E-Infrastructures

H2020-EINFRA-2015-1

EINFRA-5-2015: Centres of Excellence for computing applications

EoCoE

Energy oriented Center of Excellence for computing applications

Grant Agreement Number: EINFRA-676629

D2.2 - M.18

Ultra large meteorological ensemble

Project and Deliverable Information Sheet

EoCoE	Project Ref:	EINFRA-676629
	Project Title:	Energy oriented Centre of Excellence
	Project Web Site:	http://www.eocoe.eu
	Deliverable ID:	D2.2 - M.18
	Lead Beneficiary:	FZ Jülich
	Contact:	Hendrik Elbern
	Contact's e-mail:	h.elbern@fz-juelich.de
	Deliverable Nature:	Report
	Dissemination Level:	PU*
	Contractual Date of Delivery:	M18 31/03/2017
	Actual Date of Delivery:	M18 28/03/2017
	EC Project Officer:	Carlos Morais-Pires

* - The dissemination level are indicated as follows: PU – Public, CO – Confidential, only for members of the consortium (including the Commission Services) CL – Classified, as referred to in Commission Decision 2991/844/EC.

Document Control Sheet

Document	Title :	Ultra large meteorological ensemble
	ID :	D2.2 - M.18
	Available at:	http://www.eocoe.eu
	Software tool:	LaTeX
Authorship	Written by:	Jonas Berndt, Hendrik Elbern
	Contributors:	Charlotte Hoppe
	Reviewed by:	Edouard Audit, Nathalie Girard, Matthieu Haefele, Charlotte Hoppe

Contents

1 Introduction **5**

2 Model description **6**

3 Uncertainty representation **8**

3.1 Stochastically Perturbed Energy Backscatter Scheme 9

3.2 Stochastically Perturbed Parameterization Tendency Scheme 9

3.3 Perturbed Surface Parameter Scheme 10

3.4 Multi-Parameter Parameterization Scheme 11

4 Computational performance analysis and improvements **13**

5 Parallel ultra large ensemble controlling system **16**

6 A case study **18**

A Examples of the extended log-normal distribution **21**

B Extended log-normal variables for the Surface Parameter Perturbation Scheme **22**

List of Figures

1 Domain nesting configuration 7

2 SKEBS' instantaneous forcing pattern 10

3 Validation of ensemble spread for various model error representation schemes 12

4 I/O characterization of WRF on JUQUEEN 14

5 Parallel scaling behavior of WRF on JUQUEEN 15

6 Visualization of the MPI approach to modify the WRF code to a stand-alone ensemble controlling system. 17

7 Computational performance of the stand-alone ensemble controlling system. 18

8 Ensemble diversity for a case study 19

9 Time series of wind speed quantiles at measurement towers for a case study 20

10 Examples of the extended log-normal distribution for the Surface Parameter Perturbation Scheme 21

List of Tables

1 Spatial extent of the domain configuration 8

2 Comparison of computational performance of the WRF code with default and current compiler optimization. 14

3 Comparison of computational performance of the original and computational improved SKEBS scheme. 15

4 Extended log-normal variables for the uncertainty representation of albedo (annual minimum). 22

5 Extended log-normal variables for the uncertainty representation of roughness length (annual minimum). 23

Abstract

This work is motivated by the need to avoid exceptionally high costs of large errors in energy forecasts by developing warning capabilities for cases of low probability. This implies the processing of ensembles, the sizes of which are drastically increased. Therefore, a novel approach of an ultra large ensemble control system is developed. The computational performance of the underlying atmospheric model is monitored on JUQUEEN and a speedup of 100 % is achieved via compiler optimization. Code development is conducted to make existing model error schemes applicable and complemented by two multi-parameter approaches. The system is designed to efficiently realize particle filter assimilation on a HPC architecture. Multiple assimilation cycles can be performed within a single application, enabling communication among numerous ensemble members. The parallelization is demonstrated with up to 4096 ensemble members running on 262'144 cores.

1. Introduction

The main scientific objective to be addressed in Task 2.1 is the short term predictability of wind and solar power with focus on low-probability extreme-error events by ultra-large meteorological ensemble forecasts. Such events are caused by failed forecasts of weather phenomena as well as ramping events. In this context, predictions of single weather forecasts are insufficient and have fundamentally limited usefulness. Firstly, they are optimized to average statistical error values. Secondly, since Numerical Weather Predictions contain errors due to uncertainties of the initial conditions as well as the model formulation itself, the predictions should be furnished with likelihood, leading to a probabilistic approach. The only promising, yet novel path is a transition to stochastic extensions, leading to the integration of ensemble models ([18],[36]).

So far, the migration of ensemble models into the field of power forecasting is yet limited. Only ensembles from meteorological institutions serve as input so far [24, 10, 19, 22] and are therefore restricted in size and resolution. However, higher resolution advances the capture of smaller scale atmospheric and geographic features that differ from site to site. With increasing ensemble size, a wide range of forecast uncertainties of the short-range can be tackled. Most notably, convective scale instabilities, boundary layer dynamics, cloud induced modulation of insolation, and the various mechanisms to trigger or influence these processes, must be accounted for in the parameterization of the physical processes by various perturbations of parameters.

The actual limits of ensemble forecasting shall be ascertained in the frame of a demonstrator which serves as a warning system for low probability high impact events in the energy sector. The aim is to overcome restricted spatial resolution and limited ensemble member number (from ∼50 to ∼1000) to attain an improved approximation of the model state's probability density function. The selection of suitable ensemble members during runtime shall be done using a Sequential Importance Resampling Filter (SIRF, [38]), a fully nonlinear data assimilation technique proven to be applicable to atmospheric systems and most valuable in cases of non-linear observation mapping ([2], [26]). The continuous ensemble updates to improve the model accuracy by not loosing ensemble spread is a main innovation goal of the present project.

The deliverable D2.2 is intended to demonstrate appropriate software developments

to realize an ultra large ensemble of a full fledged regional numerical forecast model, as part of Ensembles for Stochastic Integration of Atmospheric Systems (ESIAS). The newly developed system is suitable for particle filtering and operating with compute cores of the order of 10^5. The underlying atmospheric model is the Weather Research and Forecasting Model (WRF), it's setup is described in Section 2. Various approaches of uncertainty representation within the ensemble are outlined in Section 3, with a focus on physical deficiencies, realized by different stochastic parameterization techniques. A performance analysis of WRF on JUQUEEN is given in Section 4 and principle actions of improvement within this project are summarized. As the major workload of this deliverable, a computational efficient setup controlling various WRF ensemble realizations is introduced in Section 5. This setup is in particular designed for the feasibility of particle filtering on a HPC system. The evaluation of the ultra large meteorological ensemble in the frame of extreme error events is scheduled to deliverable D2.4. The demonstration of the model's full potential in the realm of nonlinear data assimilation conjoined with Big Data Analytics approaches (e.g. Support Vector Machines [6], in coop. with WP1) will be presented in the final deliverable D2.7.

2. Model description

1) Ensemble for Stochastic Integration of Atmospheric Systems

The ultra large ensemble controlling system is part of the the newly developed software Ensemble for Stochastic Integration of Atmospheric Systems (ESIAS) at the Institute of Energy and Climate Research (IEK-8), Forschungszentrum Jülich. The development of ESIAS started from scratch within the EoCoE project, with ongoing support from WP1. ESIAS is designed to later include a chemistry transport model (EURAD–IM, [8]) for aerosol influenced solar radiation forecasts for PV power generation. The coupling realizes enhanced uncertainty representation of energy meteorology related matters or special events with the feasibility of meteorological and chemical data assimilation. Both systems are coupled in the sense that the meteorological ensemble system is used as input for the EURAD–IM ensemble system and vice versa.

2) Weather Research and Forecasting Model

The underlying atmospheric model of the ultra large ensemble controlling system is the Weather Research and Forecasting Model (WRF, see [32] and *http://www.wrf-model.org* for detailed information). WRF is a mesoscale numerical weather prediction system which is used extensively for research and operational real-time forecasting at numerous public research organizations and the private sector throughout the world. The model was developed collaboratively by the National Center for Atmospheric Research (NCAR), the National Centers for Environmental Prediction (NCEP), the Air Force Weather Agency (AFWA), the Naval Research Laboratory, the University of Oklahoma, and the Federal Aviation Administration (FAA), all in the USA. WRF offers ongoing development of various sophisticated physics and dynamics options by leading researchers of the scientific community and is fully open to the public. Within the Advanced Research WRF (ARW) Solver, sophisticated numerical methods are employed. The fully compressible nonhydrostatic Euler equations are solved, which includes the Navier-Stokes equations, the continuity equation, the equation of state for a perfect gas, the thermodynamic energy equation and the equation for conservation of water vapor mixing ratio. Spatial discretization is realized by finite difference scheme using an Arakawa C-grid staggering. As the vertical

coordinate, a terrain-following dry hydrostatic pressure coordinate is chosen and vertical grid stretching is possible. For time integration, a third-order Runge-Kutta method is used. A time-split integration scheme is applied for the high-frequency acoustic and gravity-wave modes. Up to 6th-order centered and upwind-biased advection schemes are available in horizontal and vertical direction. One-way and two-way nesting is possible with multiple nesting levels. The model undergoes updating to the current version during the project period.

