SafeWind

Collaborative project funded by the European Commission under the 7th Framework Program, Theme 2007-2.3.2: Energy

"Multi-scale data assimilation, advanced wind modelling & forecasting with emphasis to extreme weather situations for a safe large-scale wind power integration"

Grant Agreement N°: 213740

Deliverable Dc-5.11a

"Progress on ECMWF experiments and their impact on near-surface wind speed"

DOCUMENT TYPE	Deliverable
DOCUMENT NAME:	swind_deliverable_Dc-5.11a_ProgressECMWF_v2.0.pdf
VERSION:	V2.0
DATE:	2011.02.11
CLASSIFICATION:	R0: General public
STATUS:	Approved/Released

Abstract: This report reviews the experiments and operational upgrades at ECMWF which are relevant to the EU project SafeWind and to wind power prediction in general. They include the providing of 100-metre wind forecasts (deterministic and ensembles) and corresponding analysis, the upgrade of horizontal resolution of the whole system, research and development based on the TIGGE dataset, as well as the revision of the methodologies for representating initial and model uncertainties when generating ensemble predictions. Owing to its interim nature, the report does not go into details with the description of the methods and experiments, only giving an overview of already obtained and expected results.

AUTHORS[1], REVIEWERS				
MAIN AUTHOR/EDITOR:	P. Pinson[1,2]			
AFFILIATION:	[1] European Centre for Medium-range Weather Forecasts [2] Technical University of Denmark, DTU Informatics			
ADDRESS:	[1] Shinfield Park, Reading, RG2 9AX, United Kingdom [2] Richard Petersens Plads 305(212), 2800 Kgs. Lyngby, Denmark.			
TEL.:	[1] +44 (0)118 9499664 / [2] +45 4525 3428			
EMAIL:	[1] pierre.pinson@ecmwf.int / [2] pp@imm.dtu.dk			
FURTHER AUTHORS:	T. Petroliagis (ForWind, Germany), Renate Hagedorn (ECMWF)			
PEER REVIEWERS:	Robin Girard (Armines, France)			
REVIEW APPROVAL:	Approved :	X	Rejected (improve as indicated below) :	
SUGGESTED IMPROVEMENTS:	*For a long list of remarks make reference to another document*			

VERSION HISTORY			
VERSION[2]:	DATE:	COMMENTS, CHANGES, STATUS:	PERSON(S):
v1.0	2010.11.10	First version for review	P. Pinson
v2.0	2011.02.11	Final version after review (by R. Girard, Armines)	P. PInson

Acknowledgements

The work presented has been partly supported by the European Commission under the SafeWind project (ENK7-CT2008-213740), which is hereby acknowledged. Acknowledgments to due to a significant number of ECMWF staff members for their contribution to the developments covered by this report, including Martin Leutbecher, Martin Steinheimer and Roberto Buizza among others. Robin Girard is finally acknowledged for his review of the present document.

STATUS, CONFIDENTIALITY, ACCESSIBILITY										
STATUS:				CONFIDENTIALITY:				ACCESSIBILITY:		
S0	Approved/Released	X		**R0**	General public	X		Private web site	X
S1	Reviewed			**R1**	Restricted to project members			Public web site	X
S2	Pending for review			**R2**	Restricted to European Commission			Paper copy	X
S3	Draft for comments			**R3**	Restricted to WP members + PL				
S4	Under preparation			**R4**	Restricted to Task members +WPL+PL				

PL: *Project leader* **WPL:** *Work package leader* **TL:** *Task leader*

[1] **The authors of this document are solely responsible for its content, which does not represent the opinion of the European Community and the European Community is not responsible for any use that might be made of data appearing therein.**

[2] **VERSION NAMING : V0.x** draft before peer-review approval, **V1.0** at the approval, **V1.x** minor revisions, **V2.0** major revision

1 Introduction

The European Centre for Medium-range Weather Forecasts (ECMWF) has initiated a number of changes to its operational forecasting system which may be of particular relevance to the EU-funded project SafeWind (FP-7) and to wind power prediction in general. This is while a number of experiments are undergoing in order to prepare the future operational upgrades. Some research and development activities are also specifically focusing on wind-related problems. Obviously, since the primary focus of ECMWF cannot be on near-surface winds only, the relevance of these upgrades and related benefits are studied and assessed in terms of general scores for certain reference variables, not obviously for near-surface wind speed and for Europe only. These various changes and experiments are reviewed in the present document, while the resulting benefits (already observed or expected) are discussed. Most of the discussion developed below and results presented directly link to other publications, which are either technical reports at ECMWF, contents of the ECMWF Newsletter, or articles published (or submitted) in the scientific literature. All developments considered directly relate to Task 5.2 in the Description of Work for the EU-project SafeWind. The present document is an interim report for the 1st half of the project. It will be complemented by a final report, due at the end of the involvement of ECMWF in SafeWind. This report also contributes to the achievement of the defined milestone M-5.3 ("Quantification of improvements in WPF using data from new EPS configurations") by documenting and discussing the improvement in wind forecasts observed at ECMWF.

