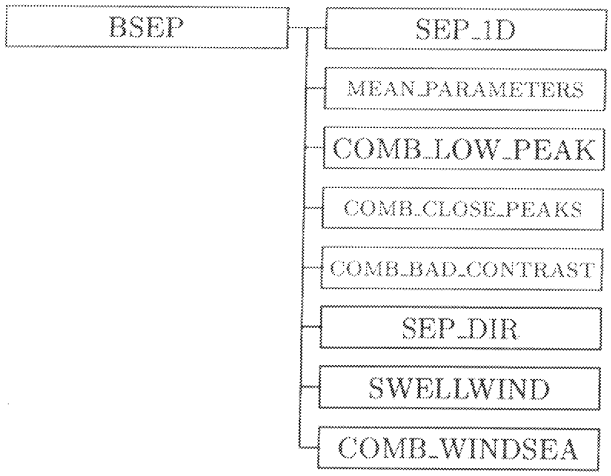

Diagram 3: Tree of the subroutines of BSEP.

partition at every spectral bin. Gaps are smoothly filled by interpolation in routine FILLGAPS. Its subroutines MAKEFRAMES and AVOIDPEAKS identify the gaps, and GAPINTERPOL makes a quadratic interpolation over the gap. In the interpolation, a least-squares solution of a set of equations is calculated by routine LSQ, which calls the Numerical Recipes routines SVDCMP and SVDKSB (Press et al, 1992). TRANSMEANS, finally, recalculates the mean parameters of the analyzed spectra which are obtained by TUSTRE.

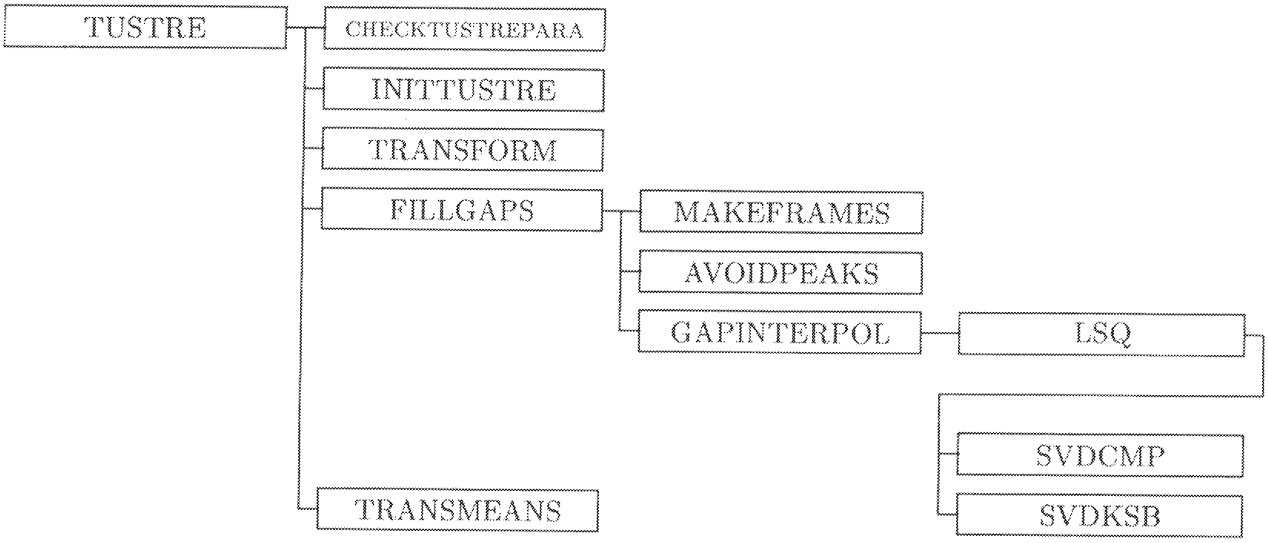

Diagram 4: Tree of the subroutines of TUSTRE.

A.5 Include files

The ISARAS package uses the following include files:

- **dim_ass.h** This include file contains the dimensions for arrays in the ISARAS routines. It includes itself the WAM include file, **parcrs.h**, and a file in which the WAM dimension variables are declared, **decl_parcrs.h**. The latter is necessary, since all ISARAS routines use IMPLICIT NONE.

- **common_ass.h** This file contains those WAM common blocks which are used by ISARAS. The only routines of ISARAS which include this file are the interface routine (ISARAS), the reading routines READSPEC and WAMIN, and the writing routine WAMOUT. This minimizes the risk of interfering with WAM by accident.

- **waveclds.h** This file contains arrays which contain information on the wave spectra from the LDS observation files.

- **waveclfd.h** This file contains arrays which contain information on the wave spectra from the LFD observation files.

A.6 Output

The main output from ISARAS are the analyzed wave spectra and wind vectors, which are copied to the WAM spectra. Apart from this, for every assimilation step a short observation logfile OBS*dtg*.out is written, where $dtg = yymmddhh$ is the date-time group of the assimilation date. Also, ISARAS writes to the main WAM logfile; the detail of the comments depends on the parameter ITEST which is defined in SARAS.

B Comparison of operational and semi-operational runs without assimilation

The performance of the operational NEDWAM runs and the semi-operational NOASS NEDWAM4 (WAM Cycle 4) runs has been compared for the period November 1995 until May 1996, by validation against the buoy observations at NOC, AUK, K13, and EPF. Only the results at analysis time have been compared. The results are summarized in tables 2, 3, 4 and 5.

The first obvious difference between the results is the bias in wind speed at K13 and EPF. These two platforms lie in the southern North Sea, where in the operational NEDWAM run the wind speed is systematically increased by 10 % to 15% (the so-called "Kok" correction). This correction is applied in order to compensate roughly the underestimation of wind speed at sea close to the coast due to the rather coarse wind model resolution. In the semi-operational NEDWAM4 runs, this correction is not applied. For EPF, the Kok-correction is clearly justified

by the observations: it should even be augmented. For K13, it seems that the Kok correction is too large: the positive bias of the NEDWAM winds is almost as large as the negative bias of the NEDWAM4 winds. In contrast to the bias, the standard deviation of the NEDWAM4 winds is smaller (probably caused by the higher resolution of the NEDWAM4 grid), and the RMS errors of NEDWAM4 and NEDWAM balance each other more or less. Note that wind speed measurements for North Cormorant were not available for most of the period, so the rather large bias in wind speed here is based on a much smaller number of observations than all the other statistical results which are presented.

The second clear difference is the negative bias in significant wave height of the NEDWAM4 runs, which is higher than the bias of NEDWAM for all four buoys. For K13 and EPF, and to some extent for AUK, this negative bias can be explained by the missing Kok-correction. This correction seems to compensate partly for the apparent underprediction of waves in the shallow southern North Sea.

But clearly, at least one extra effect determines the differences between NEDWAM4 and NEDWAM. First, at NOC NEDWAM4 has again a larger negative bias in wave height than NEDWAM, although the Kok correction plays no role here. Second, NEDWAM has a larger negative bias in mean wave period at all four buoys. This is the contrary of what one would expect for waves generated by higher wind speeds. Clearly, waves or on average steeper in the NEDWAM runs. These differences are probably due to the replacement and retuning of the source functions in the WAM model which took place when Cycle 4 was constructed (Günther et al, 1992).

In the southern North Sea, the standard deviation of the significant wave height errors is smaller for NEDWAM4. This is probably due to the higher spatial resolution of NEDWAM4. Overall, the mean wave period results are better for NEDWAM4.

In summary, the quality of the NEDWAM and the NEDWAM4 analyses runs trade off: NEDWAM is generally better in wave heights, NEDWAM4 better in wave periods (and wave steepness). For the performance of the wave model, a run with NEDWAM4 and Kok corrections would be preferable, however. The negative biases, both in wave height and in mean wave period, in the southern North Sea would be reduced by the corrections.

The main conclusion of this comparison is that the quality of the semi-operational and the operational runs are comparable, and thus the expected impact of assimilation in the operational suite can be expected to be comparable to the results described in section 3 of this report.

	U_{10} (m/s)		H_s (m)		H_{10} (m)		T_m (s)	
N	65		490		490		486	
Mean obs	10.07		3.28		1.83		8.57	
	NEDWAM4	NEDWAM	NEDWAM4	NEDWAM	NEDWAM4	NEDWAM	NEDWAM4	NEDWAM
bias	-1.53	-1.68	-0.89	-0.57	-0.88	-0.78	-1.26	-1.37
std	1.51	1.69	0.52	0.49	0.57	0.60	0.93	0.91
rms	2.16	2.38	1.03	0.75	1.05	0.99	1.56	1.65

Table 2: Statistics for the oeprational NEDWAM and the semi-operational NEDWAM4 NOASS run, for North Cormorant. U_{10}, wind speed; H_s, significant wave height; H_{10}, low-frequency wave height (frequencies below 0.1 Hz); T_m, mean wave period.

	U_{10} (m/s)		H_s (m)		H_{10} (m)		T_m (s)	
N	133		638		638		529	
Mean obs	7.46		2.29		0.83		7.41	
	NEDWAM4	NEDWAM	NEDWAM4	NEDWAM	NEDWAM4	NEDWAM	NEDWAM4	NEDWAM
bias	-0.46	-0.46	-0.37	-0.14	-0.29	-0.21	-0.62	-0.74
std	1.14	1.15	0.38	0.39	0.35	0.39	0.56	0.66
rms	1.24	1.24	0.53	0.41	0.46	0.44	0.83	0.99

Table 3: Statistics for NEDWAM and the NEDWAM4 NOASS run, for Auk Alpha.

	U_{10} (m/s)		H_s (m)		H_{10} (m)		T_m (s)	
N	660		664		664		436	
Mean obs	8.15		1.50		0.32		6.25	
	NEDWAM4	NEDWAM	NEDWAM4	NEDWAM	NEDWAM4	NEDWAM	NEDWAM4	NEDWAM
bias	-0.65	0.39	-0.19	0.05	-0.09	-0.11	-0.28	-0.55
std	1.49	1.68	0.27	0.39	0.21	0.27	0.47	0.64
rms	1.63	1.72	0.34	0.39	0.23	0.29	0.55	0.85

Table 4: Statistics for NEDWAM and the NEDWAM4 NOASS run, for K13.

	U_{10} (m/s)		H_s (m)		H_{10} (m)		T_m (s)	
N	785		786		786		421	
Mean obs	7.88		1.18		0.17		5.57	
	NEDWAM4	NEDWAM	NEDWAM4	NEDWAM	NEDWAM4	NEDWAM	NEDWAM4	NEDWAM
bias	-1.35	-0.72	-0.24	-0.07	-0.05	-0.06	-0.27	-0.33
std	1.57	1.76	0.24	0.32	0.11	0.16	0.54	0.53
rms	2.07	1.90	0.34	0.33	0.12	0.17	0.61	0.62

Table 5: Statistics for NEDWAM and the NEDWAM4 NOASS run, for Euro Platform.

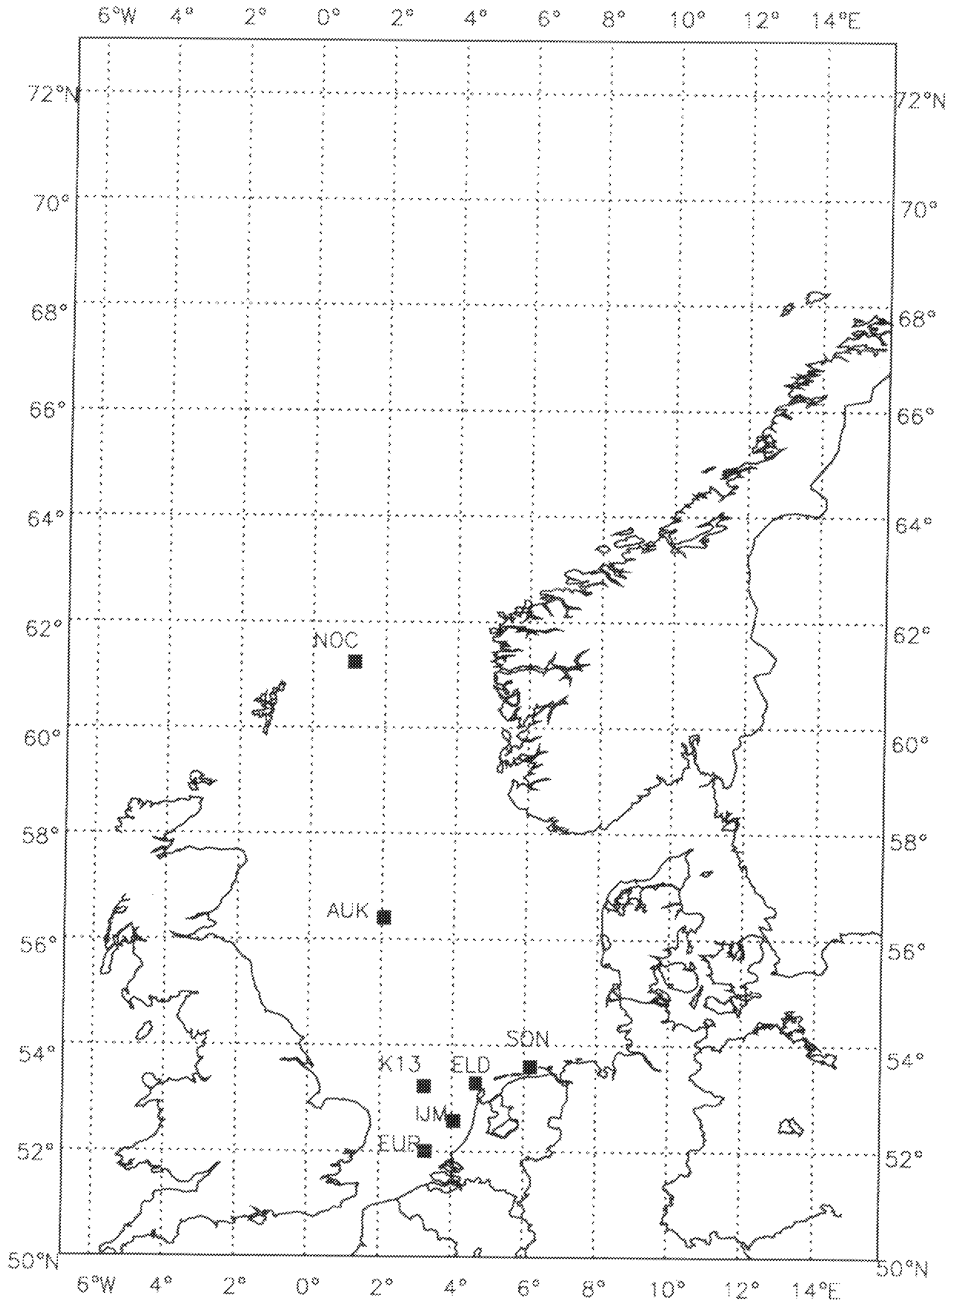

Figure 1: The model region. Indicated are the following Wavec buoy locations: *NOC*, North Cormorant; *AUK*, Auk Alpha; *K13*; *EUR*, Euro Platform; *IJM*, IJmuiden; *ELD*, Eierland; *SON*, Schiermonnikoog Noord.