3) Domain Configuration

All results stated in this deliverable have been obtained on the nesting domain configuration depicted in Figure 1. Although the target ensemble size of the order of 1000 member limits the outer domain extent to a certain degree, it is still large enough to realize ensemble perturbations on synoptic scales covering most parts of Europe. The first nest covers most of Central Europe and is the target domain for any power predictions aggregated over Germany. The innermost nest with a horizontal resolution of 1.33 km is located over geographical regions with high density power productions for sensitivity experiments. The vertical grid spacing is reduced in the boundary layer, with 4 layers within 100 m above the ground layer. Table 1 summarizes the dimensions of the domain configuration. Spatial resolutions are chosen according to state-of-the-art limited area operational models. A horizontal resolution between 10 to 5 km is intentionally avoided due to poor performances of certain parameterization schemes within this range.

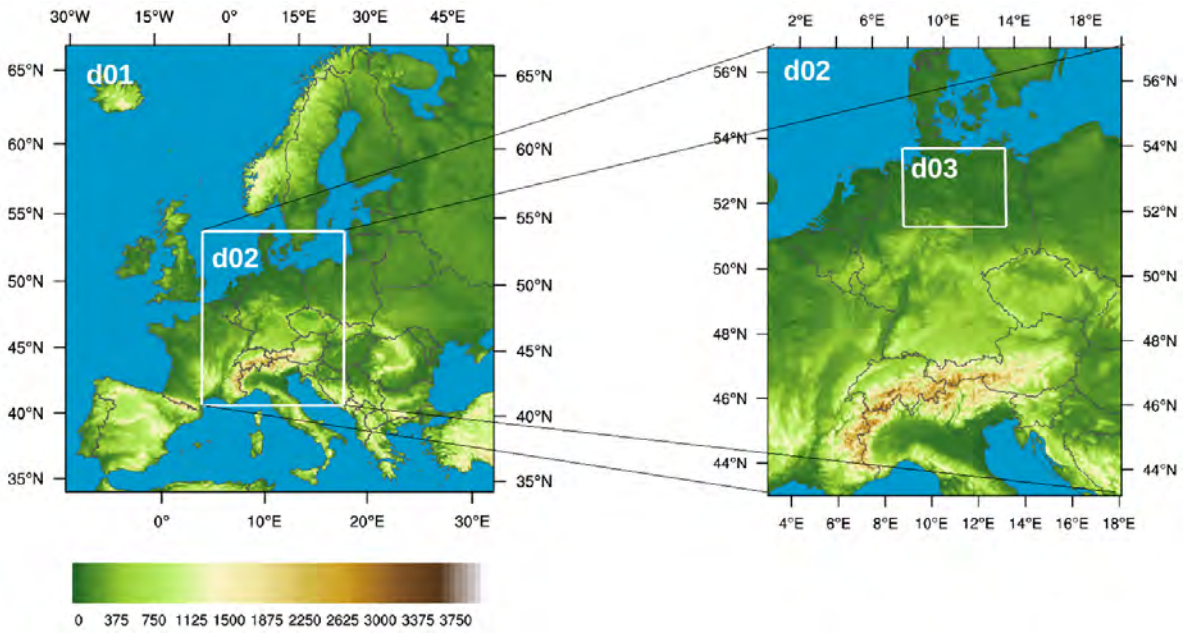

Figure 1: Surface elevation of the domain configuration with it's nesting procedure. The innermost nest is placed in areas of high density wind and/or solar power production.

Table 1: Spatial extent of the domain configuration.

Domain	Horizontal gridpoints	Horizontal resolution	Vertical layers
Mother domain (d01)	330×330	12 km	50
First nest (d02)	301×361	4 km	50
Second nest (d03)	361×301	1.33 km	50

4) Model setup

Each ensemble member uses the same model framework and differs solely by the schemes to represent model uncertainty. The following parameterizations are used: The WRF Single-moment 6-class microphysics scheme [12], the Yonsei University boundary layer scheme [13], the Rapid Radiative Transfer Model for Global Climate Change for shortwave and longwave radiation [14] and the Noah land surface model [35]. On the outer domain, the Kain–Fritsch cumulus scheme is used [15], whereas both nests explicitly resolve convection.

3. Uncertainty representation

Sources of uncertainty in atmospheric modeling can be divided into two categories: imperfect initial conditions and model formulation. The identification of optimal initial conditions links data assimilation and observability with nonlinear error growth. Manifold approaches exist at the major operational weather centres. In this project, the uncertainty in initial conditions of the limited area (at least for the first initialization) are downscaled from coarser-resolution ensemble systems: the Ensemble Prediction System (EPS) from ECMWF (51 members) and the Global Ensemble Forecasting System GEFS from NCEP (21 members). The ECMWF Ensemble Prediction System constructs initial perturbations by a combination of the singular value technique [17] and perturbations from an ensemble data assimilation system [3]. The NCEP Global Ensemble Forecasting system uses the technique of vector breeding [37]. Although such approaches are mature, choosing optimal initial conditions perturbations remains an ad hoc challenge, and yet ensembles with initial condition uncertainty alone appear to be underdispersive, which especially holds for mesoscale convection permitting ensembles in the short range (e.g. [29]).

Model error representation acts complementary to initial condition uncertainty, in which the major part is believed to originate from deficiencies in the representation of physical processes. Manifold approaches of model uncertainty representation exist, whereas multi-physics parameterizations (creating a different model framework by selecting different physical parameterizations) approaches are abandoned to preserve the advantageous of a single model framework. Thus, the focus is set on stochastic-dynamic and multi-parameter schemes. Stochastic-dynamic schemes represent uncertainty by spatial and temporal dynamic stochastic processes and result in indistinguishable ensemble members, whereas multi-parameter schemes modify certain model parameters which results in different climatologies. Since the representation of model error is far too complex to be

described by a single scheme [28], to achieve best results, multiple schemes are applied. Figure 3 exemplary illustrates the complementarity of the different schemes, which are discussed in the following.

3.1 Stochastically Perturbed Energy Backscatter Scheme

WRF offers the possibility to account for model deficiencies by applying the Stochastic Kinetic Energy Backscatter Scheme (SKEBS, [1], [31]). The key notion is that the physical process of energy upscale from unresolved to resolved scales is missing in conventional parameterization schemes. Explicitly, the upscale energy transfer from unbalanced motions associated with convection and gravity waves as well as losses due to numerical diffusion is mimiced via small-amplitude perturbations added to the rotational component of the horizontal wind and the potential temperature to stochastically correct the turbulent energy cascade. For example, the streamfunction tendency perturbations are expressed in 2-dimensional spectral space by

$$\Psi'(x,y,t) = \sum_{k=-K/2}^{K/2} \sum_{l=-L/2}^{L/2} \psi'_{k,l}(t) e^{2\pi i(kx/X + ly/Y)}, \tag{1}$$

where $\psi'_{k,l}(t)$ denote the spectral coefficients, k and l the wavenumber components and t the time. Temporal correlation is introduced by letting each spectral coefficient evolve by a first-order autoregressive process:

$$\psi'_{k,l}(t + \Delta t) = (1 - \alpha)\psi'_{k,l}(t) + g_{k,l}\sqrt{\alpha}\epsilon_{k,l}(t). \tag{2}$$

This equation describes a Markov process and $1 - \alpha$ is the autoregressive parameter, $g_{k,l}$ is a wavenumber-dependent noise amplitude and $\epsilon_{k,l}$ is a complex Gaussian white noise process with zero mean. The temporal and spatial correlation of the random forcing ensures error growth to larger scales with a prescribed spectral slope and domain averaged rate of backscattered energy. Instantaneous forcing patterns are shown in Figure 2(a) for the horizontal wind component and in Figure 2(b) for the potential temperature. To apply the random forcing pattern to the inner high-resolution domain, the decorrelation time has been reduced and no forcing is applied to the lowest two wavenumbers to amplify perturbations with smaller spatial correlation. The algorithm to compute the random forcing has been partly parallelized within this project period to make it applicable on the JUQUEEN's architecture (see Section 4).

