Numerical Weather Prediction

Objective verification of SSFM coupled with MOSES II and the non-local BL mixing scheme

Forecasting Research Technical Report No. 398

Clare Bryden

email: nwp_publications@metoffice.com

Forecasting Research
Technical Report No.398

Objective verification of the SSFM coupled with MOSES II and the non-local BL mixing scheme: winter 1997/98 trial

by

C. Bryden

January 2003

Local Forecasting R&D
Development Process
The Met Office
London Road
Bracknell
RG12 2SZ
United Kingdom

Document Description

Name of document:	Objective verification of the SSFM coupled with MOSES II and the non-local BL mixing scheme: winter 1997/98 trial
Document identifier:	Forecasting Research Technical Report No.398
Issue/version:	Issue 1.0
Date of issue:	7 January 2003
Author of current version:	Clare Bryden
Maintainer of document:	Clare Bryden
Reviewer:	Phil Hopwood
Approver:	Phil Hopwood

Document History

Version number	Date issue	Author	Reviewer/approver
Issue 1.0 Draft A	3 September 2001	Clare Bryden	Phil Hopwood
Issue 1.0 Draft B	6 March 2002	Clare Bryden	Phil Hopwood
Issue 1.0	7 January 2003	Clare Bryden	Phil Hopwood

Objective verification of the SSFM coupled with MOSES II and the non-local BL mixing scheme: winter 1997/98 trial

C. Bryden

Abstract

The purpose of the Site-Specific Forecasting Model (SSFM) is to add automatic, physically self-consistent site-specific detail to NWP products at low cost. The SSFM has been continually under development, incorporating new dynamics and parametrizations, and the winter 1997/98 trial has become the Local Forecasting R&D Group's standard for testing new configurations.

As part of a project to assess the feasibility of an ensemble low cloud prediction system, the SSFM was extended with the latest non-local physical process schemes and an ensemble approach. This configuration was trialled against the current operational version of the SSFM, the previous version without the Edwards-Slingo radiation scheme, and a number of other configurations involving observations assimilation and soil moisture nudging (both undergoing separate trials).

Overall, the results indicate a problem with the formulation of surface layer blending heights in all versions of the SSFM from 4.5 onwards. This consideration apart, the configuration under trial is warmer and drier at screen level than the current operational version. In some cases, such as for the hillier and more urban sites, the forecast shows greater skill, but the overall conclusion of the winter 1997/98 trial is that more work is needed before it could replace the current operational version.

It is recommended that the impact of observations assimilation on the new configuration be studied, and a summer trial run to look at the impact of soil moisture nudging. In addition, there are two threads of SSFM development planned or underway, to make it capable of running over the sea and to complete the post-processing from UM 5.0. The latter will generate SSFM 5.0, which is likely to incorporate MOSES II and the 6A boundary layer scheme.

Contents

1. Introduction

The purpose of the Site-Specific Forecasting Model (SSFM) is to add automatic, physically self-consistent site-specific detail to NWP products at low cost.

Before they can become operational, new models or model versions must be tested to ensure that they add skill to the old model or version. Accordingly, the initial version of the SSFM was run routinely for 14 synoptic stations during autumn 1996/7, and it showed significant improvement over the mesoscale model (MES). This first trial is described in Clark et al (1997). The tile version (version 2) was initially tested using the same data, then trialled from 1st November 1997 to 31st January 1998 at these sites plus 5 OpenRoad sites.

The SSFM has been continually under development, incorporating new dynamics and parametrizations, and the winter 1997/98 trial has become the Local Forecasting R&D Group's standard for testing new configurations.

As part of a project to assess the feasibility of an ensemble low cloud prediction system (the LCP project, see Bryden, 2001), the SSFM was extended with the latest non-local physical process schemes and an ensemble approach. This report describes the results of the standard winter 1997/98 trial.

2. Model Description

2.1. The Operational Version of the SSFM

The SSFM is a 1D model, closely based on the single column version of the Unified Model (UM). It includes a detailed treatment of all the physical processes occurring in a column, from sub-surface to the top of the atmosphere, including sub-surface soil moisture and soil temperature processes, surface exchange, boundary layer mixing, cloud and precipitation processes, and radiative exchange. Synoptic and mesoscale forcing must be provided from NWP data (or constructed using simple physical assumptions).

Not all of the characteristics of a site that impact its local meteorology can be treated in a 1D framework. The SSFM concentrates on improving forecasts through three mechanisms:

- improved accuracy, with higher resolution in soil and the boundary layer and improvements to the physics
- simple corrections to forcing due to local variations in orography
- sophisticated treatment of surface exchange processes in the local, upwind fetch to take account of the local environment

The current operational version of the SSFM is based on version 4.5 of the UM.

The SSFM uses the UM soil/surface scheme (MOSES) which includes a sophisticated treatment of plant physiology to predict surface resistance to evaporation, a multi-layer soil hydrology and thermal model including soil-moisture freezing. Treatment of a 'radiative canopy', to improve treatment of heat exchange over vegetated surfaces, and an 'urban' canopy have been added to this. MOSES is run with nine soil levels in contrast with the four used in climate simulations and operational forecasts. Surface exchange is treated using Monin-Obukhov similarity theory, and the boundary layer is treated with four times the vertical resolution of the MES.

The exchange of heat, moisture and momentum with the surface over the upwind fetch determines the local near-surface profiles of temperature, humidity and wind. The SSFM modifies MOSES by treating the inhomogeneous terrain in terms of equilibrium surface exchange over a number of different 'tiles', each with homogeneous characteristics. In the SSFM the seven available tiles are: deciduous trees, coniferous trees, C3 grass, C4 grass, bare soil, urban and open water. The proportions of each tile are derived by multiplying each element of land-use in the upwind fetch by a weighting function derived from a Source Area Model. Over the UK, 25m resolution land-use data from ITE (based on LANDSAT images) are used.

2.2. The 8A Test Configuration

At the start of the LCP project, the operational version of the SSFM incorporated the 5A boundary layer scheme and 7-tile surface exchange scheme (MOSES I), and did not include the Edwards-Slingo radiation scheme (ES). ES was incorporated in the operational version (SSFM 4.5) during the course of the project.

In the UM, development of the 5A scheme had taken two paths: the 6A scheme was developed by Adrian Lock (Lock et al, 2000) based on non-local mixing and six boundary layer morphologies; the 7A scheme was based on 5A, but incorporated the nine-tile surface exchange scheme (MOSES II). These two new schemes were brought together in 8A, which exists as a set of modifications on the 7A scheme.

The nine tiles available in MOSES II are: deciduous trees, coniferous trees, C3 grass, C4 grass, shrubs, urban, open water, bare soil and ice. Surface temperature, net radiation, heat fluxes, canopy moisture, snow mass and snowmelt are calculated for each tile within a grid box. Air temperature, humidity and windspeed on model levels above the surface, and soil temperature and moisture content below the surface are treated as homogeneous across a grid box.

The science behind the 6A boundary layer scheme is described in Lock et al (2000). It includes a representation of non-local mixing (driven by both surface fluxes and cloud-top processes) in unstable layers, either coupled to or decoupled from the surface, and an explicit entrainment parametrization. The scheme is formulated in moist conserved variables so that it can treat both dry and cloudy layers.

The 6A boundary layer scheme and MOSES II, together with ES, were incorporated into a test configuration of the SSFM (SSFM 8A). SSFM 8A was tested against a number of other SSFM configurations:

- 45std, the old operational version without Edwards-Slingo
- 45esr, the current operational version with Edwards-Slingo
- 45esy, the current version enhanced with assimilation of synoptic screen-level observations (also under trial)
- 45esm, the current version enhanced with soil moisture nudging (also under trial)
- 45abw, the current version with both assimilation and soil moisture nudging

Assimilation

NWP data is used to constrain the SSFM to the general synoptic evolution through the course of the model run and this includes in it an element of data assimilation. Work is underway to assimilate observational data into the operational SSFM more explicitly for a particular site. The assimilation used in the 45esy and 45abw runs was a test version, and is still under development. Rather than initialising the model from the T+6 dump from the previous run, the model is initialised from the T+0 dump. The current run is then 'spun up', incorporating observations from T-6 to T+0. The effect of the observations is washed out after only a few hours of the run, and this is the area in which development is concentrating.

Soil moisture nudging

Soil moisture is controlled by the driving data from the NWP product. For the MES, this is calculated using the MORECS system. A scheme (Best and Maisey, 2002) has been developed to determine errors in forecast temperature and specific humidity in comparison to observations and this data is then used to calculate an increment to the soil moisture, which is applied to the model.

3. Trial Description

The primary objective of the trial was to show that the incorporation of MOSES II and the non-local boundary layer scheme improves the performance of the SSFM over the current operational version, and preferably any of the other configuration under test. Some of the errors derive from the MES forcing, but all configurations were run using the same forcing, so these errors would be discounted in a comparative study.

Table 1 : Winter 1997/98 trial sites.

Site	*Site Code*	*Location*	
Glasgow Airport	3140	55.87°N	4.43°W
Manchester Airport	3334	53.36°N	2.28°W
Waddington	3377	53.17°N	0.52°W
Cranwell	3379	53.03°N	0.50°W
Coningsby	3391	53.09°N	0.17°W
Wittering	3462	52.61°N	0.46°W
Aberporth	3502	52.14°N	4.57°W
Birmingham Airport	3534	52.45°N	1.75°W
Church Lawford	3544	52.36°N	1.33°W
Northolt	3672	51.55°N	0.41°W
Odiham	3761	51.24°N	0.94°W
Beaufort Park	3763	51.39°N	0.78°W
Heathrow	3772	51.48°N	0.45°W
Gatwick Airport	3776	51.15°N	0.19°W

The model was tested for the fourteen sites listed in Table 1. However, there are a lot of results for fourteen sites, so four sites were chosen as representing different orography and land-use (see Figure 1), and the analysis focuses on these:

- Glasgow Airport – hilly, urban
- Coningsby – very flat, crops or grass
- Beaufort Park – rolling, mixed land-use

- Heathrow – fairly flat, very urban

The model was run over the 3 months from November 1997 to January 1998. According to the Weather Log of the Royal Meteorological Society, the weather for this period was:

- November 1997 – Warm and wet
- December 1997 – Changeable and rather mild.
- January 1998 – Mild and stormy till the 20th, then drier and colder

3.1. Running the Trial

The model trial was run using the standard set-up on the T3E, with routine transfer of the model output to the HP. The Fortran was originally developed on the HP, so it had to be slightly reconfigured for the T3E, and copied up and compiled on the T3E.

The trial was run over 22 October 1997 to 31 January 1998, with start times at 00Z, 06Z, 12Z and 18Z. The October runs acted as 'spin up', and only November-January output was analysed. The run length was 18 hours, constrained by the MES forcing data.

The first run was initialised from the MES dump, with subsequent runs initialised on the T+6 output of the previous run. Ancillary data files were not needed, as the data were available from the MES dump. MES forcing data for the trial period are held on-line on the T3E. Land-use data for each site had to be converted from 7 to 9 tiles. The 9-tile data were held on the HP.

Model output was copied back to the HP, and held on-line for the short term.

3.2. Analysing the Output

The model output was analysed and compared with observations using a standard series of Unix scripts and PV-Wave programs.

For each run hour in 00Z, 06Z, 12Z and 18Z and each month in November 1997, December 1997 and January 1998, a Unix script was run to extract and store the five verification diagnostics:

- screen temperature
- screen dew point temperature
- screen visibility
- 10m wind speed
- 10m wind direction

These diagnostics were then compared with observations, and bias and rms errors calculated. PV-Wave was used to generate objective verification statistics such as bias and rms errors, with a number of options for combining and comparing results:

- one run time (from 00Z, 06Z, 12Z or 18Z), for the whole trial period or month by month
- a combination of run times, for the whole trial period (month by month option added recently)
  - by validity time, intended to show any diurnal dependence of the errors – the output is repositioned, according to the time of initiation of the run, before it is combined
  - by forecast time, intended to show any deterioration over the forecast period – the output is combined without repositioning
- comparison of multiple sites (one configuration)

- comparison of multiple configurations (one site, or a combination of sites, but not multiple sites)

4. Results and Assessment

4.1. Screen Temperature

Comparison of Configurations: by Validity Time

Figure 2 shows the mean and rms errors in screen temperature, averaging over all sites and displaying the results by validity time. See (2.2) for an explanation of configurations and their codes used in the figures.

The mean errors in all configurations exhibit a diurnal dependency, with maxima at around 06Z, and minima at around 18Z. 8A shows a warm bias, especially over 00-12Z, though the absolute error still shows an improvement over the old operational version (45std). Furthermore, the error over 12-24Z shows an improvement over the cold bias in the other configurations.

The rms errors exhibit some diurnal dependency, with the 8A and 45esr/esm configurations showing better skill during the day than the night. There is little to choose between these configurations, but all three show greater skill than 45std during the day. The configurations incorporating assimilation (45esy/abw) show downward spikes every 6 hours; assimilation gives markedly lower errors at the beginning of forecast period (see the analysis by forecast time below), as would be expected since relaxation to the MES basis state dominates.

Comparison of Configurations: by Forecast Time

Figure 3 shows the mean and rms errors in screen temperature by forecast time.

As expected, the mean errors in all configurations increase with forecast time. In the case of 45esy/abw it is not until T+3 that the effects of assimilation are washed out, and the less-pronounced effects of initialisation are also seen in the other configurations.

All configurations show an improvement over 45std. 8A again shows the warm bias, which is close to zero earlier in the forecast period but worsens. The other configurations show a cold bias earlier, and are very close to zero over T+12 to 18.

The rms errors also broadly increase with forecast time. Again all show an improvement over 45std, but there is not much to choose, even though the assimilation runs are obviously better to T+3.

Focus on Sites

Figure 4 shows the spread across all sites of errors in configuration 8A. Figures 5-8 show the errors for the four representative sites. Table 1 and Figure 1 give an explanation of the site codes used in the figures.

For Glasgow Airport (3140), the 8A configuration shows strong diurnal variation in bias. There is an improvement in the forecast in the evening, but the warm bias in the morning represents a worsening in the forecast. The bias plunges during the middle of the day, possibly because the blending height does not represent the local orography and there is too much mixing. All configurations show strong diurnal variation in rms, with little to choose between them (excepting 45std, which is the worst performer).

For Coningsby (3391), all configurations show a consistent warm bias, apart from 45std which does not use the ES radiation scheme. 8A performs poorly in that it shows the highest rms, except in the early afternoon, and the highest bias. All configurations show a diurnal pattern, possibly owing to differences in land-use patterns. Coningsby is a flat, arable farmland region, and the land-use is primarily crops (C3 grass), with some woodland to the north. In the model, the tile fractions do not vary with season, while in winter the actual surface is more likely to be bare soil. Hence, differences in surface exchange would lead to a greater diurnal variation in temperature, i.e. a colder bias to screen temperature during the day, and a warmer bias during the night. This could also be the explanation of the diurnal variation shown for Glasgow, as C3 grass is also the primary land-use fraction.

The results for Beaufort Park (3763) again show diurnal variation in bias, possibly for a similar reason; Beaufort Park is in a forest clearing and the model would not account for leaf-drop in winter. 8A performs slightly better despite again showing a consistent warm bias. It shows the least diurnal variation, and produces the best forecast in the early afternoon, when the other configurations show cold biases. 45std is consistently too cold.

The results for Heathrow (3772) almost all exhibit a cold bias. The model's treatment of the urban canopy may not be capturing all of the urban heat island effect (Glasgow Airport also shows this tendency to cold bias, though to a lesser extent). However, the heat island is primarily a storage phenomenon, giving higher nighttime minima, and this is not really borne out by the diurnal profile in the bias in mean temperature.

