



Met Office


Forecasting Research Technical Report 566

The benefit of continuous data assimilation in a regional NWP model

by

**Clive Jones &
Bruce Macpherson**

July 2012



Abstract

Experiments to measure the benefit of the operational North Atlantic & European (NAE) model data assimilation system are reported. Three versions are compared: the operational continuous assimilation system, an intermittent assimilation system with one 6-hour NAE assimilation cycle beginning from the previous global analysis, and a 'noDA' system where each NAE forecast begins from an interpolated global analysis. Two resolutions of global analyses were available to the experiments, from versions of the global model with grid-length of either 40km or 25km at mid-latitudes.

Results from 70 forecasts over a 35 day period in winter 2010 show that continuous NAE assimilation gives significantly better screen temperature and visibility verification over the UK during the first 24 hours of the forecast. For other fields, the noDA run generally performs better, and initial conditions interpolated from the 25km global model give better results than those from the 40km global model. An intermittent NAE assimilation system performs less well, and this appears to be due to a spin-up process which follows interpolation of the global analysis at t-6 hours. The implications for the configuration of the operational NWP system are explored. Issues related to the retirement of the operational NAE model are discussed.

Contents

1. Introduction	2
2. Impact of continuous data assimilation	3
2.1 Previous work	
2.2 The present study	
3. Results	5
3.1 UK NWP Index	
3.2 Surface variables	
3.3 Upper-air fields	
3.4 Investigation of intermittent NAE assimilation	
4. Discussion, recent developments and recommendations	17
5. References	12
Figures	13

1. Introduction

The North Atlantic and European (NAE) model has been the principal model used by the Chief Forecaster in the Met Office for the issue of short range forecasts (up to t+48 hours) for the UK region since autumn 2004, when it achieved operational status (Bush et al., 2006). Since February 2006, it has been run with a grid-length of 12km. The high resolution UK models, run to supply extra detail, have had their boundary conditions supplied by the NAE.

The overall benefit of running the NAE has been measured for several years by the UK NWP Index, whose components measure the skill of forecasts for screen-level temperature, 10-metre wind, screen-level visibility, 6-hourly precipitation accumulation, total cloud amount and cloud base height. A time series of the annual mean difference between the UK NWP index derived from the NAE and that obtained from the global model is shown in Figure 1.

The 'added value' of the NAE over the global model is in the range 5-8%, with some of the variations attributable to major operational changes in either the NAE or global system. For example, the NAE benefit rises significantly over the year following implementation of 4DVAR in the NAE in spring 2006. There is a decline in NAE benefit after the two changes made together in early summer 2007, a reduction of the NAE domain size (see Figure 2) and a redefinition of the UK NWP index to include forecasts up to t+48 rather than the previous limit of t+24. The time series of NWP index by forecast time in Figure 3 clearly shows the levelling off at later forecast times after the domain reduction, whereas the contribution of short forecast periods (less influenced by the boundaries) continues to rise. The reduced domain was introduced to improve the affordability of an increase in the number of vertical levels from 38 to 70 – this eventually happened in March 2010.

Another significant change came in early summer 2008, when screen-level temperature and humidity observations were introduced to the global assimilation. Over the 40 days of Parallel Suite 18, when this change was the main difference between the global and NAE trial packages, the NAE 'added value' declined by 10-17% of its operational level at the time. A decline is not present in the following 12 months of the Figure 1 time series, which may indicate that other changes in subsequent parallel suites affected the

evolving signal. It could also be that the change in signal measured in the parallel suite period was not representative of the annual mean.

A major change was introduced to the global model in March 2010, with an increase in resolution from 40km to 25km (mid-latitude values). During Parallel Suite 23 when this was tested, the global resolution upgrade reduced the NAE's added value by 15-30% depending on verification area. The time series in Figure 1 does show a noticeable decrease over the following 12 months of operational running.

Looking within the 5 components of the UK Index, it is the visibility which contributes the largest proportion of the NAE's added value. Over a 12-month period up to spring 2010, for example, this proportion was 42% in verification area 503 covering the UK and nearby continent. This signal will come from both assimilation and model aspects: the NAE assimilates visibility observations and carries prognostic murky aerosol with emission sources (Clark et al., 2008), whereas the global model has no such assimilation and uses a fixed aerosol content in visibility diagnosis.

