Smooth Dynamical Downscaling of the Wind with the Model for Prediction Across Scales (MPAS)

Author: Marta Gil Bardají, mgilbardaji@gmail.com
Vortex Factoria de Càlculs, Marie Curie 8, 08042, Barcelona, Spain

Director: Pedro Peixoto, ppeixoto@usp.br
Institute of Mathematics and Statistics (IME) of the University of São Paulo (USP), Brazil

University Supervisor: Bernat Codina, bcodina@ub.edu
Facultat de Física, Universitat de Barcelona, Diagonal 645, 08028 Barcelona, Spain

Abstract: In the wind industry, wind time series for the past years are commonly generated using an atmospheric model to dynamically downscale large-scale reanalysis to local wind flow. Instead of relying on a nesting strategy like the Weather Research and Forecasting (WRF) model, the Model for Prediction Across Scales (MPAS) runs on variable resolution meshes that allow for smooth transitions. The goal of this study is to design MPAS meshes that are robust, accurate and computationally efficient for wind resource assessment applications.

We have designed a benchmark validation of one-year simulations in wind-energy-relevant locations representing different geographies and flow complexity scenarios. To focus on identifying real differences between modeling wind time series using regional MPAS meshes compared to using WRF nested domains, both models share the same settings whenever possible and the post-processing is analogous for MPAS and WRF output. Besides, WRF simulations were run with and without *nudging* (the assimilation of reanalysis data on all points of the outer WRF domain), which is an option known to improve the accuracy of the results, and it is not implemented in MPAS yet.

The real-data comparison shows that MPAS improves all wind speed metrics with respect to WRF simulations without nudging, but it has generally worse accuracy than the WRF simulations that do have nudging. This is a strong indication that the MPAS model indeed benefits from smooth transitions between scales, and that further developments on MPAS may turn it into a standard for wind resource applications.

I. INTRODUCTION

The correct representation of wind conditions and patterns is essential for wind resource assessment, and virtual measurements are a key tool to optimize the prospecting stages of a wind farm project (*Murthy and Rahi* 2017). Arguably the most accurate way of generating synthetic wind time series is to use an atmospheric model to downscale large-scale reanalysis to local wind flow (*Pinard et al.* 2010, *Tammelin et al.* 2013). The Weather Research and Forecasting (WRF) model is the current state of the art for wind resource assessment applications, and it relies on a nesting strategy to sequentially model higher resolution regions embedded in lower resolution ones. However, this methodology is known to create spurious numerical reflections at the transitions between nests which may contaminate the signal at the desired location (*Jiménez and Dudhia* 2013).

To avoid nesting, there has been a trend in recent years towards multi-scale models that smoothly transition from large to local scales (*Sakaguchi et al.* 2015). The original developers of WRF are now involved in the Model for Prediction Across Scales (MPAS) project, not only to provide such smooth transitions but also to couple the atmosphere with other earth-systems (*Skamarock et al.* 2012). The philosophy behind MPAS could be a game-changer for the wind industry, and although the model is still in development, the goal of this study is to present a first operative MPAS-powered downscaling solution for wind resource assessment.

The main challenge of running MPAS is to generate a mesh that covers the globe with cells of the desired resolutions which is computationally efficient and does not create undesired numerical errors (*Yang et al.* 2018). This study has focused on designing a mesh for first-stage wind prospecting: long-term simulations at relatively high resolutions over the wind farm area fed by reanalysis boundary conditions. The innovative ideas behind the methodology of building reliable meshes, as well as the limitations still in place, will be discussed in detail.

In section II we introduce the MPAS project, and outline the specifications of the MPAS atmospheric core. In section III we discuss how to set up WRF simulations that can be considered equivalent to MPAS simulations, since MPAS does not have all the schemes that are available in WRF. For example, a very relevant option for the wind industry, nudging, is explained and discussed in the subsection III B. The meshes that have been designed for this study are described in section IV, and the sensitivity study is discussed in subsection IV B. Finally, section V presents a one-year validation benchmark in 50 locations around the world. The conclusions of this paper sum up the contributions to the wind industry this research may entail.

II. MPAS MODEL OVERVIEW

MPAS stands for Model for Prediction Across Scales, and is a collaborative project that aims to couple the modeling of the atmosphere and other earth-systems for use in global and regional climate and weather studies (*Skamarock et al.* 2012). The project is being led by the National Center for Atmospheric Research (NCAR) and Los Alamos National Laboratory (LANL). NCAR, which is also the official developer of the Weather Research and Forecasting (WRF) model, has the primary responsibility for the atmospheric model MPAS-Atmosphere.