To briefly list the various topics that are treated here, we start by presenting in Section 2 the aspects related to the recent operational providing of 100-metre wind forecasts. Section 3 concentrates on the increase of horizontal resolution of the operational forecasting system. Section 4 looks at the work carried out based on the TIGGE dataset including global ensemble forecasts from various research and operational weather forecasting centres, and comparison with suitable recalibration of ECMWF ensemble forecasts. In Section 5, we then describe the new methodology employed at ECMWF for the generation of the initial perturbations of the ensemble forecasts, and for the representation of model uncertainties. We finally gather conclusions and perspectives for further work in Section 6.

2 Forecasts of 100-metre winds

Traditionally, the near-surface wind forecasts that are made available by ECMWF are the u and v components of the 10-metre winds, then permitting to derive wind speed and direction. These variables are highly relevant to a number of applications and decision-making problems, including wind power prediction. In view of the ever-increasing size of wind turbines, however, forecasts of the whole wind profile, or at least at vertical levels close to hub height, are necessary. Acknowledging this requirement, ECMWF started to issue operationally (and also to archive) forecasts of 100-metre winds. They similarly consists of the u and v wind components. Both deterministic and ensemble forecasts are available. The first archiving of 100-metre wind forecasts was carried out on the 26th 2010 (for research purposes), while this 'product' was officially added to the catalogue of variables available to ECMWF forecast users at the beginning of August 2010. A week after that only, these aspects were mentioned in an article that appeared in the Weather Eye column of The Times (Plester 2010) (even though the SafeWind project was not directly mentioned).

Evaluation of wind forecasts may be carried out against the analysis or against observations, potentially yielding fairly different verification results (Pinson and Hagedorn 2010). Already when focusing on

10-metre winds, verification against observations may be a subtle task. This is because observational uncertainty may somehow have a significant impact on verification results, while the manner such an uncertainty is taken into account may also affect various scores in different ways (Candille and Talagrand 2008). When looking at 100-metre wind forecasts instead, is clear that the quantity of observations will be rather limited, since the number of sites where continuous meteorological measurements are obtained at these heights is very small. One of the European sites that can be considered for the purpose is the FINO platform(s) located in the German bight. Verification against observations has been carried out at this site, and recently reported by Petroliagis (2010).

In the first stage, Petroliagis (2010) looked at the discrepancy between the 100-metre analysis of wind speed with observations at 103 metres at the location of the FINO1 measurement platform. No correction is performed, hence assuming that this small difference in height does not introduce any significant change in the level or variability of wind speed. An example comparison of analysis and observations is depicted in Figure 1 for a period ranging from the 19$^{\text{th}}$ of November and until the 31$^{\text{st}}$ of December, with a temporal resolution of 12 hours. At a first glance, there is no apparent bias in the analysis, while the variability of the observations is obviously of larger magnitude than that in the analysis. This is due to the spatial and temporal scales at which the analysis is computed (i.e. on a grid with horizontal resolution of ∼75 kms, and time steps of 6 hours).

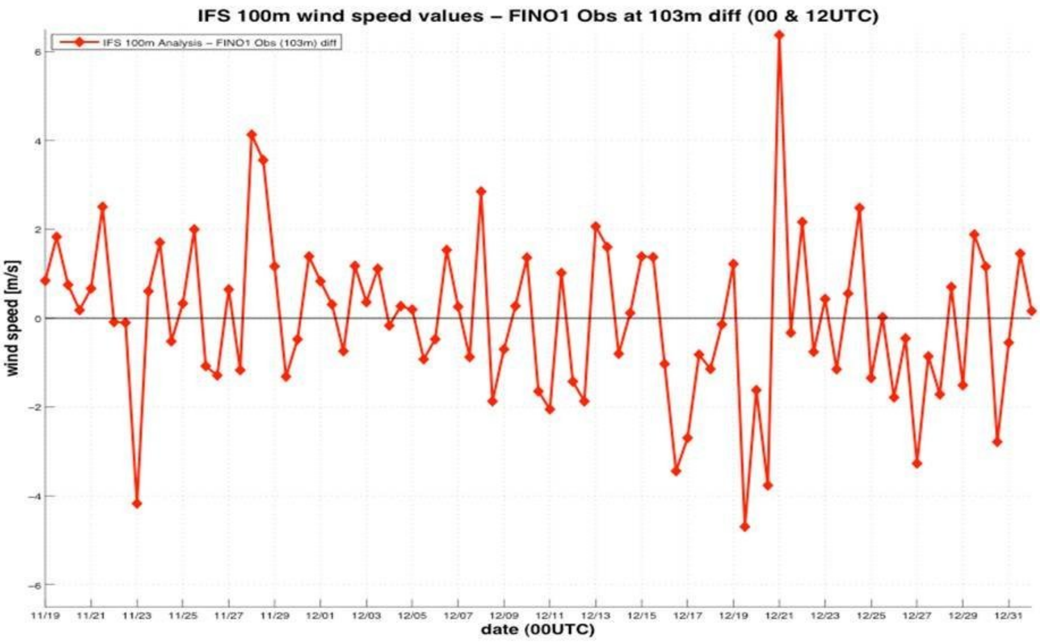

FIGURE 1: *Comparison of the 100m analysis of wind speed with observations at the FINO measurement platforms, over the period of November-December 2009.*

In the second stage, emphasis was placed on the verification against observations of the high-resolution deterministic forecasts of wind speed at FINO (for the two measurement platforms FINO1 and FINO2). Example results are gathered in Figure 2 for the same period as in the above, while more extensive

verification results can be found in Petroliagis (2010). The verification results consists of the Mean Average Error (MAE), Root Mean Square Error (RMSE) and bias of the forecasts for horizons between 0 (analysis) and 168 hours ahead (7 days). The bias, comprising the systematic component of the prediction error, is stable over the set of horizons. It could easily be removed thanks to simple post-processing techniques. Such a systematic error component is not surprising for wind speed forecasts relying on Numerical Weather Prediction (NWP) models. It also has a magnitude that can observed for 10-metre wind speed forecasts, and in onshore environments.