Heathrow is the only site for which 8A shows a cold bias. Because 8A is warmer than all the other configurations, it produces much the best forecast. However, the bias is particularly small in the early morning, whereas the rms is lowest in the afternoon. It is possible that in the morning, the forecast is correct for the wrong reason, while in the afternoon the correct physical processes are too efficient at reducing the screen temperature. This may be related to the canopy and surface exchange, or the choice of blending height.

Interim Conclusions

Broadly speaking, the incorporation of the Edwards-Slingo radiation scheme into the operational version of the SSFM has tended to warm the forecast, and the introduction of the 6A boundary layer scheme and MOSES II has tended to warm it further. Assimilation and soil moisture nudging affect the forecast far less. The assimilation of observations from T-6 to T+0 is washed out after 3 hours by the MES forcing and atmospheric mixing. Soil moisture nudging is more significant in the summer. Hence 45std shows the coldest bias, 8A the warmest, and the other configurations are much closer together.

Overall 8A tends to perform well in the afternoon and evening, and less well than the other configurations excluding 45std in the morning. The performance does seem to depend on site characteristics, giving an improved forecast for urban and hilly areas, and a worse for flat vegetated areas.

4.2. Screen Dew Point Temperature

Comparison of Configurations

Figures 9 and 10 show the mean and rms errors in dew point, averaging over all sites. The first thing to note is that all configurations show a positive bias, which in most cases is combined with a zero or negative bias in screen temperature, i.e. the forecasts are moister than actual.

Considering validity time, 8A is unique in showing no diurnal pattern. The other configurations show a cooler bias at night and warmer during the day, with 45std performing the best. 8A and 45std are roughly comparable over the whole 24 hours. The bias in screen dew point in 8A is close to the bias in screen temperature, which implies that 8A, alone among the configurations, is diagnosing moisture quite well. Although 8A is too warm, the new surface exchange and boundary layer formulations improve the latent / sensible heat partition of energy.

There is a spike in rms at 12Z for the non-assimilation runs. Further investigation revealed that this appears in the analysis for Manchester Airport in all configurations, and is caused by a problem with January observations.

The comparison by forecast time (Figure 10) shows a strong deterioration in forecast performance by all configurations. 8A and 45std are very close and substantially more accurate than the other configurations. The picture is less clear for rms, although 8A and 45std still perform relatively well.

Focus on Sites

Figure 11 shows the spread of sites for the 8A configuration, and Figures 12-15 focus on the four representative sites for a comparison of configurations.

For Glasgow Airport, all configurations except 8A show a warm, fairly flat bias throughout the day. 45std is the best of these forecasts. The profile for 8A is quite different from all the other sites, showing a sharp drop in bias during the day, as in screen temperature. Nevertheless, 8A is probably the best forecast overall, borne out by the rms error.

For Coningsby, there is a diurnal pattern in all configurations, although 8A shows less of an increase in warm bias during the day. 8A is relatively accurate, contrasting its performance for screen temperature, but 45std shows the least bias of all configurations.

Beaufort Park and Heathrow show a similar pattern of diurnal variation, with an increase in warm bias during the day except in 8A. In Beaufort Park's case 8A shows the least bias, while 45std still outperforms the rest. For Heathrow, 8A shows the least bias during the day, and 45std during the night, and 8A shows the lowest rms errors (discounting the effects of assimilation before T+3).

Interim Conclusions

Comparing 45esr with 45std, the Edwards-Slingo radiation scheme lessens the cold bias and improves the forecast of temperature, but increases the warm bias and worsens the forecast of dew point. However, ES should have little effect on moisture, and the dew point depression (the difference between temperature and dew point) increases only slightly.

Under the 8A configuration, however, there is a warmer bias in screen temperature and a colder bias in dew point than under 45esr, implying that the boundary layer is drier. This should imply improved visibility, showing up as a greater bias.

4.3. Screen Visibility

Comparison of Configurations

Figures 16 and 17 show the mean and rms errors in visibility, averaging over all sites. Clearly none of the configurations perform particularly well. Biases are of the order of ±2 km, while the rms error goes as high as 14 km; but it notoriously difficult to forecast visibility.

As expected, given that it produces a drier boundary layer, 8A shows a greater bias in visibility than 45esr. The bias under 8A is positive throughout the day, possibly slightly higher at night. The bias under 45esr is around zero at night, but is negative during the day, and is an improvement on 45std. It is open to debate which of 8A and 45esr produces a better forecast. For example, when forecasting fog it is probably safer to have a higher false alarm rate (likely under 45esr) than a higher miss rate (likely under 8A).

The comparisons by forecast time (Figure 17) show how initialisation of moisture is washed out. The bias drops sharply until T+3, and then increases slowly over the rest of the forecast period.

Focus on Sites

Figure 18 shows the spread of sites for the 8A configuration. The forecast for Aberporth, with a bias of 10 km, is particularly bad. A strong negative bias in the 10m wind also indicates that there is too much drag and too much mixing. Aberporth is located on a west-facing cliff top, and the choice of blending height is likely to be the underlying factor in these forecasts.

The forecasts for Heathrow, Manchester and Glasgow Airports, with biases of –5 km, are of less concern. Compared with the other configurations (Figures 19-22 show the plots for the four representative sites) 8A is relatively good for Heathrow and Glasgow Airport. The biases are less negative, and the rms errors lower.

For all four sites, in fact, the 8A bias is higher than the other configurations, consistent with the average across sites. For Coningsby and Beaufort Park, this means that the 8A forecast has a greater positive bias, and therefore lower skill. The plots of rms errors bear this out.

4.4. 10m Wind Speed and Direction

Comparison of Configurations

Figures 23-26 show the mean and rms errors in wind speed and direction.

Biases in wind speed are positive for all configurations. That is, there is too little drag compared with actual. 45std shows the lowest biases, followed by 8A, but 8A shows the highest rms errors. The bias in 8A shows only a slight diurnal variability, but the other configurations, especially 45std, show higher biases at night and lower biases during the day. In the treatment of nocturnal stable layers in the SSFM, an asymptotic profile has been imposed in order to avoid a cut-off in mixing. This may be the cause of the increase in nighttime bias. Alternatively, as has already been mentioned, it could be the result of the choice of blending height.

There is little difference between the errors in wind direction across the configurations. All show a positive bias. That is, the model wind tends to veer compared with actual.

This is consistent with lower drag and higher wind speed; the wind is closer to geostrophic.

Focus on Sites

Figure 27 shows the spread of sites for the 8A configuration, with Aberporth's strong negative bias and poor rms error standing out strongly. Beaufort Park and Church Lawton are the other sites showing high rms errors, though positive bias.

Figures 28-31 focus on the four representative sites for a comparison of configurations.

Set against the other configurations, the 8A forecast for Beaufort Park is comparatively good, and 45std also outperforms the others. But all show too little drag, probably because the site is well sheltered in a forest clearing.

However, 8A forecasts wind speed at Heathrow and Coningsby badly, showing substantially higher bias and rms error than all other configurations. 45std again performs relatively well. At Glasgow Airport, where the terrain is much more hilly, 8A performs much better owing to changes in the treatment of roughness length and blending height in the 6A boundary layer scheme.

5. Conclusions and Recommendations

Overall, the results indicate a problem with the formulation of surface layer blending heights in all versions of the SSFM from 4.5 onwards. There is a need to investigate the possibility that the orographic and level 1 blending heights have been transposed.

This consideration apart, 8A is warmer and drier at screen level than the current operational version of the SSFM, 45esr. In some cases, such as for the hillier and more urban sites, the forecast produced is an improvement. 8A performed particularly well for Glasgow Airport. However, 8A does not cope well with flat, vegetated sites, such as Coningsby.

In terms of visibility, which is badly predicted by the current model, 8A tends to worsen the forecast overall by increasing the bias. For some sites, where 45esr underestimates visibility, this does mean that 8A performs better. But the worst forecasts for visibility under 8A are for those flatter sites where temperature is poorly predicted.

The overall conclusion of the winter 1997/98 trial, therefore, is that SSFM 8A as it stands at present could not replace the current operational version. However, it would be interesting to look at the impact of observations assimilation on 8A, and soil moisture nudging in a summer trial. In addition, there are two threads of SSFM development currently planned or underway, to make it capable of running over the sea and to complete the post-processing from UM 5.0. The latter will generate SSFM 5.0, which is likely to incorporate MOSES II and the 6A boundary layer scheme.

Finally, there are two issues, arising from the model development phase rather than analysis of the winter 1997/98 trial, which will need to be resolved in SSFM 8A/5.0.

Firstly, the implementation of MP Section's version of MOSES raised concerns over the way in which the scheme does the calculations over tiles. In order to economise on processor time, MOSES II ignores tiles that are not in the fetch, as the fetch is the grid-box aggregate and constant with respect to wind direction. This becomes a problem in a site-specific model as the fetch is allowed to change with wind direction. Some variables (e.g. surface temperature by tile) depend on their lagged value. This problem was solved for 8A in the short term by cutting out the economy and performing calculations over all tiles.

Secondly, some of the results for MOSES diagnostics are poor. The results for leaf area index and canopy height depend largely on initialisation data, but there may be a problem with canopy capacity and other diagnostics. This was given a low priority for the ensemble work but should nevertheless be addressed in an operational version of the model.

6. References

Best, MJ and Maisey, PE, 2002: 'A physically based soil moisture nudging scheme', Hadley Centre Technical Note (HCTN) No.35.

Bryden, C, 2001: 'The Feasibility of an Ensemble Low Cloud Prediction System', Forecasting Research Technical Report No.361, August 2001.

Clark, PA, WP Hopwood, MJ Best, CC Dunlop and PE Maisey, 1997: 'Assessment of the first version of the Site Specific Forecast Model', Forecasting Research Division Technical Report No.213, June 1997.

Lock, AP, AR Brown, MR Bush, GM Martin and RNB Smith, 2000 : 'A New Boundary Layer Mixing Scheme. Part I: Scheme Description and Single-Column Model Tests', Monthly Weather Review, 3187-99, September 2000.

Figure 1: Characteristics of the four representative sites.

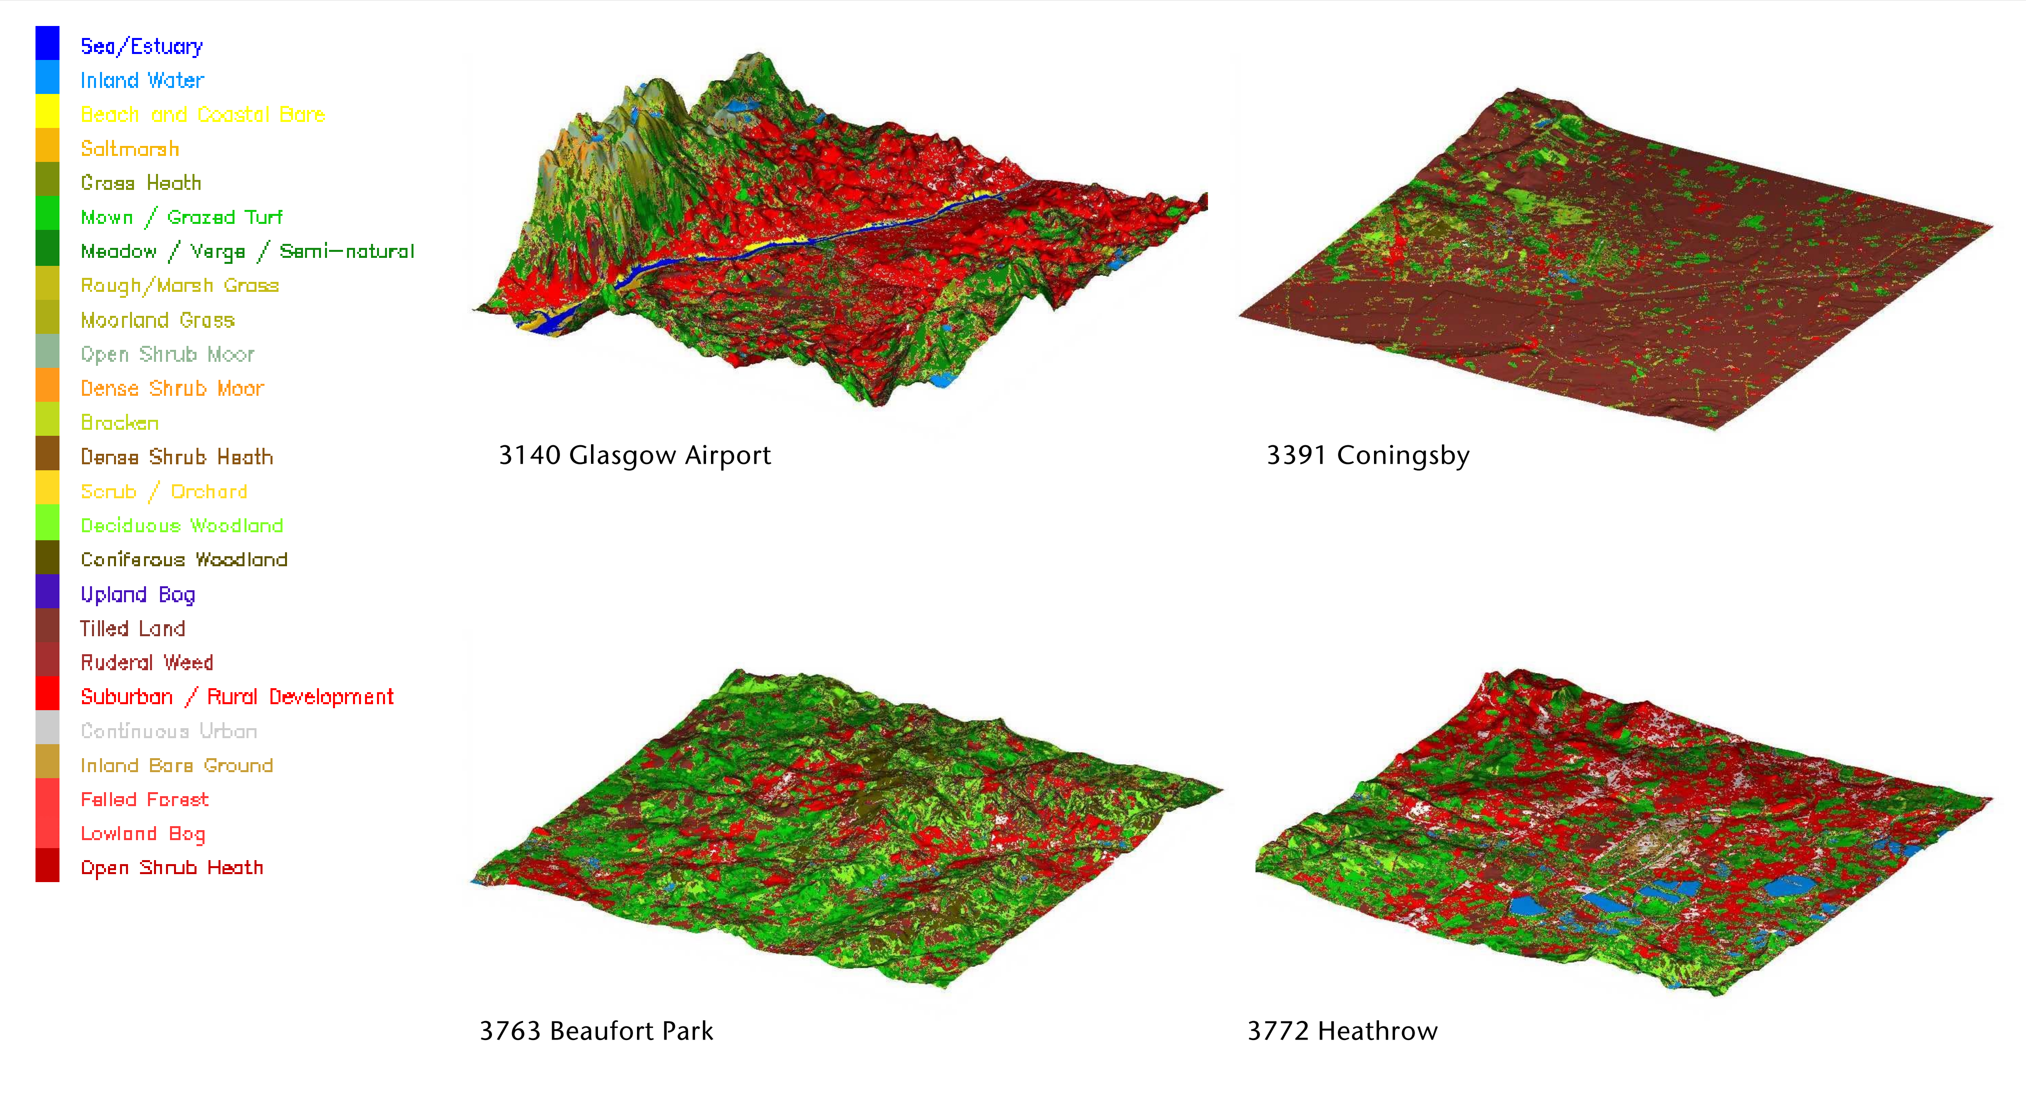

Figure 2: Screen temperature, comparison of configurations, validity time, sites combined