2. Impact of continuous data assimilation

It is of interest to isolate the separate contributions of model and assimilation to the marginal benefit of a regional NWP system relative to its driving global model. Although the regional assimilation ought to have an advantage from finer resolution, for example in analysis of features related to orographic detail, there are several reasons why a regional assimilation is at a potential disadvantage compared with a global one. Firstly, the global analysis inside the regional domain may be improved by observations outside the regional domain, which are not available to the regional analysis. Secondly, regional analysis boundary conditions enforce zero increments for scalar variables at the boundary, meaning that observations close to the boundary are not optimally used. Thirdly, the 'main run' NAE analysis is run with a shorter data cut-off (90 minutes) than the global (120 minutes) and so receives fewer observations.

2.1 Previous work

Two earlier studies have looked at the impact of continuous NAE assimilation. In 2005, Bush (*personal communication*) conducted a one month study. The NAE version used 3DVAR and covariance statistics derived from the old 12-km UK Mesoscale model, while

the global model already had a 4DVAR assimilation scheme. In the experiment, the NAE was run from an interpolated global analysis and an initial murky aerosol field set to climatology. It was found that continuous NAE assimilation gave better screen temperature and visibility forecasts, while runs from interpolated global analyses gave better 10-metre wind, cloud and precipitation verification.

In 2006, Renshaw (*personal communication*) carried out a shorter follow-up study for a 12-day period. This time the NAE assimilation used 4DVAR and covariances taken from the operational global system. It also employed long cut-off 'update' assimilation cycles as well as the shorter cut-off 'main run' cycles from which the long forecasts were started. As well as continuous NAE assimilation, and NAE runs from global analyses, a third option tested was an intermittent NAE assimilation. This began each cycle from an interpolated global update analysis at t-6 hours, followed by a single cycle of NAE assimilation and forecast. At this time, the global model still did not assimilate screen temperature and humidity observations.

Renshaw's study again confirmed that continuous NAE assimilation gave better screen temperature (Figure 4) and visibility forecasts than runs from interpolated global analyses. This benefit in visibility was matched by the intermittent NAE assimilation suite but the continuous NAE assimilation was still better for screen temperature. Intermittent assimilation was slightly better than continuous for 10-metre wind and precipitation, while impact on cloud was mixed.

2.2 The present study

The new results reported here come from a 35 day period in February/March 2010 during which Parallel Suite 23 was running with a 25km global model ready to replace the operational 40km global model. The PS23 NAE version adopted the same 70 levels which had been operational in the global model since PS22. This meant that the NAE assimilation could once more be run with covariances derived from the same set used in the global model, a consistency important when trying to assess the benefit of other aspects of the NAE's regional analysis capability. Other differences from the previous study by Renshaw were that the NAE domain was now smaller as in Figure 2 and the global assimilation now included screen temperature and humidity observations. This was anticipated to reduce the NAE's advantage in screen-level verification.

In terms of conventional and satellite observational input, the global and NAE systems were a fairly close match - the main differences being extra cloud (Renshaw and Francis 2011), precipitation (Macpherson, 2001), visibility (Clark et al., 2008) and Zenith Total Delay (total column water vapour observations) (Bennitt and Jupp, 2011) assimilated by the NAE. Of these, the ZTD data have since gone into the operational global assimilation, in March 2012.

To summarise, the various experiments run were:

CONTROL – continuous NAE assimilation

noDA25 - NAE forecast run from interpolated 25km global analysis

noDA40 - NAE forecast run from interpolated 40km global analysis

intNAE – intermittent single cycle NAE assimilation starting from interpolated t-6 25km global update analysis

Forecasts were run twice daily from 00 and 12 UTC.

3. Results

3.1 UK NWP Index

We begin by reporting (Table 1) the overall impact of the different configurations on the 5-component UK NWP Index computed over the period t+0-t+48, where the figure quoted is the difference from the value for the continuous NAE CONTROL.

	Whole NAE domain	UK
noDA40	-0.1%	-1.56%
noDA25	+0.63%	-0.45%
intNAE	+0.31%	-1.68%

Table 1 – UK Index values relative to CONTROL

The first observation from Table 1 is that improving the resolution of the global model does improve the quality of its analysis, since noDA25 outperforms noDA40. Over the UK, this is true for all 5 components of the UK index. If we compare noDA25 with

CONTROL, for verification over the UK, the extra resolution of the 12km NAE assimilation does indeed provide a modest benefit over the 25km global analysis. However, for verification over the whole NAE domain, noDA25 slightly outperforms CONTROL, which tells us that the drawbacks of a limited area analysis relative to a global one outweigh any advantage of the finer 12km resolution of the NAE when we assess performance for the whole model domain.