3.2 Stochastically Perturbed Parameterization Tendency Scheme

The Stochastically Perturbed Parameterization Tendency Scheme (SPPT, [4],[21]) rests on the notion, that in particular with increasing model resolution, subgrid scale processes can no longer be represented by an equilibrium mean and have to be sampled. The deterministic outcomes of various parameterization schemes are therefore reformulated as a sample from a probability density function by adding multiplicative noise to the several parameterized physics tendencies:

$$X_p = (1 + r_x)X_c, \tag{3}$$

where X_c and X_p are the unperturbed and perturbed tendencies, respectively, and r_x is a random number. This formulation ensures flow dependent perturbation amplitudes: regions with small tendencies are less effected and vice versa. To associate smaller spatial

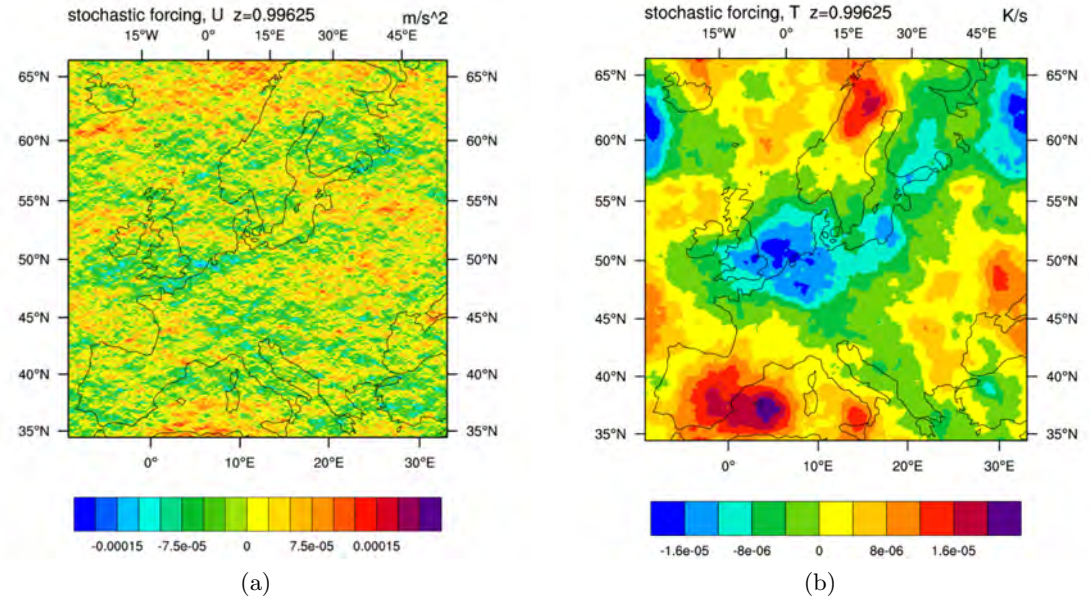

(a) (b)

Figure 2: Examples of instantaneous forcing pattern for (a) the horizontal wind component and (b) the potential temperature of the SKEBS scheme over the Mother domain. Perturbations are interpolated on the nests.

perturbation patterns with short timescales and larger spatial perturbation patterns with long timescales, a spectral forcing pattern is used with the same design as used by the SKEBS scheme (Equation 2), but with different properties. To apply the SPPT random forcing pattern to the inner high-resolution domain, the decorrelation time has been reduced to 2 hours and the gridpoint standard deviation of random perturbations to 0.25.

3.3 Perturbed Surface Parameter Scheme

Within this project, a Perturbed Surface Parameter Scheme has been implemented to represent the uncertainty of albedo and roughness length. Although the perturbations are confined to the surface, the model solution is expected to be sensitive to changes in these parameters since they indirectly effect phenomena such as surface winds, lower atmosphere temperature and cloud height. Moreover, they may effect the free troposphere, most prominent by Ekman pumping. [34] have shown, that atmospheric models may react sensitive to surface parameters even in the short range, though the magnitude of error growth is expected to be small in comparison to perturbations on the synoptic scale. Nevertheless, in this work the effort has been undertaken, since we seek for the most comprehensive variety of model perturbations. In addition, appropriate studies with high resolution are rare. The perturbation strategy follows the idea of [34] and [11] and was modified to fit the current WRF model version.

As the Noah model is used for land-surface parameterization, annual minimum and maximum of albedo and roughness length from the U.S. Geological Survey (USGS) 30-second global dataset are used for 27 land-use classifications. If applicable, values are linearly interpolated between the minimum and maximum values depending on the vegetation fraction. Although in principle there is full flexibility in perturbing such values, we have chosen to use temporally and spatially fixed perturbations. By doing so, we account

for different model climatologies to create distinguishable ensemble members for the particle filter step. The ranges of uncertainty for both parameters are quite distinct, values are taken from [20], [33], [27], [30] and [23]. We added additional uncertainty to the values to account for the gridded land-use representation. For each minimum, a unique probability density function has been designed, which has it's peak at the USGS's original value. For this purpose, a log-normal distribution has been chosen, due to its ability to take a variety of shapes, and two additional parameters have been added for more flexibility:

$$f(x;\mu,\sigma,t,r) = \begin{cases} \frac{1}{(x-t)\sigma\sqrt{2\pi}}\exp\left[\frac{-(\ln(x-t)-\mu)^2}{2\sigma^2}\right] & r = false \\ \frac{1}{(2\exp(\mu)-x-t)\sigma\sqrt{2\pi}}\exp\left[\frac{-(\ln(2\exp(\mu)-x-t)-\mu)^2}{2\sigma^2}\right] & r = true \end{cases} \quad (4)$$

where x is the random variable (albedo or roughness length) and μ and σ are the location and scale parameter on a logarithmic scale, respectively, with mean $e^{(\mu+\sigma^2)/2}$ and variance $(e^{\sigma^2}/2-1)e^2\mu+\sigma^2$. t is a translation variable and r a logical variable for reversing the distribution. The uncertainty in the seasonal variability is provided by multiplying the difference between the original minimum and maximum USGS values with a sample from a Gaussian with zero mean. Figure 10 shows examples of PDFs, Table 4 and 5 list all 184 variables of the extended log-normal distribution.
Despite of the high resolution, results in general indicate a relatively small sensitivity to surface parameters, confined to the boundary layer in the short range. Nevertheless, several case studies suggest there are distinct exceptions:

- Pronounced baroclinicity (frontogenesis, cyclogenesis etc.) tends to carry ensemble spread into the free troposphere
- Convective events in the warm season lift and enhance ensemble spread into the free troposphere

Since the ultra large ensemble is applied to events of low predictability like listed above, the inclusion of surface parameters is still valuable. Exemplary, Figure 3(c) shows the ensemble spread of solely surface parameter perturbations of a 256 member ensemble after 42h forecast time. Ensemble spread is not necessarily confined over regions of larger uncertainty in surface parameters, but may be strongly triggered due to instabilities of the underlying flow or simply transported and enhanced in the direction of the flow.

3.4 Multi-Parameter Parameterization Scheme

The Multi-Parameter Parameterization Scheme modifies parameter within the physical parameterizations to a certain degree of uncertainty. Again, we have chosen to use temporally and spatially fixed perturbations to account for different model climatologies to create distinguishable ensemble members for the particle filter step. The parameterizations applied in this study are given in Section 2, uncertain parameters have been identified by the corresponding literature and sampled within a realistic range of uncertainty (using Equation 4). Perturbed values include among others

- scattering of short wave radiation
- intercept parameter in the exponential rain drop size distribution
- scaling of the boundary layer entrainment rate
- maximum turbulent kinetic energy in the sub-cloud layer