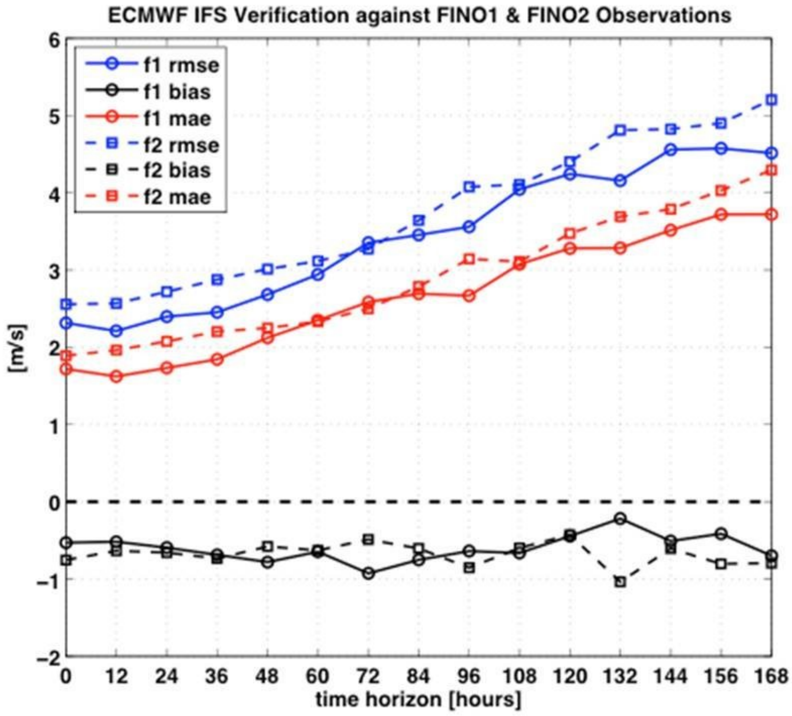

FIGURE 2: *Example results from the verification of the 100m wind speed forecasts against observations at the FINO measurement platforms over the period of November-December 2009.*

The evaluation period is fairly short, covering 42 days only. The results obtained are as one would expect, however, with forecast accuracy (as quantified by MAE and RMSE criteria) degrading with further lead times. The MAE increases from 2 to 4 m.s^{-1} while the RMSE ranges between 2.5 and 5 m.s^{-1}. These values are slightly higher than what would be usually observed for 10-metre wind forecasts and for onshore environments. These results are consistent with the fact that higher heights results in higher wind speeds, which then typically lead to forecasts of lower accuracy. Further verification results for 100-metre winds are expected to be produced by ECMWF, by partners of the SafeWind consortium, as well as by other users of ECMWF forecasts, in the coming future. They may be reported in the updated version of that report, i.e. the SafeWind Deliverable Report Dc5.11b.

3 Increase in the horizontal resolution

ECMWF upgraded the horizontal resolution of its operational forecasts from the 26th of January. This is the case for both the high-resolution deterministic forecasts and for the ensemble forecasts. In the case

of the former type of forecasts, the resolution was changed from $\sim$ 25 kms to $\sim$ 16 kms, while for the latter ones, it was increased to $\sim$ 32 kms (from $\sim$ 50 kms) for lead times up to 10 days. These changes in horizontal resolution are the result of a global upgrade that also concerned the data assimilation component or the wave model. A number of aspects related to this upgrade are extensively discussed by Miller *et al.* (2010). A more theoretical discussion on the impact of horizontal resolution on forecast skill for the ECMWF prediction system is given by Buizza (2010).

In order to evaluate the impact of such a change on the quality of ECMWF forecasts two versions of the ECMWF operational forecasting system have been running in parallel for a targeted experiment over a period of almost 3 months. This experiment yielded 187 forecast series issued over a period starting from the 3rd October 2009 and ending on the 26th of January 2010. Their starting times are 00UTC and 12UTC. No forecasts are available between the 4th and 23rd November 2009. This type of experiment is usually employed at ECMWF for assessing potential improvements brought by the new version of the system before its actual start of operation. Such improvements are usually looked at with the analysis as a reference, and by focusing on upper-air variables (e.g. Z500). We have concentrated instead on 10-metre wind speed forecasts, while seeing observations as the reference. All aspects of this study are extensively described and discussed in Pinson and Hagedorn (2010).

Maybe the most important aspect is the improvement of overall scores, calculated for all stations, hence giving an overview of potential improvements over Europe. They are given in Figure 3 as a function of lead time, and expressed as a percentage of the scores obtained for the coarser resolution. The improvements we look at are based on the bias, MAE, RMSE of point forecasts extracted from the ensembles, as well as on the CRPS of ensemble forecasts (CRPS standing for Continuous Ranked Probability Score). The potential effect of observational uncertainty is not considered, firstly owing to computational costs, and also since for an average over such a large number of stations it is expected to be negligible.

