Grid-based modeling in "Wissensnetz Energiemeteorologie"*

Jan Ploski[1], Thomas Petroliagis[2], Detlev Heinemann[2,3], Thomas Scheidsteger[3], and Wilhelm Hasselbring[1,4]

[1] OFFIS, Escherweg 2, 26121 Oldenburg, Germany,
jan.ploski@offis.de,
[2] ForWind - Center for Wind Energy Research,
Marie-Curie-Str. 1, 26129 Oldenburg, Germany,
{thomas.petroliagis,detlev.heinemann@forwind.de},
[3] University of Oldenburg, Institute of Physics, 26111 Oldenburg, Germany
thomas.scheidsteger@uni-oldenburg.de,
[4] University of Oldenburg, Software Engineering Group, 26111 Oldenburg, Germany,
hasselbring@informatik.uni-oldenburg.de

Abstract. The research project WISENT (Wissensnetz Energiemeteorologie) combines efforts of computer and atmospheric scientists to optimize the cooperation of scientific organizations in the field of energy meteorology. One of the goals is to provide easy access to distributed computing resources of the German Grid (D-Grid) for running numerical weather prediction models. Based on experience with our introductory implementation of the Weather Research and Forecasting (WRF) model in the computing cluster at OFFIS, the applicability of Grid computing approaches in this context along with their benefits and technical challenges in the areas of usability, multi-user operation, software configuration and data management are examined. Results of verification tests of different WRF configurations are presented. Daily WRF runs for March, April & May (MAM) 2007 period have been performed in an "off-the-shelf" configuration with a forecast horizon of up to 48 hours, based on the NCEP analysis cycle of 00 UTC. The initial results showed that, meteorologically, the WRF forecasts did not perform as well as expected. Our future work will focus on tuning WRF to improve its prediction quality and on overcoming the technical challenges to establish numerical weather prediction as one of the scientific applications in D-Grid.

1 Introduction

In our research project WISENT [1], funded by the German Federal Ministry of Education and Research [2], computer scientists and meteorologists work together to optimize the cooperation of scientific organizations in the field of energy meteorology. One of the work packages of the project focuses on providing easy access to distributed computing resources of the German Grid [3] for running NWP (Numerical Weather Prediction) models.

This paper is structured as follows. In Section 2 we introduce our application, a NWP model in the domain of energy meteorology. In Section 3 we briefly explain the general concept of computational Grids[1]. This concept is realized by the German Grid, described in Section 4. In Section 5 we first discuss the expected benefits of Grid-based modeling. We also present the technological challenges that need to be addressed for the envisioned benefits to become true and our initial implementation steps. A meteorological study which served to verify the WRF model using our established technical infrastructure follows in Section 6. We summarize and conclude the paper in Section 7.

2 Numerical Weather Prediction with WRF

The German name of our project, WISENT, is an acronym for "knowledge net energy meteorology". The main objective of energy meteorology is obtaining the information needed to characterize the fluctuating

* This work is supported by the German Federal Ministry of Education and Research (BMBF) under grant No. 01C5968.

[1] Note that the term *grid*, with a lowercase *g*, is also used in this paper and refers to the data representation within a model.

generation of solar and wind energy. Fulfilling this objective fundamentally depends on our ability to model and predict the weather. Our knowledge acquisition is based on interdisciplinary cooperation: physical and meteorological methods, transformation of wind and solar energy (physics), energy supply structures (electrical engineering, economics), monitoring methods (computer science, sensor technology) and efficient, flexible distributed systems (computer science). Not surprisingly, one of the primary means of representing knowledge related to energy meteorology is by installing, implementing, and optimizing weather prediction models.

The Weather Research and Forecasting (WRF) model [4], on which we focused our initial efforts in Grid context, is a mesoscale NWP model, suitable for research and operations, capable of running on a variety of platforms, either serially or in parallel, with or without multi-threading.

The current WRF Version 2.2 (December 2006) was installed, which is available to the general community for download at http://www.wrf-model.org. WRF makes use of the ARW (Advanced Research WRF) solver, which is composed of several initialization programs for idealized and real-data simulations, and the numerical integration program.