The defining characteristic of MPAS is the use of unstructured horizontal meshes (formally SCVTs: Spherical Centroidal Voronoi Tessellations) (*Ringler et al.* 2008), that allow for both quasi-uniform cell sizes and customizable refinements, as shown in Figure 1. The choice of the mesh is a key point in the design of an MPAS simulation, and can allow refinement over personalized regions where finer grid-spacing helps better model the desired event. For example, there have been studies that use topographical-refined meshes with smaller grid cells over complex regions like the Andes (*Santos and Peixoto* 2021).

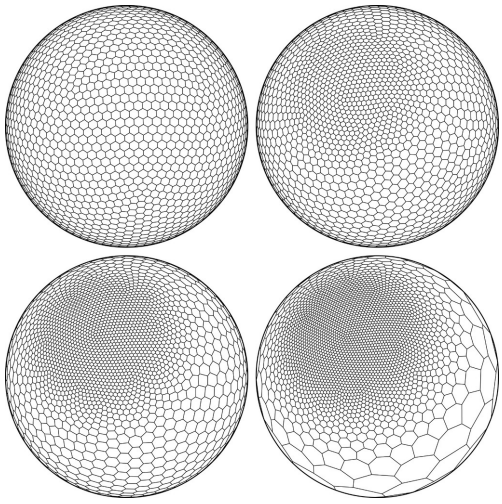

Figure 1: *Example of four MPAS meshes (SCVTs) with 2562 cells created by density functions that lead to different refinements. From Ringler et al. (2011).*

For many years, the focus of MPAS has been in global simulations that use the possibility to define coarse grid cells outside the area of interest (*Imberger et al.* 2021). The ability to run limited-area (regional) simulations was not available until version 7, released in mid 2019 (*Skamarock et al.* 2018). In wind resource assessment, mesoscale models are used to dynamically downscale global reanalysis to the area of interest, so the atmospheric model chosen needs to be able to feed the reanalysis as boundary conditions to the simulation. That is why the MPAS model was not a candidate to supersede WRF for wind industry applications until two years ago.

MPAS has not been designed as a substitute for WRF (*Powers et al.* 2017), meaning that their developers do not plan to include everything which is available in WRF in the new model. Mostly this aims to reduce the complexity of the model and target the most relevant add-ons. However, it is also very important to note that MPAS regional capabilities have still not been extensively studied and compared. If regional MPAS showed advantages over WRF, at least for some of its applications, the development of MPAS would surely move towards making the most of this edge. That is why this study aims to establish a benchmark of the current possibilities MPAS offers in the context of wind resource assessment.

A. Solver formulation

MPAS is a global fully compressible nonhydrostatic model that uses finite-volume numerical methods on an unstructured mesh and a C-grid discretization. The formulation of the equations in these grids is described in, amongst others, *Skamarock et al.* (2012), *Ringler et al.* (2010) and *Thuburn et al.* (2009). Some of its key characteristics are:

1. *Spatial discretization*: The centers of the cells are mathematically the generators of a spherical centroidal Voronoi Tessellation (SCVT), which means they are the centroids of their cells (taking into account customized density functions on the sphere), and all points inside a cell are closer to their center than to any of the other cells centers. This results in cells that are mostly hexagonal, but include some pentagons and heptagons. The most common algorithm used to generate SCVTs is Lloyd's algorithm (*Yang et al.* 2018).

2. *Vertical coordinate*: Height-based terrain-following vertical coordinate, that rapidly removes terrain influences in the coordinate surfaces with increasing height to avoid spurious errors due to steep terrain (*Klemp* 2011).

3. *C-staggering*: some of the variables (including the wind speed) are computed at the edges of the cells.

III. MPAS & WRF EQUIVALENT SIMULATIONS

Each case study will be simulated using:

- *MPAS:* Model MPAS (version 7.3, regional).
- *WRF:* Model WRF (version 3.9) without nudging.
- *WRF-WN:* Model WRF (version 3.9) with nudging applied on the outer domain.

where we distinguish between two WRF possibilities depending on the nudging option, which is explained in subsection III B.

In order to make a fair comparison of MPAS regional simulations with respect to WRF simulations, it is important to configure both models with the same settings, or at least, the most similar settings possible. In that way, differences in the results can be mainly attributed to the different horizontal structure used (variable resolution in MPAS versus nesting in WRF), or to the nudging scheme.

A. Model settings in WRF and MPAS

The reanalysis used for the initial and boundary conditions of the simulations is ERA5, generated by the European Centre for Medium-Range Weather Forecasts (*ECMWF*) and distributed by the Copernicus Climate Data Store (*Copernicus Climate Change Service (C3S)* 2017). ERA5 provides hourly estimates of a large number of atmospheric, land and oceanic climate variables in a 30 km grid and 137 vertical levels (*Hersbach et al.* 2020).