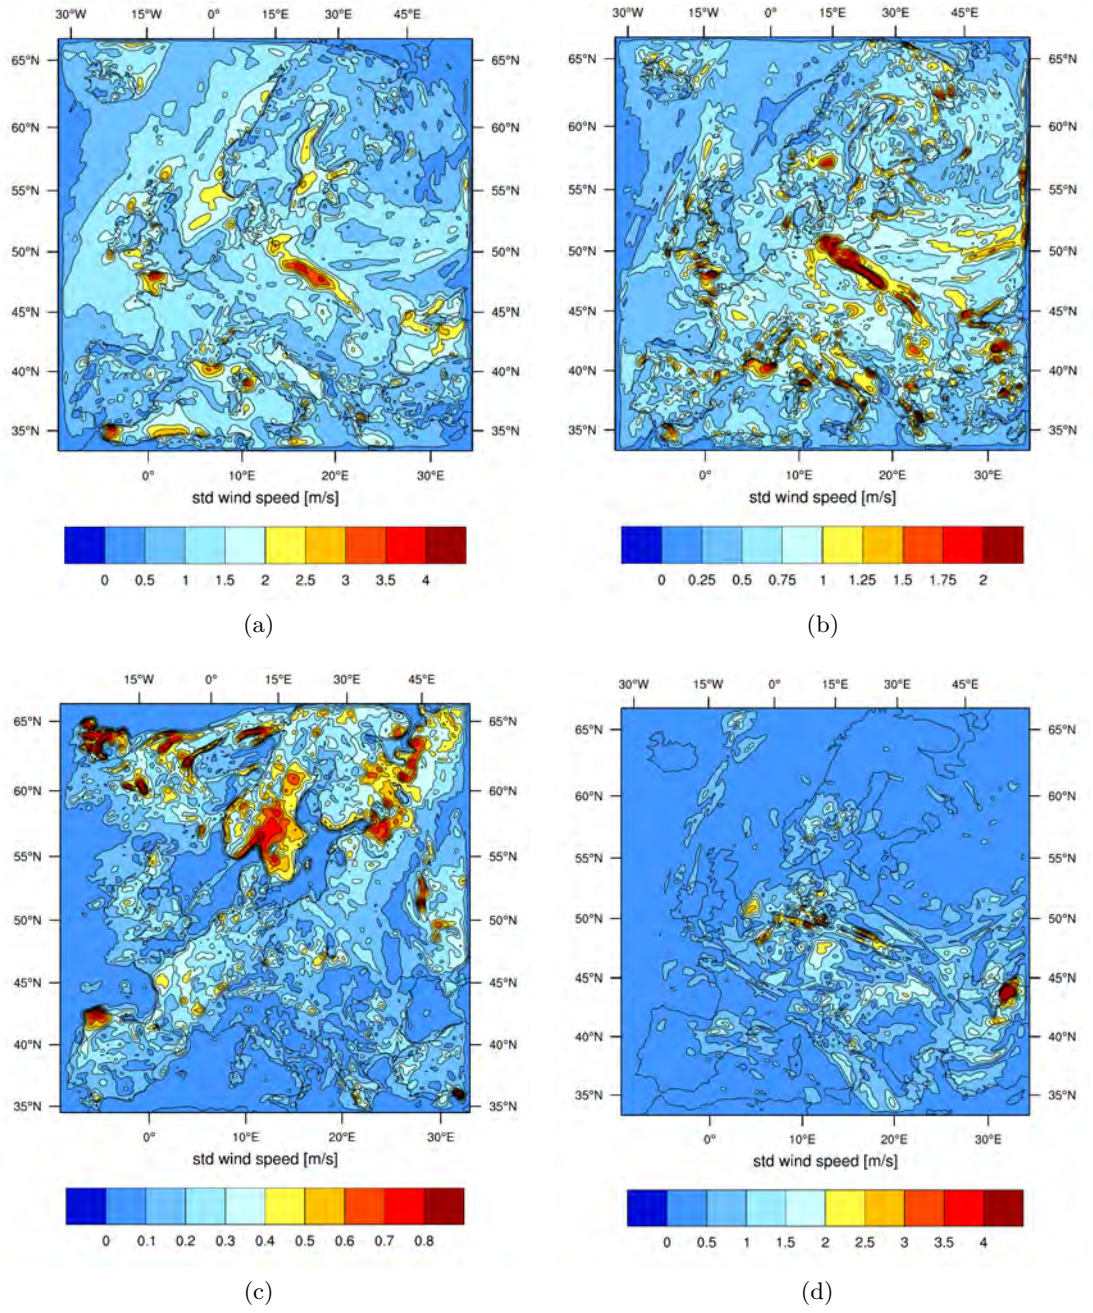

Figure 3: Validation of the ensemble spread of wind speed at 100 meter hub height for a 256 member ensemble after 42h model integration at 24-07-2014 18UTC for (a) SKEBS scheme, (b) SPPT scheme, (c) Perturbed Surface Parameter Scheme and (d) Multi-Parameter Parameterization Scheme. Note the different scale.

- downdraft mass flux of convective motion

Results indicate strong non-linear behavior between model diversity and perturbation technique when multiple parameters are perturbed at the same time. This finding suggests that parameter estimation is likely to be very difficult and a single-parameter perturbation strategy for the future application within the particle filter environment may be more suitable.

4. Computational performance analysis and improvements

1) WP1 interactions

Interactions with WP1 related to monitoring and optimization of the code included so far:

- Compute performance analysis with the *Scalasca* tool [9]
- I/O performance analysis with the *Darshan* tool [5]
- Adaptation of WRF's metadata output strategy to satisfy the GPFS File System (Initialization problems removed)

The WRF code was built on the Blue Gene/Q architecture of JUQUEEN with pure MPI support as well as a hybrid OpenMP/MPI configuration. The hybrid configuration may occasionally perform superior to pure MPI, but only slightly and depending on the grid extensions. Therefore, a pure MPI parallelization was chosen to demonstrate the code performance and the hybrid OpenMP/MPI is not discussed in this report. A number of 32 ranks per node has been put on the IBM PowerPC® A2 processors. Further SMT is not limited by memory (typically 200 MB per core), but does not lead to an overall improvement in computational time due to increased I/O expenses. As an I/O strategy, Parallel NetCDF is used.

The following performance metrics are computed for a single ensemble member with a model integration time of 3 hours on the outer domain (see Figure 1), but is also transferable to the complete ultra large ensemble thanks to the ensemble parallelization procedure (see Section 5). Total CPU time is computation dominated with 40% MPI processes. Point-to-point MPI calls prevail over collective calls due to the local spatial discretization scheme. The average message size is 16 kB. Load imbalance is acceptable and makes up to 5 % of the total computation time. I/O costs are about 10 % of the total CPU time with mostly collective output. An average of 130 MB/sec of I/O throughput can be estimated. Figure 5 shows the parallel scaling behavior for a single ensemble member with a model integration time of 48 hours on the outer domain (see Figure 1), the corresponding CPU times are given in Table 2. One can simply extrapolate the corresponding scaling behavior of the ensemble controlling system according to Figure 6.

2) General performance improvements

The compilation was built with the IBM® XL compilers and gnu compilers, the latter were rejected due to poor performance. The finding of the most suitable compiler settings for code optimization is a crucial task and requires detailed sensitivity tests and considerable effort has been invested at the beginning of this project. An optimization level of *-O3* is in principle applicable, yet higher order transformations of loops has to be

disabled due to an impact on the result's accuracy. The *-qsimd=noauto* option disables the conversion of loop array operations into vector instructions and has a positive impact. The function level of WRF is dominated by slow math intrinsic functions and linking with the IBM MASS library is not effective since the library contains different names of entry points. By adding the option *-qstrict=nolibrary*, the XL compiler does not change the names of floating point routines and faster alternatives can be used. The suboption *-qnohot=noarraypad:level=2:novector:fastmath* enables the fast scalar versions of math functions instead of the default. Reciprocal and square root functions need special considerations on Blue Gene/Q. Scalar and SIMD estimate instructions (e.g. Newton's approximation method) with high throughput exist, and their utilization has to be forced by adding *-qdebug=recipf:forcesqrt*. The final compiler optimization instructions read:

-O3 -qnohot=noarraypad:level=2:novector:fastmath -qstrict=nolibrary
-qdebug=recipf:forcesqrt -qsimd=noauto -qarch=qp -qtune=qp.

Table 2 shows a comparison of computational performance of the WRF code with default and current compiler optimization flags with varying number of cores. The reduction in CPU time is almost 100 %.

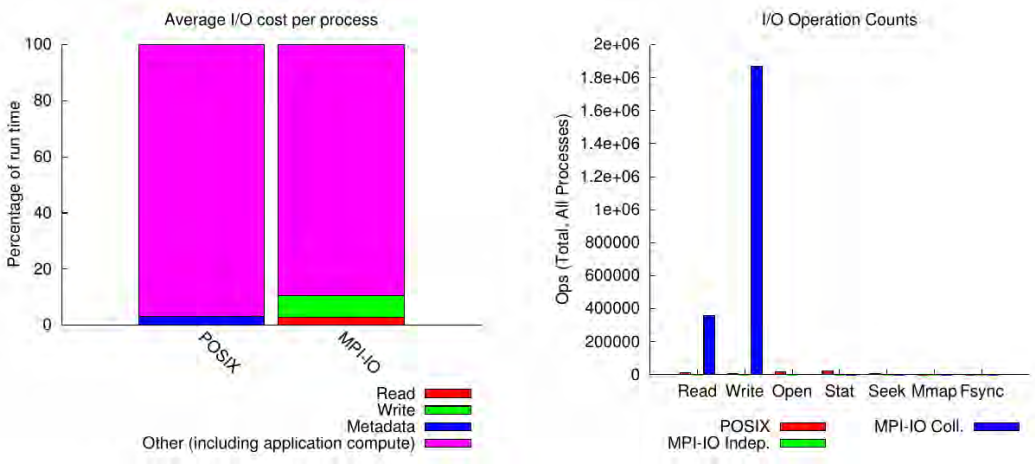

Figure 4: I/O characterization of WRF on the Blue Gene/Q architecture of JUQUEEN with Parallel NetCDF using 512 cores.

Table 2: Computational time in seconds for a 6h simulation on the mother domain (see Figure 1) with default and current compiler optimization for varying number of cores. 32 ranks per node have been used.