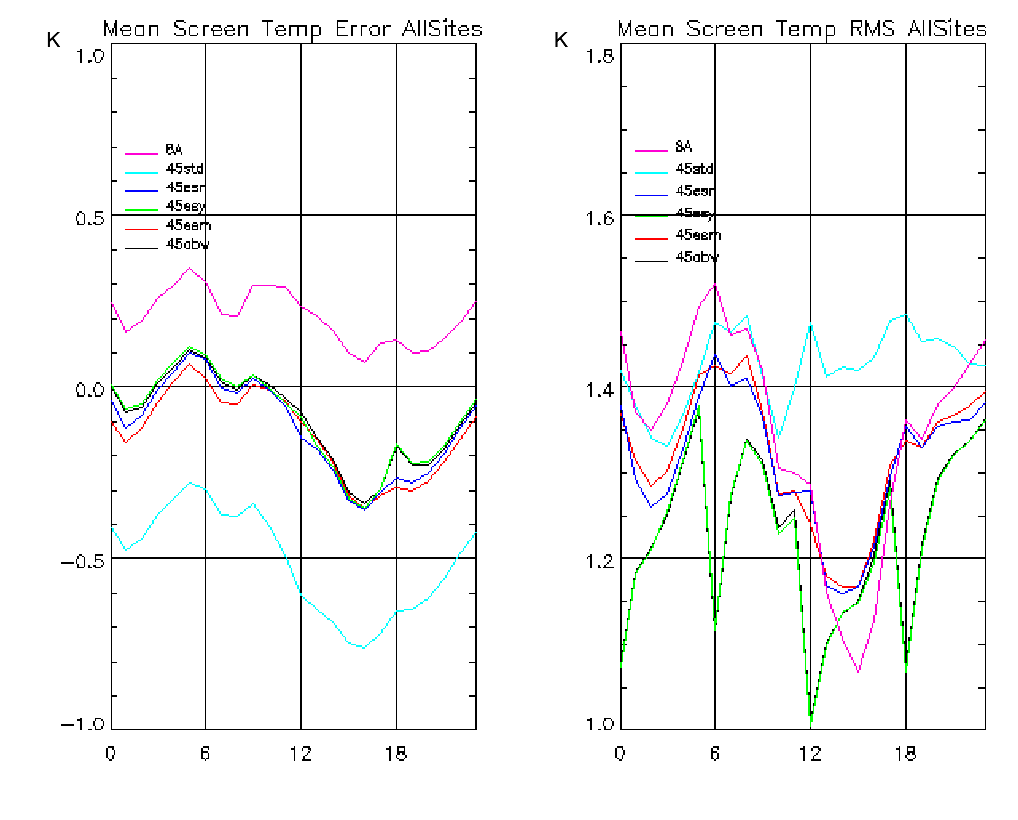

Figure 3: Screen temperature, comparison of configurations, forecast time, sites combined