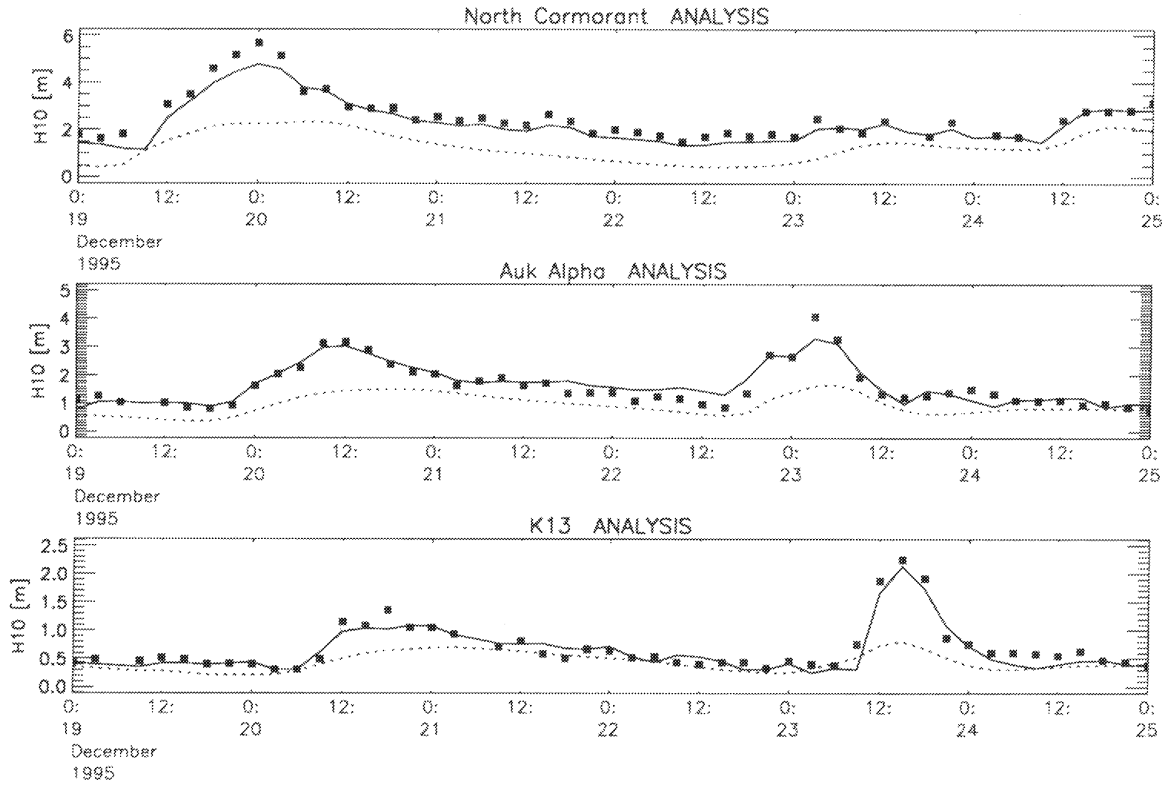

Figure 2: Time series of H_{10} from December 19, 1995 until December 25, 1995, for North Cormorant, Auk Alpha, and K13. Boxes indicate observations, solid lines ASS analysis results, dashed lines NOASS analysis results.

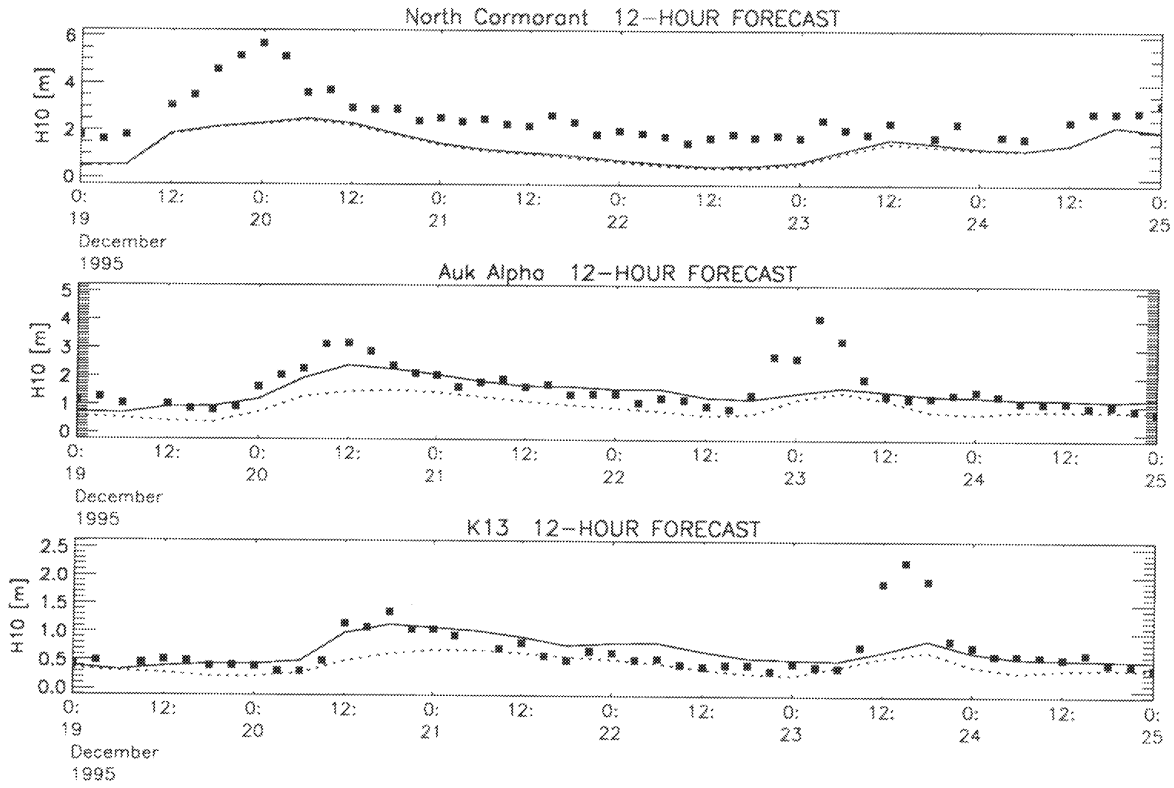

Figure 3: Time series of H_{10} from December 19, 1995 until December 25, 1995, for North Cormorant, Auk Alpha, and K13. Boxes indicate observations, solid lines ASS 12-hour forecasts, dashed lines NOASS 12-hour forecasts.

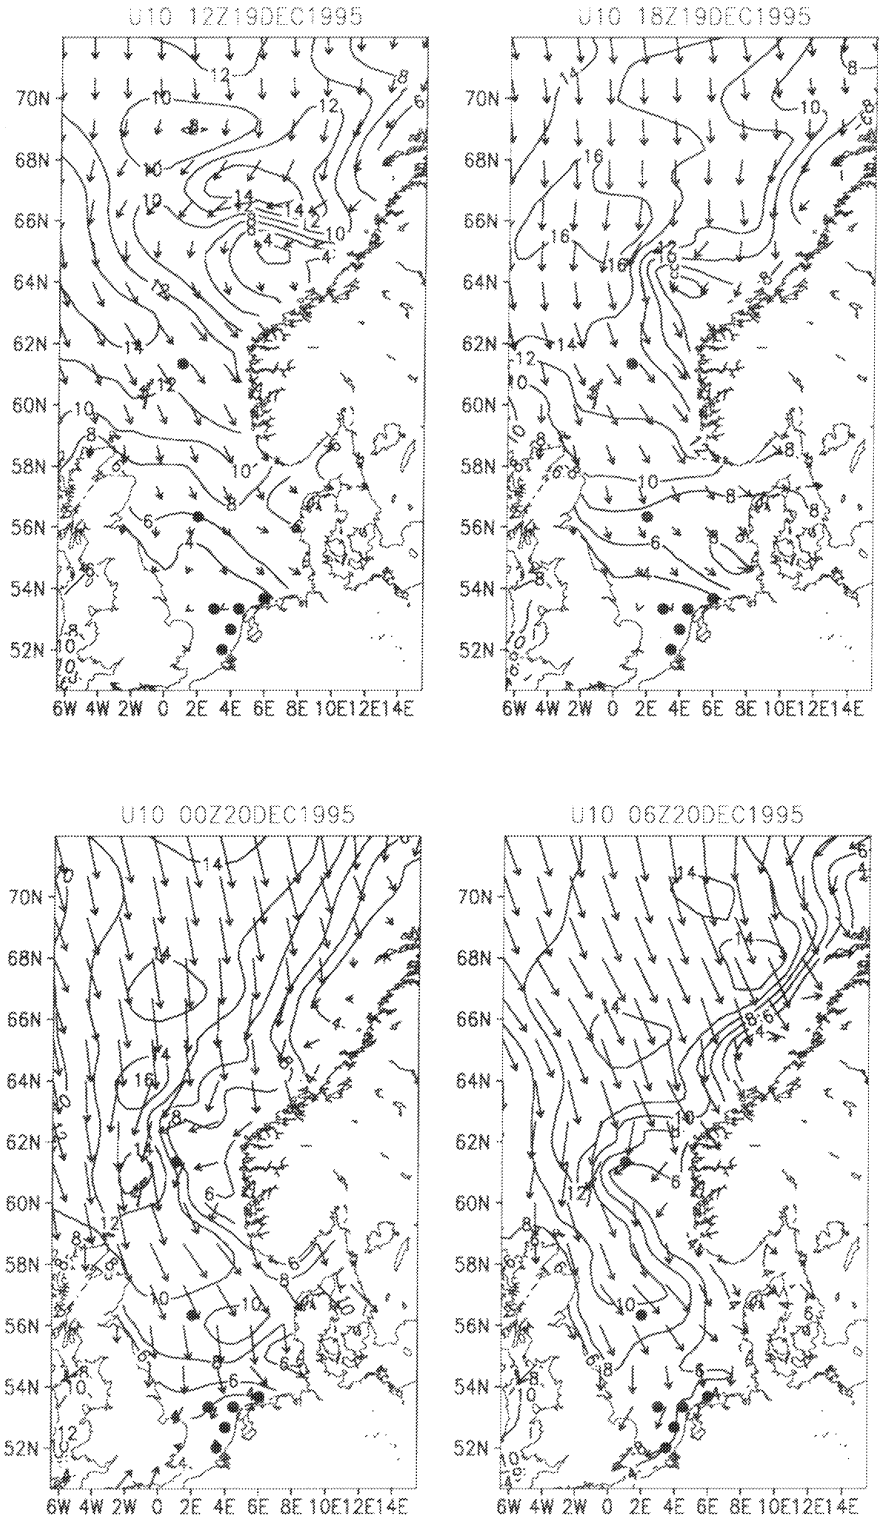

Figure 4: Six-hourly HIRLAM wind fields from 19 December 1995, 1200 GMT until 20 December, 0600 GMT.

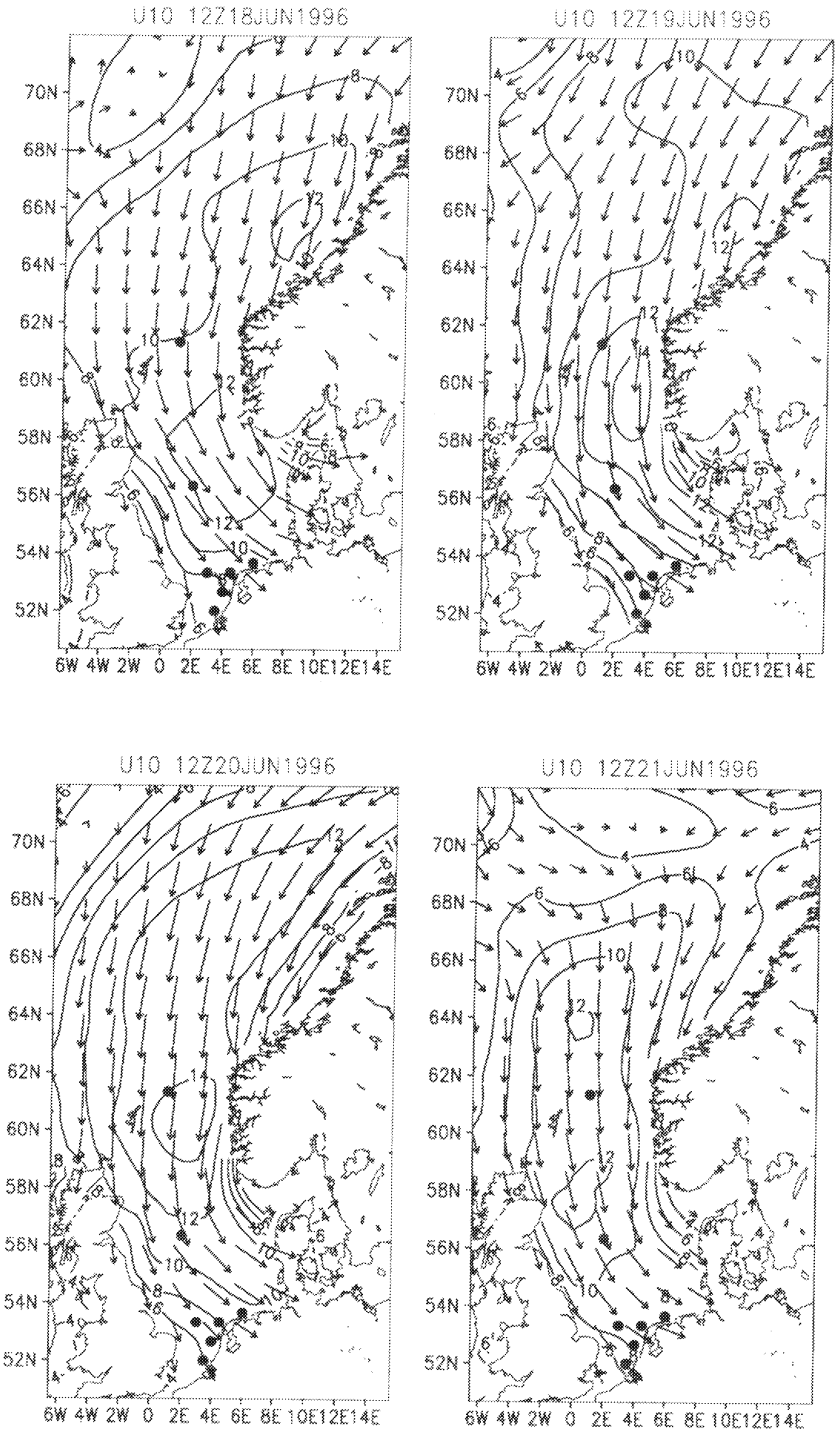

Figure 5: Daily HIRLAM wind fields from 18 June 1996, 1200 GMT until 21 June, 1200 GMT.