# cores	Default optimization level	Current optimization level
512	119.2 s	66.2 s
256	192.0 s	107.7 s
128	321.8 s	205.0 s
64	593.0 s	362.9 s
32	1125.22 s	686.4 s

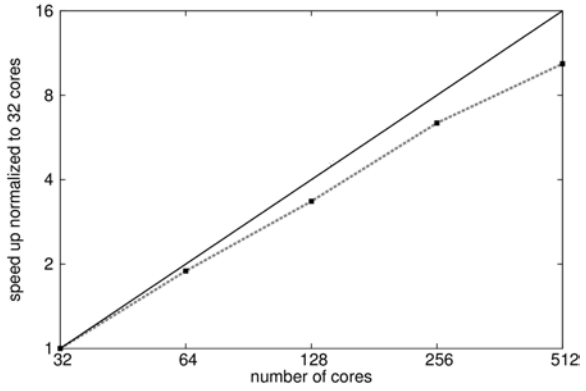

Figure 5: Parallel scaling behavior of WRF on the Blue Gene/Q architecture of JUQUEEN. CPU times are given in Table 2.

Table 3: Comparison of CPU times for model integration with the SKEBS scheme for the original and computational improved STOCH code for varying number of cores.

#cores	Original STOCH code	Computational improved STOCH code (global MPI)	Computational improved STOCH code (point-to-point MPI)
512	924.9 s	96.1 s	86.5 s
256	973.3 s	158.64 s	145.3 s
128	1093.6 s	272.66 s	250.8 s
64	1301.5 s	505.1 s	470.0 s
32	1780.4 s	968.4 s	920.7 s

3) Improvement of stochastic pattern computation

The performance analysis discussed above no longer holds if either SKEBS or SPPT is used for model perturbation. The computation time increases up to a factor of ~ 10 (see Table 3) compared to an unperturbed run. The largest portion of the additional computing time is due to the calculation of the complex Gaussian white noise process $\epsilon_{k,l}$ of the stochastic pattern (Equation (2), STOCH code hereafter, which is used by both schemes). Within the officially released WRF code, the calculation of $\epsilon_{k,l}$ uses global indexes to provide necessary symmetries and anti-symmetries of the random forcing to ensure real-valued Fourier back transformation. Hence, random numbers are calculated over the whole domain by each single processor. This leads to an unacceptably high CPU time on JUQUEEN's architecture with its IBM PowerPC®A2 (1.6 GHz) processors.

Therefore, the computation of $\epsilon_{k,l}(t)$ has been parallelized, letting each processor computing its own random numbers. Initialization of the seed for random number generation is done by each processor's rank in the *MPI_COMM_WORLD* communicator. The implementation of point-to-point MPI communication is non trivial since the rotation centre of symmetry is never in the domain centre. A comparison of CPU times for the model integration with the SKEBS scheme for the original and computational improved STOCH code (global and point-to-point communication) is shown in Table 3. The parallelization of the code makes SKEBS and SPPT finally applicable on JUQUEEN with an acceptable additional CPU time. However, Table 3 shows that even with the improved STOCH scheme, a high ratio of gridpoints per processor is still to be avoided when using SKEBS or SPPT. Point-to-point communication performs only slightly superior to global communication due to a high number of single MPI requests. Results with the SPPT scheme look

analogous, since the same stochastic pattern is used. Within the improved scheme, every processor uses a different seed and simulation results are not comparable anymore to the original STOCH code. However, in the case of an infinite large ensemble size, one would expect the statistical properties to be identical. This constrain is assumed to be fulfilled by comparing the spatial and temporal evolution of a 256 member ensemble's standard deviation of several test runs under different atmospheric conditions.

5. Parallel ultra large ensemble controlling system

The application of a Sequential Importance Resampling Filter (SIR-PF) and its smoothing variant (SIR-PS) as a fully non-linear ensemble data assimilation technique is the main innovation of Task 2.1. From an HPC perspective, particle filtering suffers the pitfall that dynamical updates with inter-member communication between the ensemble members during the run is required. This is explained in the following.

As other data assimilation strategies (e.g. ENKF, 4D-Var), the Sequential Importance Resampling Filter is based on Bayes' Theorem

$$p(x|y) = \frac{p(y|x)p(x)}{\int p(y|\hat{x})p(\hat{x})d\hat{x}} \ , \tag{5}$$

combining the information of the model state x and the observations y optimally. The a prior PDF is represented by the sum of Dirac delta functions, each centered around the ensemble member state x_i

$$p(x) = \frac{1}{N}\sum_{i=1}^{N}\delta(x - x_i) \ , \qquad \delta(x - x_i) = \begin{cases} 1 \text{ for } x = x_i, \\ 0 \text{ otherwise,} \end{cases} \tag{6}$$

where N is the ensemble size.
Substituting (6) into Bayes' Theorem yields

$$p(x|y) = \sum_{i=1}^{N} w_i \delta(x - x_i) \ , \tag{7}$$

with weights

$$w_i = \frac{p(y|x_i)}{\sum_{j=1}^{N} p(y|x_j)} \ . \tag{8}$$

Depending on the weight w_i, an ensemble member gets resampled ([16],[7]), i.e. ensemble members with low weights are rejected, whereas ensemble members with higher weights are kept in the system and multiple copies are generated until the total number N is restored, giving each ensemble member an equal weight of $1/N$ again (Sequential Importance Resampling). If each ensemble member runs autonomous, as it is done traditionally (e.g. script based ensemble versions running in a cloud), unacceptably high computational times would arise during a resampling step due to (1) I/O costs of writing restart files, (2) model shut-down, (3) resampling by copy procedures of restart files and (4) model initialization.

Thus, WRF has been successfully modified towards a stand-alone ensemble version, realizing inter-member communication within a single application and controlling thousand of ensemble members in a computationally efficient way.

For this purpose, two competitive approaches have been followed. The first approach is based on the intrinsic MPI-2 routine *MPI_COMM_SPAWN*, which allows to execute an arbitrary number of MPI programs in parallel and offers the possibility of communication

via an intercommunicator. Although the simplicity of this approach is quite appealing, due to poor performance this approach has been rejected.

Instead, the more advantageous approach of integrating a second stage of MPI parallelism within the WRF code was adopted. Immediately after MPI initialization, the global MPI communicator *MPI_COMM_WORLD* is split into n new communicators, with n the number of ensemble members. Each of the new communicators is associated with a certain ensemble member and the number of processors is equally distributed among them. The obsolete *MPI_COMM_WORLD* communicator is consequently replaced by the new ensemble communicator throughout the WRF code. An additional communicator is created which groups all processors with identical ranks within the new ensemble communicators. Thereby, every ensemble member for itself runs on multiple processors as well as parallel to the remaining ensemble members. This novel parallelization strategy of the WRF code is illustrated in Figure 6. Due to the second parallelization stage, initial values and boundary values as well as a full model state can be send from one member to the other, realizing multiple particle filter resampling steps within one application.

There do not exist any restrictions on model execution with the novel ensemble setup. The approach is among others compatible with 1-way and 2-way nesting or different time-stepping. I/O strategies are not conflicting since every member has its own input and output files. Complimentary, the entire WRF Preprocessing System (WPS) has been modified towards a stand-alone ensemble version with the same approach as the main solver. Figure 7 proofs the computational performance of the approach. Without the expanse of communication between the members, the loss in computation time with increasing number of ensemble members is of the order of 1 - 3% and therefore neglectable. The performance of the ensemble setup has been tested on JUQUEEN with up to 262'144 cores and 4096 members.

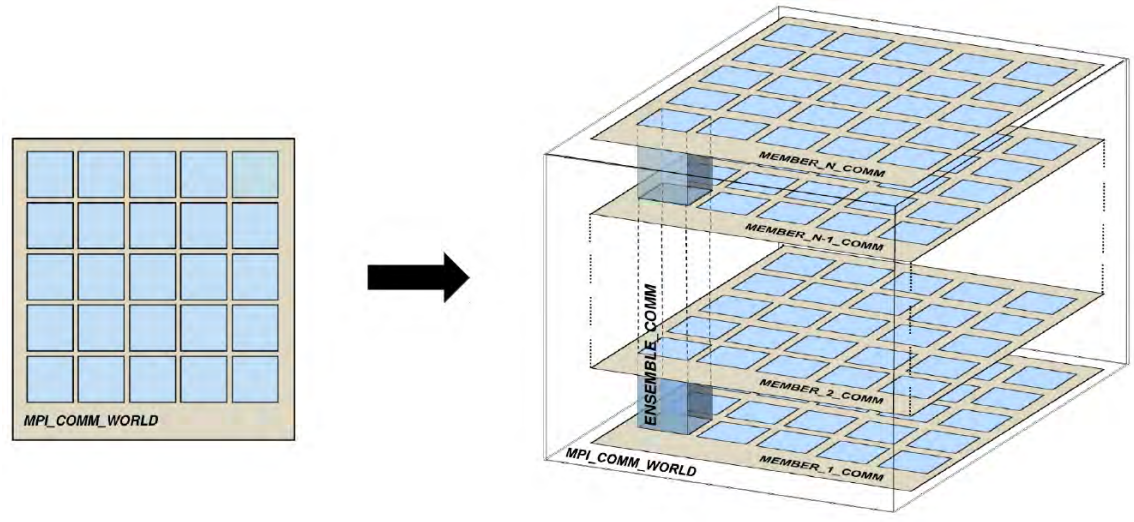

Figure 6: Visualization of the MPI approach to modify the WRF code to a stand-alone ensemble controlling system (as part of ESIAS). The modification has its signification in computational efficiency of particle filter experiments on a HPC architecture.