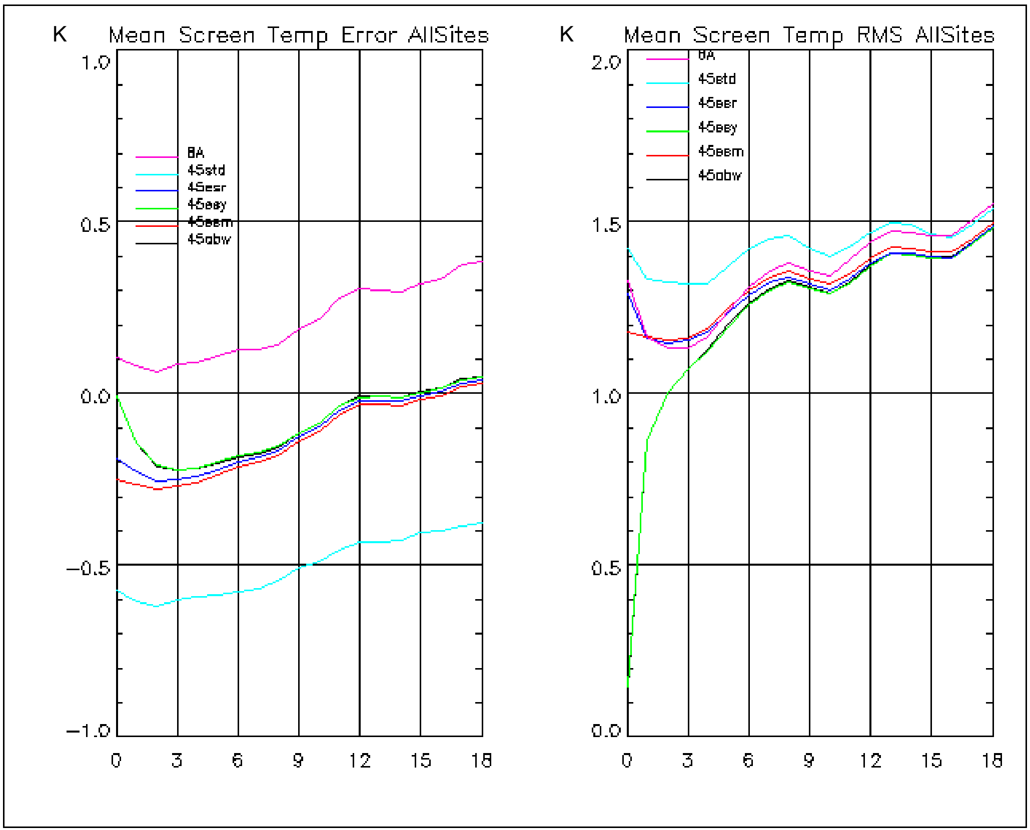

Figure 4: Screen temperature, 8A configuration, validity time

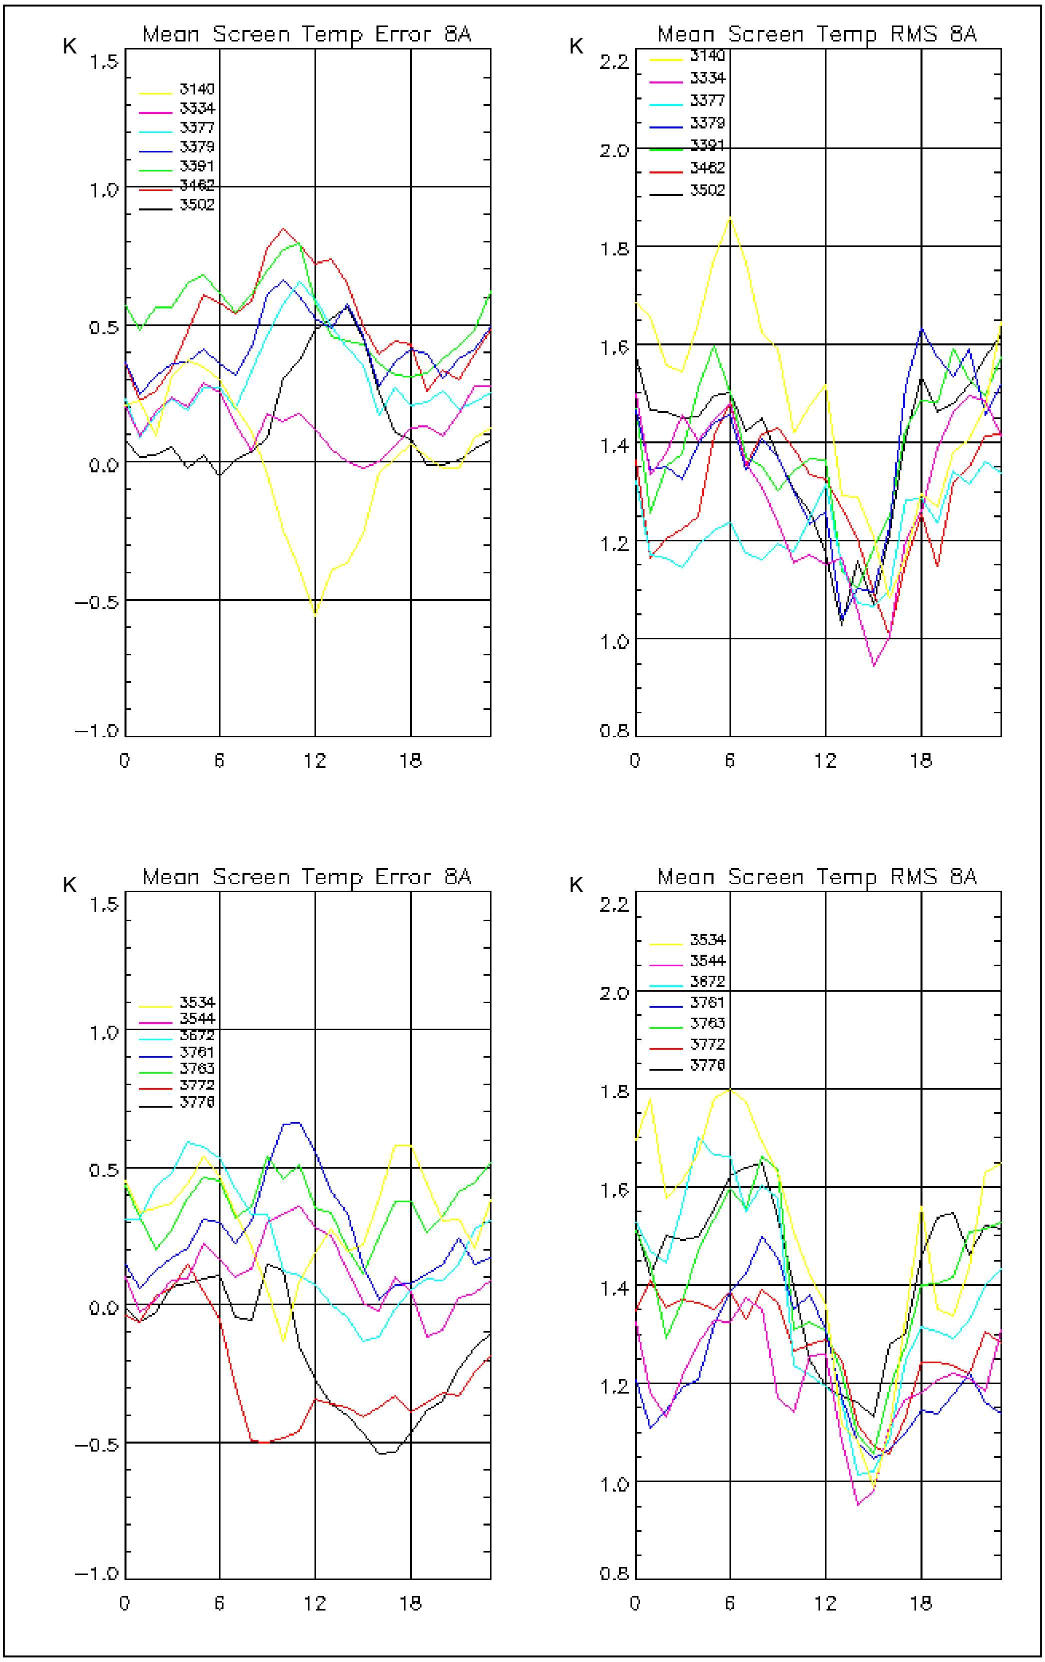

Figure 5: Screen temperature at Glasgow Airport, validity time

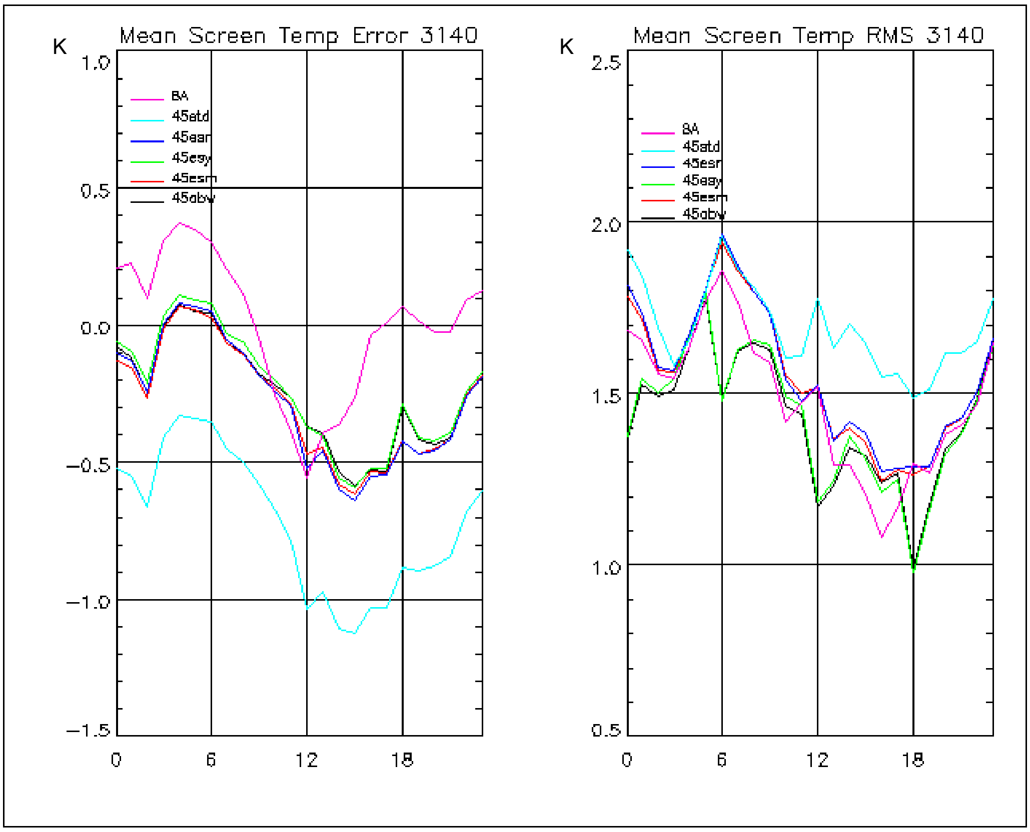

Figure 6: Screen temperature at Coningsby, validity time

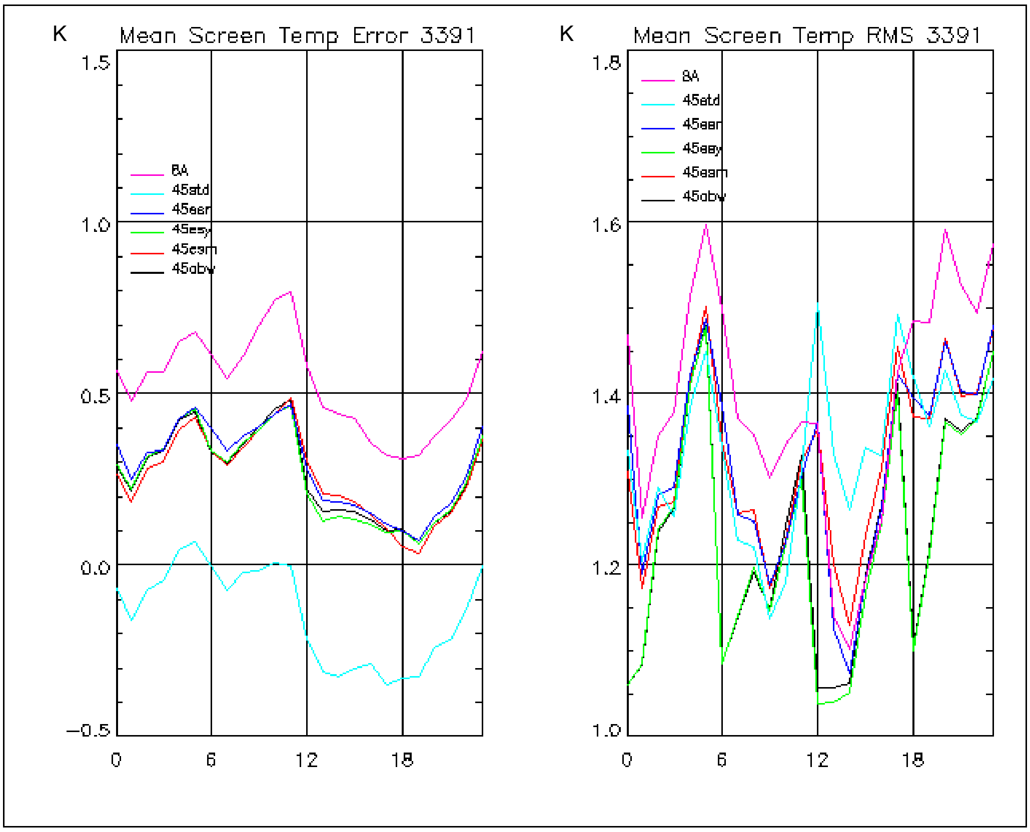

Figure 7: Screen temperature at Beaufort Park, validity time

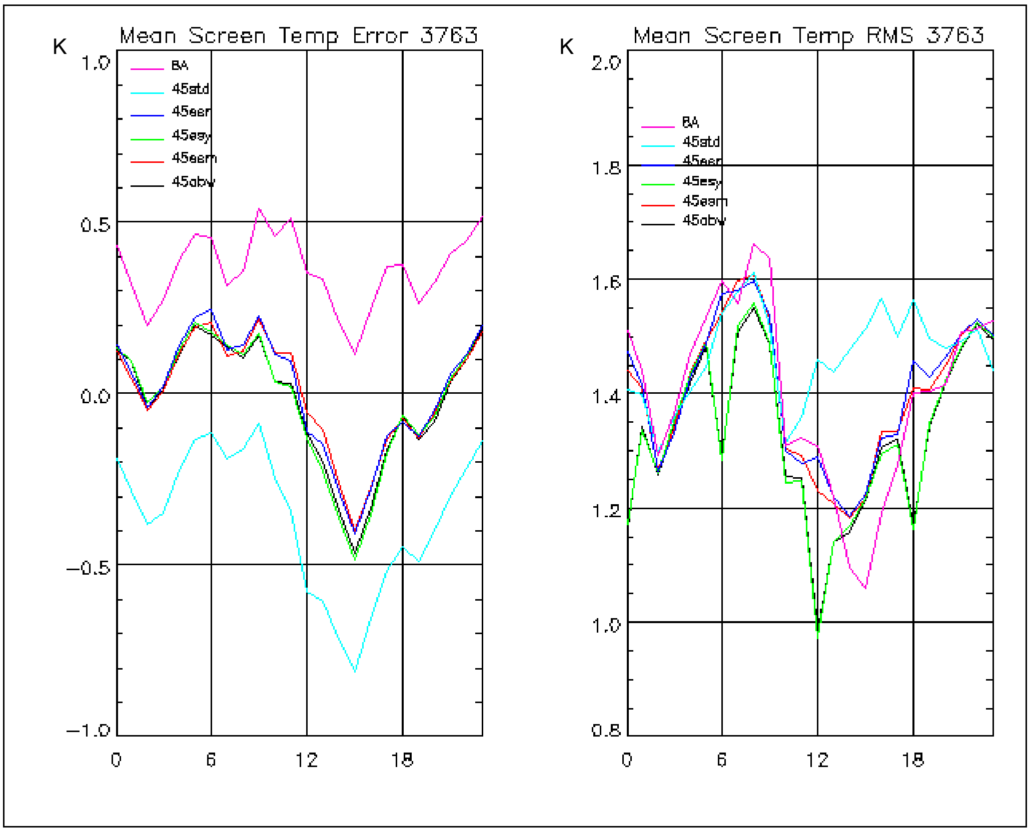

Figure 8: Screen temperature at Heathrow, validity time

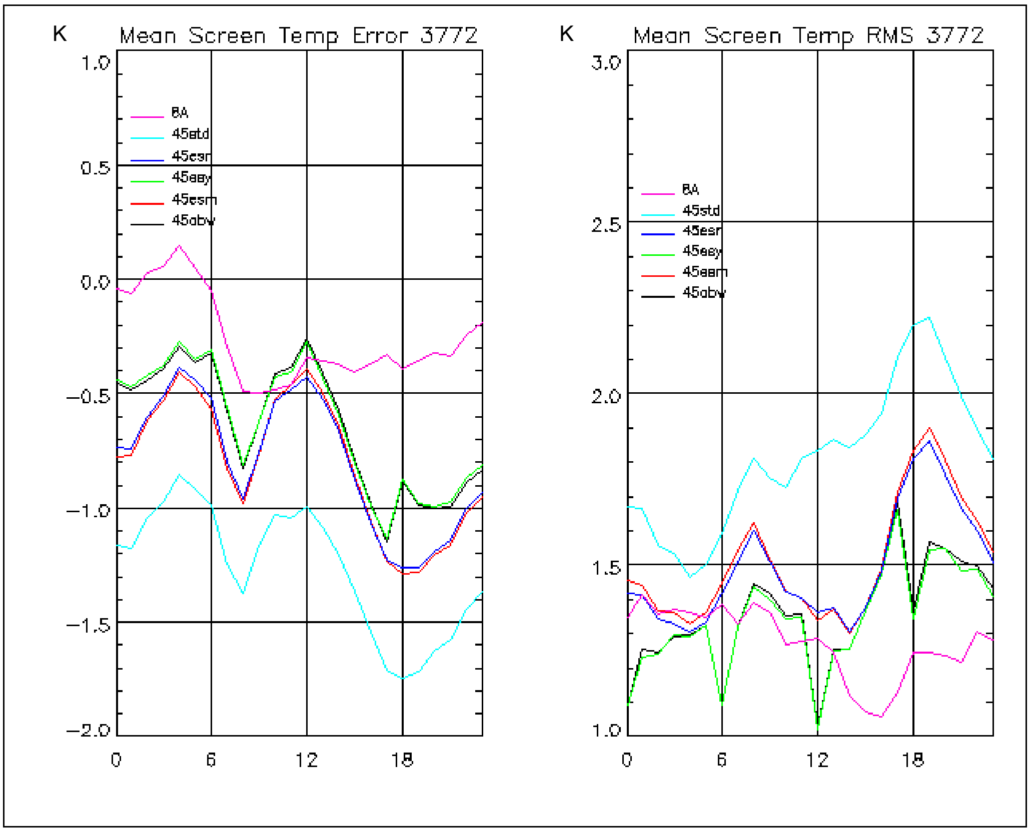

Figure 9: Screen dew point, comparison of configurations, validity time, sites combined

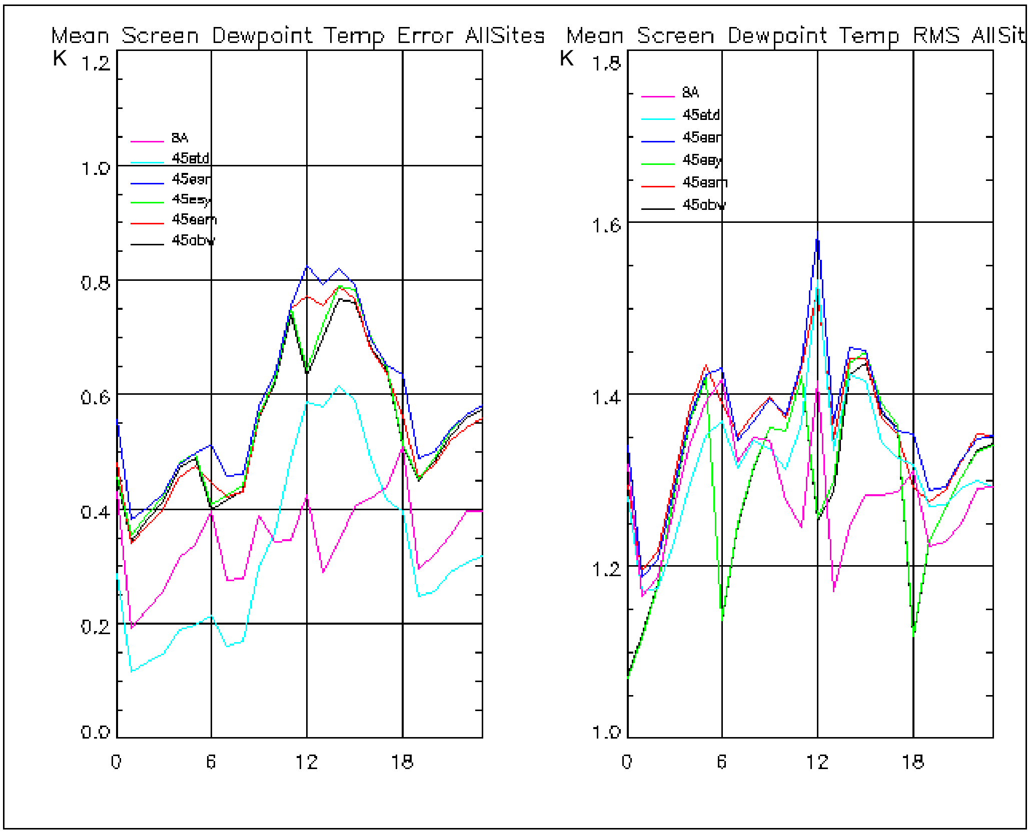

Figure 10: Screen dew point, comparison of configurations, forecast time, sites combined