WRF is fully compressible, Euler non-hydrostatic with a run-time hydrostatic option available and it is conservative for scalar variables. Its prognostic variables are: the velocity components u and v in Cartesian coordinates, the vertical velocity w, the perturbation potential temperature, the perturbation geopotential, and the perturbation surface pressure of dry air. Optionally, the turbulent kinetic energy and any number of scalars such as water vapor mixing ratio, rain/snow mixing ratio, and cloud water/ice mixing ratio can be also predicted. The model's vertical coordinate is the terrain-following hydrostatic-pressure, with vertical grid stretching permitted. The top of the model is a constant pressure surface. Its horizontal grid is the Arakawa C-grid staggering. Its time integration scheme is a time-split integration using a 3rd order Runge-Kutta scheme with smaller time step for acoustic and gravity-wave modes. Its spatial discretization comprises 2nd to 6th order advection options in horizontal and vertical. WRF provides the user with the capability to use both turbulent mixing and model filters (WRF tuning parameters). An explicit filter option is also available.

3 Grid Computing

Numerical weather prediction, like many other scientific applications, requires significant computing resources. As new and more powerful computer hardware becomes available, the employed meteorological models such as WRF can be easily refined by increasing their time and space resolution, which further drives the demand for CPU cycles, memory and disk space. Moreover, as the amount of (market) interest in meteorological data and the diversity of end-user applications grows, so does the number of executed model runs. Overall, these challenges can be referred to as *vertical* and *horizontal* scalability: the former means to make existing models run faster and improved models run fast enough; the latter to make it possible to run more model instances concurrently.

Traditionally, the scalability of scientific applications has been addressed by building bigger supercomputers with special high-performance hardware. The advent of cluster computing in the mid-90s made it possible to replace some of these supercomputers with less expensive commodity servers connected through fast local area networks. Computing clusters are especially attractive for improving the horizontal scalability, as they can accomodate many users and can be extended relatively easily. However, in practice not every organization interested in weather prediction can afford buying and administering an own cluster. For this reason, access to a computing cluster is typically offered to scientific users by a computing center.

For many users, having access to a single computing cluster, together with a service-level agreement about its availability, would be just enough to satisfy their operational forecasting needs. However, in contrast to operational use, research tends to generate high workloads at unpredictable times. A single computing center may not be able to match these ad-hoc needs timely. On the other hand, if one could pool the computing power (along with RAM, disk space) of multiple computing centers, it might well turn out that sufficient resources are available. This idea lies at heart of *Grid computing*, which is more abstractly defined as *coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations* [5]. The term *resource* in this definition in most direct sense refers to hardware. However, other kinds of resources are also required for successful distributed computing, such as the correctly programmed and installed software and the available expertise. The need for *virtual organizations* arises from the simple fact that the persons engaged in providing and consuming computing resources usually belong to many different institutions. These cannot individually keep track of the identities and privileges of each others' participants, but must find ways to exchange such crucial information.

4 The German Grid

While it is not unusual to hear or read about "the Grid", in reality there are many. A Grid is created when several computing centers agree to collaborate in sharing workloads generated by their users and install some appropriate supporting software. The German Grid (D-Grid [3]) is an example of such an initiative, funded by the federal government.

At present, 20 different organizations contribute their computing clusters and 11 organizations contribute dedicated storage space to D-Grid. A contributed cluster (or a part thereof) is also known as a "Grid site". The hardware in each D-Grid cluster varies (from 4 to 32 CPUs per nodes, 4 to 128 GB per node). The largest such cluster, installed at FZK (Karlsruhe), makes 2700 CPU cores available to D-Grid users. The cumulative amount of disk space managed by D-Grid is approximately 2 petabytes. Each participating organization is responsible for administering their own Grid site, but the administration is coordinated through a central D-Grid integration project.