We note also that the intermittent system intNAE does not perform as well as either the CONTROL or the noDA25 system over the UK, although it does match the CONTROL over the whole domain. As we might expect intNAE to get the benefit of both a good large scale analysis from the global t-6 field and a good fine scale initialisation over its single cycle at NAE resolution, the performance of intNAE is a little disappointing. We investigate this a little more in section 3.4.

3.2 Surface variables

Previous studies found benefit of regional assimilation for screen temperature and visibility, so we begin by examining these variables. Figures 5a and 5b show the mean and rms **screen temperature** verification over the UK area and NAE domain respectively.

Over the UK, the continuous NAE DA has a significantly lower rms error than other configurations run from global analyses, as far as about t+24. All NAE runs have lower rms errors and smaller bias than the full global 25km system shown for completeness. It is noticeable that intNAE shows a modest improvement over noDA25. (There were actually two variants of intNAE run, one with and one without MOPS cloud data. The one with MOPS cloud has slightly lower rms errors in Figure 5a).

Over the whole NAE domain, continuous NAE DA loses its advantage at analysis time even by t+6 hours, with both noDA25 and intNAE performing better thereafter.

For **visibility**, Figure 6a shows the Equitable Threat Score (ETS) averaged over 3 thresholds of 200m, 1000m and 4000m which are part of the UK Index. Continuous NAE DA gives improved skill over the first 24 hours compared to noDA25, especially so at t+6. The intNAE DA experiment shows improvement over noDA25 for the first 18 hours. This confirms that the NAE assimilation of visibility observations is having a positive impact on the subsequent visibility forecast. In Figure 6b we examine verification of

log10 (visibility) which is dominated by the aerosol field rather than by humidity. This shows that continuous NAE DA and intNAE, both with a visibility analysis, outperform noDA25 which starts from an aerosol climatology. The benefit lasts until around t+24 over the whole NAE domain. The intNAE results show that a single assimilation cycle is enough to improve the aerosol analysis significantly.

For **10-metre wind**, there is a short-lived detrimental impact of interpolation from coarser 40/25km global grids to the 12km NAE grid as seen in the fit at analysis time. All NAE configurations verify similarly from t+6 onwards and better than the 25km global model (Figure 7). The lowest errors in day 2 of the forecast come from NAE forecasts beginning from interpolated global analyses, noDA25 and noDA40. The results shown are for the UK area, but the same is true for verification over the whole model domain (not shown).

For **fractional total cloud cover** (Figure 8), the global model clearly performs worse than the NAE and there exists a spin-up error when not running continuous NAE DA. (Although incidental to the present investigation, two flavours of intNAE were run, one with and one without assimilation of MOPS cloud cover, the former giving slightly smaller rms errors). It is also worth mentioning that the large negative bias characteristic of the operational global model in Figure 8 has changed since these experiments with the introduction of the PC2 cloud scheme in the global model. Cloud fraction bias is typically positive now.

For **6-hourly precipitation accumulation**, the global model appears to be the best performer over the whole NAE model area with light and moderate precipitation but not so good with heavy precipitation. When precipitation is heavy, noDA25 is the best performer although continuous NAE DA is not far behind. For light precipitation over the UK area (Figure 9), noDA25 and noDA40 perform best at t+6 and beyond t+24, though continuous NAE DA does best at t+12. The situation is mixed for heavier precipitation. In Figure 9, the graph of most interest is the probability of false detection in the bottom right. This measures the probability that an event forecast to occur will not actually do so. The intNAE experiment clearly shows a spin-up problem in light precipitation, with a greater over-forecasting of spurious light precipitation early in the forecast (that is reduced slightly in the run assimilating MOPS cloud data).

For **mean sea-level pressure** (not part of the UK NWP Index), the global model verifies better than the various NAE versions over the whole NAE area, with smaller negative bias and lower rms error (Figure 10). The effect of interpolation from a coarser global analysis is clearly seen at t+0, but makes negligible difference from t+6 onwards. Results for the UK area (not shown) show less difference between the full global system and the NAE experiments, although runs beginning from global analyses verify better during day 2 of the forecast.

3.3 Upper-air fields

Although the primary purpose of running a regional model is to improve prediction of *surface* weather elements, it is worth comparing the NAE's upper-air performance with the global model. We summarise this in Figures 11a and 11b, which give the vertical profile of t+24 temperature and wind errors respectively, over the whole NAE area. The 25km global model does best, with noDA40 and noDA25 runs of the NAE performing less well but better than the continuous NAE CONTROL or intNAE.