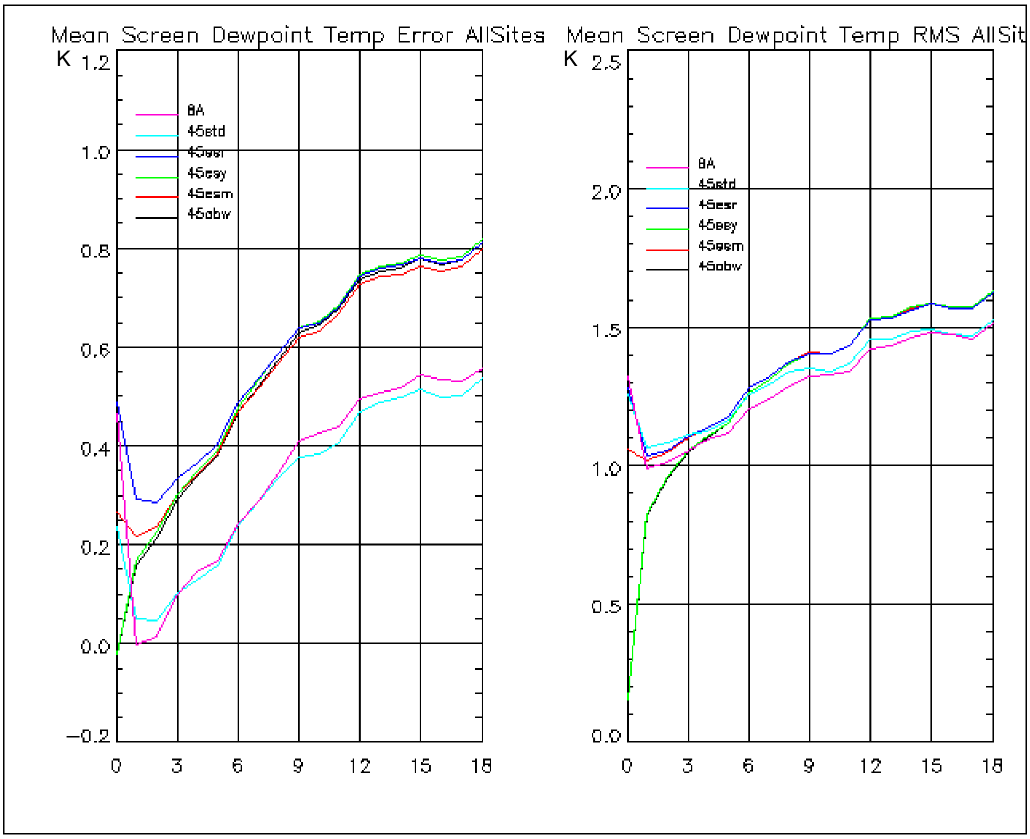

Figure 11: Screen dew point, 8A configuration, validity time

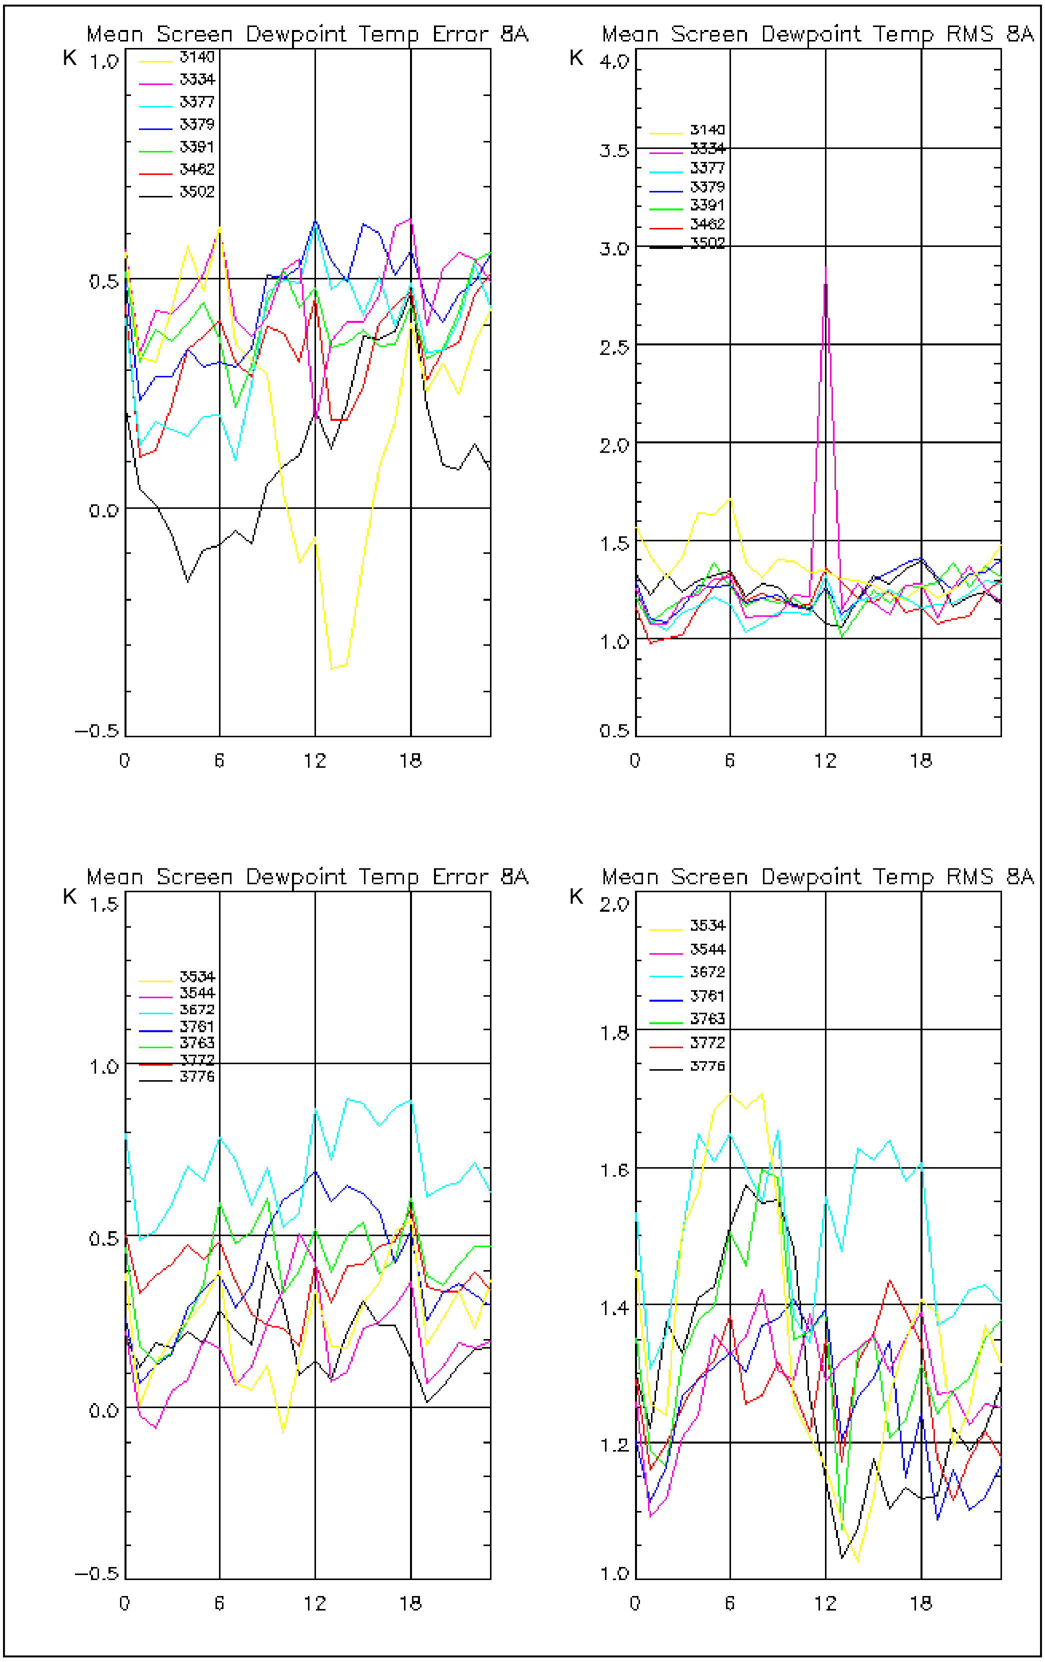

Figure 12: Screen dew point at Glasgow Airport, validity time

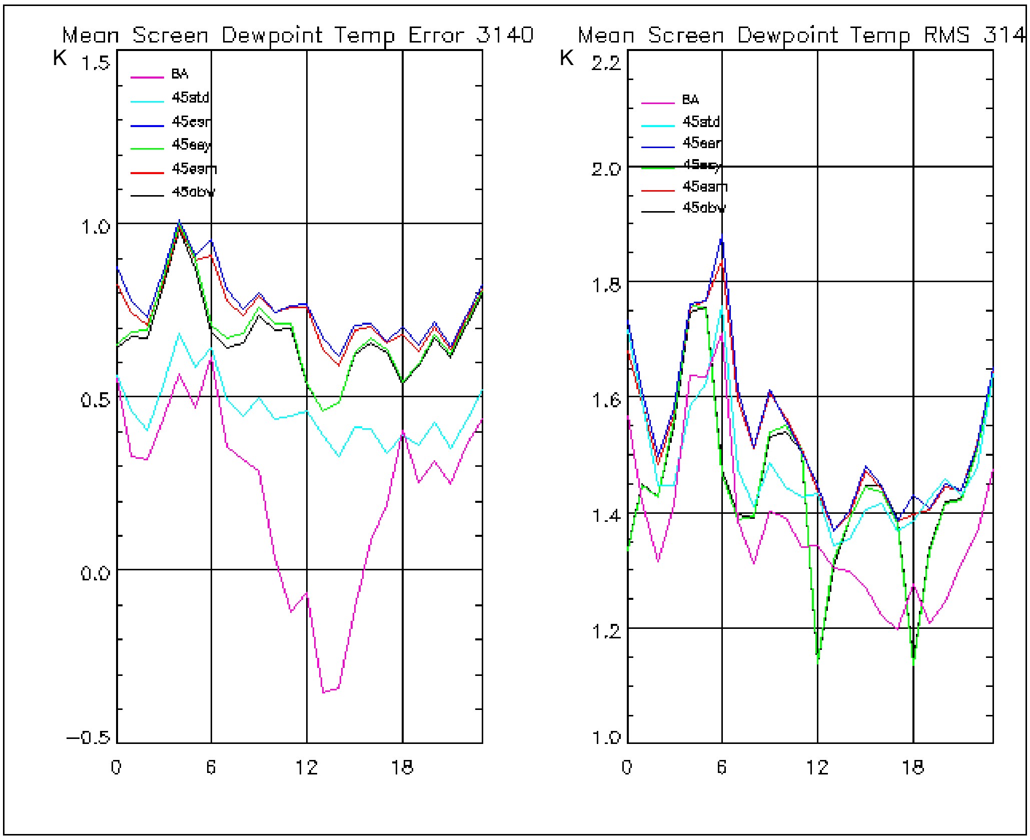

Figure 13: Screen dew point at Coningsby, validity time

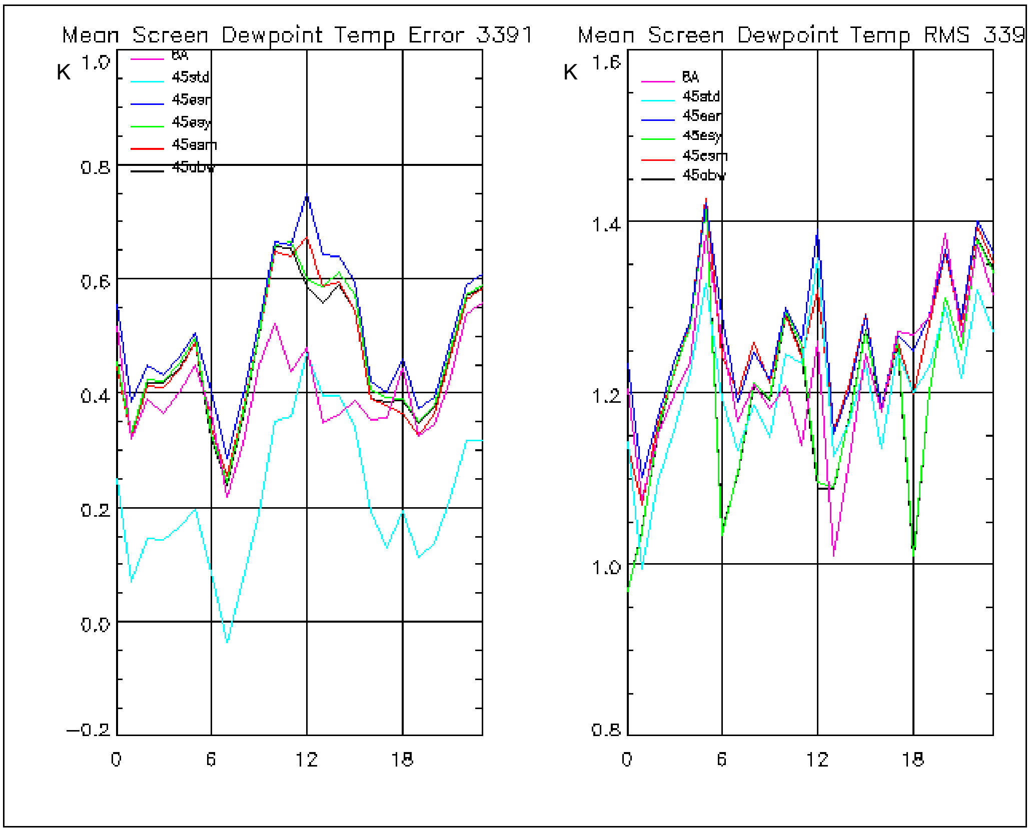

Figure 14: Screen dew point at Beaufort Park, validity time

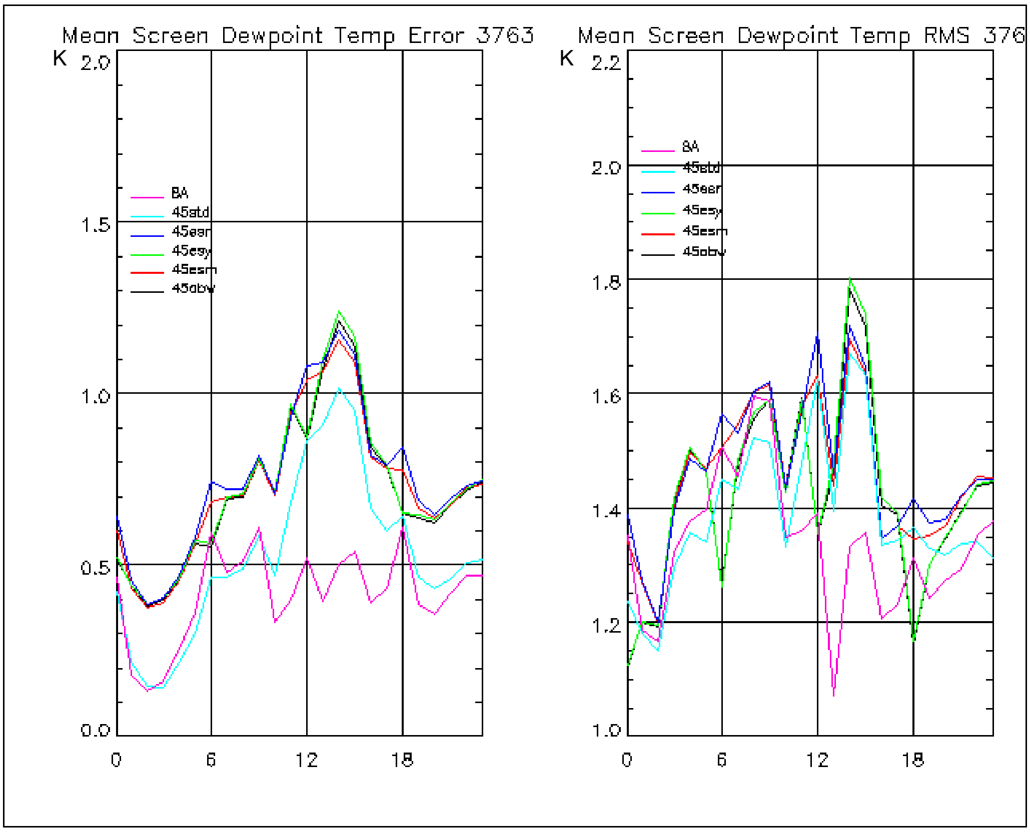

Figure 15: Screen dew point at Heathrow, validity time

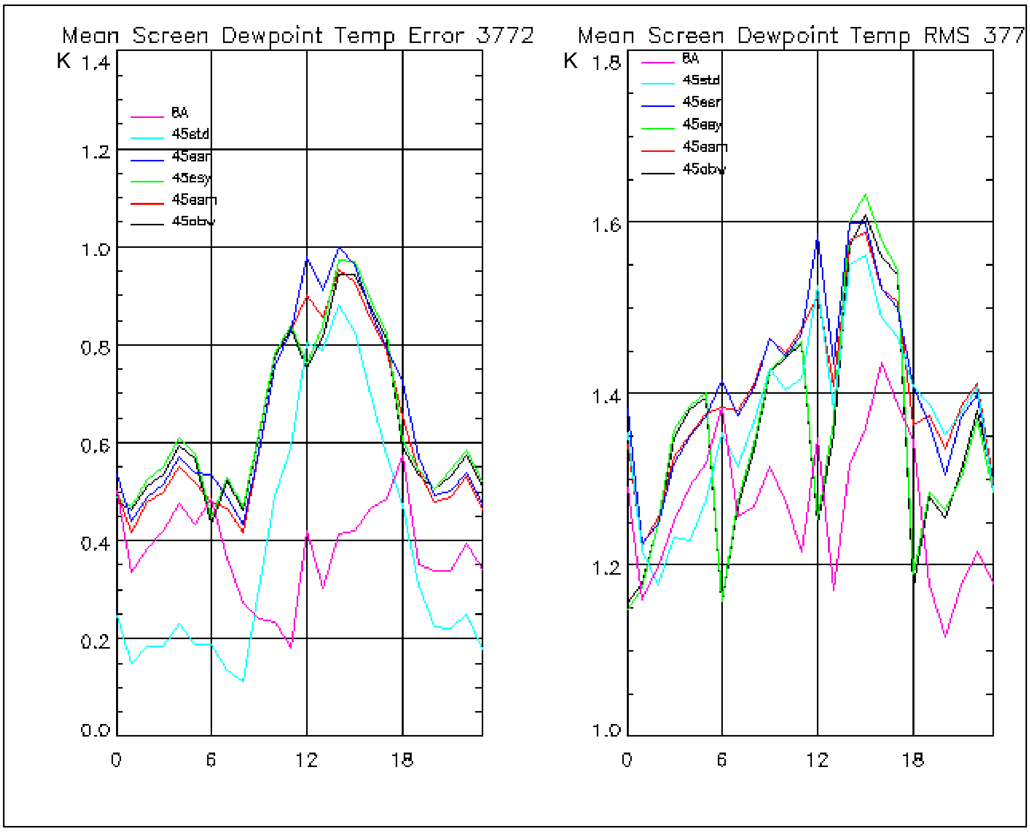

Figure 16: Screen visibility, comparison of configurations, validity time, sites combined