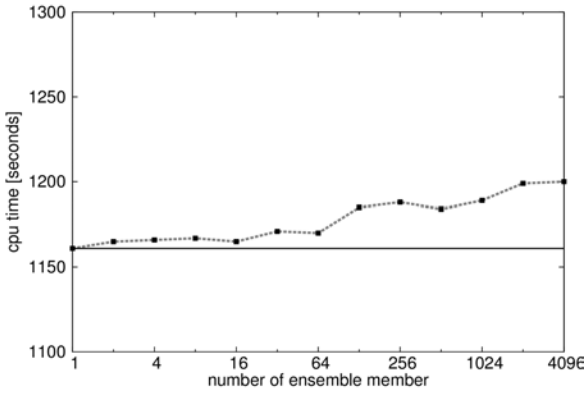

Figure 7: Computation time of the stand-alone ensemble controlling system with increasing member number for a 24h simulation on the largest domain (dashed line). The number of processors per ensemble member is kept fixed (64 processors each) with increasing ensemble size. Deviations from a perfectly parallelized ensemble (solid line) are neglectable and due to maximum bandwidth during I/O processes.

6. A case study

Within the frame of a case study, the principle realization of the ultra large ensemble is shown in the following. A detailed evaluation of the ultra large ensemble's impact on the energy forecast is scheduled to deliverable 2.4. A case study for the 09-08-2014 is chosen since it has been identified by Fraunhofer IWES as one of the major error events of 2014. At 09-08-2014 12:00 UTC, the day-ahead forecast of the Transmission System Operators (TSOs) underestimated the wind power production by 7.8 GW [25]. Since the power forecast is produced by combining different meteorological forecast models, it is assumed that the majority failed to properly predict the weather situation. The event is characterized by the low predictability of pronounced cyclogenesis and frontogenesis, with the challenging aspect of forecasting the correct location and shape of a surface low in Numerical Weather Prediction.

Between 08-08-2014 and 09-08-2014, the weather situation in Germany was dominated by strong cyclogenesis: An upper air flow over the North Atlantic developed its low pressure system close to Iceland, which originated in a secondary cut-off low in the progressing development phase. The cut-off low system reached a core pressure minimum of below 985 hPa at 09-08-2014 12:00 UTC located in the North Sea between Scotland and South Norway. It's associated frontal system passed north and central Germany with wind speeds up to 90 km/h in coastal areas.

A 512 member ensemble has been initialized at 08-08-2014 00:00 UTC with GEFS initial and boundary conditions via a cold start, yet no local data assimilation with the Sequential Importance Resampling Filter has been performed. Model error is represented by the schemes described in Section 3. Figure 8(a) shows the isolines of geopotential at the 850 hPa pressure level, indicating severe model diversity over the centre of the secondary low pressure system by distinct differences in the isoline's shape over the North Sea. This result is even more apparent given the corresponding ensemble spread of the wind speed at 100 m hub height shown in Figure 8(b). A pronounced maximum is located over the area of the core pressure east of Scotland. The maximum does not primarily arise due to an overestimation of the core pressure but rather indicating different extends and position, resulting in different pressure gradients and too low wind speeds on average. A secondary local maximum is formed over the Baltic Sea which extends south east of Germany, indi-

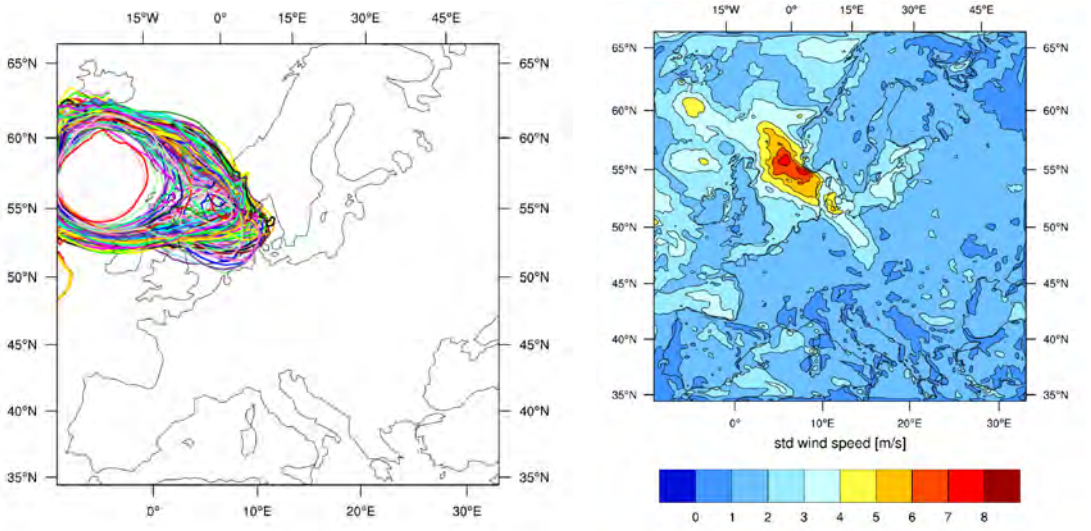

(a) (b)

Figure 8: Diversity of 512 ensmble members over the Mother domain at 09/08/2014 12UTC, generated by the ESIAS system. All ensemble members have been initialized at 08-08-2014 00:00 UTC. (a) Isolines of geopotential at 850 hPa pressure level (134 gpdam). (b) Ensemble spread of wind spead interpolated to 100m hub height.

cating the frontal system's movement. The stationarity of the low pressure system south to Iceland is represented by less ensemble spread, indicating higher predictability.

Figure 9 shows the time series of wind speed quantiles at several measurement towers at around 100 m hub height. The ensemble median indicates a clear underestimation of wind speed in off-shore regions (Figure 9(a), 9(b), 9(c)) and does not capture well enough strong ramping events in coastal areas (Figure 9(d)). Wind speeds are better represented further inland (Figure 9(e) and 9(f)), though with less impact on the wind power production due to less installed capacities. Nevertheless, observations are captured within the quantiles of the ultra large ensemble, except for a few peak values. However, there is no single member who adequately represents the spatial and temporal evolution of the wind speed. Hence we draw the conclusion, that the ultra large ensemble is yet too much confined to the initial and boundary conditions of the global ensemble members, pointing out the need for nonlinear data assimilation with the potential of encapsulating low probability events.