The basic access to D-Grid resources is officially provided through three different kinds of Grid software: Globus, gLite and UNICORE, although not every participating site supports each software. This software provides to users two basic capabilities:

1. submit computational jobs; that is, run a Unix executable or script on one or more of the Grid site's available machines ("cluster nodes")
2. transfer data files; either from a local machine to a Grid site or between Grid sites

In theory, every D-Grid participant may access every available computing resource, as well as the amount of storage space allocated to their virtual organization (such as WISENT). To become a D-Grid participant, one must meet the following requirements:

1. belong to or at least be affiliated with one of the D-Grid projects
2. apply for a "user certificate" from your organization's own Grid Registration Authority[2]; in exchange for signing an acceptable use policy and presenting your personal identity card to an administrator, you receive a pair of files which are installed on your computer to prove your identity to Grid sites
3. visit a web site to register with one of the "virtual organizations" set up in D-Grid (for example, with the WISENT-VO)
4. become familiar with the above mentioned software

5 Benefits and Challenges of Grid-based Modeling

To introduce the basic concept of Grid computing in this paper, we focused on the most obvious scalability benefits. In this section, we list some expected benefits that are less apparent. Finally, we report our practical experiences with NWP in context of D-Grid and the additional requirements which are specific to running NWP models.

5.1 Benefits

Successfully implementing Grid computing means not only acquiring uniform access to multiple computing centers, but also *re-thinking existing scientific workflows* to take advantage of such access. Based on our experiences, we recognize potential benefits in the following areas:

Availability. Increase robustness of workflows by depending less on individual resources. Use one Grid site as a backup for another.

[2] In practice, such an authority must first become established by an organization. Although not technically difficult, it is a major bureaucratic hurdle for many Grid participants.

Separation of concerns. Improve productivity by better division of labor. Delegate solving technical problems to Grid site maintainers and application developers, while focusing on meteorological applications.

Quality of service. Let multiple Grid sites compete for workloads based on the quality of technical assistance, availability and performance. Expect overall better service than from a single exclusive provider.

Comparability of results. Compare results produced by different models, or even by differently compiled versions of the same model, running in different hardware environments. Benchmark commercial compilers against each other.

Community building. Learn modern technology to support efforts of colleagues and educate students. Transfer knowledge about Grid computing from IT experts to meteorologists. Transfer knowledge about meteorological applications in the opposite direction.

Business models. Approach companies with new offers that previously could not be made because of lacking computational resources.

Data access. Search for data more efficiently, describe data more precisely, transfer large amounts of data quicker than before.

Traceability of results. Publish scientific results in a way which makes it easier for other researchers (as well as the original authors) to reproduce them.

5.2 Challenges

Despite the above described benefits and being part of D-Grid with an own cluster located in OFFIS (Oldenburg), our implementation of WRF in the WISENT project started in a very "traditional" way. First, the software and the supporting libraries and tools needed for parallel execution of the model were installed. This process was slowed down by the need to acquire a working knowledge of cluster computing, Grid technology and to familiarize ourselves with WRF tools more or less at the same time. Second, the necessary Global Forecast System (GFS) data was downloaded from the NCDC NOMADS [6] server to support the model's execution. In order to speed up downloading, the real-time and historical data was downloaded in parallel to the cluster in Oldenburg and to another D-Grid cluster in Jülich. Our users, whose main priority was getting access to the computing power and storage space, were granted direct Unix shell access to the computing cluster. Furthermore, they were provided with technical support related to running their jobs in the cluster. In summary, while we got the model up and running in Oldenburg (and are close to a similar working configuration in Jülich), little of the genuine Grid technology or Grid-related workflows have entered the picture so far. Why?

The WRF software is normally installed on a single machine or on a single computing cluster by an end-user or for a limited group of users. Its configuration and installation requires substantial effort and technical knowledge. The software can be configured and compiled in a great number of variants with tradeoffs in runtime performance and in precision of results. To obtain useful results (such as wind speed predictions at specified locations), post-processing must be applied, which means installing, configuring and understanding additional software. The involved FORTRAN programs tend to be very sensitive to user mistakes and they produce invalid output or difficult to comprehend error messages upon failure. As a result, time-consuming configuration work has to be performed, even before the first Grid job is submitted. On a technical level, the WRF configuration consists of editing multiple text files with different semantics, compiling various source codes, and arranging the file system's contents in a manner expected (but not always fully documented) by the software components.