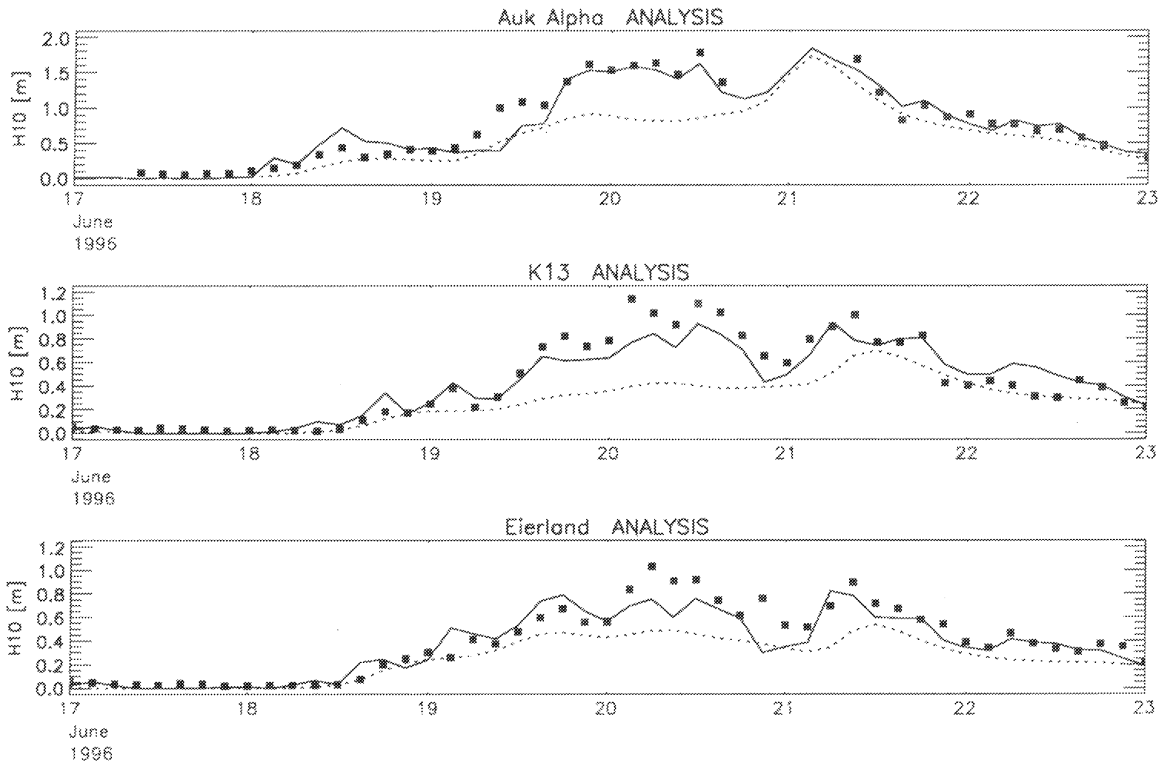

Figure 6: Time series of H_{10} from June 17, 1996 until June 23, for Auk Alpha, K13, and Eierland. Boxes indicate observations, solid lines ASS analysis results, dashed lines NOASS analysis results.

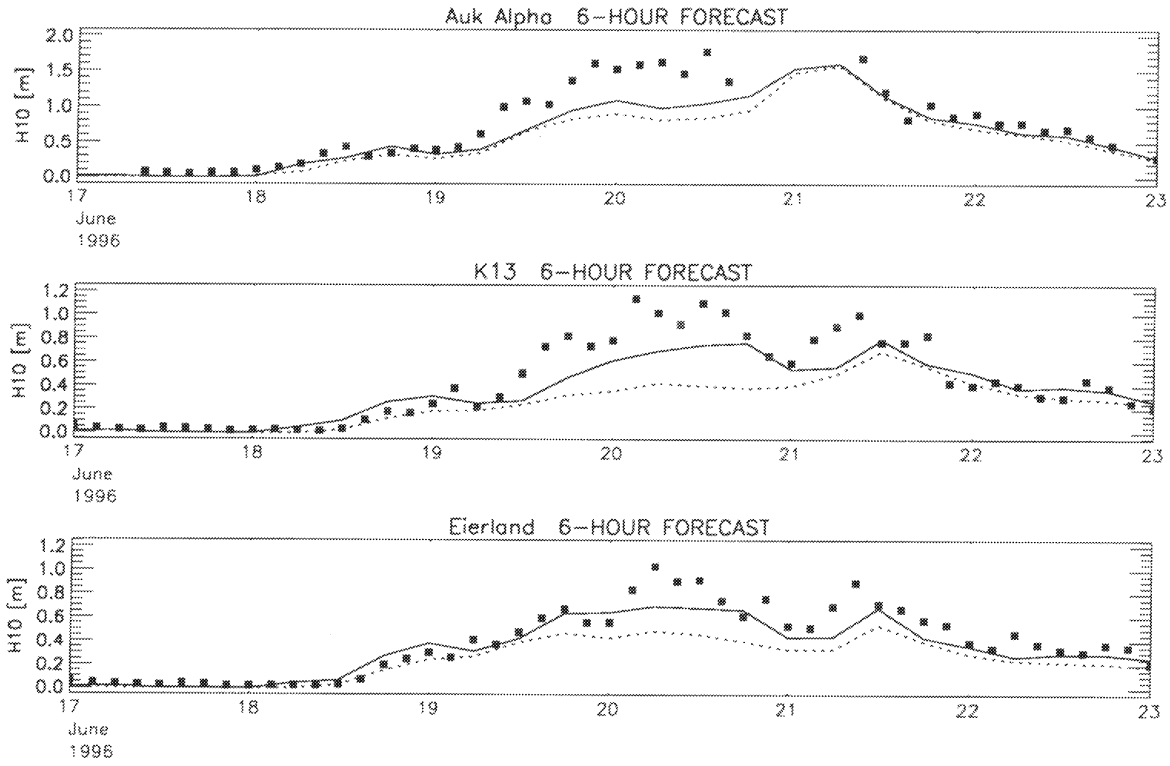

Figure 7: Time series of H_{10} from June 17, 1996 until June 23, for Auk Alpha, K13, and Eierland. Boxes indicate observations, solid lines ASS 6-hour forecasts, dashed lines NOASS 6-hour forecasts.

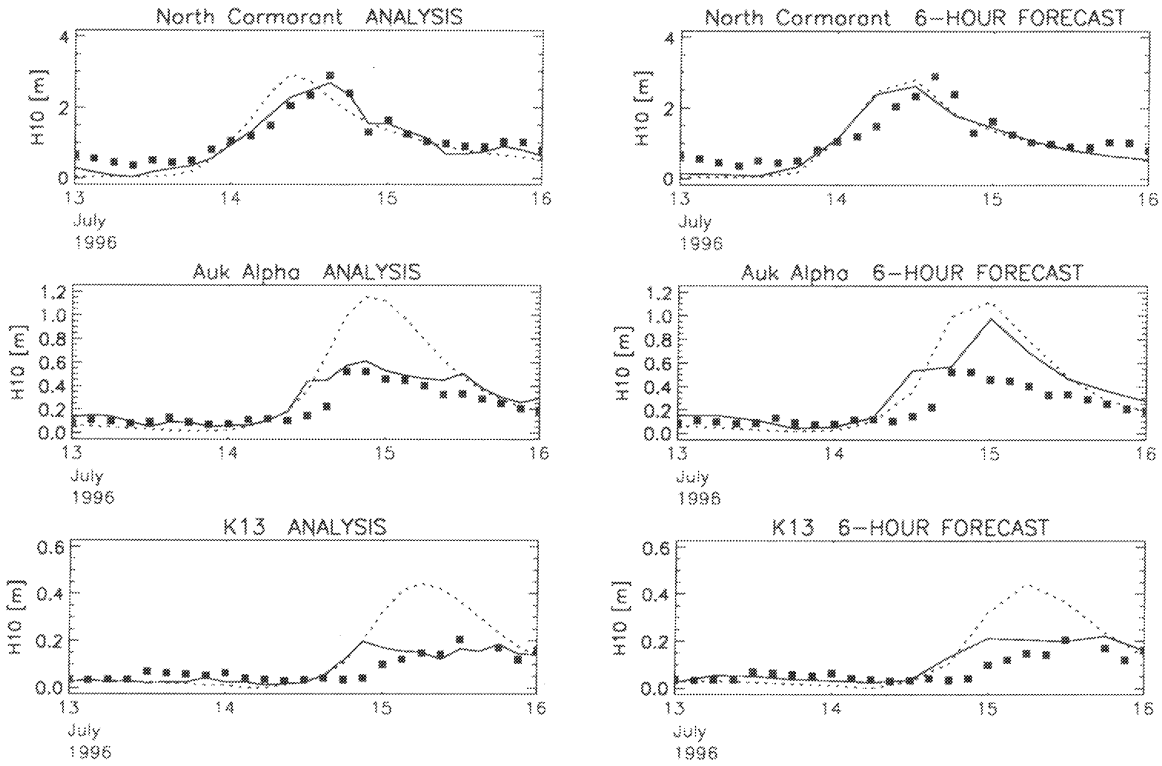

Figure 8: Time series of H_{10} from July 13, 1996 until July 26, for North Cormorant, Auk Alpha, and K13. Left panels: model analyses. Right panels: 6-hour forecasts. Boxes indicate observations, solid lines ASS results, dashed lines NOASS results.

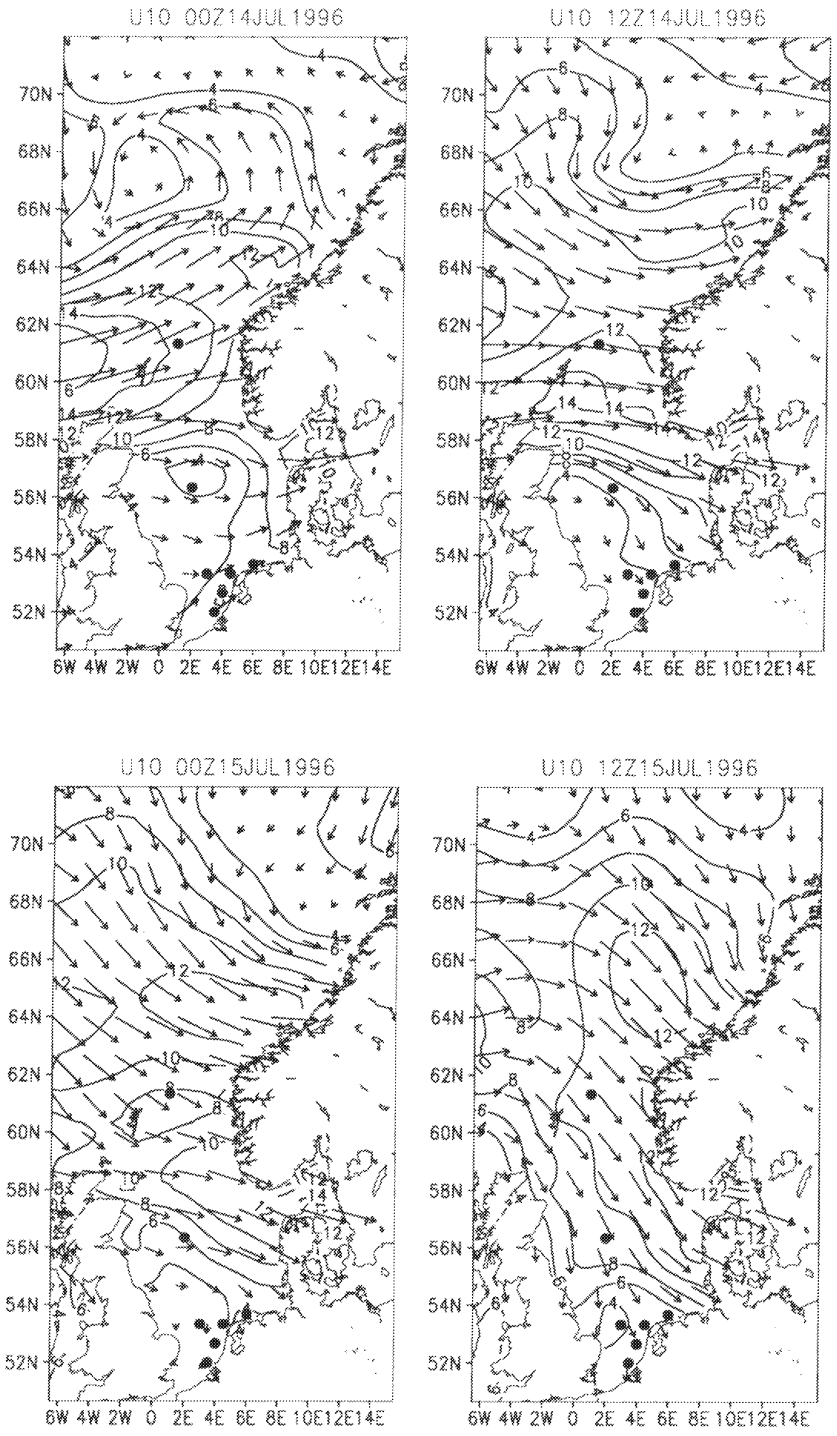

Figure 9: 12-hourly HIRLAM wind fields from 14 July 1996, 0000 GMT until 15 July, 1200 MT.