3.4 Investigation of intermittent NAE assimilation

We now look further into why the intermittent system intNAE does not perform as well as either the CONTROL or the noDA25 system over the UK, even though it does match the CONTROL over the whole domain. When running a single cycle of NAE assimilation for intNAE, one first runs a NAE forecast from a reconfigured global analysis at t-6 to produce background fields for the observation processing system and 4D-Var. These fields are valid from as early as three hours into the run and so may still be affected by a spin-up process following the interpolation from global to NAE grids. To assess this we show the total observation penalty at the start and end of the 4DVAR process for both the continuous NAE CONTROL and the intermittent intNAE (Figure 12).

It is clear that the starting total observation penalty, which measures the fit of the background to observations, is consistently higher for intNAE than in the continuous data assimilation run. This indicates that 'spin-up' errors are still present in the VAR background fields. A less accurate background field will on average give rise to worse quality control decisions and indeed it was found for example that intNAE rejected more SYNOP temperature and humidity data than CONTROL. Although the total observation penalty at the end of the 4DVAR assimilation process, which measures the fit of the

analysis to the observations, is very similar for intNAE and CONTROL, these diagnostics are consistent with a worse intNAE analysis as measured by the forecast verification.

4. Discussion, recent developments and recommendations

The overall results in terms of UK NWP Index show that model resolution is much more important than data assimilation as an explanation of the NAE's added value relative to the global model. Over the UK, the NAE assimilation contributes ~0.5% benefit, whereas the full NAE model system gives benefit of 5-8% in total and 2.5-4% if the visibility component is excluded. Indeed, over the full NAE model domain, interpolation from a 25km global analysis outperforms continuous NAE assimilation overall by ~0.5% and improves on interpolation from a 40km global analysis. These results support the strategy of **increasing global model resolution**.

Within the overall UK NWP index results, it was found that continuous NAE assimilation did still supply significant benefit for UK forecasts of temperature and visibility during day one of the forecast. In this respect the present study agrees with the earlier ones of Bush and Renshaw. An attempt to retain these NAE assimilation benefits and introduce the benefits of global analysis via an intermittent NAE assimilation was unsuccessful. Performance of the intermittent system was found to be hindered by the after-effects of interpolation from the global 25km grid to the 12km NAE grid, leading to a detriment in the background field for the single cycle of NAE assimilation.

The superior upper-air verification of the global model (or of NAE forecasts started from global analyses) relative to the continuous NAE assimilation, suggests a role for the global model as preferred source of **lateral boundary conditions for high resolution UK models**. Since these experiments were performed, encouraging results were obtained from a 4km UK 'downscaler' system starting from interpolated global analyses and driven by global model lateral boundary data (Weeks, *personal communication*). Subsequently, the full 4km UK assimilation system was tested with global boundary data and trials over about three months gave an overall 0.7% benefit in terms of UK NWP index relative to the system forced by NAE boundary data. Operational implementation of the switch from NAE to global lateral boundary forcing took place in both the 4km and 1.5km UK models on 17th January 2012.

We have highlighted the benefit of continuous NAE assimilation on a 12km grid for short-period **screen temperature** forecasts, relative to NAE forecasts from global analyses of

25km or 40km resolution. This supports the strategy of continuing to assimilate observations which are related to orographic detail when running high resolution models. So when the source of 'best UK forecast data' was switched from the NAE to the 4km UK model in April 2010, and to the 1.5km UK model in January 2012, it was recognised that enhancement of the surface observing network and improvement of fine-scale forecast error covariances would be worth pursuing. In January 2012, the UK models began to assimilate screen-level temperature and humidity from a subset of the Highways Agency roadside sensor network, and there are plans to expand the coverage of this new data source. Also, tests have begun with a new forecast error covariance generation technique which will allow shorter horizontal length scales to be imposed on the analysis increments.

The current results have confirmed once more the benefit of continuous assimilation for visibility forecasts in a model with a prognostic aerosol variable like the NAE. Running the NAE from a global analysis of atmospheric fields and an initial aerosol derived from climatology did not perform so well. This supports the retention of **visibility assimilation** within the system delivering 'best UK forecast data' which has moved from the NAE to the 4km UK to the 1.5km UK model.