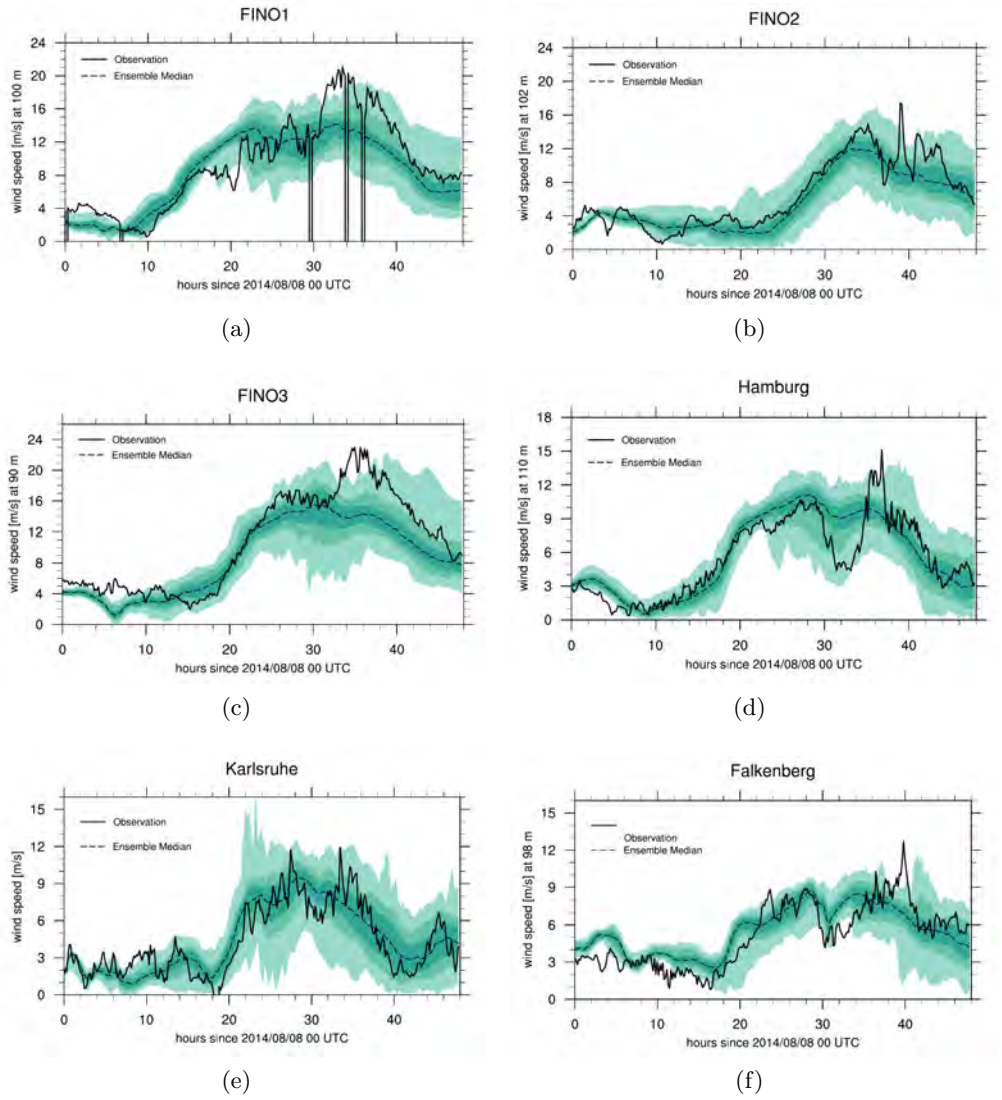

Figure 9: Time series of the wind speed quantiles at different measurement towers located over Germany from 08/08/2014 00UTC to 09/08/2014. FINO1 and FINO3 are placed in the North Sea, FINO2 in the Baltic Sea. Unphysical zero values in case of FINO1 observations indicate data gaps.

A. Examples of the extended log-normal distribution

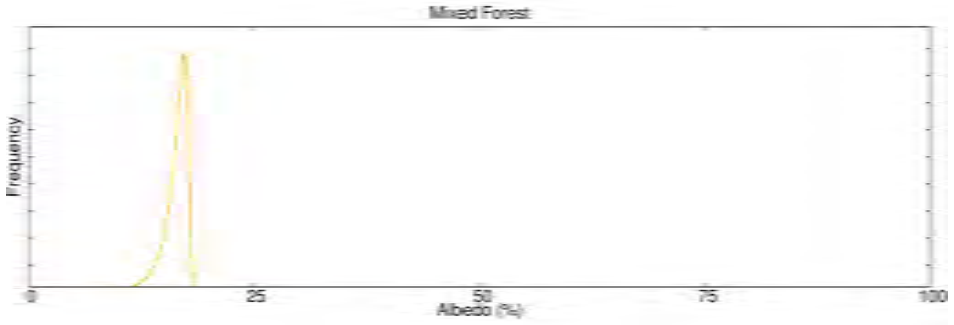

(a)

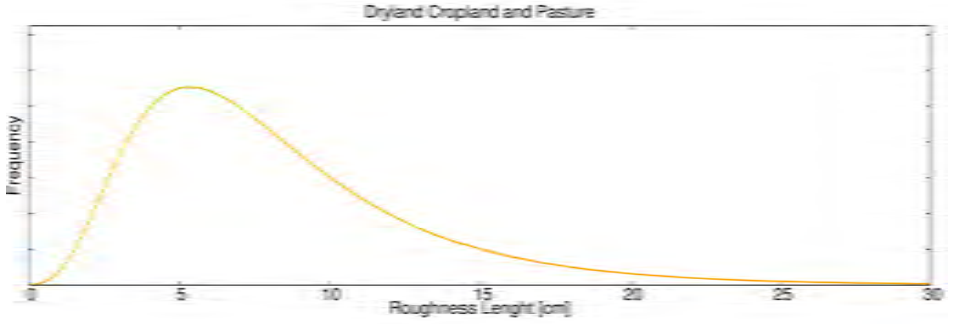

(b)

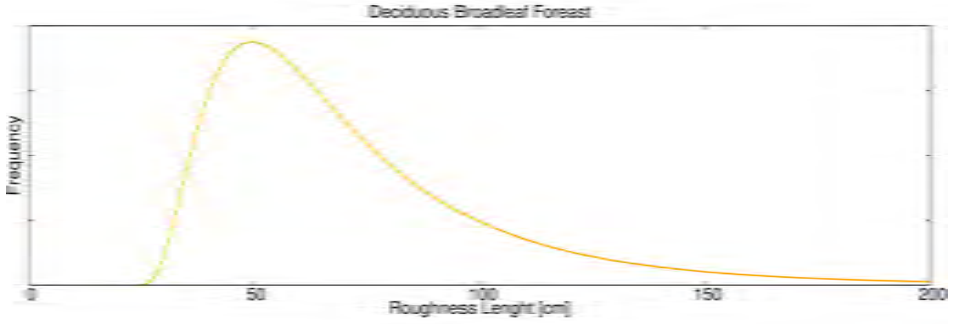

(c)

Figure 10: Examples of the extended log-normal distribution (Equation (4)) used for the Surface Parameter Perturbation Scheme. The associated parameter of the PDF are given in Table 4 and 5.

B. Extended log-normal variables for the Surface Parameter Perturbation Scheme

Table 4: Extended log-normal variables for the uncertainty representation of albedo (annual minimum). Values for "Water Bodies", "Snow and Ice" and "Lava" are not perturbed. μ denotes the location parameter, σ the scale parameter, t the translation parameter and r the reversing parameter.

Land Use Category	σ	μ	t	r
Urban and Built-Up Land	0.13	2.7	0.4	*false*
Dryland Cropland and Pasture	0.35	1.5	13.0	*false*
Irrigated Cropland and Pasture	0.3	1.0	17.0	*false*
Mixed Dryland \Irrigated Cropland and Pasture	0.35	1.5	15.0	*false*
Cropland \Grassland Mosaic	0.42	1.25	15.0	*false*
Cropland \Woodland Mosaic	0.3	1.8	9.0	*false*
Grassland	0.2	2.0	12.0	*false*
Shrubland	0.3	0.8	22.0	*false*
Mixed Shrubland \Grassland	0.7	0.6	17.0	*false*
Savanna	0.7	0.6	17.0	*true*
Deciduous Broadleaf Forest	0.7	0.2	13.5	*true*
Deciduous Needleleaf Forest	0.1	2.0	6.7	*true*
Evergreen Broadleaf Forest	0.1	2.0	4.8	*true*
Evergreen Needleleaf Forest	0.4	1.0	9.0	*false*
Mixed Forest	0.7	0.5	15.0	*true*
Herbaceous Wetland	0.1	2.0	6.5	*true*
Wooded Wetland	0.1	2.0	6.5	*true*
Barren or Sparsely Vegetated	0.6	0.7	36.0	*true*
Herbaceous Tundra	0.7	0.7	13.0	*false*
Wooded Tundra	0.7	0.7	13.0	*false*
Mixed Tundra	0.7	0.7	13.0	*false*
Bare Ground Tundra	0.2	1.2	21.5	*false*
Playa	0.2	2.7	15.0	*true*
White Sand	0.45	2.0	51.0	*true*

Table 5: Extended log-normal variables for the uncertainty representation of roughness length (annual minimum). Values for "Water Bodies", "Snow and Ice", "Lava", Playa" and "White Sand" are not perturbed. μ denotes the location parameter, σ the scale parameter, t the translation parameter and r the reversing parameter.

Land Use	σ	μ	t	r
Urban and Built-Up Land	0.5	3.5	24.0	*false*
Dryland Cropland and Pasture	0.085	2.9	-11.3	*false*
Irrigated Cropland and Pasture	0.7	0.5	1.0	*false*
Mixed Dryland \Irrigated Cropland and Pasture	0.085	2.8	-11.3	*false*
Cropland \Grassland Mosaic	0.085	2.8	-11.3	*false*
Cropland \Woodland Mosaic	0.11	3.1	-2.0	*false*
Grassland	0.095	3.1	-12.1	*true*
Shrubland	0.6	0.5	0.1	*false*
Mixed Shrubland \Grassland	0.18	2.2	12.5	*false*
Savanna	0.06	2.6	1.6	*false*
Deciduous Broadleaf Forest	0.65	3.7	23.0	*false*
Deciduous Needleleaf Forest	0.65	3.7	23.0	*false*
Evergreen Broadleaf Forest	0.65	3.7	23.0	*false*
Evergreen Needleleaf Forest	0.65	3.7	23.0	*false*
Mixed Forest	0.4	2.7	7.0	*false*
Herbaceous Wetland	0.06	4.2	-46.0	*false*
Wooded Wetland	0.4	3.7	6.0	*false*
Barren or Sparsely Vegetated	0.5	0.1	0.15	*false*
Herbaceous Tundra	0.04	4.2	-56.5	*false*
Wooded Tundra	0.3	3.7	-6.5	*false*
Mixed Tundra	0.25	3.0	-3.8	*false*
Bare Ground Tundra	0.1	2.6	-8.3	*false*

References

[1] J. S. Berner, S.-Y. Ha, J. P. Hacker, A. Fournier, and C. M. Snyder. Model uncertainty in a mesoscale ensemble prediction system: Stochastic versus Multiphysics Representations. . *Monthly Weather Review*, 139:1972–1994, 2010.