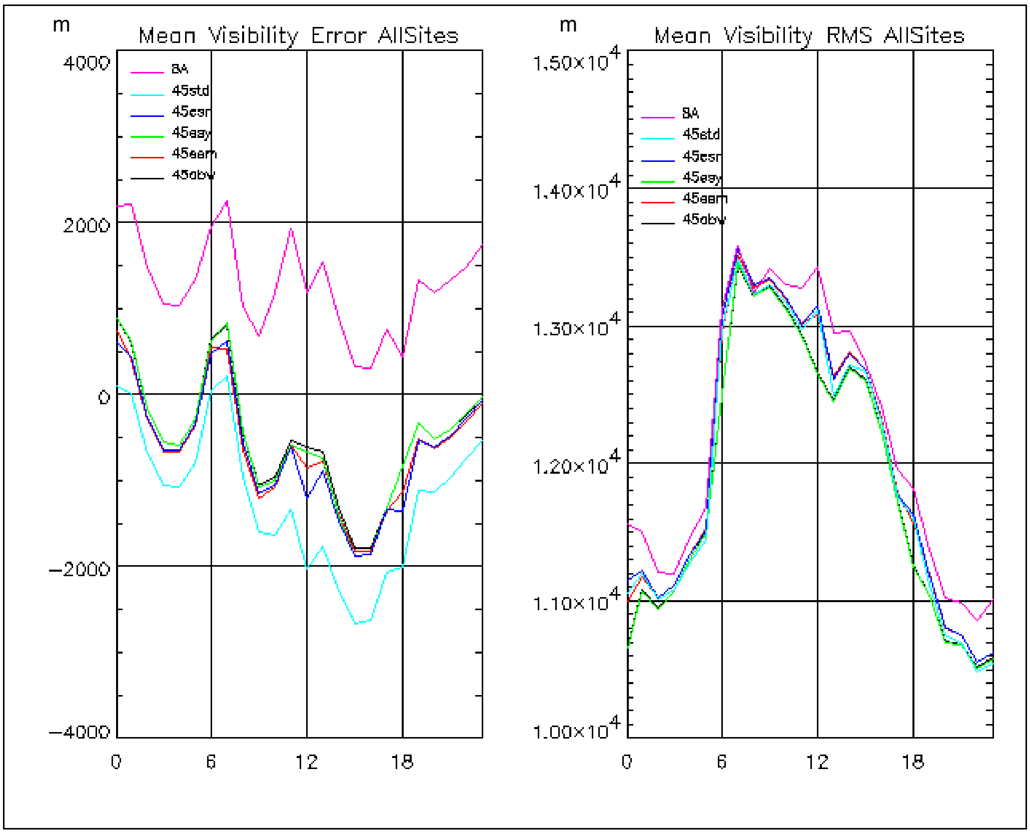

Figure 17: Screen visibility, comparison of configurations, forecast time, sites combined