The set of physics schemes defined as part of the *convection-permitting suite* are not the same for WRF-3.9 and MPAS-7.3. Generally speaking, MPAS has less parametrization options than WRF, so in choosing the selected schemes the limiting model has usually been MPAS. In Table I there is a summary of the schemes and options used.

Physics Parameterization	**Scheme selected**
Microphysics	None
Radiation (Long Wave)	RRTMG (*Iacono et al.* 2008)
Radiation (Short Wave)	RRTMG (*Iacono et al.* 2008)
Land Surface	Noah (*Chen et al.* 1997)
Planetary Boundary Layer	MYNN (*Olson et al.* 2019)
Surface Boundary Layer	MYNN (*Olson et al.* 2019)
Convection/Cumulus	Grell-Freitas with shallow convection (*Grell and Freitas* 2014)
Cloud Fraction for Radiation	Xu-Randall (*Xu and Randall* 1996)
Gravity Wave Drag by Orography	None

Table I: *Physics schemes selected for the simulations.*

The vertical levels are fixed in both models to be quasi-analogous. In WRF, they are specified by pressure levels, whereas in MPAS they are given in vertical height units. We define 38 vertical levels, with the highest level being 30 km (in MPAS) and 5 hPa (in WRF). The specific locations of the levels is passed to the simulations as a list of approximately equivalent pressure and height levels, so that the first hundreds of meters of the atmosphere are well represented.

B. Nudging

Nudging is a four-dimensional data assimilation (FDDA) scheme that uses external information to adjust the tendencies of the simulation and prevent it from deviating too much from a reference (*David R. Stauffer and Nelson L. Seaman* 1990, *Liu et al.* 2008). There are different nudging implementations used in WRF for dynamical downscaling: analysis nudging, observational nudging and spectral nudging. In the context of wind industry prospecting studies, analysis nudging is the most widely used flavour of the scheme and it uses the reanalysis as the large-scale reference towards which the regional-model solution is relaxed. It has been seen that the wind flow is captured more accurately when analysis nudging (from now on, called simply nudging) is activated in WRF (*Ma et al.* 2016).

Nudging has been implemented and tested in global MPAS (*Bullock et al.* 2018), but it is neither available in the official version nor for MPAS version 7 (and therefore for regional simulations). One of the goals of this study is to evaluate whether nudging should be included in regional MPAS. For that purpose, in this study we will run two kinds of WRF simulations: with and without nudging. The WRF simulations without nudging can be considered equivalent to MPAS simulations in terms of model settings, whereas the WRF simulations with nudging show how can this technique improve the results.

Specifically, in this study the analysis nudging in WRF is performed relaxing temperature, humidity and wind variables to the reanalysis values for all vertical levels including those in the Planetary Boundary Layer, and the nudging coefficients are kept at their default values.

IV. GENERATING MPAS MESHES FOR WIND RESOURCE ASSESSMENT

For wind resource assessment, the main goal is to obtain a long-term description of a specific location or small area around the prospective location of the wind farm. Unavoidably, there is a trade-off between computational cost and horizontal resolution and period of the simulation. In this analysis we focus on 1-year simulations at relatively low resolutions, aiming for a grid-spacing of 3 km. This is the finest grid-spacing that has been extensively validated in the bibliography, and increasing it might mean encountering parametrization and double-counting issues. Since the reanalysis resolution is around 30 km, we want a variable resolution mesh that has cells of that approximate size on the region borders. To keep the number of cells as small as possible, the transition zone from low to high resolution should be as narrow as possible while conserving the quality of the mesh.

In the official MPAS-Atmosphere Release there is no mesh generation utility included. Instead, there is a web page (*NCAR* 2019) with several global quasi-uniform and variable-resolution meshes available for download. To ob-

tain a regional mesh from a global one, they do provide the `MPAS-Limited-Area` utility. However, none of the available global meshes can be used to generate a regional mesh that fulfills the wind industry needs. The high resolution areas are simply too large and the transition zones too wide, producing regional meshes with hundreds of thousands of cells.

To make the generation of regional meshes for the wind industry accessible, we created a GitHub open-source repository `marta-gil/vtx-mpas-meshes` that includes tools and examples to generate an MPAS mesh (ie, a SCVT) from a personalized global grid-spacing map. The process relies heavily on the software JIGSAW-GEO (*Engwirda* 2017) and on the MPAS-Tools package, which have been extensively used in the MPAS-Ocean community (*Engwirda* 2018, *Hoch et al.* 2020).

The grid-size or grid-spacing of a Voronoi cell is defined as the average distance from its center to its neighbours (*Santos and Peixoto* 2021). The mesh generation procedure creates MPAS cells with a grid-size as close to the requested one as possible, while maintaining the quality of the cells and the mesh as a whole.

In subsection IV A we describe the family of *doughnut* meshes and the parameters that define them: highest-grid spacing, central region size, transition area, etc. Changing these parameters allows us to create similar meshes and analyze the sensitivity of the simulated time series to the mesh used. This sensitivity test is explained in subsection IV B.