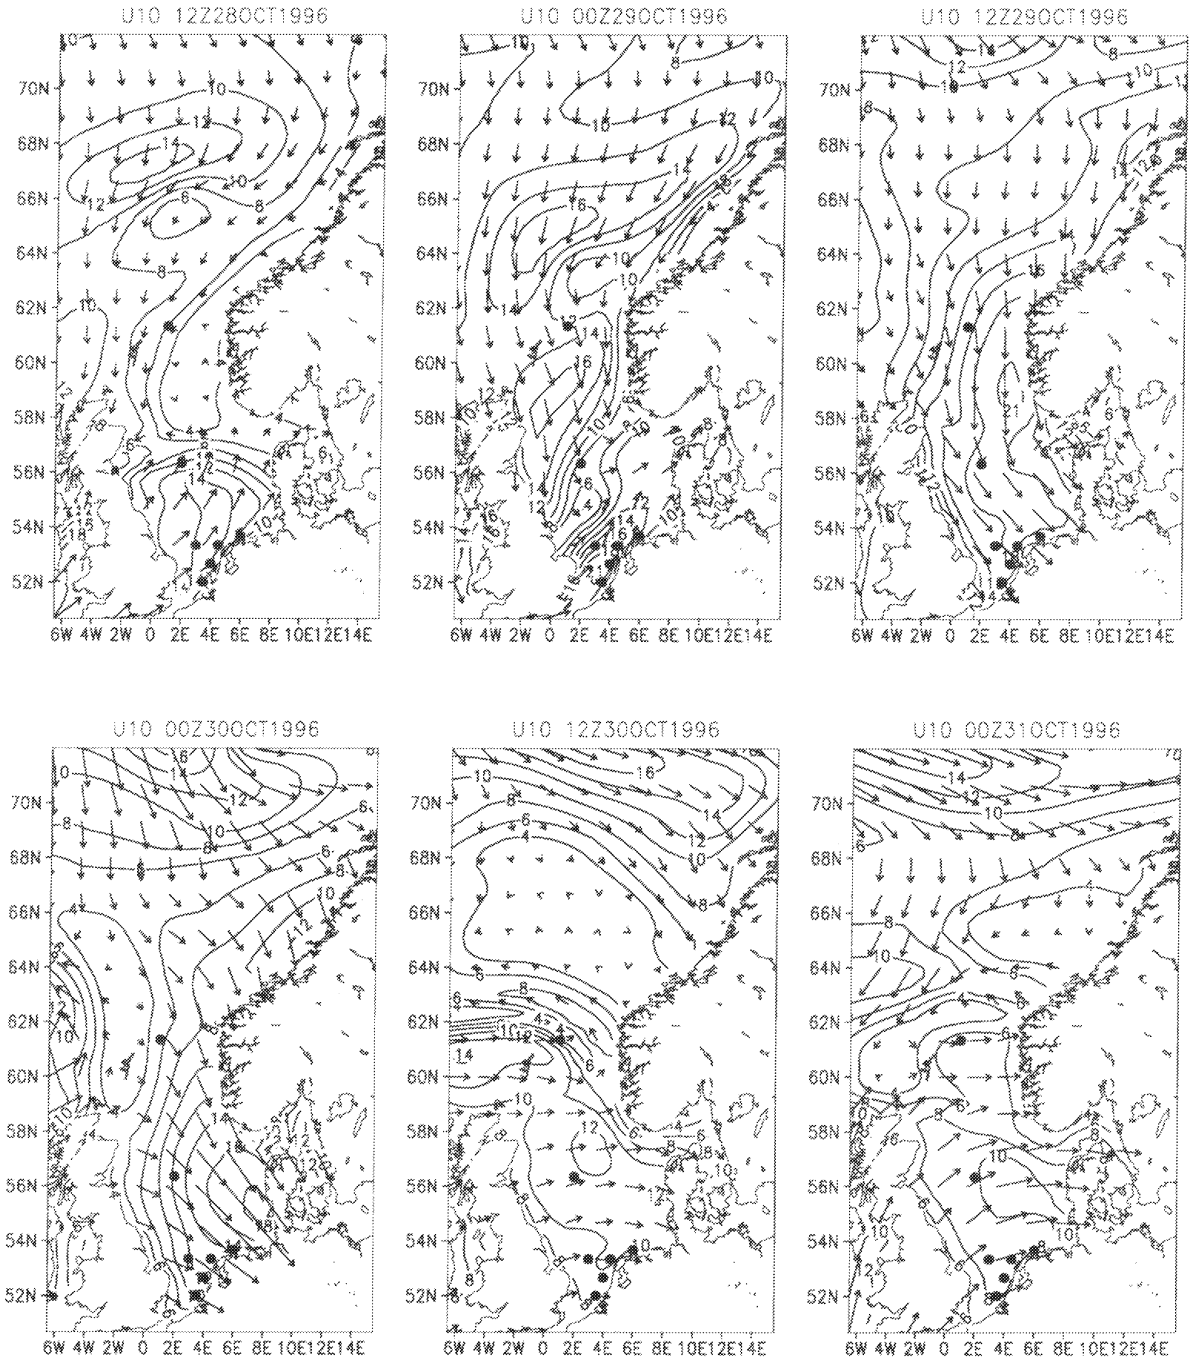

Figure 10: 12-hourly HIRLAM wind fields from 28 October 1996, 1200 GMT, until 31 October, 0000 GMT.

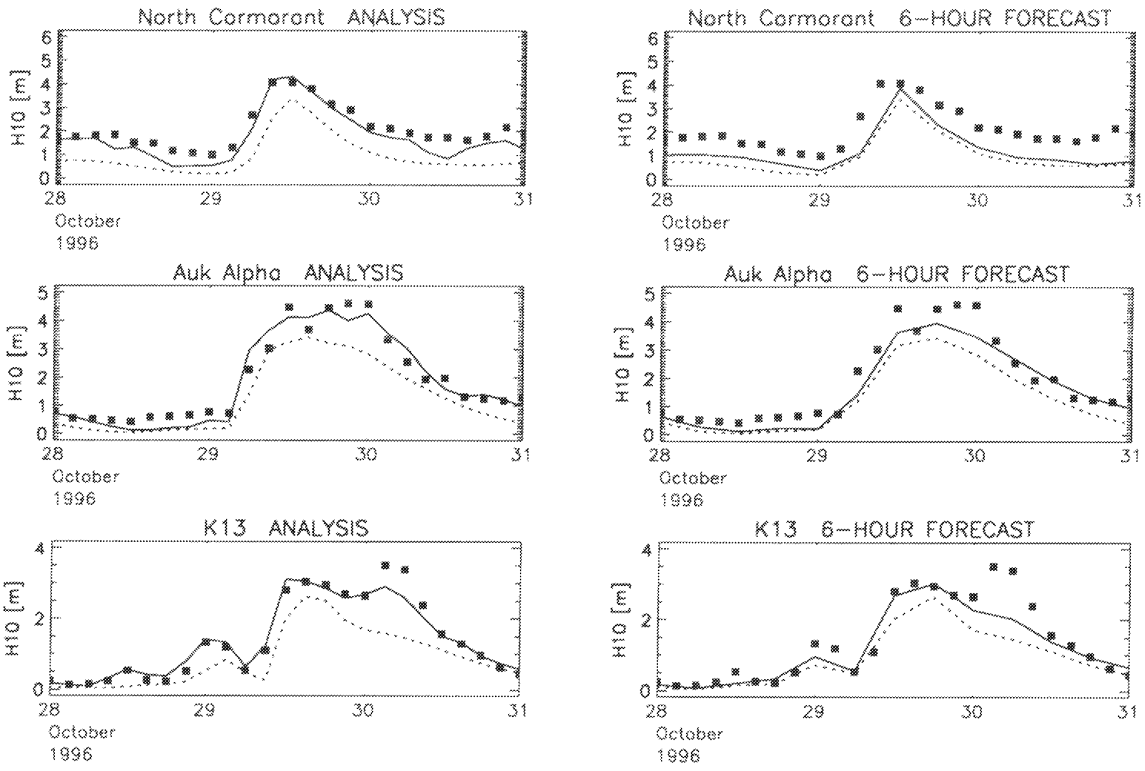

Figure 11: Time series of H_{10} from October 28, 1996 until October 31, for North Cormorant, Auk Alpha, and K13. Left panels: model analyses. Right panels: 6-hour forecasts. Boxes indicate observations, solid lines ASS results, dashed lines NOASS results.

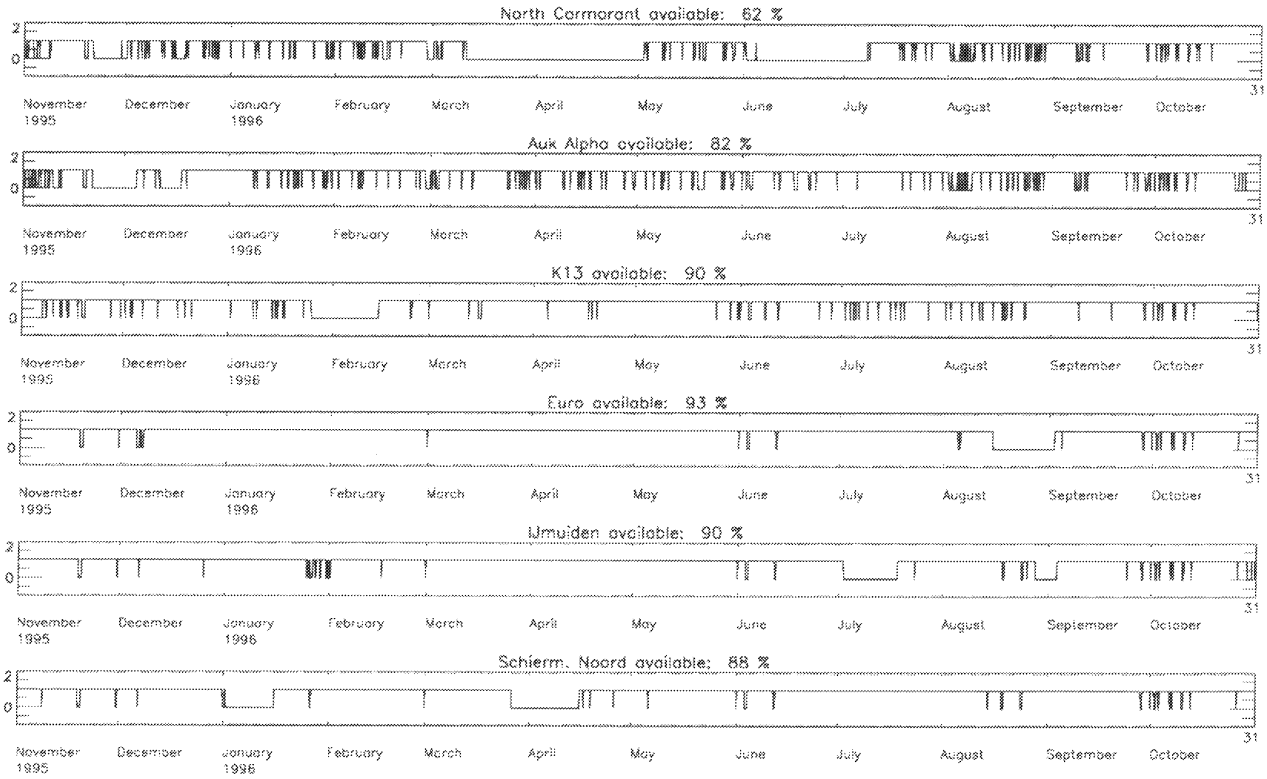

Figure 12: Availability of wave data of the six buoys which were used for the assimilation. One means available, zero means not available.