One consequence of the switch from NAE to global boundary forcing for the UK models is that '**murk**' **aerosol boundary data** is no longer available, since this variable is not prognostic in the global model. The advantage of the NAE over the global model in visibility verification implies that the loss of aerosol boundary data could degrade the UK model forecasts, especially when the inflow direction is from the south-east bringing more polluted air from industrial parts of Europe. This was confirmed in one test period by Bornemann (*personal communication*), although it was not enough of a problem to degrade the visibility forecasts overall for prevailing wind direction regimes. Without aerosol in the boundary data, the UK model simply imports values at inflow boundaries which were advected there by previous outflow. This is not a satisfactory formulation. Instead, therefore, an aerosol climatology will be assembled from a long history of NAE forecasts and this climatology will be used at inflow boundary points. This should at least give improved behaviour on average.

In the light of the developments discussed above, the operational NWP system is almost ready for the **retirement of the NAE model**, scheduled to take place in 2013. The NAE's role as supplier of 'best data' for the UK NWP index has passed to the UK1.5 model and its function to supply lateral boundary data to the UK models has been taken

over by the global model. One remaining connection between the NAE and UK models comes from the use of NAE profiles within the AUTOSAT Cloud Top Height diagnosis scheme. This supplies CTH data to the UKPP cloud analysis, from which MOPS cloud data are derived for assimilation in the UK models. Work is in hand to use either global or UK1.5 model profiles within AUTOSAT instead.

As regards the **distinctive capabilities of the NAE assimilation system** relative to the global model (assimilation of precipitation rate, cloud fraction and visibility), these already form part of the UK model assimilation system. The potential to add a precipitation assimilation capability to the global model has been identified, and this would be based on variational assimilation of accumulations rather than Latent Heat Nudging of instantaneous rates. Cloud assimilation could also be added to the global model via the assimilation of GeoCloud CTH data already used by the NAE. In future, direct assimilation of surface cloud reports should be possible in models of all resolutions without the need to process them first by a retrieval of the cloud fraction profile, as in the MOPS system.

Visibility assimilation is not feasible in the global model because of the cost of 'murk' aerosol as a prognostic variable, the lack of global coverage of aerosol emissions data, and the significant extra cost of visibility assimilation (~30% in NAE 4DVAR). For UK forecast products, visibility assimilation within the UK model provides adequate benefit as discussed earlier. To meet any particular visibility forecast requirements within the NAE domain but outside the UK area, an option is to run a high resolution model (say 4km) from interpolated global analyses but with prognostic murk aerosol and emissions data. The higher model resolution relative to the NAE would help compensate for the absence of visibility assimilation.

5. References

Bennitt, G.V. and A. Jupp, 2012. Operational Assimilation of Zenith Total Delay observations into the UK Met Office Numerical Weather Prediction models, accepted by *Monthly Weather Review*. doi: 10.1175/MWR-D-11-00156.1

Bush, M., Bell, R.S., Christidis, N., Renshaw, R.J., Macpherson, B. & C.A. Wilson, 2006. Development of the North Atlantic European Model (NAE) into an operational model. Forecasting Research Technical Report No. 470.

Clark P.A., Harcourt, S.A., Macpherson, B., Mathison, C.T., Cusack, S., Naylor, M. 2008. Prediction of visibility and aerosol within the operational Met Office Unified Model. Part 1: model formulation and variational assimilation. *Q. J. R. Meteorol. Soc.*, **134**, 1801-1816.

Macpherson, B. 2001: Operational Experience with Assimilation of Rainfall Data in the Met. Office Mesoscale Model, *Meteorol. Atmos. Phys.*, **76**(1-2), 3-8.

Renshaw, R.J. and P. Francis, 2011: Variational assimilation of cloud fraction in the operational Met Office unified model. *Q. J. R. Meteorol. Soc.* **137**: 1963–1974




Figure 1 – Time series of the percentage difference between the UK NWP Index for the NAE model and the global model (NAE-global). The points plotted are a 12-month running mean. Red arrows indicate timing of a significant operational change in either the NAE or global (GL) model as discussed in the text.




Figure 2 – the original larger NAE domain and its smaller successor from spring 2007




Figure 3 – time series of the 12-month running mean UK NWP Index from the NAE model, with break down by forecast time in range $t+6 \rightarrow t+48$ as identified in the key. The point is indicated at which the smaller NAE domain in Figure 2 was introduced.

Equalized and Meaned from 10/9/2006 00Z to 27/9/2006 18Z

Cases: —+— PS12 NAE —x— Global —*— Recon —o— Recon + DA




Figure 4 – Results from the study of Renshaw (*personal communication*). The plots show mean (left) and rms (right) errors in screen-temperature as a function of forecast time over a period of about two weeks for: NAE from interpolated global analysis (green), continuous NAE assimilation (red), intermittent NAE assimilation (yellow), continuous global assimilation (blue).