A. Mesh definition

The meshes are always created as global meshes centered at $(0, 0)$, and they are rotated to the desired location and cut to regional meshes at simulation time. That means that we can save a lot of space and speed in the generation of the meshes up by accepting huge MPAS cells outside the region of interest. Another consideration is that we will define meshes that have circular symmetry around the central point with respect to spherical distance. In operative applications, it could be interesting to rethink the shape of the mesh and adapt it to specific flow conditions at each site. However, in the scope of this study, it is more relevant to have the same mesh validated at different sites, and circular meshes are the most general choice.

A mesh of the *doughnut* family is fully characterized by the 4 parameters `cgs`, `mgs`, `s` and `m`. The central region of the mesh has radius `s` km (size) and uniform grid-spacing `cgs` km (central grid-spacing). Then, there is a linear increase of grid-spacing up to `mgs` km (mid grid-spacing) over a distance of `m` km (margin). After reaching the distance `s+m`, we keep the grid-spacing constant for 10·`mgs` km. This is done because `s+m` is going to be the nominal `radius` of the circular region that will be cut using MPAS-Limited-Area. However, the region will include 8 layers of cells *outside* the requested nominal radius. We

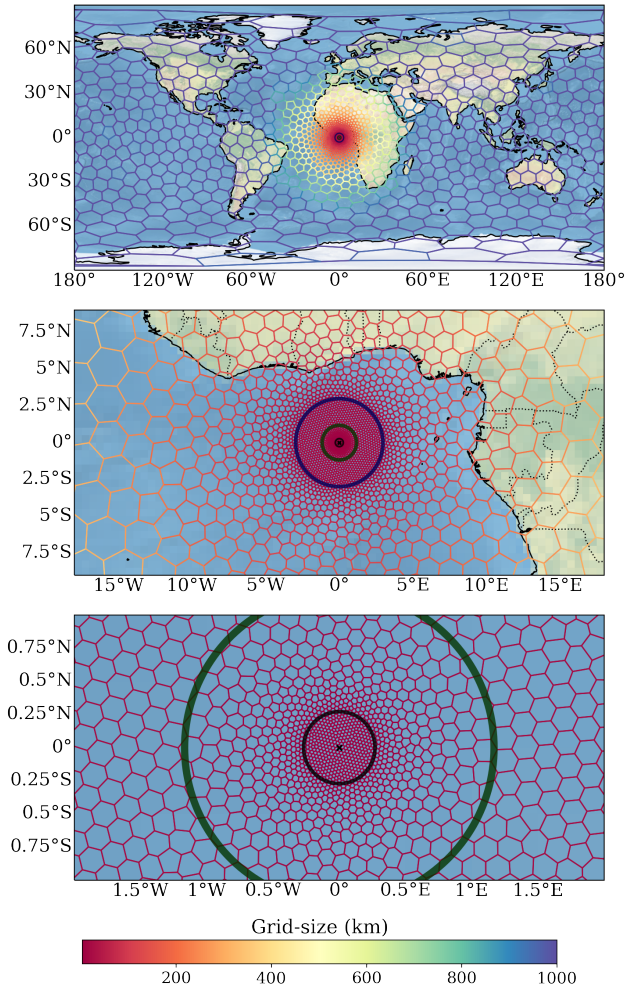

Figure 2: *Different zooms for a doughnut global mesh with `cgs`=3 km, `mgs`=20 km, `s`=30 km and `m`=100 km centered at* $(0, 0)$*. The thick blue ring marks the distance* α *where the grid-spacing is higher than `mgs`, the green ring the nominal radius of the region, and the black ring the central region with the smallest cells.*

therefore create this constant grid-spacing region so that these extra boundary cells all have grid-spacing `mgs`. After this hiatus in grid-spacing, we resume the linear increase of resolution until a very large grid-spacing (that was set to 1000 km for this study).

Mathematically, for a certain point in the globe $x = (lat, lon)$, the expected grid-spacing of the MPAS mesh $s(x)$ is obtained by the formula:

$$s(x) = \begin{cases} \texttt{cgs} & d(x) \leq \texttt{s} \\ \texttt{mgs} + \lambda\,(d(x) - \texttt{s}) & \texttt{s} \leq d(x) \leq \texttt{s+m} \\ \texttt{mgs} & \texttt{s+m} \leq d(x) \leq \alpha \\ \texttt{mgs} + \lambda\,(d(x) - \texttt{s}) & \alpha \leq d(x) \leq \beta \\ 1000 & \beta \leq d(x) \end{cases}$$

where $d(x)$ is the spherical distance of the point x to the center in km, $\lambda \equiv \frac{\texttt{mgs}-\texttt{cgs}}{\texttt{m}}$ is the slope of the tran-

sition regions, $\alpha \equiv \texttt{s} + \texttt{m} + 10 \cdot \texttt{mgs}$ is the radius of the maximum buffer area around the central region, and β is the distance from the center where $s(x) = 1000$ km, ie, $\beta \equiv \alpha + \frac{1000 - \texttt{mgs}}{\lambda}$.