We discovered that our model users are goal-oriented, pressed for time, independent, and inventive enough to help themselves with all their prior accumulated experience. There are few users who have the freedom and interest in learning new technologies, software tools, changing their established work procedures, and yet still coping with daily routine tasks at the same time. Technological innovations such as Grid rarely directly address a scientific community's most immediate problems. To become accepted, they have to be ready-to-use and easy-to-learn. Ideally, the users should not need to understand technical aspects of Grid computing at all, but nonetheless benefit from the new technology through the improved quality of process steps: models running faster and reliably, post-processing tools becoming easier to understand and apply, input data being available in an instant, and so on. The bad news is that the present Grid technology is bound to miss such high expectations. A lot of joint work and patient communication among meteorologists and computer scientists is necessary. Our main immediate areas of concern can be summarized as follows:

Usability. A key question for Grid-based modeling is: what kind of user interface(s) should the NWP models and their pre- and post-processing tools expose to users? To even begin answering this question, we must consider what solutions users are comfortable with today, and we must understand their working style. Traditionally, NWP users rely on interactive Unix shell access, programs with command-line parameters and/or manually edited configuration files, self-authored utility programs and commercial data processing applications (such as MATLAB [7], Hugin [8]). In face of such diversity, any Grid project immediately becomes an IT integration project. Furthermore, a significant risk exists that the technology will be rejected because of limited compatibility with technically mature commercial alternatives. Unfortunately, there is currently no single dominant approach to Grid-based integration. Desktop-based front-ends (such as gEclipse [9]) and web portals (such as GridSphere [10]) both promise to remove some of the technical details from the end users' view, potentially saving time and preventing mistakes by automating often needed sequences of operational (batch processing) steps. However, they lack in power with regard to detailed model customization, ad-hoc scripting capabilities, integration with proprietary tools, operational transparency, and interactive data manipulation capabilities. Overall, it is easier to implement an individual, well-specified workflow on the Grid than to provide a single working environment which balances the need for flexibility and ease-of-use. However, to successfully compete with older, proprietary technology for automated execution of NWP models (whose maintenance can be afforded only by larger institutions) and proprietary software for manual data post-processing (with associated software license costs), Grid technology must form the basis for a general, open problem solving framework rather than become an amalgamate of point solutions.

Multi-user operation and access. The WRF model as well as the pre- and post-processing tools make strong assumptions about the relative locations of different files on the execution machine. It is a common approach among users to copy and compile an entire model in a separate directory tree when preparing a new run, or to selectively reuse parts of a previously installed model by creating symbolic links in the file system. However, this simple procedure does not translate well into the distributed world of Grid computing where a single installed model instance should be accessed by multiple concurrent users. Individual model runs hardly translate into traditional "Grid jobs" that copy the required executables and input data to the user's home directory on some target machine inside of a cluster (a so-called "stage-in phase") and transfer data out after completion. An installed model occupies approximately 300 MB of disk space. The geographical data, which is routinely reused across runs, occupies 9.6 GB. Thus, due to bandwidth limitations, models must be pre-installed in their target execution environment. This is another aspect to which the current Grid technology does not provide ready answers.

Software configuration management. Software configuration management is a subfield of software engineering concerned with the tools and procedures for successfully maintaining multiple versions of a software product. The differences in software that one must keep track of are primarily caused by the process of software evolution, customization and the heterogeneity of existing software platforms. Software configuration management is typically not an issue for small-sized projects (with respect to the number of users, developers and supported execution environments), but becomes crucial when projects expand. The notion of distributed numerical weather prediction on the Grid naturally includes software configuration management challenges. Different Grid sites already run different versions of the operating system, libraries, compilers, utilities. Thus, installing software such as WRF on the Grid scale is difficult. Keeping it up-to-date and well-performing despite the possibility of uncontrollable changes in the environment (such as patches, system upgrades) is another issue that must be considered up-front. In addition to general software configuration management, model configuration management is also necessary to provide the ability to reproduce historical runs and re-generate past results.

Data management. NWP models both consume and produce considerable amounts of data. For example, the input data processed to obtain the verification results presented in Section 6 approaches 300 GB, while the output sum of the model approaches 360 GB. Not only must adequate disk space be provided in the target execution environment to accomodate the running models, but (some of) the results must be also archived to make it possible to reproduce research results. Furthermore, near-real time data must often be collected in advance, as ordering archived data from the original vendors introduces unacceptable delays. Ideally, the whole set of possible input data should be available at each Grid site where a model can execute. However, Grid sites have limited storage capability (e.g., 4 TB in the OFFIS cluster) and in practice data has to be transferred between sites at some time before model execution. Besides of space management, the issue of metadata management (both input- and process-related) is also important to ensure the traceability of results.