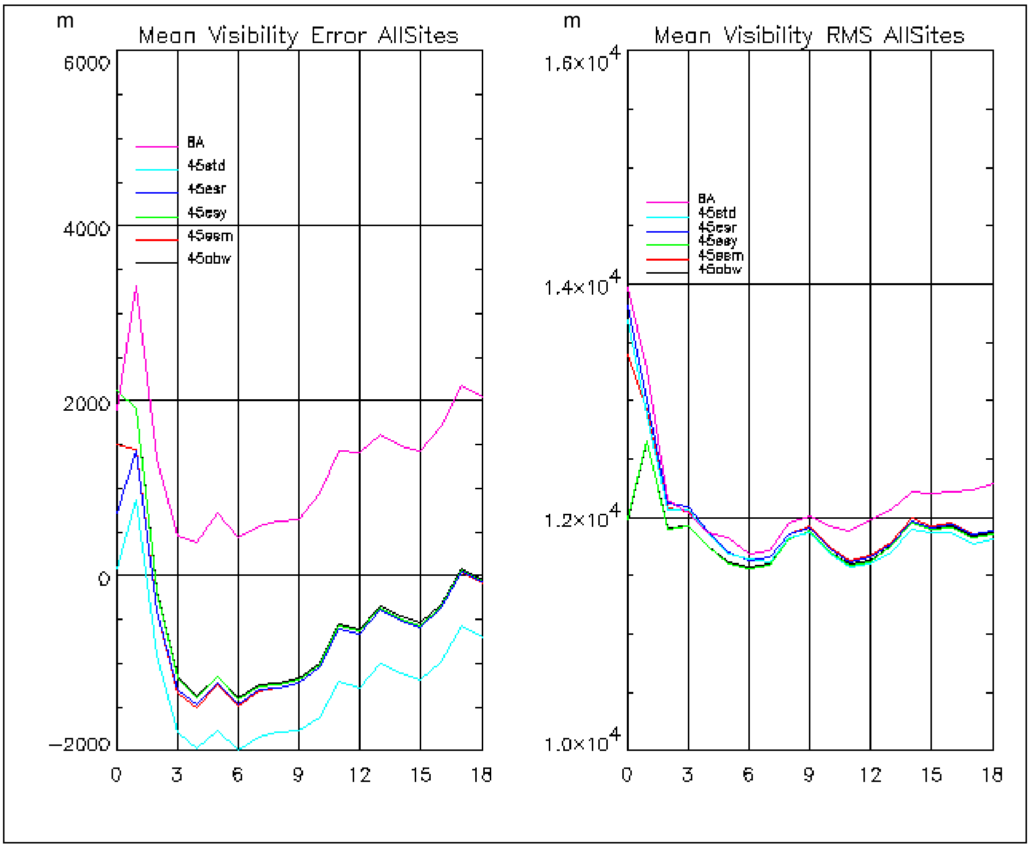

Figure 18: Screen visibility, 8A configuration, validity time

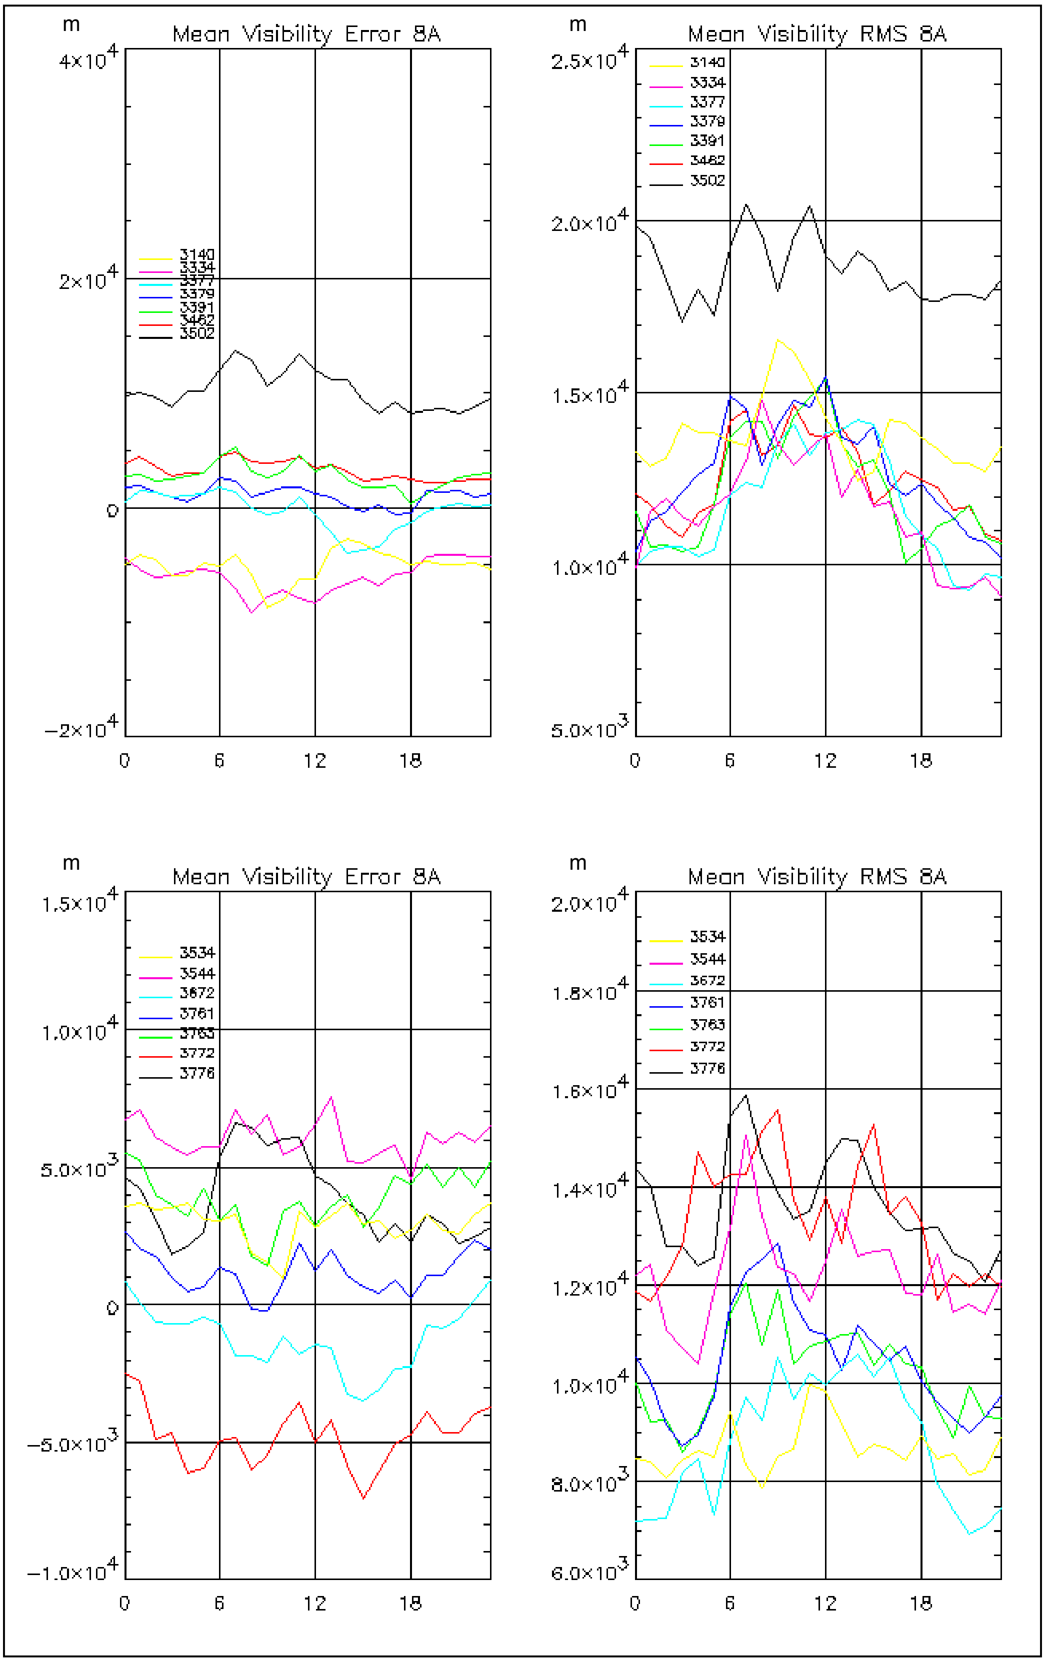

Figure 19: Screen visibility at Glasgow Airport, validity time

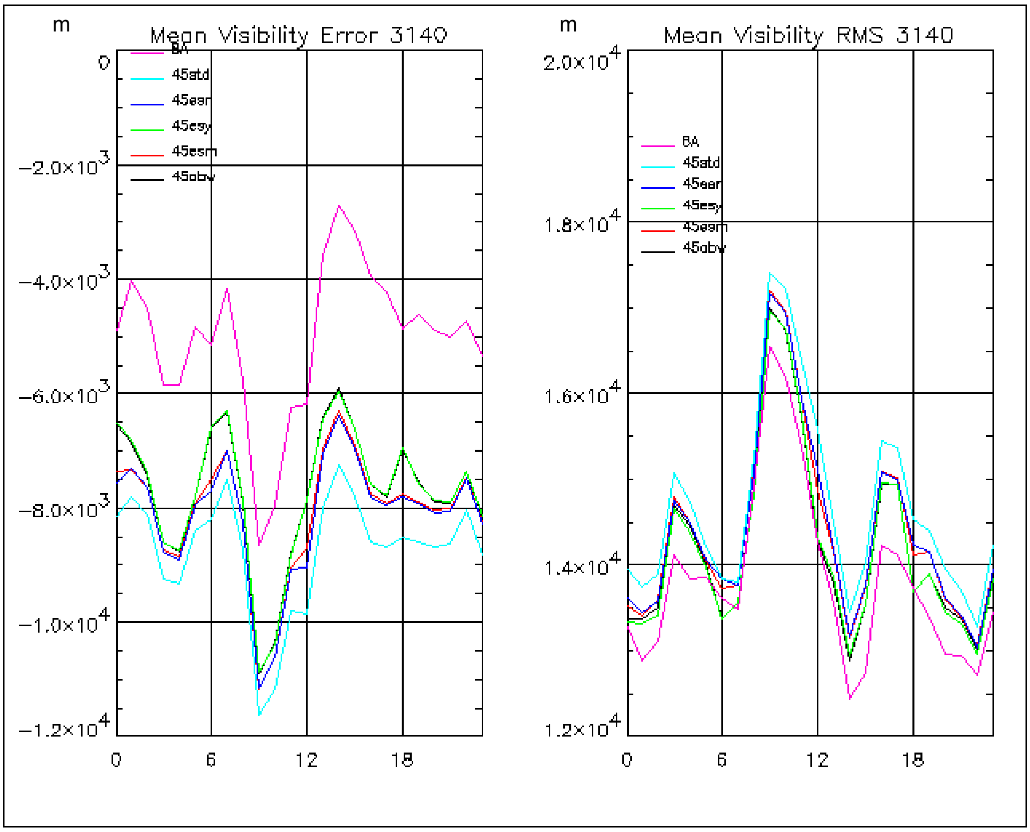

Figure 20: Screen visibility at Coningsby, validity time

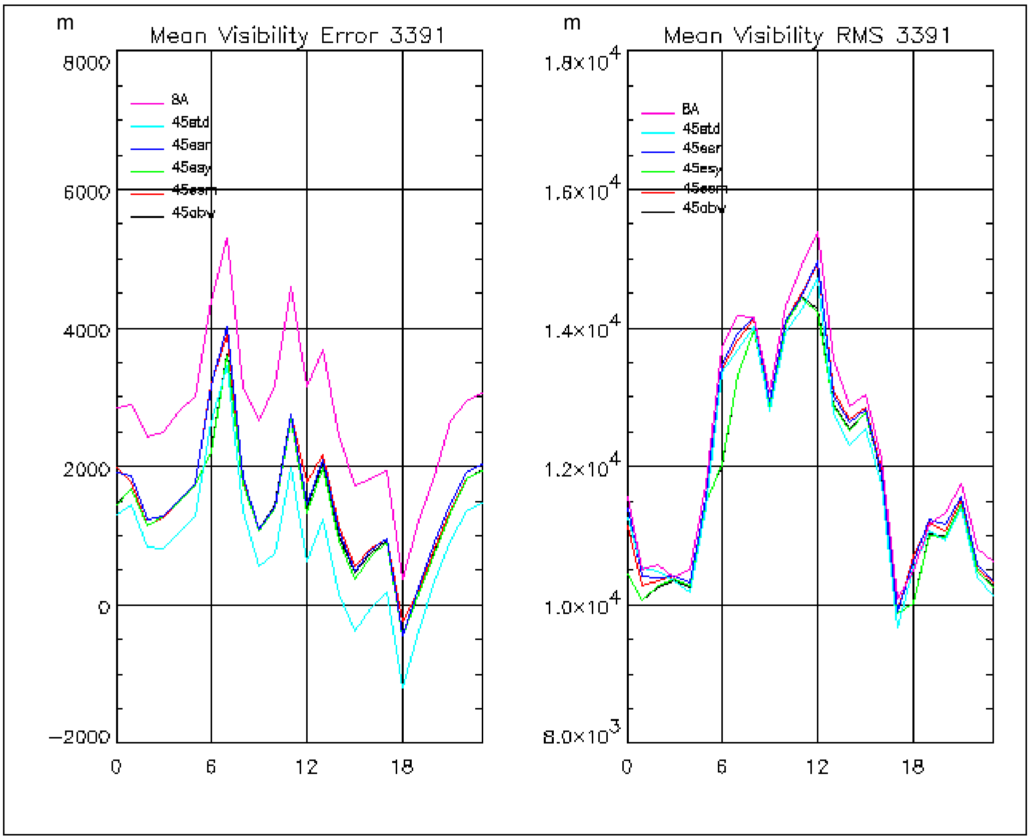

Figure 21: Screen visibility at Beaufort Park, validity time

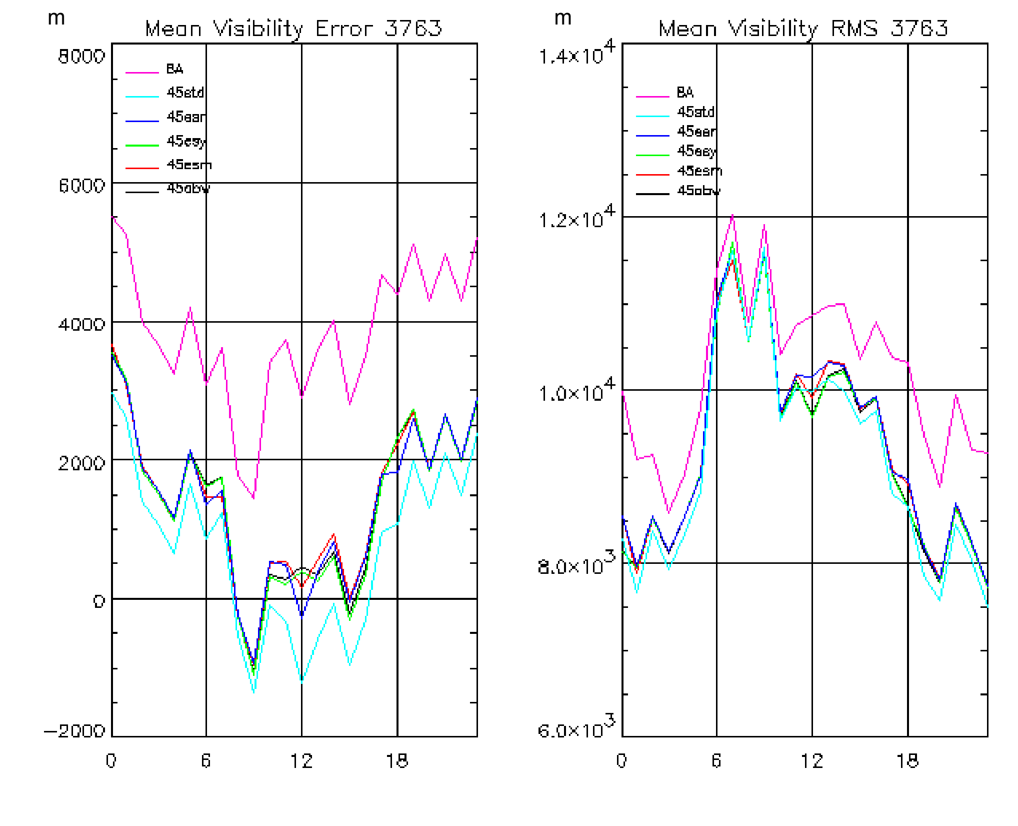

Figure 22: Screen visibility at Heathrow, validity time

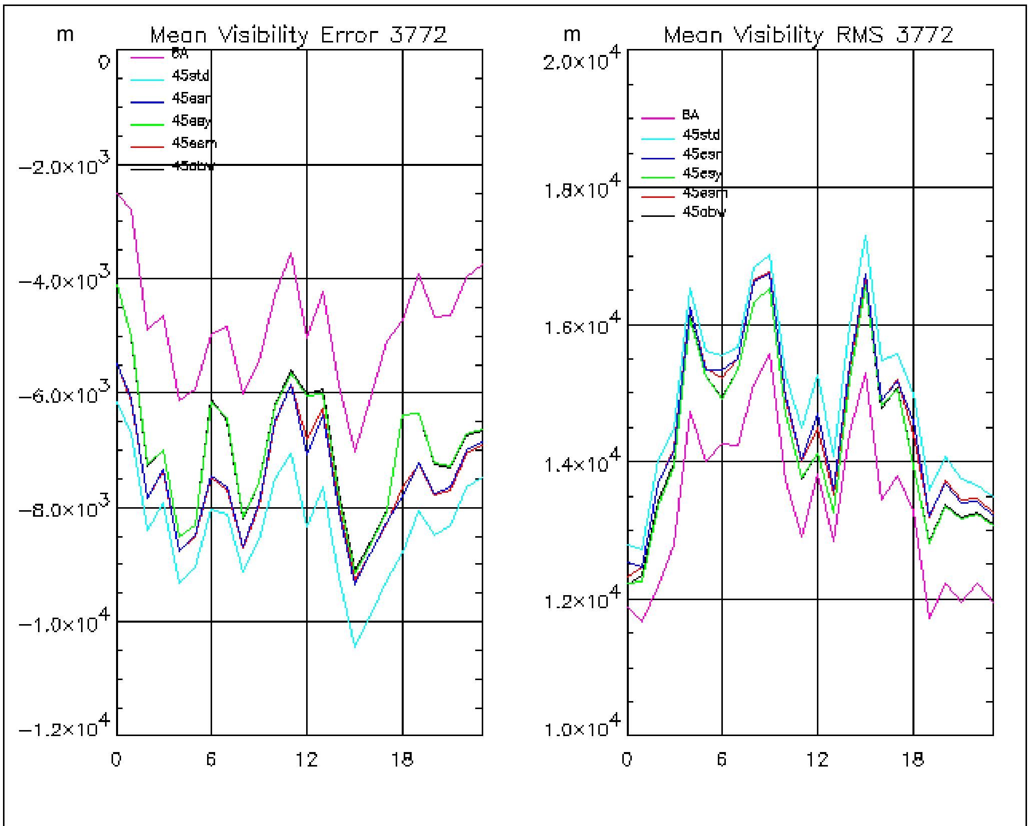

Figure 23: 10m wind speed, comparison of configurations, validity time, sites combined

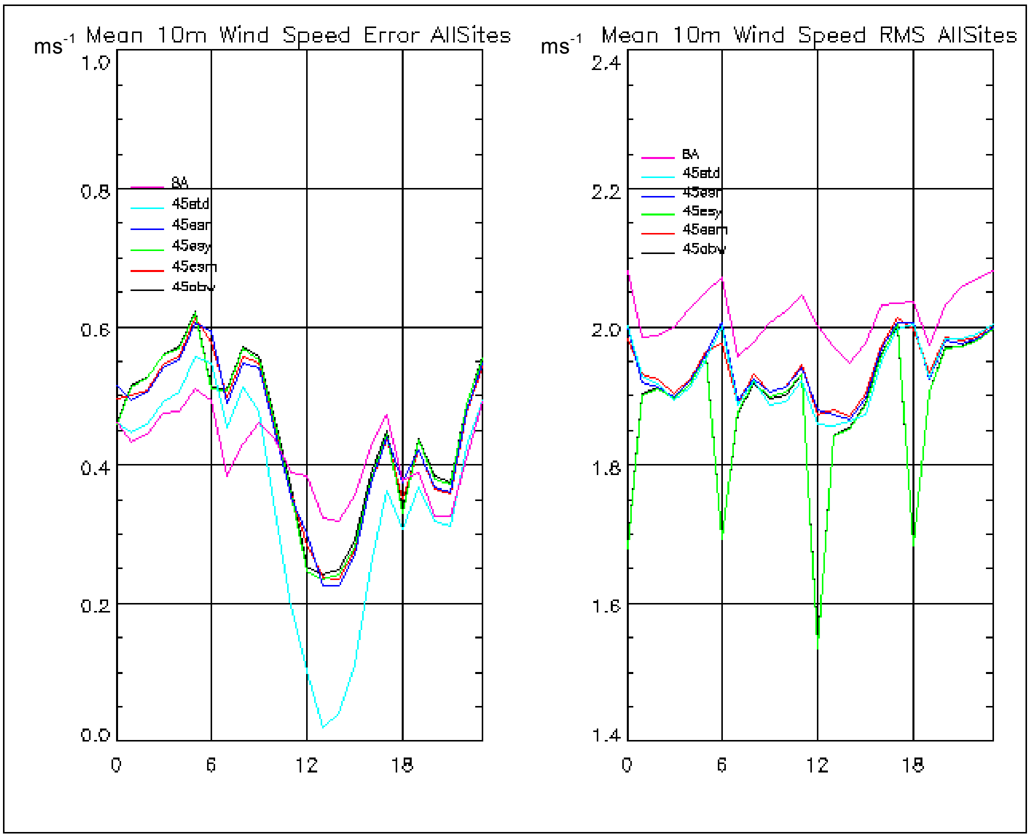

Figure 24: 10m wind speed, comparison of configurations, forecast time, sites combined