[2] P. Browne and P. Van Leeuwen. Twin experiments with the equivalent weights particle filter. *Quarterly Journal of the Royal Meteorological Society*, 141:3399–3414, 2015.

[3] R. Buizza, M. Leutbacher, and L. Isaksen. Potential use of an ensemble of analyses in the ECMWF Ensemble Prediction System. *Quarterly Journal of the Royal Meteorology Society*, 134:2051–2066, 2008.

[4] R. Buizza, M. Miller, and T.N. Palmer. Stochastic representation of model uncertainties in the ECMWF Ensemble Prediction System. *Quarterly Journal of the Royal Meteorology Society*, 125:2887–2908, 1999.

[5] P. Carns, K. Harms, W. Allcock, C. Bacon, S. Land, R. Latham, and R. Ross. Understanding and improving computational science storage access through continuous characterization. *ACM Transactions on Storage*, pages 1–8, 2011.

[6] G. Cavallaro, M. Riedel, and J.A. Benediktsson. On Understanding Big Data Impacts in Remotely Sensed Image Classification Using Support Vector Machine Methods, journal = IEEE Journal of Selected Topics in Applied Earth Observation and Remote Sensing. pages 4634–4646, 2015.

[7] R. Douc, O. Cappé, and E. Moulines. Image and Signal Processing and Analysis, 2005. ISPA 2005. Proceedings of the 4th International Symposium on Image and Signal Processung and Analysis. pages 64–69, 2005.

[8] H. Elbern, A. Strunk, H. Schmidt, and O. Talagrand. Emission rate and chemical state estimation by 4-dimensional variational inversion. *Atmospheric Chemestry and Physics*, 7:1–59, 2007.

[9] M. Geimer, F. Wolf, B. J. N. Wylie, E. Ábrahám, D. Becker, and B. Mohr. The Scalasca performance toolset architecture. *Concurrency and Computation: Practice and Experience*, 22:702–719, 2010.

[10] G. Giebel, L. Landberg, J. Badger, K. Sattler, H. Feddersen, T. S. Nielsen, H.Aa. Nielsen, and H. Madsen. Using ensemble Forecasting for Wind Power. *Paper presented on the European Wind Energy Conference and Exhibition, Madrid, Spain*, 2003.

[11] J. P. Hacker, S.-Y. Ha, C. Snyder, J. Berner, F. A. Eckel, E. Kuchera, M. Pocernich, S. Rugg, J. Schramm, and X.Wang. The U.S Air Force Weather Agency's mesoscale ensemble: scietific description and performance results. *Tellus*, 63:625–641, 2011.

[12] S. Y. Hong and J.O Lim. The WRF single-moment 6-class microphysics scheme (WSM6). *Monthly Weather Review*, 42:129–151, 2006.

[13] S. Y. Hong, Y. Noh, and J. Dudhia. A new vertical diffusion package with an explicit treatment of entrainment processes. *Monthly Weather Review*, 134:2318–2341, 2006.

[14] M.J. Iacono, J. S. Delamere, E. J. Mlawer, M. W. Shephard, S. A. Clough, and W. D. Collins. Radiative forcing by long-lived greenhouse gases: Calculatons with the AER radiative transfer models. *Journal of Geophysical Research*, 113:D13103, 2008.

[15] J. S. Kain. The Kain-Fritsch convective parameterization: An update. *Journal of Applied Meteorology*, 43:170–181, 2004.

[16] N. Metropolis and S. Ulam. The Monte Carlo method. *Jounral of the American Statistical Association*, 44:335–341, 1944.

[17] F. Molteni, R. Buizza, T.N. Palmer, and T. Petroliagis. The ECMWF Ensemble Prediction System: Methodology and validation. *Quarterly Journal of the Royal Meteorological Society*, 122:73–119, 1996.

[18] F. Molteni and T.N. Palmer. Predictability and finite-time instability of the Northern winter circulation. *Quarterly Journal of the Royal Meteorological Society*, 119:269–298, 1993.

[19] H.Aa. Nielsen, H. Madsen, G. Giebel, J. Badger, L. Landberg, K. Sattler, L. Voulund, and J. Tofting. From wind ensembles to probabilistic information about future wind power production – results from an actual application. *9th International Conference on Probabilistic Methods Applied to Power Systems KTH, Stockholm, Sweden, Invited Paper*, 2006.

[20] T. R. Oke. Boundary Layer Climates, 2nd Edition. *Routledge*, 1987.

[21] T. N. Palmer, R. Buizza, F. Doblas-Reyes, T. Jung, M. Leutbacher, G. Shutts, M. Steinheimer, and A. Weisheimer. Stochastic Parameterization and Model Uncertainty. ECMWF Technical Memordandum. *available at http://www.ecmwf.in/publications/*, 598, 2009.

[22] R. Perez, E. Lorenz, S. Pelland, M. Beauharnois, G. Van Knowe, K. Hemker Jr., D. Heinemann, J. Remund, S. C. Müeller, W. Traumüller, G. Steinmauer, D. Pozo, J. A. Ruiz-Arias, V. Lara-Fanego, L. Ramirez-Santigosa, M. Gaston-Romero, and L. M. Pomares. Comparison of numerical weather prediction solar irradiance forecasts in the US, Canada and Europe. *Solar Energy*, 94:305–326, 2013.

[23] R. A. Pielke. Mesoscale Meteorological Modeling. *Academic Press*, 2013.

[24] P. Pinson, H.Aa. Nielsen, H. Madsen, and G. Kariniotakis. Skill forecasting from ensemble prediction of wind power. *Applied Energy*, 86:1326–1334, 2009.

[25] EEX Transparency Platform. *http://www.eex-transparency.com*.

[26] P. Poterjoy and J.L. Anderson. Efficient assimilation of simulated observations in a high-dimensional geophysical system using a localized particle filter. *Monthly Weather Review*, 144:2007–2020, 2016.

[27] C. H. B. Pristly. Turbulent Transfer in the Lower Atmosphere. *The University of Chicago Press*, 1959.

[28] G. Romine, C.S. Schwartz, J. Berner, R. Fosseell, C.M. Snyder, J.L. Anderson, and M.L. Weisman. Representing forecast error in a convection-permnitting ensemble system. *Monthly Weather Review*, 134:2051–2066, 2008.

[29] C.S. Schwartz, G. Romine, K. Smith, and M. Weisman. Characterizing and optimizing precipitation forecasts from a convection-permitting ensemble initialized by a mesoscale ensemble Kalman Filter. *Weather and Forecasting*, 29:1295–1318, 2014.

[30] W. D. Sellers. Physical Climatology. *The University of Chicago Press*, 1965.

[31] G. J. Shutts. A kinetic energy backscatter algorithm for use in ensemble prediction systems. *Quarterly Journal of the Royal Meteorology Society*, 131:3079–3102, 2005.

[32] William C. Skamarock, Joseph B. Klemp, Jimy Dudhia, David O. Gill, Dale M. Barker, Wei Wang, and Jordan G. Powers. A Description of the Advanced Research WRF Version 3. *AVAILABLE FROM NCAR; P.O. BOX 3000; BOULDER, CO*, 2008.

[33] J. W. Snow. Wind Power Assessment along the Atlantic and Gulf Coasts of the U.S. *Ph.D. dissertation, Dept. of Environmental Science, University of Virginia, Charlottesville*, 1981.

[34] D. J. Stensrud, H.E Brooks, J. Du, M.S Tracton, and E. Rogers. Using ensembles for short-range forecasting. *Journal of the Korean Meteorological Society*, 127:433–446, 2000.

[35] M. Tewari, F. Chen, W. Wang, J. Dudhia, M.A. LeMone, K-Mitchell, M. Ekk, G. Gayno, J. Wegiel, and R.H. Cuenca. Implementation and verification of the unified NOAH lad surface model in the WRF model. *20th conference on weather analysis and forecasting/16th conference on numerical weather prediction*, pages 11–15, 2004.

[36] Z. Toth and E. Kalnay. Ensemble forecasting at NMC: The generation of perturbations. *Bulletin of the American Meteorological Society*, 74:2317–2330, 1993.

[37] Z. Toth and E. Kalnay. Ensemble Forecasting at NCEP and the Breeding Method. *Monthly Weather Review*, 125:3297, 1997.

[38] P. Van Leeuwen. Particle filtering in geophysical systems. *Monthly Weather Review*, 137:4089–4114, 2009.