6 Verification of WRF

Verification of a model refers to the processes and techniques used to assure that the model is correct and matches certain specifications and assumptions. During our initial verification tests different model configurations were examined. Nested one- or two-way formulations of WRF with its inner grid to take values as low as 1 km were configured and run in a single or multi-threading environment. For a typical two-way WRF two domain nested formulation, with an inner grid of 5 km covering the greater area of Germany of 241 x 241 grid points, and a vertical resolution comprising 35 levels, the completion of a 48 hour forecast took 23.1 hours on a single processor. Running WRF (with exactly the same configuration setup) in a parallel environment using 16 CPUs and MVAPICH [11], the integration time dropped down to 2.6 hours, showing the significance, potential and capabilities of parallelization techniques.

For our initial verification purposes, daily WRF runs for March, April & May (MAM) 2007 period were performed with a forecast horizon of up to 48 hours, based on the NCEP analysis cycle of 00 UTC (initial & boundary fields). A two-way nesting was employed, utilizing two domains. The outer (coarse) domain was configured with a 45 km grid resolution, while the inner domain (covering Germany) had a resolution of 15 km (ratio parent to child 3:1). WRF uses an Arakawa-C stagger scheme allowing an equal number of points in the grid to be distributed across the 16 processors used on the computational platform. The Runge-Kutta solver used allowed a long time step of 90 seconds to be implemented despite the 15 km (inner) grid spacing. In the vertical, 35 full levels (34 computational layers for the mass variables) were used. This formulation was chosen to serve as our "basis" version, for future intercomparisons. The physical parametrizations used in our basis version are summarized in Table 1.

Parameter	**Value**
Land-surface model	Thermal diffusion scheme
Microphysics	WRF Single-Moment (WSM) 3-class simple ice scheme. A simple efficient scheme with ice and snow processes suitable for mesoscale grid sizes. [12]
Cumulus parametrization	Kain-Fritsch (new ETA) scheme
Planetary Boundary Layer (PBL)	Yonsei University Scheme (YSU): Next generation MRF-PBL. Non-local-K scheme with an explicit entrainment layer and parabolic K profile in unstable mixed layer. [4]
Shortwave radiation	Dudhia scheme. A simple downward integration allowing for efficient cloud and clear-sky absorption and scattering. [13]
Longwave radiation	Rapid Radiative Transfer Model (RRTM) scheme. An accurate scheme using look-up tables for efficiency. Accounts for multiple bands, trace gases, and microphysics species. [14]

Table 1. Physical parameterizations used during MAM WRF model integrations

Using objective verification scores such as Bias (Mean Error), MAE (Mean Absolute Error) and RMSE (Root Mean Square Error), WRF 24 and 48 hour forecasts were verified against corresponding NCEP analysis fields. A similar set of scores was used for NCEP 24 and 48 hour forecasts, as shown in Figures 1, 2, and 3.

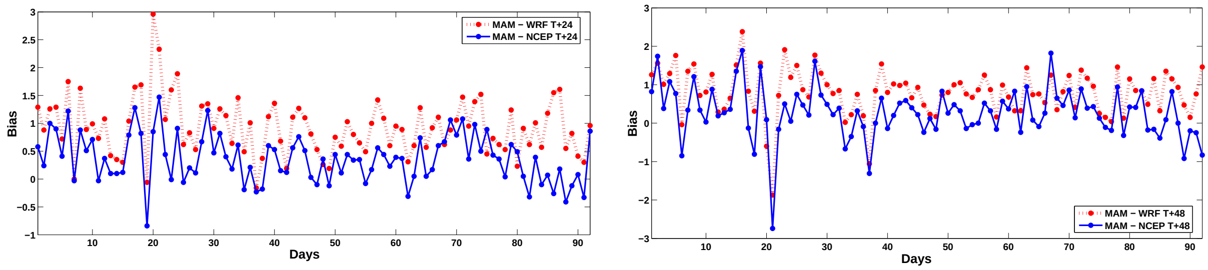

(a) (b)

Fig. 1. Bias (Mean Error) for WRF and NCEP 24 (a) and 48 (b) hour forecasts for wind speed at 10 meters valid for MAM period