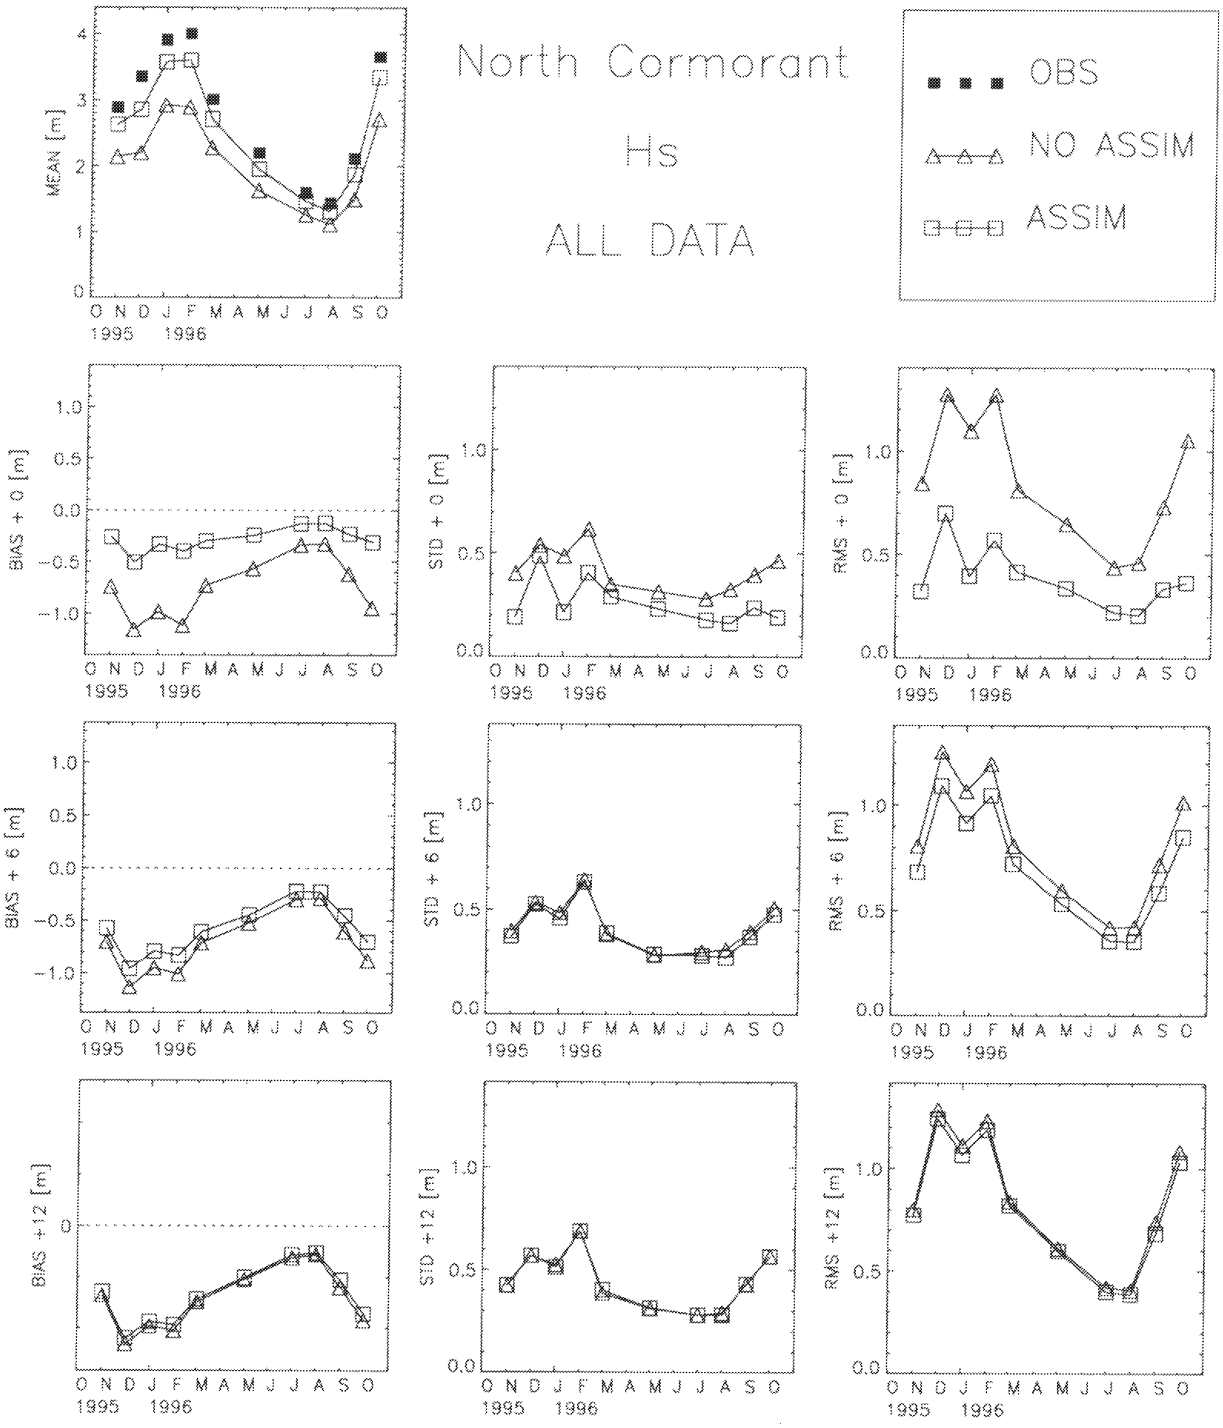

Figure 13: Monthly averages of Hs for North Cormorant. Top left panel: monthly means for observations (filled boxes), NOASS analysis results (triangles) and ASS analysis results (open boxes). Lower three rows: NOASS and ASS monthly bias (left), standard deviation (middle), and RMS error (right) at analysis time (second row), 6-hour forecast (third row) and 12-hour forecast (bottom row).

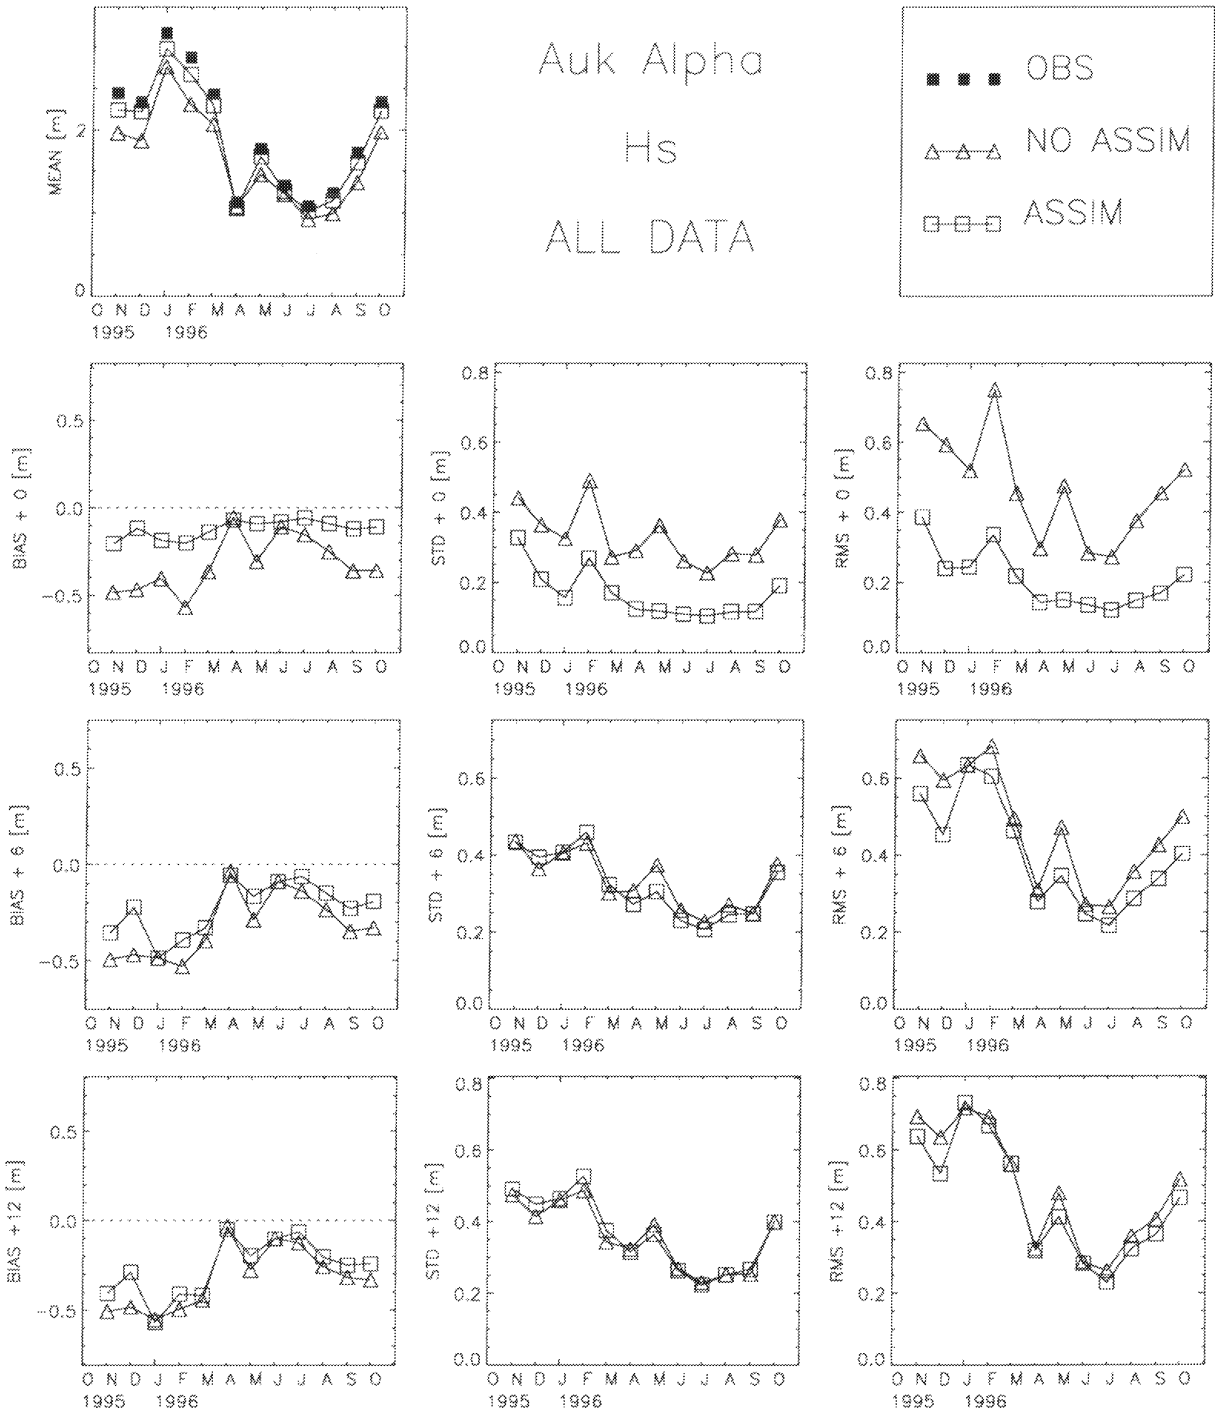

Figure 14: Monthly averages of Hs for Auk Alpha. Panels and symbols as in figure 13.

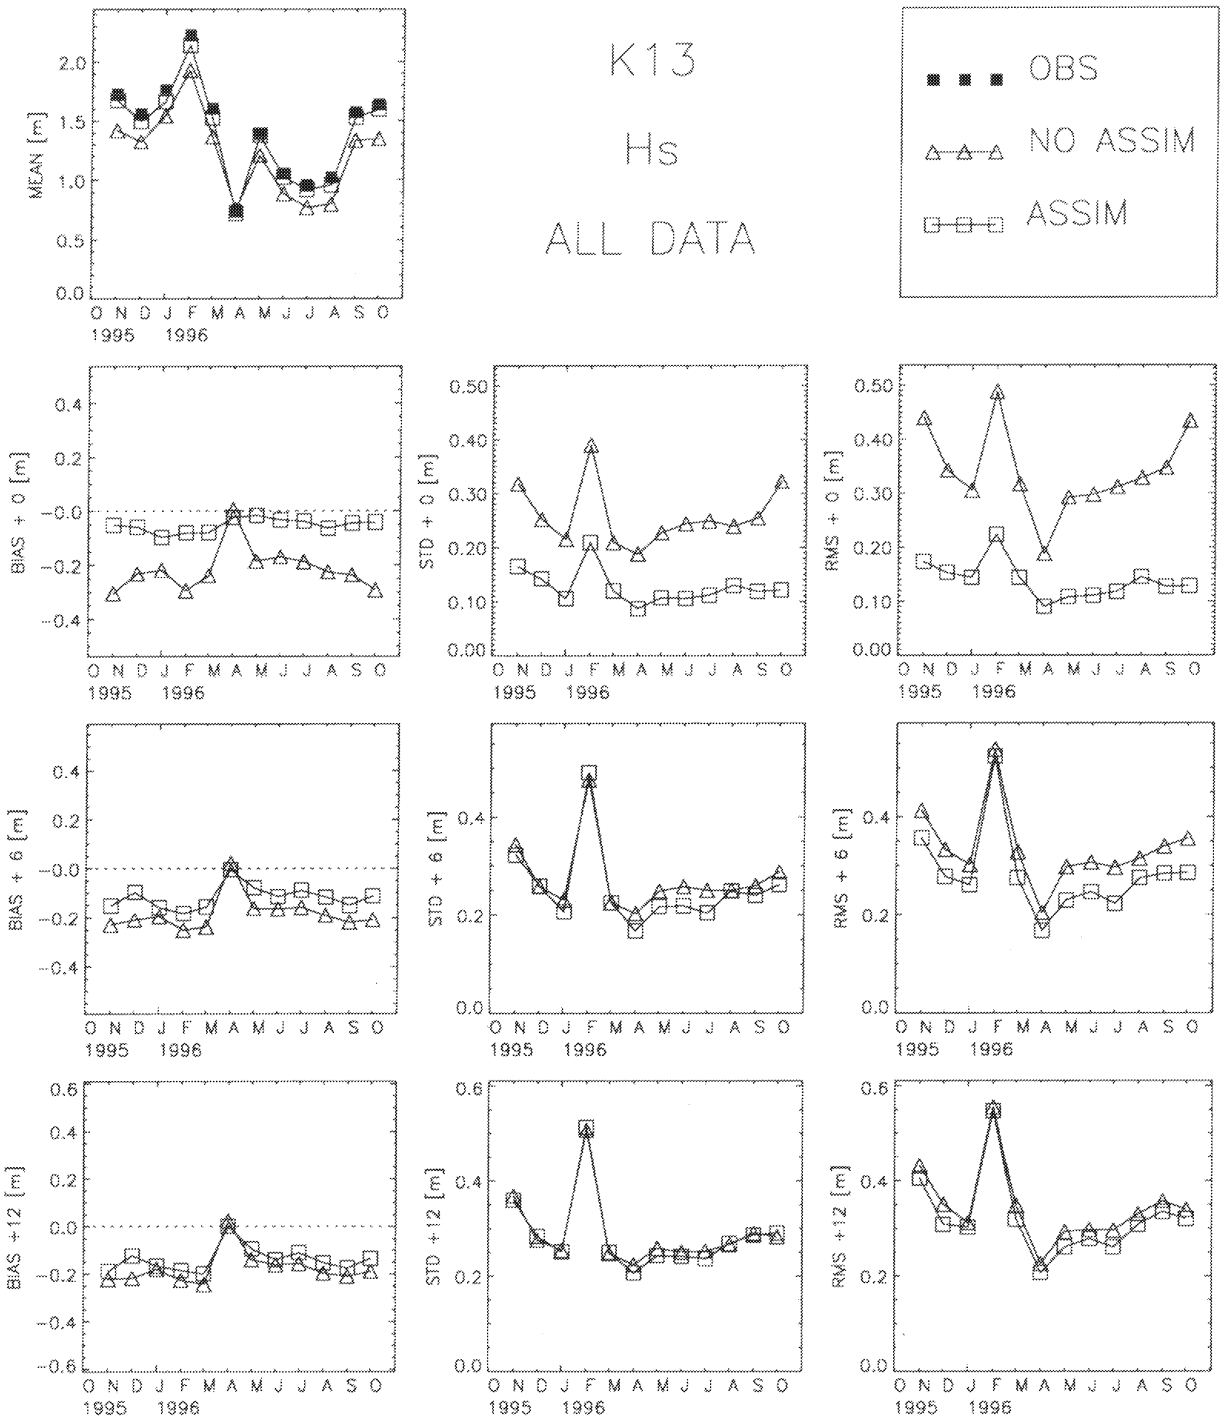

Figure 15: Monthly averages of Hs for K13. Panels and symbols as in figure 13.