All improvements are positive over the forecast range considered, up to 6 days ahead. They are between 2 and 4% for the MAE and RMSE scores, while ranging from 3 to 5.5% for the CRPS. In view of the number of forecast series and stations involved, these improvements can be seen as significant. They are even more significant for the bias, being up to 22% for 3-day ahead forecasts. In parallel, the seasonalities present for all scores (though especially for the bias, which then affect other scores) show that the change of resolution also impacted the way local diurnal effects are captured by the models. The maximum improvements for all scores are reached in the early medium range, that is, between 2 and 3 days ahead. Finally, it is interesting to see that improvements in the CRPS are larger than improvements for the more deterministic scores MAE and RMSE (since relying on point forecasts only). A potential explanation can be that the forecast quality improvements are not only related to the better ability of ensemble forecasts to target observations, and to a higher sharpness, but also originates from a better calibration.

Consequently, we have investigated that point by plotting reliability diagrams for the various stations and groups of stations in order to see how calibration is affected by the change of horizontal resolution. Example results are gathered in Figure 4, where all horizons are considered indifferently. For a large number of stations, the situation is similar to that shown for Amsterdam Schipol and Cork airports. It consists of a significant improvement of probabilistic calibration for the finer resolution forecasts. For some other stations e.g. Thyboron in Denmark, however, probabilistic calibration actually seems to be worse for the forecasts with finer horizontal resolution. This is also the case for some of the stations where the worst calibration results can usually be observed, like Cap Béar in the South of France. Such a result is counter-intuitive since one would expect that more local regimes e.g. coastal effects may be better captured by increasing resolution. This may well also depend upon the physics behind the models instead. When looking at all stations altogether, the improvements in probabilistic calibration seem to

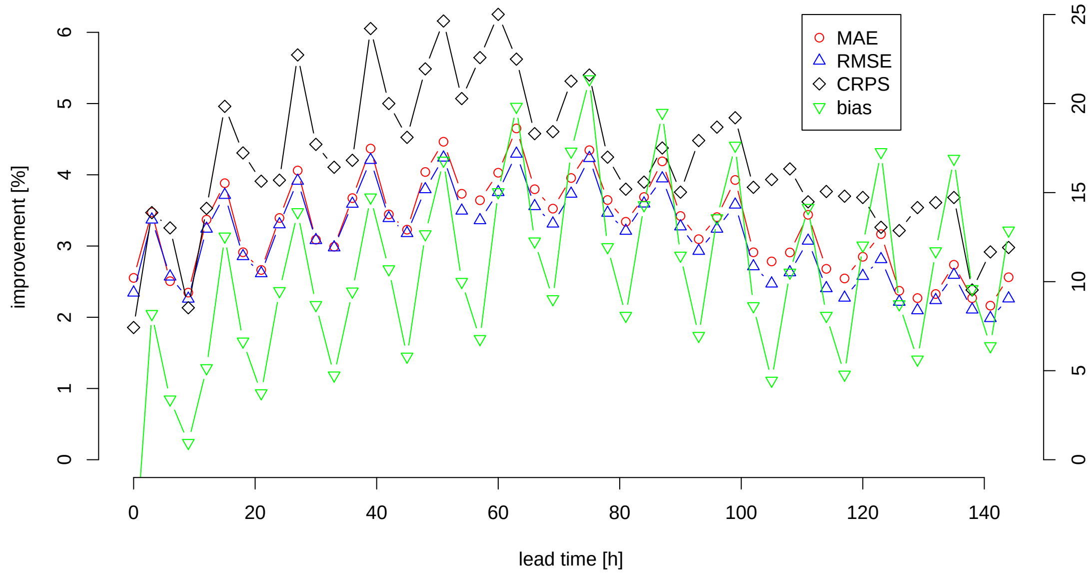

FIGURE 3: *Evaluation of the impact of the change of horizontal resolution on the quality of ensemble forecasts of wind speed over Europe. The left axis scale is for the MAE, RMSE and CRPS scores, while the right one is for the bias.*

exist, though being small. Note that the reliability diagrams we use here may not be the best visual tools for spotting very small differences between calibration curves. An alternative presentation based on the deviation from the diagonal as proposed by Pinson *et al.* (2007b) and Marzban *et al.* (2010), or on the probability paper (Bröcker and Smith 2007), may then be more appropriate.

4 Recalibration and TIGGE dataset

The TIGGE dataset is an archive consisting of ensemble forecast data from ten global NWP centres, which has been made available for scientific research. The dataset starts from October 2006. TIGGE stands for THORPEX Interactive Grand Global Ensemble. It was initially set with the objective of accelerating the improvements in the accuracy of 1-day to 2-week high-impact weather forecasts. This dataset is widely used today in a wide range of research projects and for various applications. More information is available at the TIGGE website hosted by ECMWF (TIGGE 2010), while an overview of the system and the rationale behind the TIGGE project is presented in Bougeault *et al.* (2010).