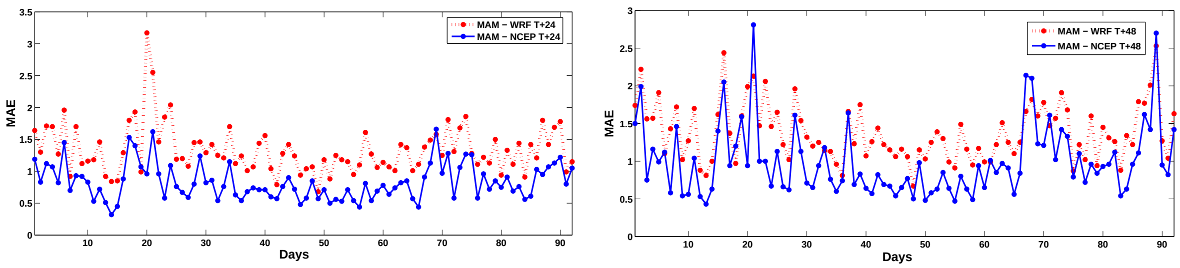

(a) (b)

Fig. 2. MAE (Mean Absolute Error) for WRF and NCEP 24 (a) and 48 (b) hour forecasts for wind speed at 10 meters valid for MAM period

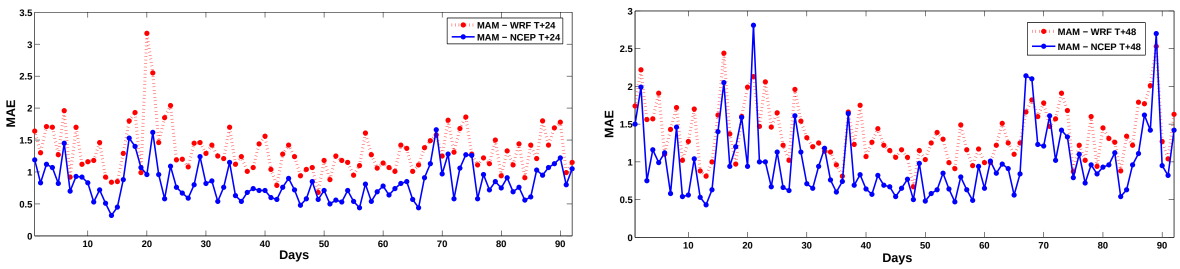

(a) (b)

Fig. 3. RMSE (Root Mean Square Error) for WRF and NCEP 24 (a) and 48 (b) hour forecasts for wind speed at 10 meters valid for MAM period

7 Conclusions and Future Work

The results of our efforts on Grid-based weather modeling in the WISENT project can be broadly described in organizational, technological and meteorological categories.

The organizational achievements lie in establishing a working relationship between the researchers from ForWind and OFFIS as well as bringing each project partner to understand the most important requirements and visions of another. Furthermore, the installed WRF model can be easily tailored to local needs and interests, serving all future potential Grid users, scientists and researchers, even university students at an undergraduate level.

The technological achievements consist of setting up a functional high-performance cluster integrated into D-Grid and capable of running the WRF model. Furthermore, thanks to the utilization of the Grid file transfer technology, once input data is downloaded initially, it can be distributed to execution sites much faster than if it had to be acquired from the original vendor (up to 10 MB/s vs. 200 kB/s).

Initial results showed that, meteorologically, the WRF forecasts did not perform as well as expected. However, running a mesoscale model capable of simulating correctly local topographic effects takes time and a lot of tuning. WRF wind forecasts exhibit a positive bias at 10 meter height level (WRF flow too strong) in both the 24 and 48 hour forecast horizons. Nevertheless, it is important to bear in mind that most of the currently operational mesoscale models (as the NCEP/ETA for example) and their associated post-processed fields (e.g., 2 meter temperature, 10 meter winds) have undergone extensive tuning since their implementations, most of them for several years, whereas the WRF model is a relatively new model and was used by us in an "off-the-shelf" configuration.