Figure 5a – mean (top) and rms (bottom) error in screen temperature for the UK area averaged over the trial period, as a function of forecast time. Experiment definitions are given in the key above the upper panel. In terms of the labels used in the text, these are (in order): CONTROL, noDA40, noDA25, intNAE without MOPS cloud, intNAE with MOPS cloud, 25km global assimilation.

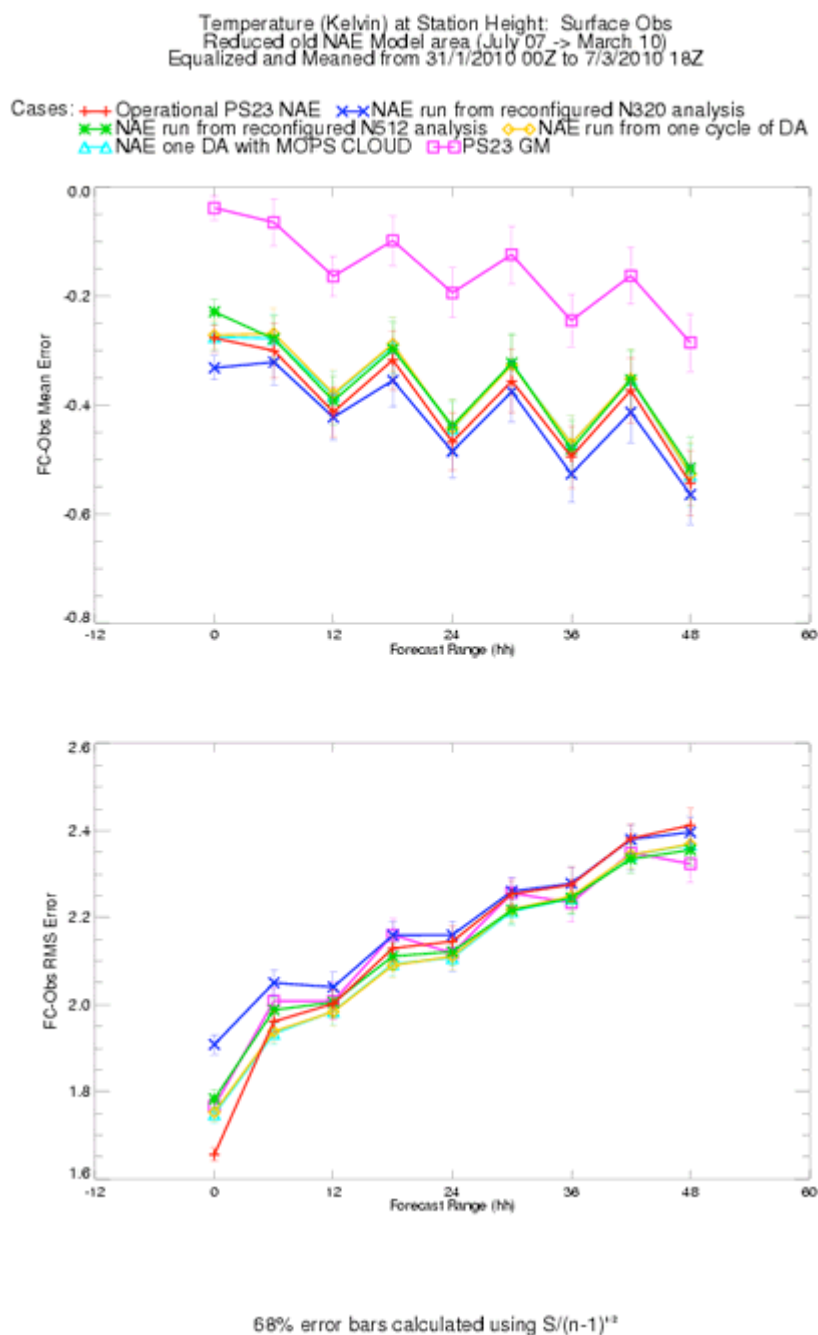


Figure 5b – as Figure 5a for the whole NAE model domain.




Figure 6a - Average Equitable Threat Score (ETS) across all visibility thresholds (200m, 1000m and 4000m) as a function of forecast time averaged over the trial period. Results shown for CONTROL ie continuous NAE assimilation (red), noDA25 ie NAE forecast from interpolated 25km global analysis (green) and intNAE ie intermittent NAE assimilation (turquoise).