Some of the early works based on the TIGGE dataset are those reported by Park *et al.* (2008), where emphasis is placed on evaluating the quality of the individual components of the TIGGE archive, over a period covering the first 3 months (October - December 2007). This study focused on 500-hPa geopotential height and 850 hPa temperature. The various individual components were evaluated against their own analysis, but also against other analyses (to study the sensitivity of verification results with the choice of the reference). One of the major results of the work of Park *et al.* (2008) is the illustration of the major differences in forecast quality of the 10 global ensemble prediction models. A clear advantage was given

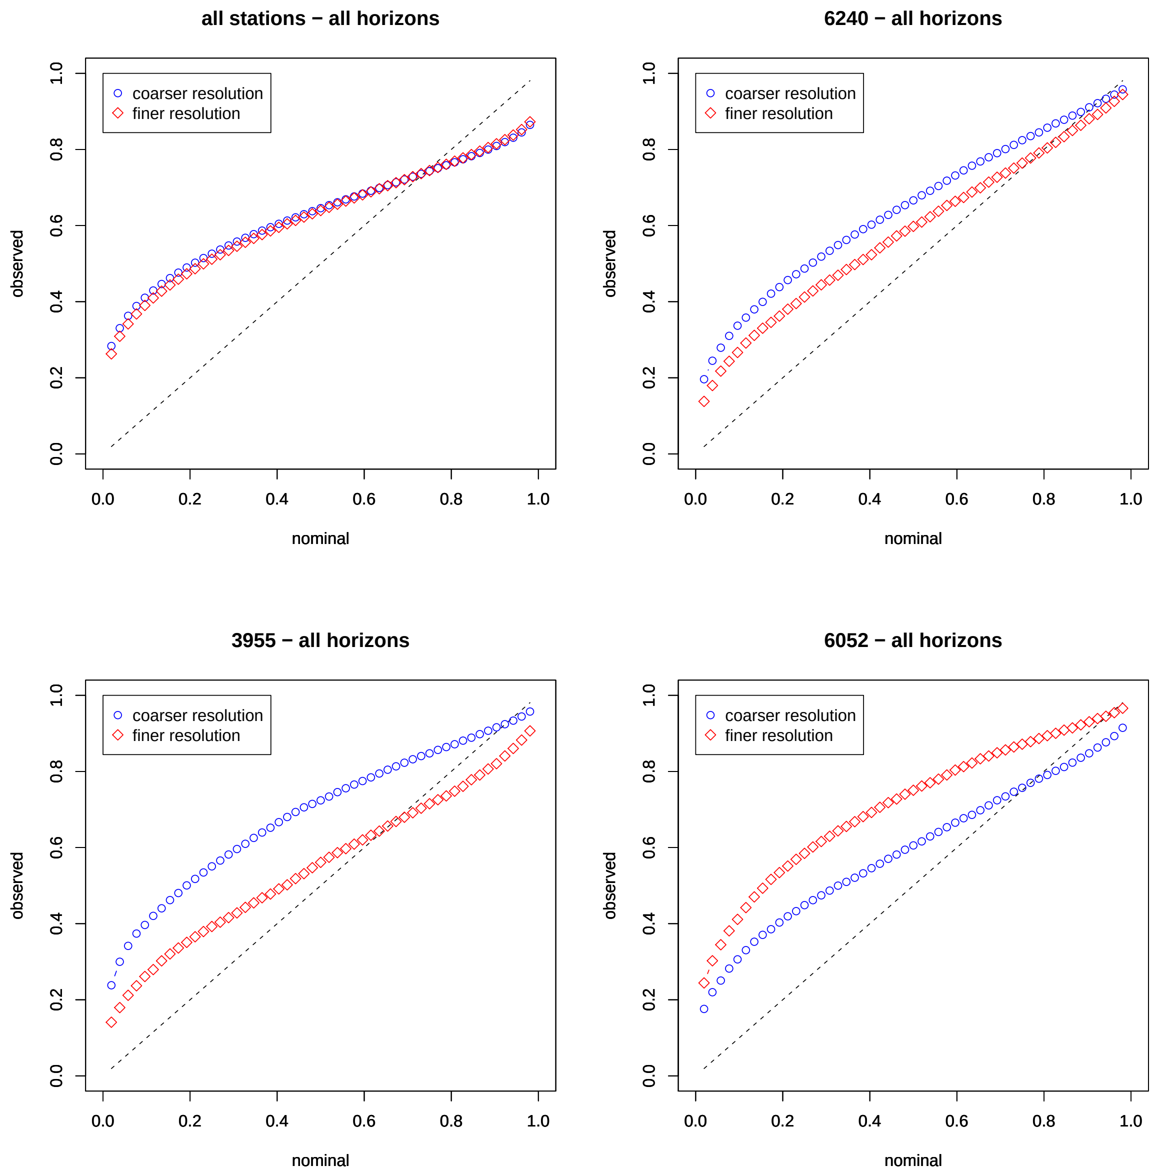

FIGURE 4: *Impact of the change of horizontal resolution on the reliability of ensemble forecasts at synoptic stations. Reliability diagrams as depicted as the cumulative form of PIT histograms. These diagrams are for all stations (top left), Amsterdam Schipol airport (station id 6240, top right), Cork airport (station id 3955, bottom left) and finally Thyboron station in Denmark (station id 6052, bottom right). All horizons are considered indifferently.*

to that of ECMWF. Park *et al.* (2008) finally hinted at the possibility of combining all these individual components in order to maximize the skill of resulting ensemble forecasts.