In Figure 2 there are three plots that show a global *doughnut* mesh with three levels of zoom, and in Figure 3 we see the same mesh once it is centered at the desired location and the circular region (of nominal radius `s+m`) has been cut. Note that the boundary layers of the region lay outside the nominal radius, and that these boundary cells are inside the constant mid grid-spacing region.

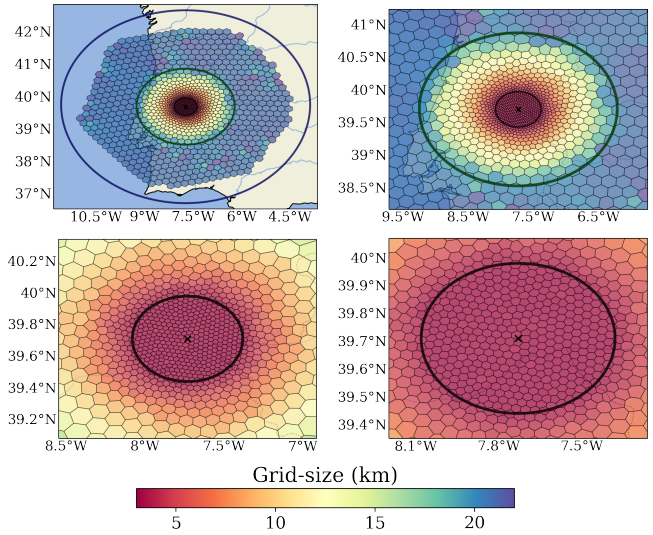

Figure 3: *Different zooms for a regional mesh obtained from the global mesh of Figure 2 centered at Perdigão. The blue ring marks the distance α where the grid-spacing is higher than `mgs`, the green ring the nominal radius of the region, and the black ring the central region.*

B. Sensitivity Analysis

The sensitivity analysis includes meshes of the *doughnut* family with different values for the parameters size `s` and margin `m`. All of the tested meshes have a central grid-spacing of `cgs`=3 km and mid grid-spacing `mgs` = 20 km (which is comparable to the reanalysis resolution). Increasing the margin parameter modifies the rate at which the grid-spacing increases in the transition section, so that higher margins give smoother meshes. The goal of this sensitivity analysis is to choose a value for the size and margin parameters that produces an accurate, robust and computationally efficient mesh.

1. *Meshes tested and WRF domains*

The sensitivity test has considered *regional doughnut* meshes with 3 different sizes `s`: 15 km, 25 km and 50 km, and 6 different margins `m`: 50 km, 75 km, 100 km, 125 km, 150 km and 250 km. That is, 18 different MPAS meshes. The smallest mesh has 617 cells and was generated in 26 seconds, whereas the largest mesh has 5199 cells and was generated in 82 seconds.

For each regional MPAS mesh, one needs to define which WRF domains will be used as comparison. In all cases, there are two one-way nested domains. The outer domain has 9 km grid-spacing and it is designed to cover most of the the MPAS region, while the inner domain has 3 km grid-spacing and covers the central region of the MPAS mesh. In Figure 4 the smallest and largest meshes are plotted next to their corresponding WRF domains, and in Figure 5 we zoom the meshes and domains.

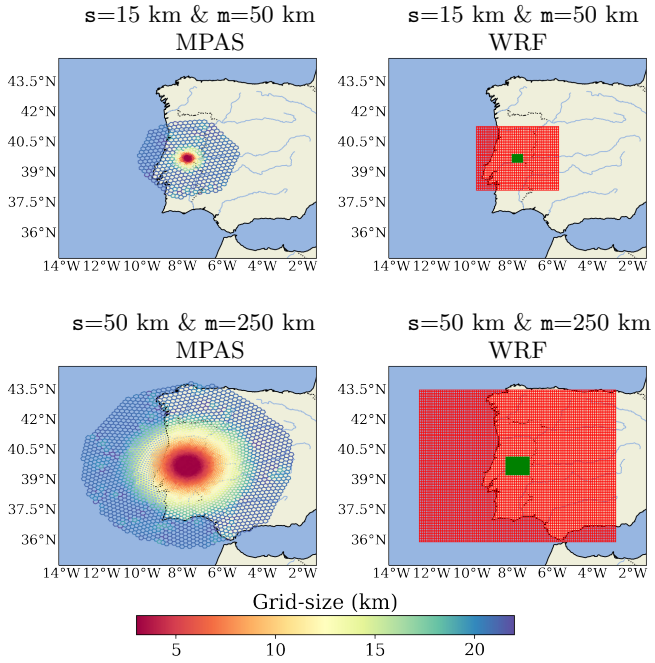

Figure 4: *The MPAS mesh and its corresponding WRF nested domains for the smallest and largest test cases of the sensitivity study. The outer WRF domain (red) has 9 km grid-spacing and the inner domain (green) 3 km.*