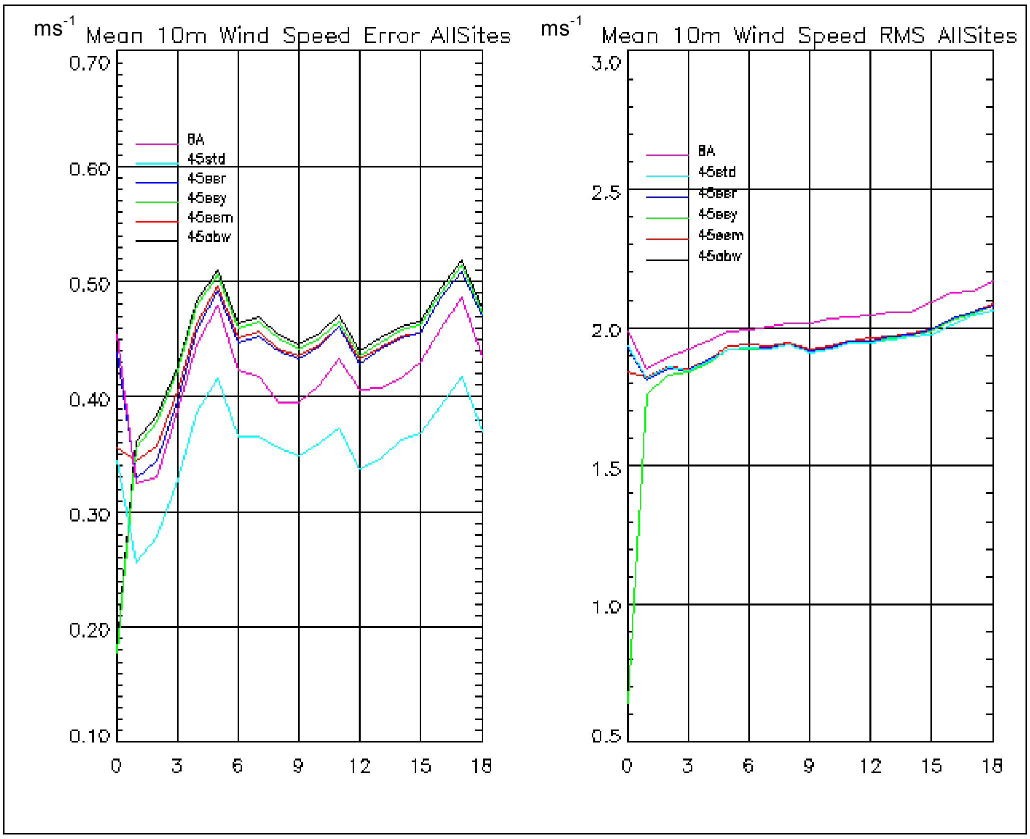

Figure 25: 10m wind speed, 8A configuration, validity time

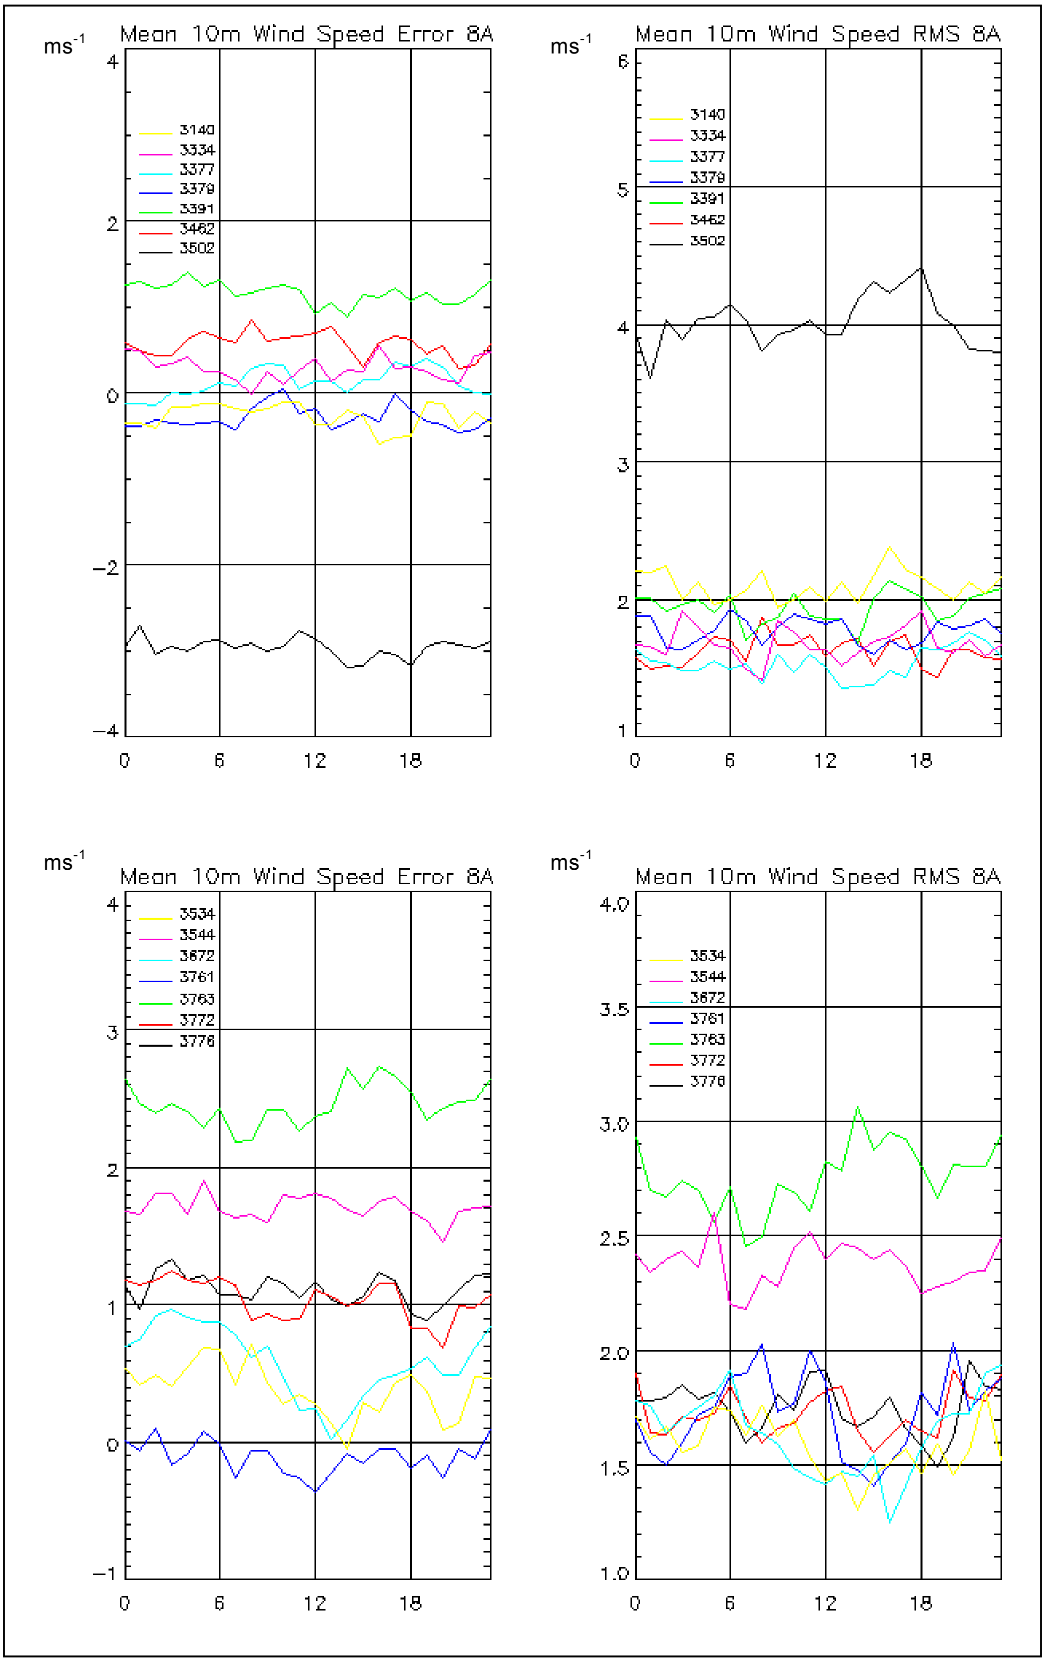

Figure 26: 10m wind speed at Glasgow Airport, validity time

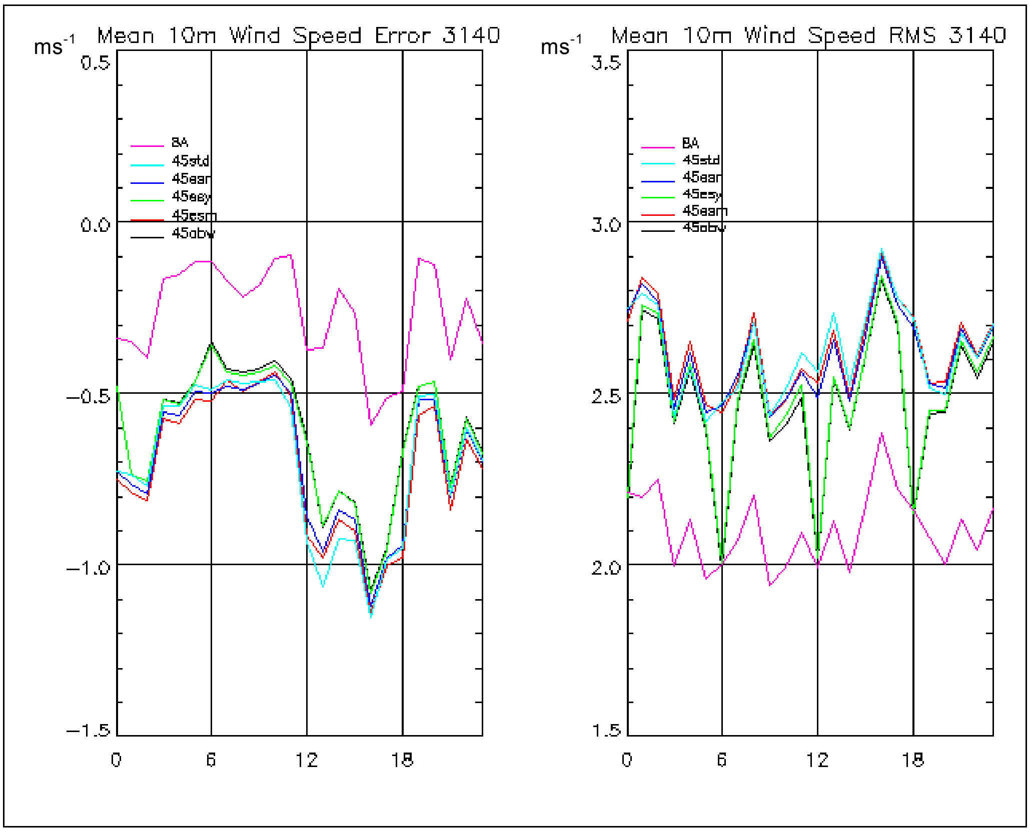

Figure 27: 10m wind speed at Coningsby, validity time

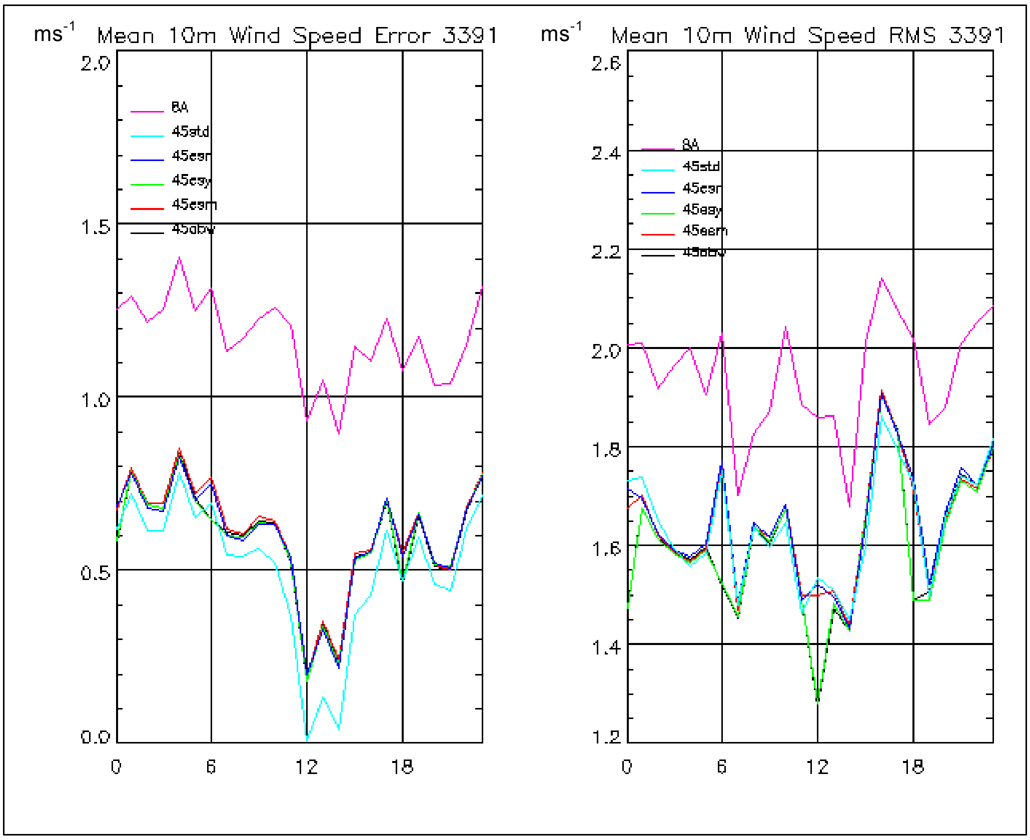

Figure 28: 10m wind speed at Beaufort Park, validity time

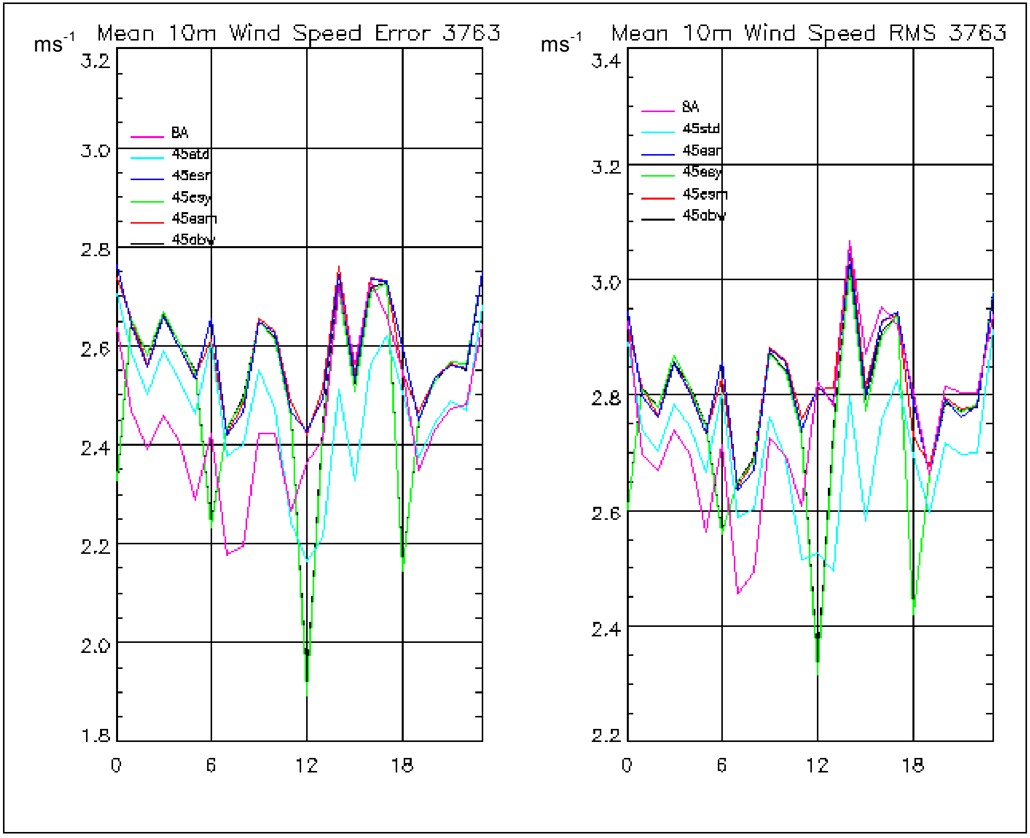

Figure 29: 10m wind speed at Heathrow, validity time

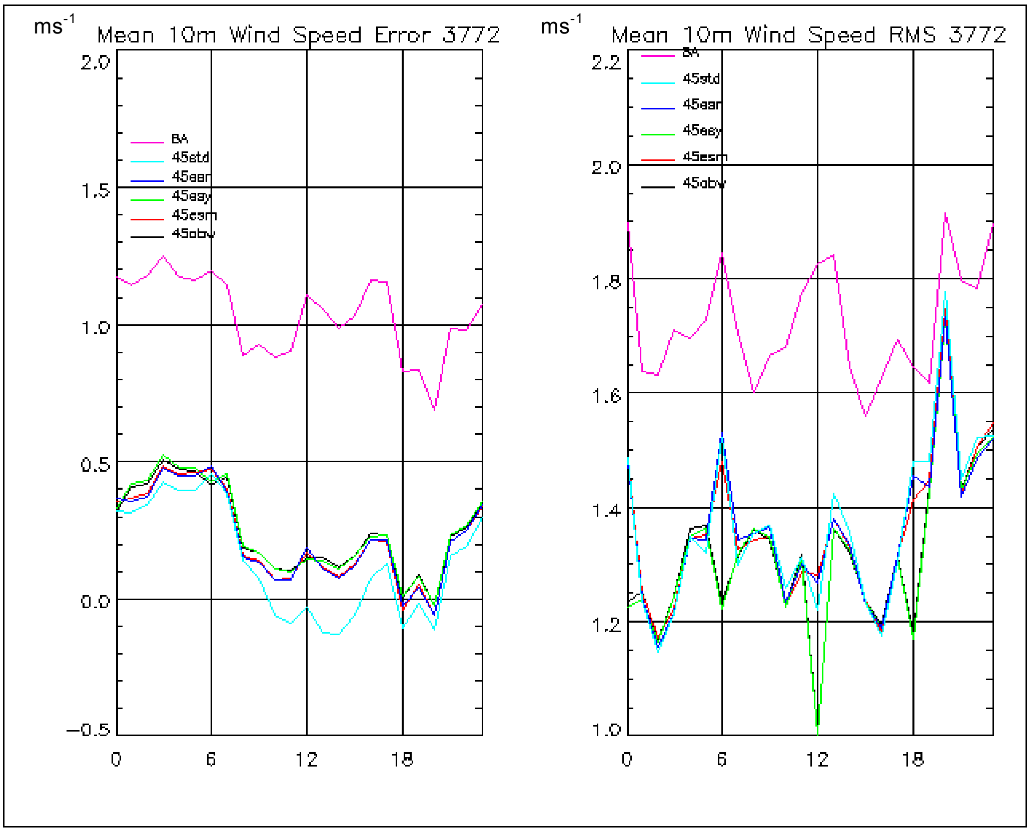

Figure 30: 10m wind direction, comparison of configurations, validity time, sites combined

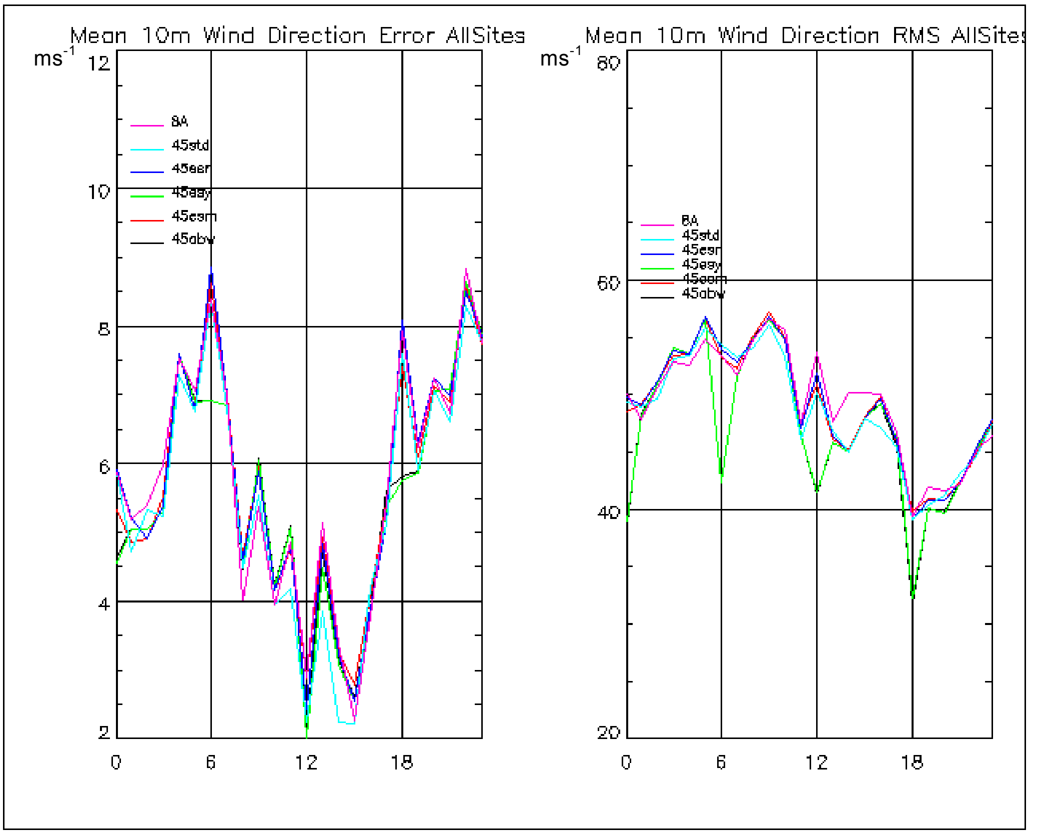

Figure 31: 10m wind direction, comparison of configurations, forecast time, sites combined