Further investigating this idea of combining the individual components of the TIGGE dataset, the results of the study performed by Hagedorn *et al.* (2010) at ECMWF has shown the respective merit of combination against recalibrating the best component (i.e. the ECMWF ensemble forecasts). A first conclusion from this study is that combination appears to be beneficial mainly if picking appropriate components

as input, and not if just aiming at combining all inputs with accounting for their respective skill and information content. In comparison, picking the best component and carrying out appropriate recalibration appears to yield forecasts of similar level of quality, if not better. This is somehow due to the fact that the various models for global ensemble forecasting are quite similar, leading to a lack of new and independent information brought in by the other systems. In such a case picking (and recalibrating) the best ensemble forecasting system may be seen as sufficient for optimizing forecast quality. A practical interest of TIGGE dataset is, however, that one can stack the various ensemble forecast components and increase the number of alternative scenarios for the coming period. These scenarios can then be used as input to various types of decision-making problems.

For the case of 10-metre wind speed, example results are discussed in the following. A first evaluation of the ability of the various models in the TIGGE dataset to provide reliable probabilistic forecasts of wind speed was based on their spread-skill relationship. In other words, this comprises an evaluation of how the uncertainty expressed by the ensemble forecasts, a priori, is consistent with the actual uncertainty of the forecasts observed a posteriori. The results from this evaluation are gathered in Figure 5 for the Northern Hemisphere and for two different lead times, 1-day and 7-day ahead. Are considered the ensemble forecasts from CMC (Canada), ECMWF, Met Office (UK), NCEP (United States), as well as the overall TIGGE forecasts (obtained by stacking all available ensembles) and the recalibrated ECMWF ensemble forecasts. Recalibration is based on the reforecasts dataset available at ECMWF, while employing Nonhomogeneous Gaussian regression (NGR) as initially introduced by Gneiting *et al.* (2005).

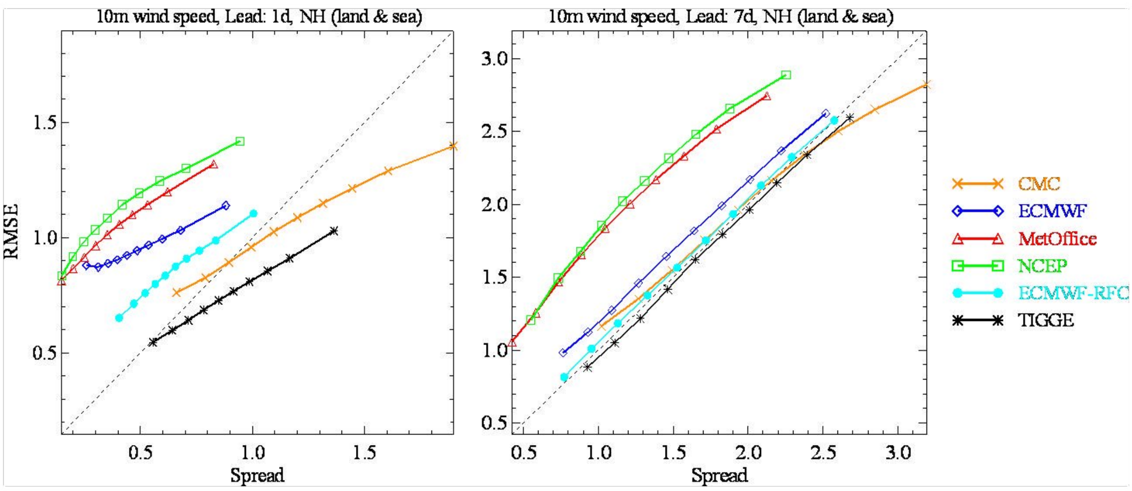

FIGURE 5: *Spread-skill TIGGE, days 1 and 7.*

The ensemble forecasts from the various centres tend to be under-dispersive, especially for the early-range, meaning that their spread is too small in comparison with actual uncertainty. This typical result can be observed for the 1-day ahead lead time in Figure 5 for most of the ensemble prediction systems. As mentioned in the above, one can directly observe one of the benefits of the TIGGE dataset which is that by stacking all ensemble forecasts one obtain ensemble forecasts with higher spread. These certainly are more realistic than the under-dispersive individual components. For the medium-range, and with more particular focus on 7-day lead times here, there is generally a better agreement between ensemble spread and skill. Ensemble forecasts from NCEP and the Met Office still seem to be under-dispersive

here though.