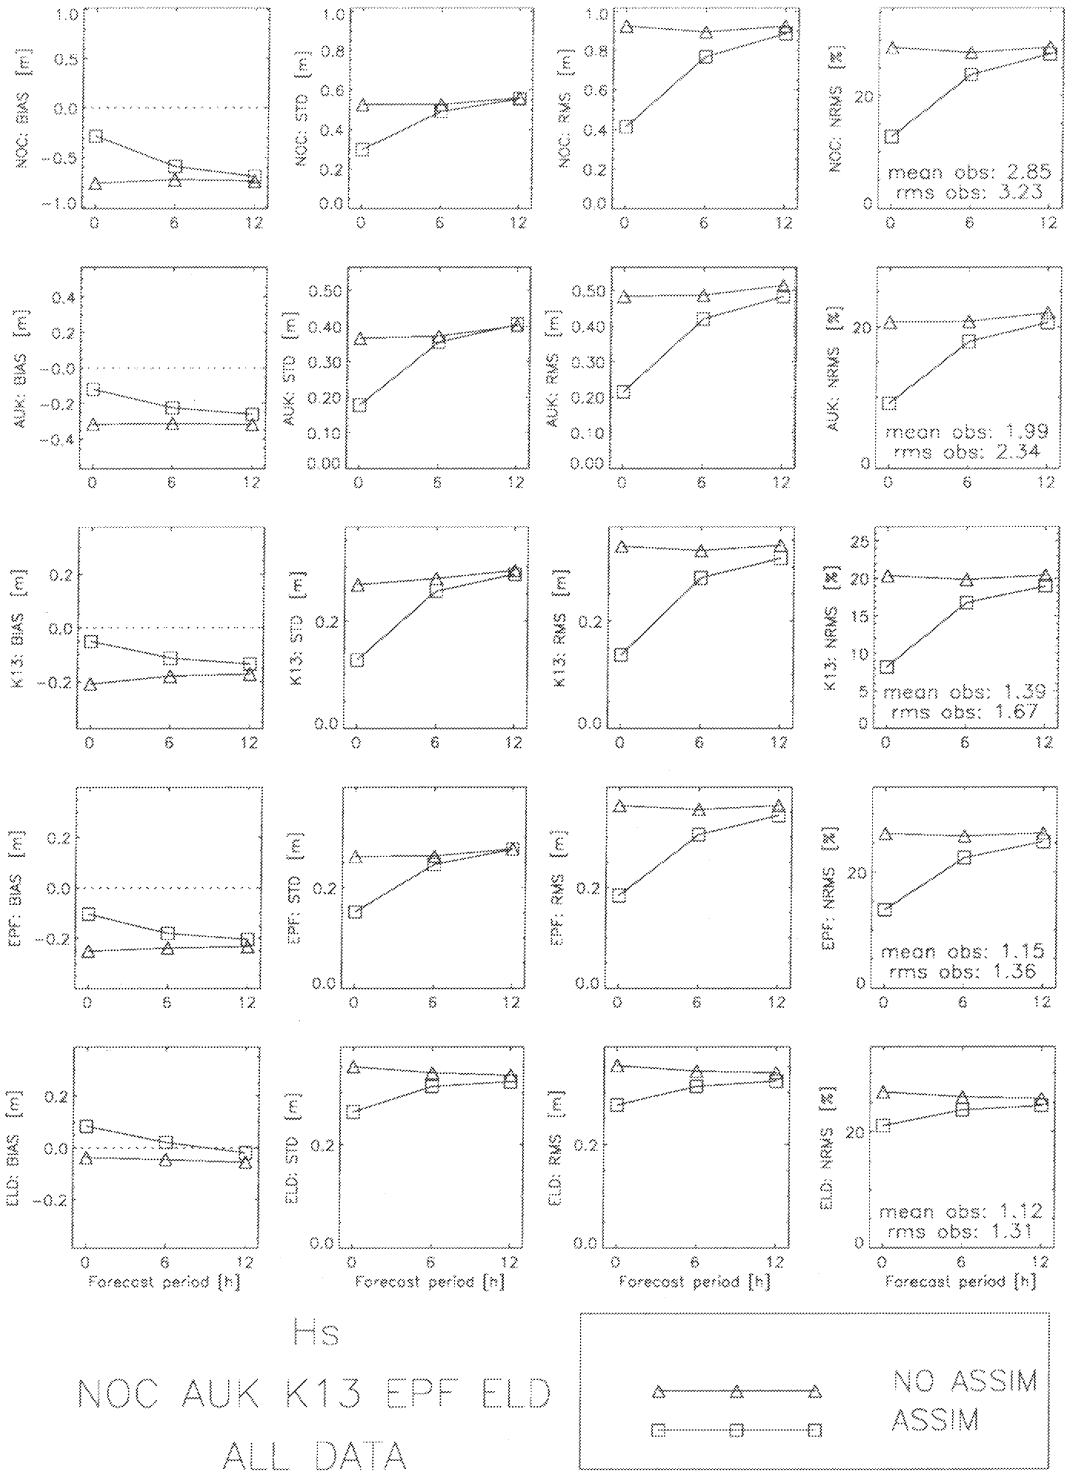

Figure 16: Statistics of all data over the period November 1995 - October 1996, as a function of forecast period, for significant wave height H_s. Boxes, runs with assimilation; triangles, runs without assimilation. From left to right: bias, standard deviation, RMS error, normalized RMS error. From top to bottom: North Cormorant, Auk, K13, Euro Platform, and Eierland.

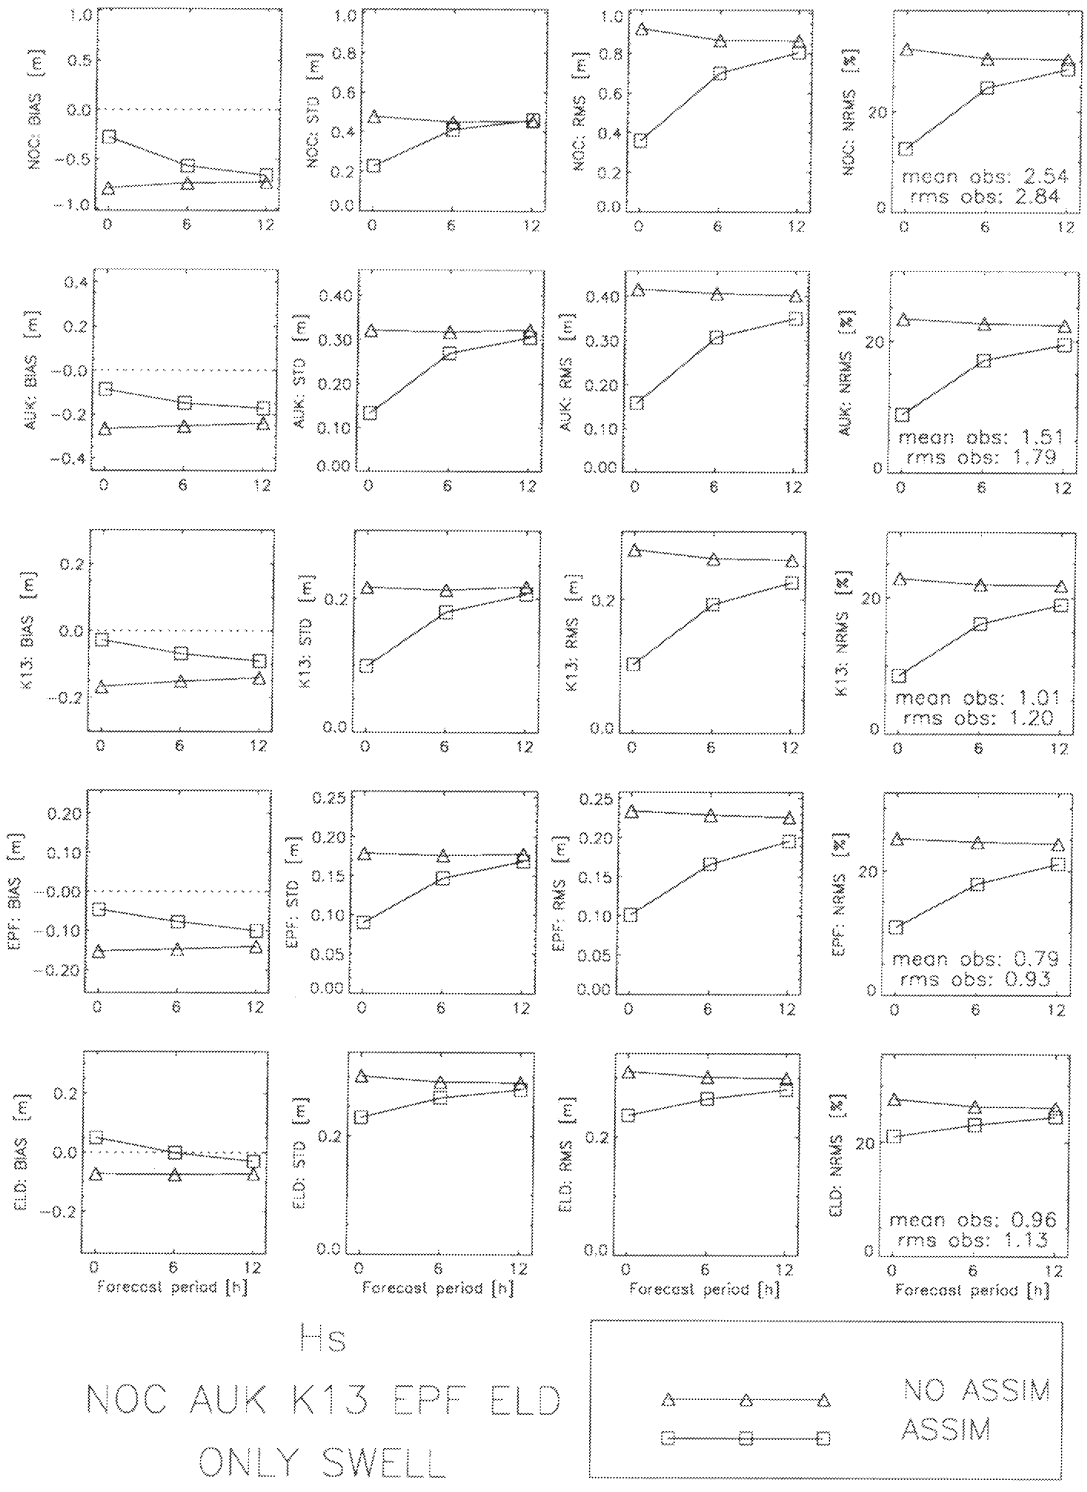

Figure 17: Statistics over swell data as a function of forecast period, for significant wave height H_s. Parameters and symbols: as in figure 16.