Figure 6b – as Figure 5b for log10 (visibility)




Figure 7 – as Figure 5a for 10-metre wind

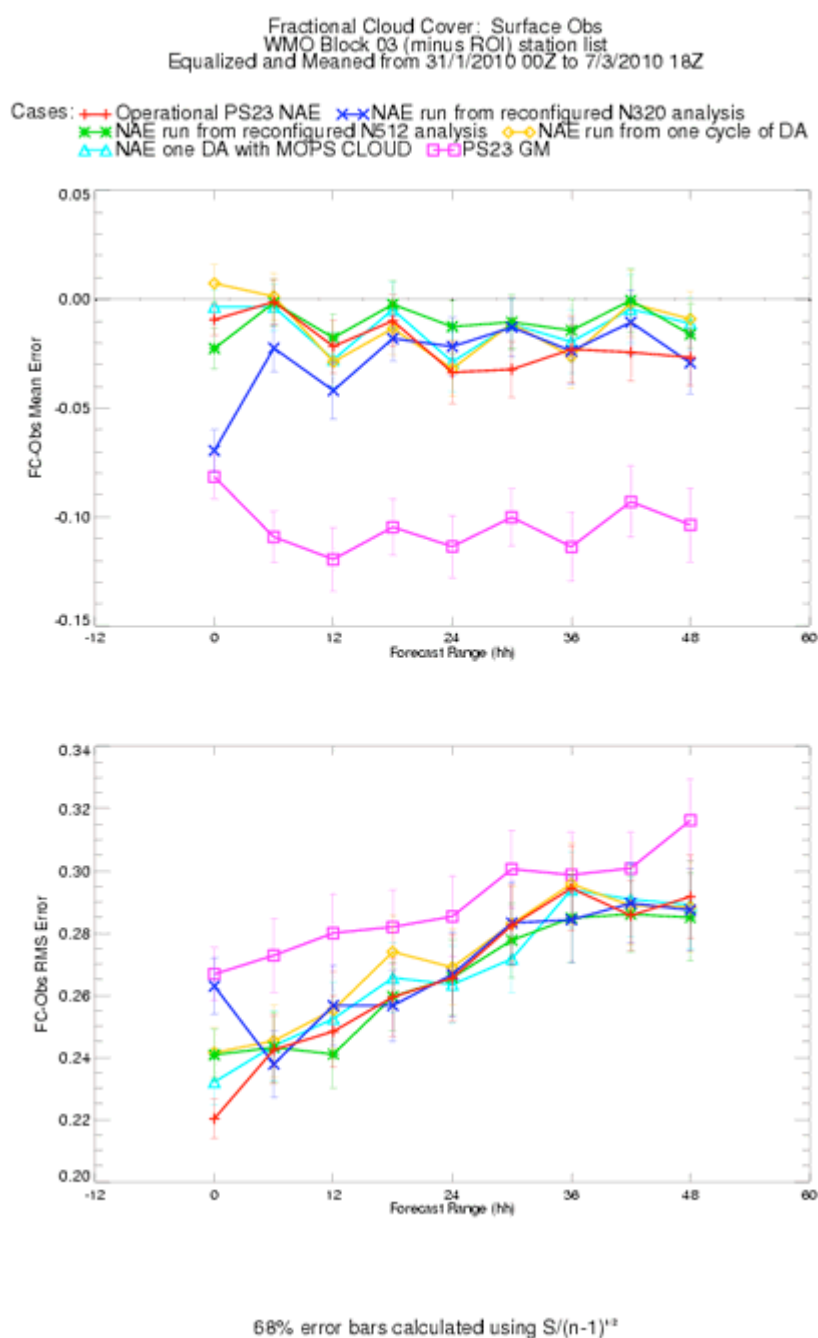


Figure 8 – as Figure 5a for total cloud fraction

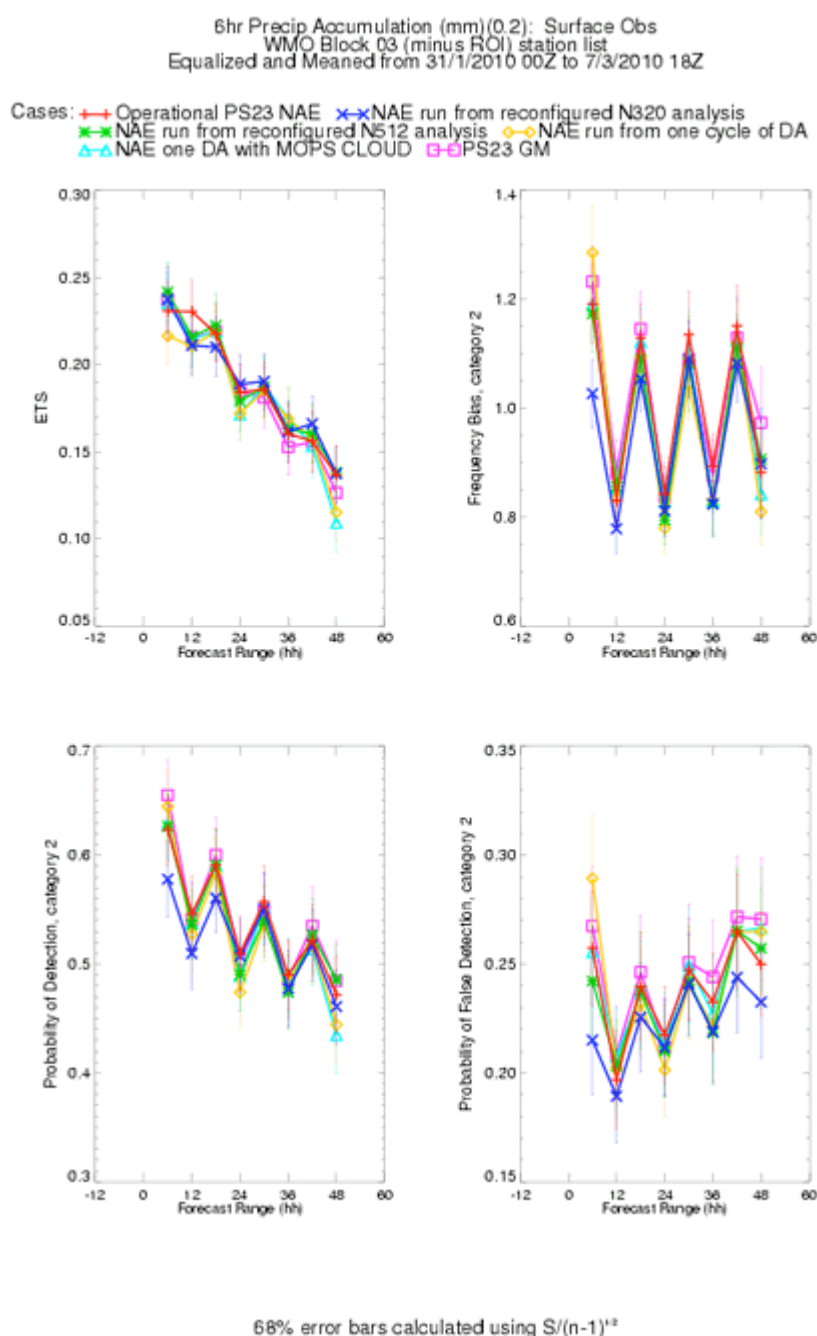


Figure 9 – skill scores as a function of forecast time for 6-hourly precipitation accumulation greater than 0.2mm. Results averaged over the trial period for UK area. Top left is Equitable Threat Score (ETS). Top right is Frequency Bias. Lower left is Probability of Detection. Lower Right is Probability of False Detection. Experiment definitions are given in the key above the upper panels. In terms of the labels used in the text, these are (in order): CONTROL, noDA40, noDA25, intNAE without MOPS cloud, intNAE with MOPS cloud, 25km global assimilation.




Figure 10 – as Figure 5b for mean sea-level pressure




Figure 11a – vertical profile of the mean (upper panel) and rms (lower panel) error in temperature with respect to radiosondes for the whole NAE area, averaged over the trial period. Experiment definitions are given in the key above the upper panel. In terms of the labels used in the text, these are (in order): CONTROL, noDA40, noDA25, intNAE without MOPS cloud, intNAE with MOPS cloud, 25km global assimilation.




Figure 11b – vertical profile of the mean wind speed error (upper panel) and rms vector wind error (lower panel) with respect to radiosondes for the whole NAE area, averaged over the trial period. Experiment definitions are given in the key above the upper panel. In terms of the labels used in the text, these are (in order): CONTROL, noDA40, noDA25, intNAE without MOPS cloud, intNAE with MOPS cloud, 25km global assimilation.




Figure 12 - Total observation penalty Jo at start and end of minimisation for 4DVAR analysis, as a function of analysis cycle number through the trial period. Results shown for continuous NAE DA (labelled PS23) and intNAE (labelled Exp).

Met Office

FitzRoy Road, Exeter
Devon EX1 3PB
United Kingdom

Tel: 0870 900 0100

Fax: 0870 900 5050

enquiries@metoffice.gov.uk

www.metoffice.gov.uk