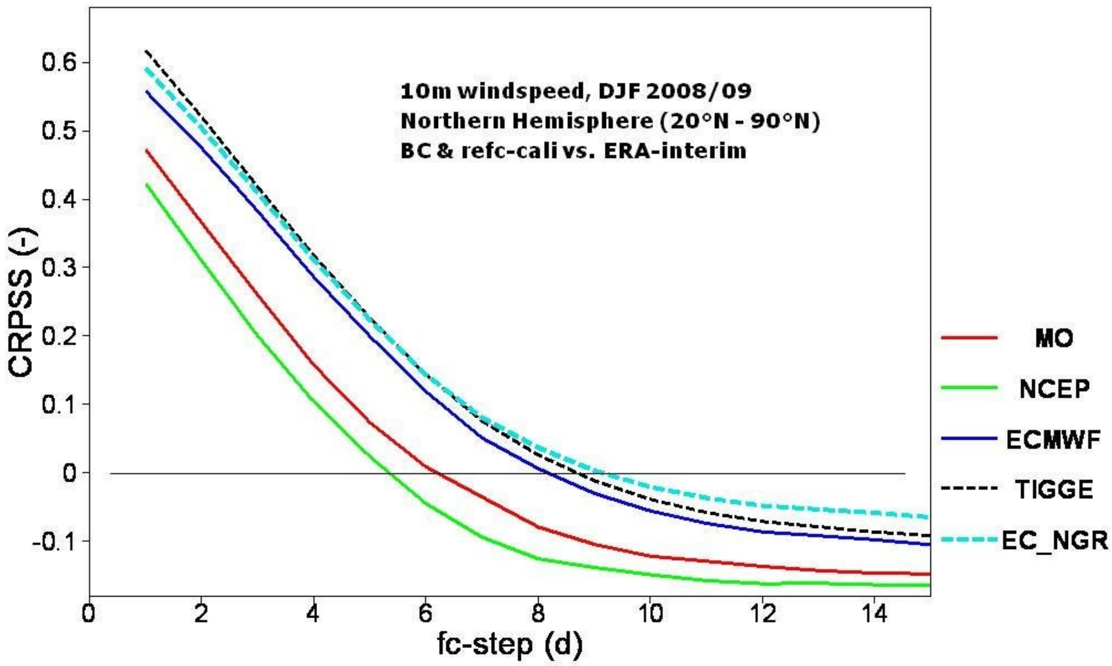

FIGURE 6: *TIGGE CRPSS.*

5 New representation of initial and model uncertainties

Originally, ensemble forecasts at ECMWF were produced based on the simulation of initial uncertainties only, with a singular-vector approach. Singular vectors form a basis for the phase-space directions of maximum growth during a time interval covering the first 48 hours of the forecast series. They also are scaled so that they have an amplitude comparable to analysis error estimates. The methodology originally developed for the computation of singular vectors is presented in Buizza and Palmer (1995). The underlying idea of employing singular vectors is that small errors in the initial conditions along these directions would amplify most rapidly, and further affect forecast accuracy.

It was later realized that non only the misspecification of the initial conditions but also the misspecification of models parameters were affecting forecast uncertainty. A very comprehensive paper on that aspect was published by Orrell (2005). AT ECMWF, this problem was tackled by implementing a stochastic parameterization of the forecasting model, which is extensively described and discussed in e.g. Buizza *et al.* (1999).

Since that time, extensive research on the topic on initial and model uncertainties has been performed at ECMWF. Recently, emphasis has been placed especially on:

- representing initial uncertainties by combining perturbations from the ensemble of perturbed 4D-

Vars (Ensemble Data Assimilation - EDA) and by perturbations based on initial singular vectors;

- revising the methodology for the representation of model uncertainties, by adding a Spectral Stochastic Backscatter Scheme (SPBS), and by revising the Stochastically Perturbed Parameterisation Tendency (SPPT) scheme.

So far a few experiments have been performed and operational implementation in under progress. The methodology developed for the EDA part is described by Buizza *et al.* (2010) and by Isaksen *et al.* (2010).

Regarding initial uncertainties, the amplitude of the singular vector perturbations has been halved, in order to address the problem of the lack of spread of the ensembles in the early to medium-range. This reduction of the amplitude of the singular vector perturbations actually makes the EPS more under-dispersive and thus leads to a reduction in probabilistic skill. To re-establish a balance, more weight is given to the representation of model uncertainties, based on both the SPBS and SPPT schemes. For the SPPT part, the changes mainly consists of considering spatial and temporal structures at three different scales. They roughly are of 6 hours, 3 days and a month for the temporal axis, while being of 500, 1000 and 2000 kms for the spatial axis. Experiments for evaluating the benefits from such new approaches to the representation of initial and model uncertainties have been carried out at the operational resolution of the ECMWF ensemble forecasting system. The results obtained indicated that the improvements in predicting upper air variables (500 hPa geopotential, 850 hPa wind and temperature) are significant in the Northern Hemisphere Extra-tropics and in the Tropics, while not degradation of forecast skill can be observed otherwise. Further experiments and evaluation need to be performed with particular focus on near-surface variables, especially wind components. A publication on that topic of the revision of the representation of model uncertainties is under preparation.

6 Conclusions and perspectives

Among the various activities of ECMWF within the SafeWind project, the work in Task 5.2 is specifically related to the development and evaluation of novel methods for ensemble forecasting, while being relevant for the wind power application. With that in mind the present report has described the interim results from this work. They are mainly related to the providing of forecasts of 100-metre winds, to the upgrade of the horizontal resolution of the forecasting model, to the TIGGE experiments (combination of global ensemble forecasting systems) and finally to the revision of the methodologies for the representation of initial and model uncertainties.

First of all, the providing of 100-metre wind forecasts and analysis opens the door to numerous verification works, and to their use as input to various wind power prediction methodologies. Few verification results (against observations) are available so far. More will be produced in the second part of the project as the quantity of data available increases with time. In parallel, the interest expressed by a large number of forecast users in various member states makes us hope that more verification results will appear thanks to the work of researchers and practitioners outside of the SafeWind project consortium. Globally, it has been observed that the accuracy and skill of 100-metre wind forecasts is at level almost comparable to that of 10-metre wind forecasts.