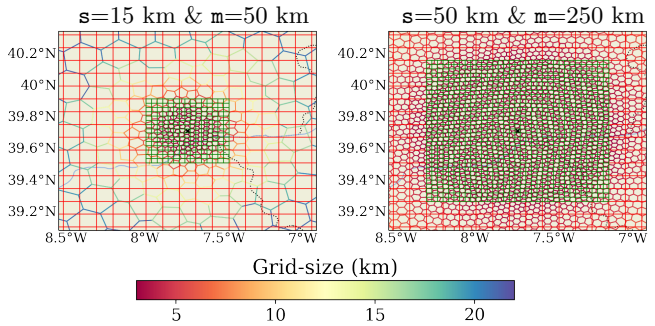

Figure 5: *Zoom on the central region of the MPAS mesh and its corresponding WRF nested domains for the smallest and largest test cases (see Figure 4).*

2. Results

The chosen test site is Perdigão, in Portugal, which has open-source anemometer measurements at 100 m for a 2-month period. The *MPAS*, *WRF* and *WRF-WN* simulations for each mesh/domain are centered at the Perdigão tower 13 (39.7136, -7.73), cover the full 2-month period, and are initialized daily with a spin-up time of 6 hours. As a reference, we will also study how the reanalysis ERA5 performs compared to measurements. We will consider two ways of extracting a time series at the desired location: considering the closest ERA5 node (*ERA5_nearest*) and interpolating the time series from the closest nodes (*ERA5_interp*).

In Figure 6 we show the evolution of the wind speed time series for one day from the different sources studied at their original frequencies, where the MPAS, WRF and WRF with nudging simulations lines show the mean wind speed and its spread across all the different meshes/domains tested.

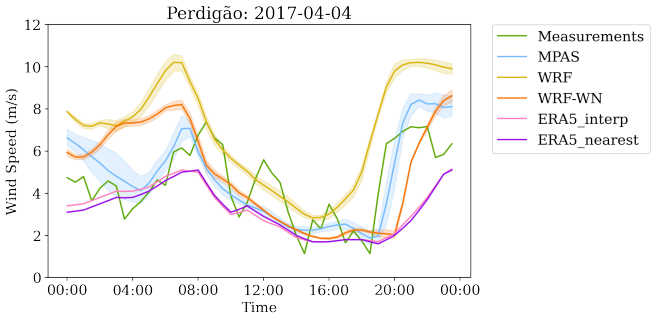

Figure 6: *Wind speed time series comparing measurements, reanalysis, and the mean across different meshes for models MPAS, WRF and WRF-WN. The shaded area represents the spread across all the different meshes/domains tested.*

After re-sampling all time series to 30-min data, we validate the wind speed against measurements. In Figure 7 we see the performance in terms of wind speed correlation and root mean squared error (RMSE) of the tests run for each model, and also the ERA5 reanalysis metrics for reference. There are several interesting observations:

- The correlation of the reanalysis is greatly improved by any of the dynamically down-scaled simulations.
- The RMSE of the reanalysis is improved by MPAS and WRF-WN simulations, but it is worsened by almost every WRF simulation that does not use nudging.
- The worst metrics are seen for WRF simulations without nudging. In fact, every WRF simulation without nudging is outperformed by any other MPAS or WRF-WN simulation.
- MPAS significantly outperforms the "equivalent" WRF simulations (without nudging), but it's not as accurate as the WRF-WN simulations, that had the nudging scheme on.
- WRF and WRF-WN show greater differences in accuracy between different meshes/domains, whereas for MPAS simulations all meshes show similar performance.

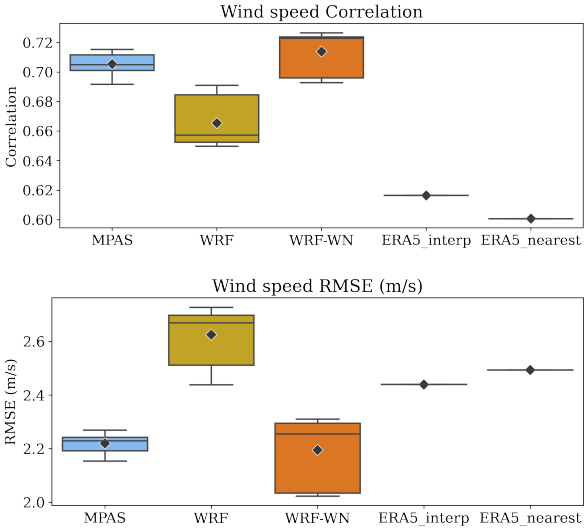

Figure 7: *Results of the sensitivity test. Wind speed validation against measurements for each simulation run using different meshes and models. The mean value of the metric across test cases is highlighted. The ERA5 reanalysis performance is also shown.*