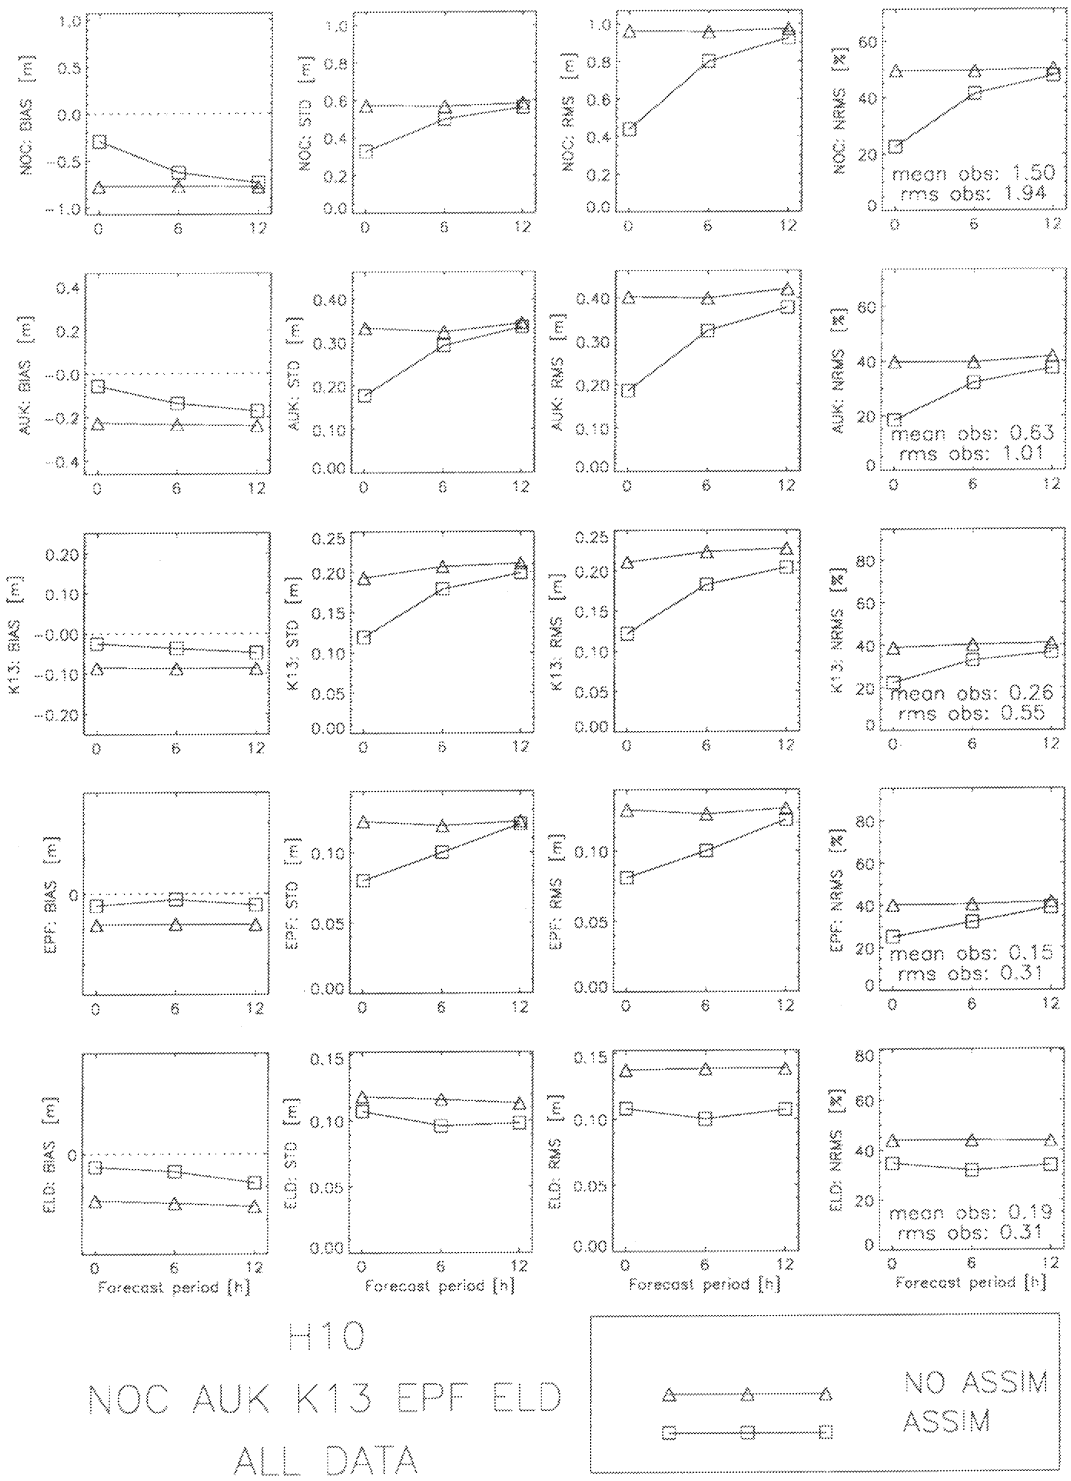

Figure 18: Statistics over all data as a function of forecast period, for low-frequency wave height H_{10}. Parameters and symbols: as in figure 16.

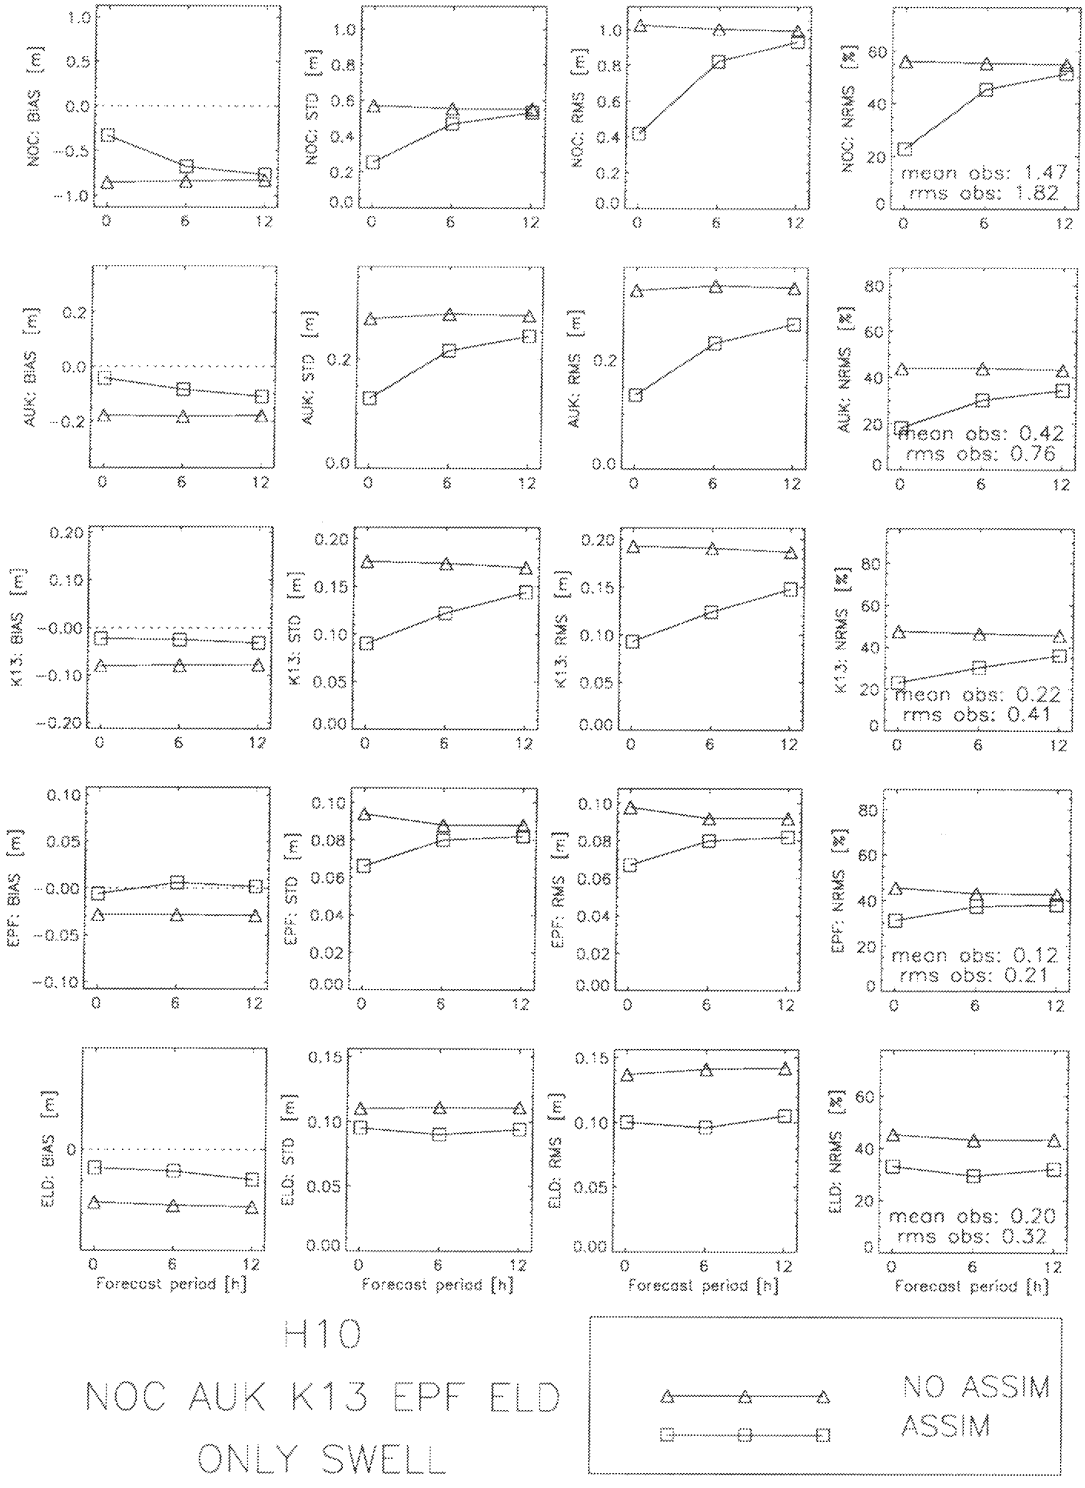

Figure 19: Statistics over swell data as a function of forecast period, for low-frequency wave height H_{10}. Parameters and symbols: as in figure 16.

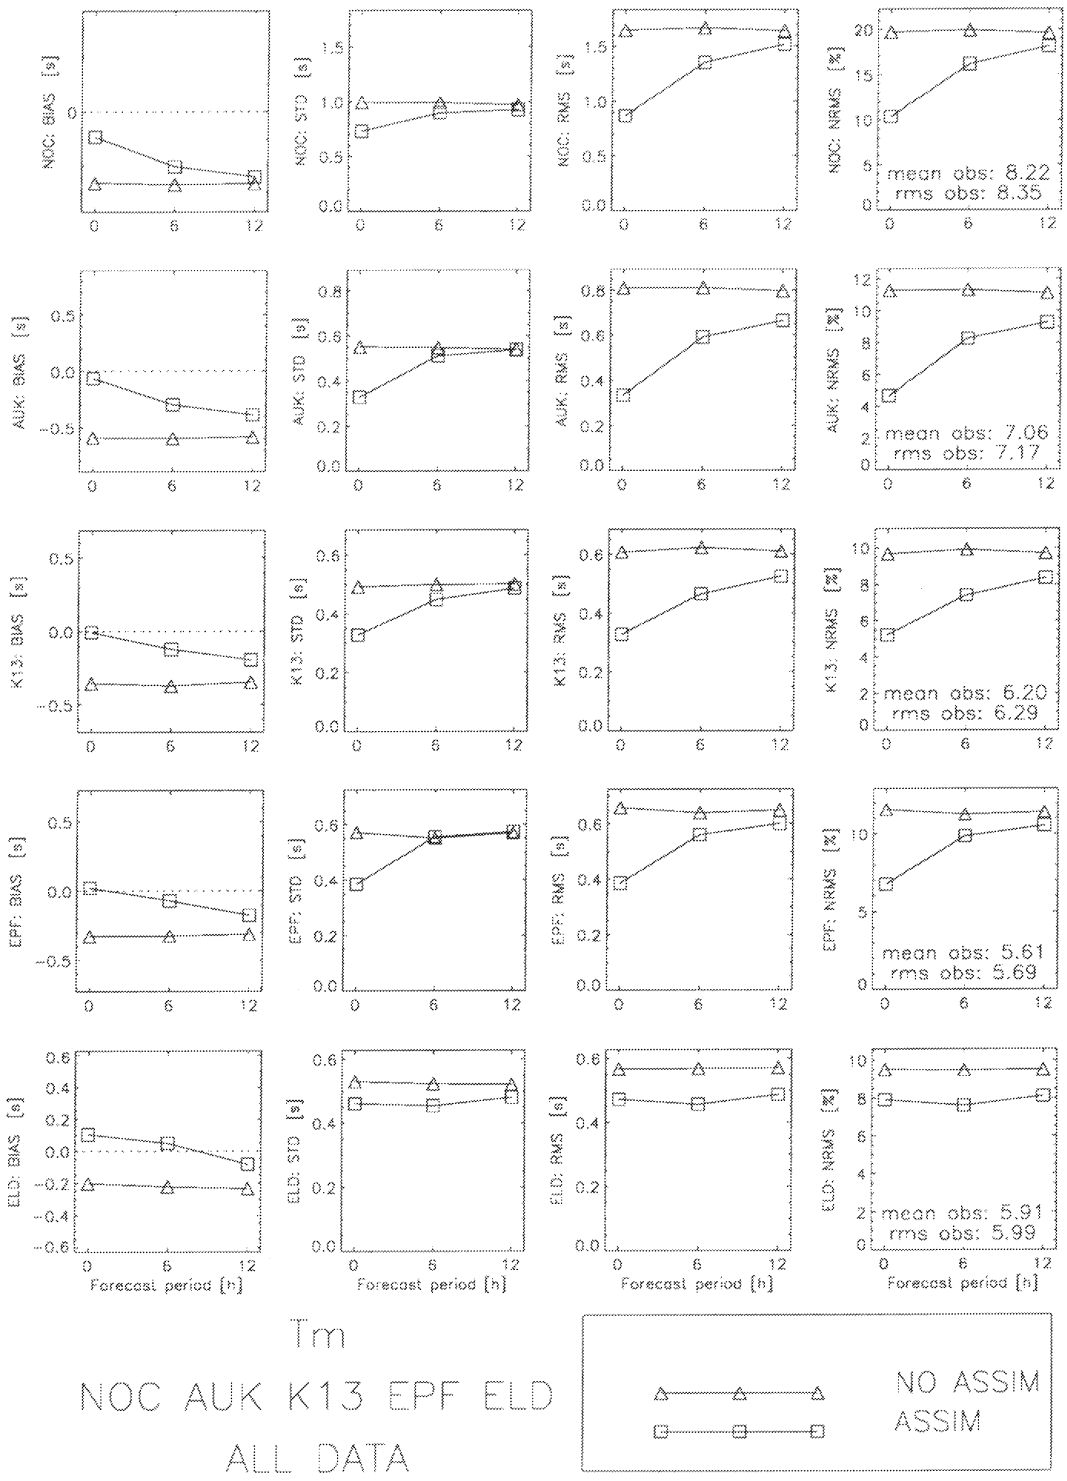

Figure 20: Statistics over all data as a function of forecast period, for mean wave period T_m. Parameters and symbols: as in figure 16.