The effect of the upgrade in the horizontal resolution of the ECMWF forecasting system has been extensively studied, and its effect on near-surface wind speeds over Europe characterized. The principal

benefits consist of a highly significant bias reduction in the point forecasts, as well as in a slight improvement of the reliability and skill of the ensemble forecasts. Still, this study concluded on the necessity of further post-processing the deterministic and ensemble forecasts issued by ECMWF over Europe based on analysis and observations data, owing to the non-negligible potential for improving the overall skill of wind forecasts.

In parallel, the TIGGE experiments showed that the leading skill of ECMWF ensemble forecasts makes them difficult to outperform when being recalibrated, even if employing advanced combination scheme. A natural benefit from combination, however, is the increase in reliability of the ensemble forecast at no cost. A second practical benefit is the increase in the number of alternative scenarios that can be used as input to various decision-making methodologies in a Monte-Carlo or Stochastic Programming framework. More work should be performed based on the TIGGE dataset, and new combination methods proposed, especially accounting for the correlation of the input components.

The revision of the methodologies for the representation of initial and model uncertainties is the result of long-lasting efforts at ECMWF for improving the reliability and skill of ensemble forecasts. The EDA approach developed is definitely the way to go if aiming at better capturing the initial uncertainties and their dependencies. More work will be performed in the future for characterizing the spatial and temporal structures that would be accounted for. In parallel, the combination of the SPPT and SPBS schemes allowed improving general scores, though further studies should particularly focus on near-surface variables such as wind components. The scope for further improvements exist, though being at the expense of a strong commitment of resources for more research, as well as for the implementation of the methods developed.

Apart from these experiments, work at ECMWF has been focused on the improvement of the reliability and skill of ensemble forecasts by appropriate combination of raw ensemble and high-resolution deterministic forecasts, or recalibration of Europe-wide ensemble forecasts against observations. The methodology involved and corresponding results will be reported in a different document.

References

Bougeault P, and co-authors. The THORPEX Interactive Grand Global Ensemble. *Bulletin of the Amercian Meteorological Society* **91**: 1059–1072.

Bröcker J, Smith LA (2007). Increasing the reliability of reliability diagrams. *Weather and Forecasting* **22**: 651–661.

Buizza R (2010). Horizontal resolution impact on short- and long-range forecast error. *Quarterly Journal of the Royal Meteorological Society* **136**: 1020-1035.

Buizza R, Leutbecher M, Isaksen L, Haseler J (2010). Combined use of EDA- and SV-based perturbations in the EPS. *ECMWF Newsletter* **123**: 22–28.

Buizza R, Miller M, Palmer T (1999). Stochastic simulation of model uncertainties in the ECMWF ensemble prediction system. *Quarterly Journal of the Royal Meteorological Society* **125**: 2887–2908

Buizza R, Palmer T (1995). The singular-vector structure of the atmospheric general circulation. *Journal of Atmospheric Sciences* **52**: 1434–1456.

Candille G, Talagrand O (2008). Impact of observational error on the validation of ensemble prediction systems. *Quarterly Journal of the Royal Meteorological Society* **134**: 957–971.

Gneiting T, Raftery AE, Westveld III AH, Goldman T (2005). Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Monthly Weather Review* **133**: 1098-1118.

Hagedorn R, Buizza R, Hamill TM, Leutbecher M, Palmer TN (2010). Comparing TIGGE multi-model forecasts with reforecast-calibrated ECMWF ensemble forecasts. *Monthly Weather Review*, submitted.

Isaksen L, Haseler J, Buizza R, Leutbecher M (2010). The new ensemble of data assimilations. *ECMWF Newsletter* **123**: 17–21.

Marzban C, Wang R, Kong F, Leyton S (2010). On the effect of correlations on rank histograms: reliability of temperature and wind-speed forecasts from fine-scale ensemble reforecasts. *Monthly Weather Review*, available online.

Miller M, Buizza R, Haseler J, Hortal M, Janssen P, Untch A (2010). Increased resolution in the ECMWF deterministic and ensemble prediction systems. *ECMWF Newsletter* **124**: 10–16.

Orrell D (2005). Ensemble forecasting in a system with model error. *Journal of Atmospheric Sciences* **62**: 1652–1659.

Park YY, Buizza R, Leutbecher M. TIGGE: Preliminary results on comparing and combining ensembles. *Quarterly Journal of the Royal Meteorological Society* **134**: 2029–2050.

Petroliagis T (2010). Wind power forecasting utilizing global and limited-area ensemble prediction. ECMWF Forecasts Products Users Meeting 2010. [Available at: `http://www.ecmwf.int/newsevents/meetings/forecast_products_user/index.html`]

Pinson P, Hagedorn R (2010). Verification of the ECMWF ensemble forecasts of wind speed against observations. ECMWF Technical memorandum.

Pinson P, Nielsen H.Aa., Møller J.K., Madsen H, Kariniotakis G (2007). Nonparametric probabilistic forecasts of wind power: required properties and evaluation. *Wind Energy* **10**:497–516.

Plester J (2010). Weather eye: Bountiful supply of wind. The Times, edition of the 19th August.

TiGGE website: `http://tigge.ecmwf.int/`; with a list of related publications at: `http://tigge.ecmwf.int/references.html`