Even though it is in the early stages of development, the performance of the WRF model is very encouraging. The problems with low level wind forecasts warrant some investigation into the implementation of the PBL, land surface, and radiation schemes and their interactions within the WRF model, which poses interesting new challenges for the potential users of WRF within the German Grid. A second evaluation period for WRF model is planned to be the last DJF period (December 2006, January and February 2007). Input data for WRF initial and lateral fields are being downloaded exploiting the German Grid capabilities of data retrieval, transferring and storage. The basic two-way nested configuration of 45 (outer) and 15 km (inner) will be applied for the DJF period as well. Verification will be performed against different weather centers' analysis fields and selected points of observations.

This evaluation of the winter period may also provide useful verification data to assess WRF performance in different weather regimes. Both MAM and DJF WRF 24 and 48 hour forecasts will be verified against observations from different wind parks. Besides the 10 meter basis level, the verification of wind characteristics will be performed for different height levels (closer to the turbine hub height) as well. At the time of writing all WRF output files are being converted to GRIB format to provide high-resolution input fields for the ForWind's IWPPP (Integrated Wind Power Prediction Platform). IWPPP focuses on wind power prediction for individual wind farms (on and offshore) and on aggregated wind power forecasts over entire Germany, utilizing different model inputs, mainly from ECMWF and NCEP. ForWind uses a continuous process which compares new model against existing model outputs. WRF higher resolution forecasts will be compared against lower resolution NCEP and ECMWF forecasts, while all model forecasts will be verified against real data (observations).

On the technical side, we are going to continue to expand our deployment of the WRF model to other D-Grid clusters, to establish better automated data management procedures, and work towards addressing the challenges described in Section 5.2. Specifically, we plan to work on constructing a modern software platform to provide easy access to NWP models to new users in the domain of energy meteorology and beyond.

References

1. WISENT: Wissensnetz Energiemeteorologie. (`http://wisent.d-grid.de`) Last accessed: 2007-08-24.
2. BMBF: German Federal Ministry of Education and Research. (`http://www.bmbf.de/`) Last acess: 2007-08-24.
3. D-Grid: e-Science-Framework für Deutschland. (`http://www.d-grid.de/index.php?id=57`) Last accessed: 2007-08-24.

4. Skamarock, W.C., Klemp, J.B., Dudhia, J., Gill, D.O., Barker, D.M., Wang, W., Powers, J.G.: A Description of the Advanced Research WRF Version 2. NCAR Tech Note 468+STR, NCAR (2005) Available online at: http://box.mmm.ucar.edu/wrf/users/docs/arw_v2.pdf.
5. Foster, I., Kesselman, C., Tuecke, S.: The anatomy of the Grid: Enabling scalable virtual organizations. Lecture Notes in Computer Science **2150** (2001) 2–13
6. National Climatic Data Center: NOAA National Operational Model Archive & Distribution System. (http://nomads.ncdc.noaa.gov) Last accessed: 2007-08-24.
7. The MathWorks, Inc.: MATLAB® - The Language of Technical Computing. (http://www.mathworks.com/products/matlab/) Last accessed: 2007-08-24.
8. Hugin Expert A/S: HUGIN Researcher™. (http://www.hugin.com/Products_Services/Products/Academic/Researcher/) Last accessed: 2007-08-24.
9. The g-Eclipse Consortium: g-Eclipse. (http://www.geclipse.org) Last accessed: 2007-08-24.
10. The GridSphere Project: GridSphere Portal Framework. (http://www.gridsphere.org) Last accessed: 2007-08-24.
11. Ohio State University, N.B.C.L.: MVAPICH: MPI over InfiniBand and iWARP. (http://mvapich.cse.ohio-state.edu) Last accessed: 2007-08-24.
12. Hong, S.Y., Dudhia, J., Chen, S.H.: A Revised Approach to Ice Microphysical Processes for the Bulk Parameterization of Clouds and Precipitation. Monthly Weather Review **132**(1) (2004) 103–120
13. Dudhia, J.: Numerical study of convection observed during the winter monsoon experiment using a mesoscale two-dimensional model. Journal of the Atmospheric Sciences **46**(20) (1989) 3077–3107
14. Mlawer, E.J., Taubman, S.J., Brown, P.D., Iacono, M.J., Clough, S.A.: Radiative transfer for inhomogeneous atmosphere: Rrtm, a validated correlated-k model for the longwave. Journal of Geophsical Research **102**(D14) (1997) 16663–16682