The detailed results for each choice of size `s` and margin `m` parameters are shown in Figure 8. The figure also shows the computational cost of each 2-months simulation in terms of a metric measured in hours that is proportional to the energy consumption of the process. With this extra information we can make some new observations:

- The cost of running the smallest mesh in MPAS is approximately the 60% than in WRF, but for the largest mesh the cost in MPAS is the 140% of the cost in WRF. The simulation cost grows faster in MPAS than in WRF when the margin or the size parameters are increased.
- WRF and WRF-WN are more sensitive to the size `s` (which is the size of the inner domain) than to the margin `m` (which is related to the size of the outer domain). In WRF-WN it is especially evident that changing the margin parameter has almost no effect.
- When there is no nudging in WRF, using small inner domains seems to lead to very low accuracy. This is what we can expect from general WRF good practices, that require having reasonable-sized domains in order to avoid having buffer cells all the way to the center of the domain.

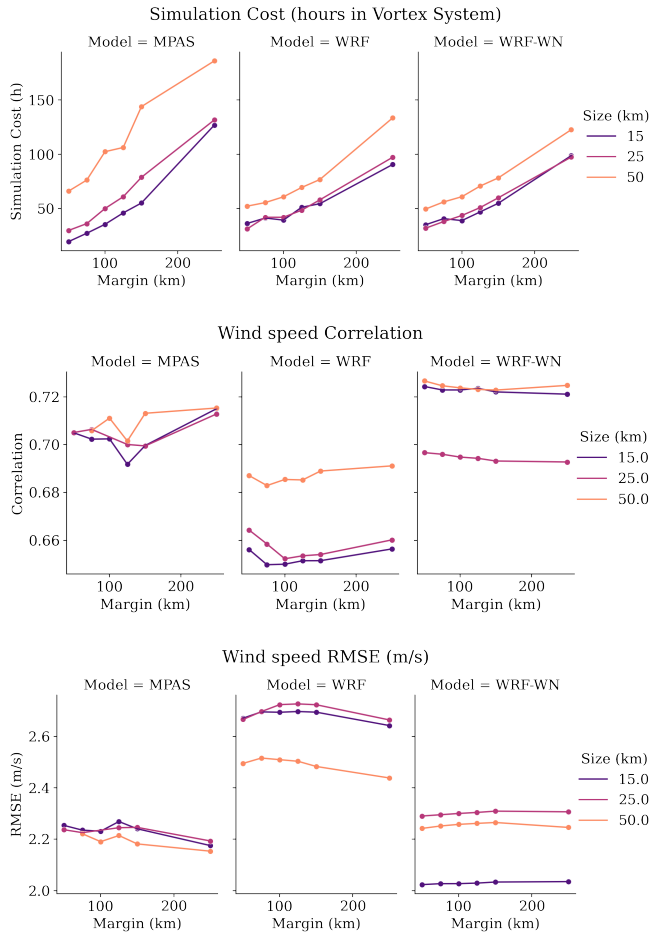

Figure 8: *Results of the sensitivity test in detail for every case, that is determined by a choice of size* `s` *and margin* `m` *parameters for the MPAS mesh.*

- For WRF-WN it is not so clear whether small or large inner domains are preferable, because not only it is the *smallest* inner domain the one that shows the best performance, but also the worst one is *middle*-sized inner domain simulation.

- In MPAS, we can see a trend for better metrics when the margin or the size parameters are increased, ie, when the transition in grid-spacing is smoother and the central region is larger.

So we see that the modification of the size and margin parameters behaves as expected in MPAS: larger size and margin lead to more accurate results, but at a higher computational cost. Surprisingly, such clear behaviour was not seen for the corresponding WRF domains.

With this results in mind, we suggest that a reasonable choice for an MPAS mesh that is accurate but also computationally-efficient is setting margin `m`=100 km and size `s`=30 km. This mesh is the one that has been shown in Figures 2 and 3, and will be used in the following section to validate real-world data.

V. VALIDATION BENCHMARK

We have selected 50 towers around the globe with wind speed and wind direction 10-min measurements for a full year, available in the private database of the company Vortex. In Figure 9, we see the locations and kind of terrain of the selected sites: flat terrain (26 sites), complex terrain (14 sites) and forest (10 sites). The height of the measurements varies from 20 to 120 meters above ground level, with the mean value being (65.7 ± 21.6) m. The period of the measurements is in the 1998-2015 range.