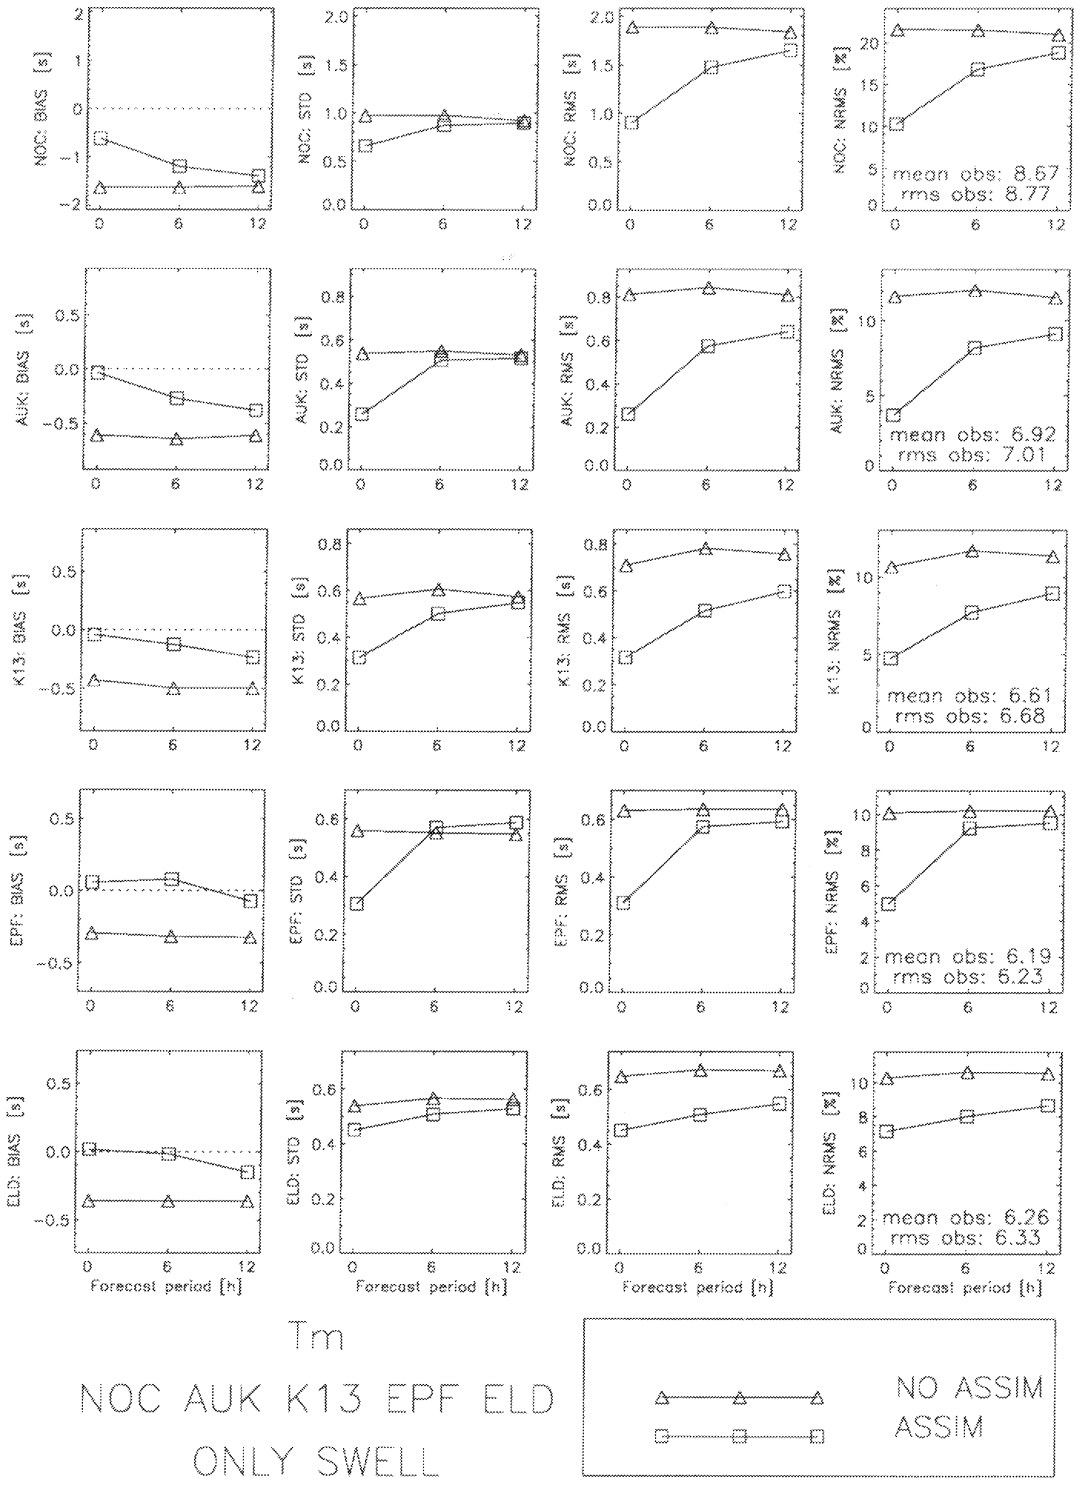

Figure 21: Statistics over swell data as a function of forecast period, for mean wave period T_m. Parameters and symbols: as in figure 16.

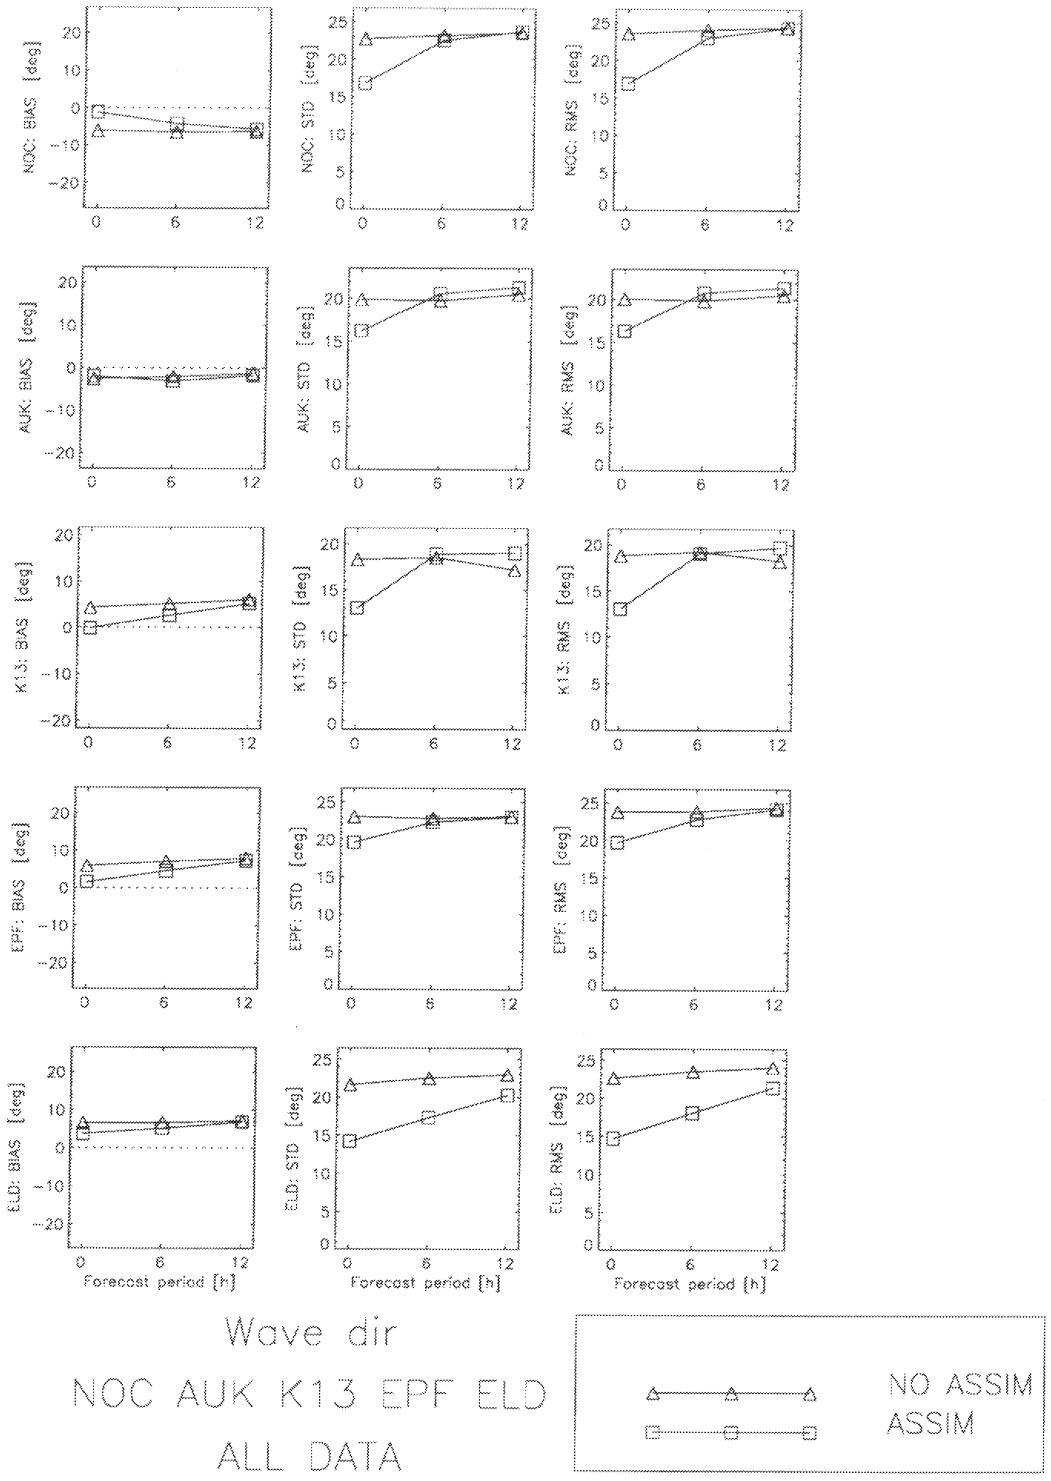

Figure 22: Statistics over all data as a function of forecast period, for mean wave direction θ_w. Boxes, runs with assimilation; triangles, runs without assimilation. From left to right: bias, standard deviation, RMS error. From top to bottom: North Cormorant, Auk, K13, Euro Platform, and Eierland.

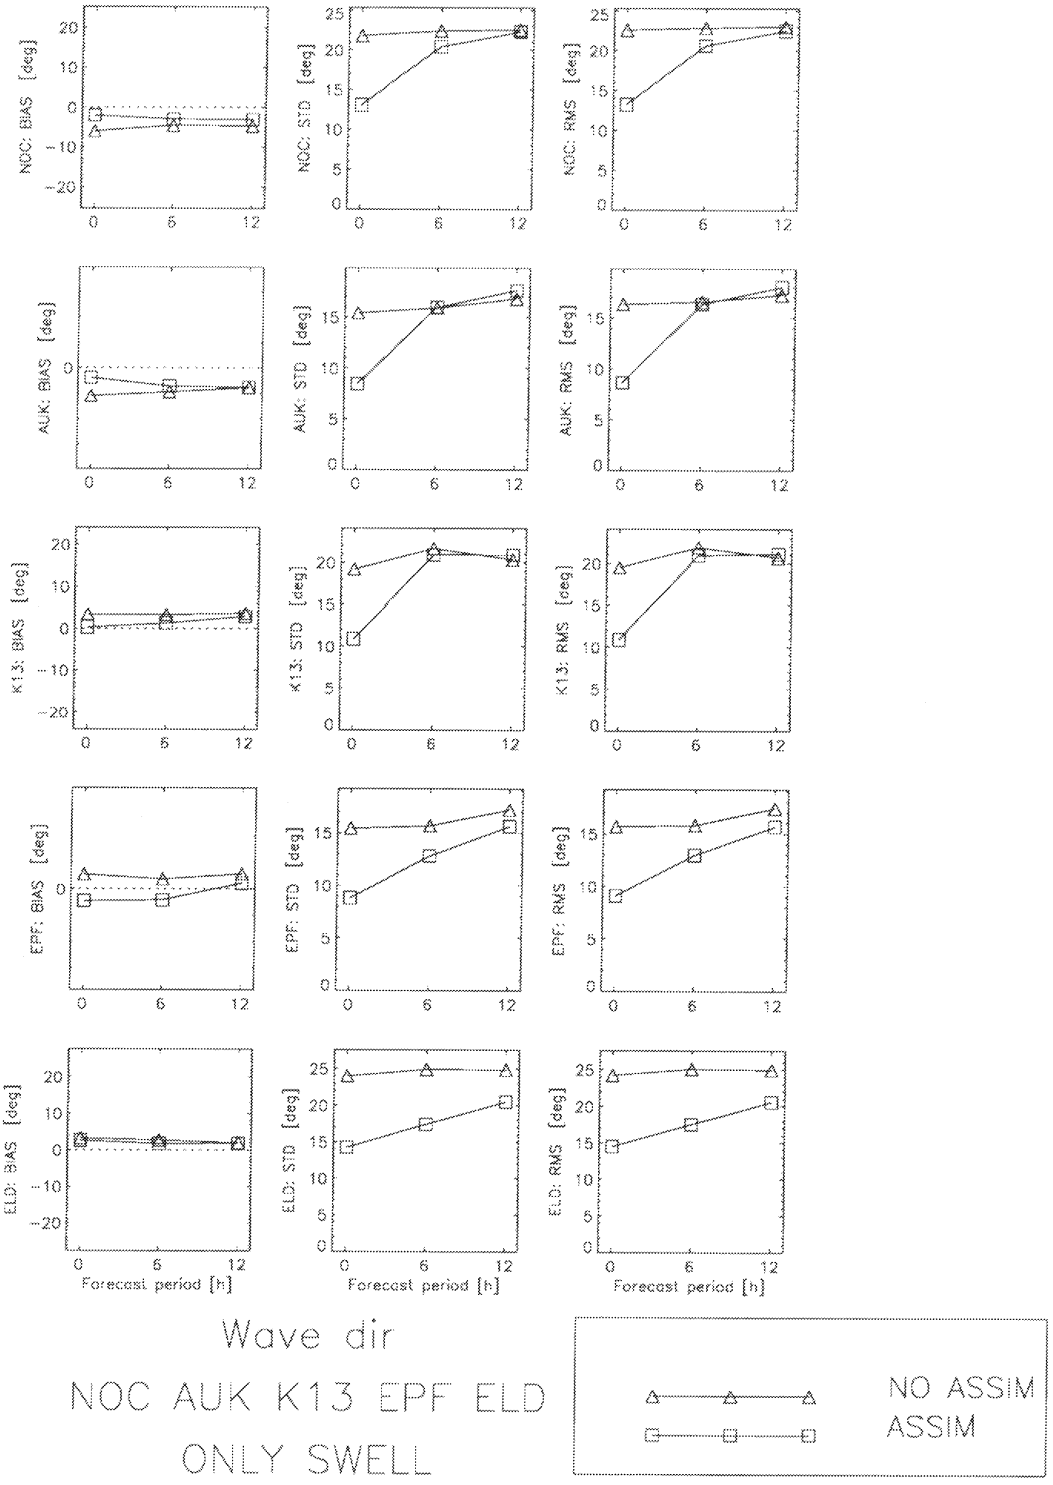

Figure 23: Statistics over swell data as a function of forecast period, for mean wave direction θ_w. Parameters and symbols: as in figure 22.