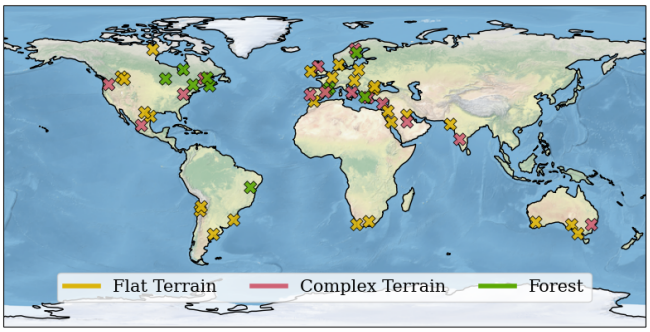

Figure 9: *Location of the 50 measurement sites, coloured according to the terrain of the area: yellow for flat terrain, red for complex terrain and green for forests.*

A. Methodology

Each site is simulated using the three configurations described in III using the MPAS mesh chosen in the sensitivity test of the previous section with size `s`=30 km and margin `m`=100 km. The one-year time series are obtained by separately modelling each day. The daily simulations take into account a spin-up time of 6 hours, so each run is 30 hours long. This allows for an efficient distribution of tasks between the cluster of computers (CPUs), since each day can be sent to a different machine.

The post-process from native MPAS and WRF formats is carried out in *python* and also using specific NCAR tools (like `convert_mpas` to convert MPAS meshes to lat-lon gridded structures). Finally, wind speed and direction time series are extracted at the coordinates and height of the measurements. Since the model output is saved every 30 minutes, measurements are re-sampled to that frequency before the validation.

B. Results and Discussion

For each site, we compare the mean wind speed of each model (bias and absolute bias in %); consider time-domain metrics (RMSE and 30-min/daily/monthly correlations); and also compute metrics that focus on the wind speed distribution, which is of great importance in

wind resource analysis. On the one hand, the wind speed histogram typically follows a *Weibull distribution* (*Bowden et al.* 1983), so we compute the absolute bias (%) of the fitted scale parameter A and shape parameter k. On the other hand, we compare the the mean error in the *wind-sectors 2d histogram* (%), which is the mean of relative error (%) between the modeled and observed number of events in specific wind speed bins and wind direction sectors. In this case, 16 wind sectors and 1 m/s wind speed bins are used.

The results have been averaged across the 50 sites and compared in Table II. Besides the validation of the three regional simulations (*MPAS*, *WRF*, *WRF-WN*), we also include the validation of the time series obtained by interpolating ERA5 nodes (*ERA5_interp*) and by using the nearest node of the reanalysis (*ERA5_nearest*). In Figure 10 some of the metrics described in the Table II are represented. In terms of computational cost, MPAS simulations were around 20% more costly than WRF simulations. MPAS could be sped up using the GPU version of the model or using parallelization.

Benchmark validation results

Wind Speed Metrics	*MPAS*	*WRF*	*WRF-WN*	*ERA5_interp*	*ERA5_nearest*
Bias (%)	4.0 ± 11.6	6.7 ± 13.0	0.0 ± 12.5	-32.6 ± 15.1	-31.2 ± 15.0
Absolute bias (%)	10.1 ± 7.0	11.4 ± 9.2	9.9 ± 7.5	32.6 ± 15.1	31.3 ± 14.8
RMSE (m/s)	2.39 ± 0.47	2.56 ± 0.50	2.23 ± 0.49	3.55 ± 1.40	3.41 ± 1.26
30-min correlation	0.64 ± 0.09	0.62 ± 0.10	0.68 ± 0.09	0.49 ± 0.22	0.51 ± 0.20
Daily correlation	0.85 ± 0.07	0.82 ± 0.08	0.86 ± 0.07	0.67 ± 0.24	0.69 ± 0.21
Monthly correlation	0.91 ± 0.06	0.87 ± 0.13	0.91 ± 0.09	0.75 ± 0.26	0.79 ± 0.20
Weibull A absolute bias (%)	10.1 ± 7.0	11.4 ± 9.2	9.8 ± 7.6	32.7 ± 15.1	31.4 ± 14.7
Weibull k absolute bias (%)	8.7 ± 8.3	9.8 ± 9.2	9.7 ± 9.6	12.9 ± 14.2	12.0 ± 13.0
Mean error wind-sectors (%)	55.3 ± 22.7	59.7 ± 24.8	59.0 ± 28.5	106.2 ± 54.6	102.7 ± 51.2

Table II: *Comparison of wind speed metrics across 50 sites for MPAS, WRF, WRF with nudging, ERA5 interpolating the nearest nodes, and ERA5 selecting the nearest node. The definition of the metrics is detailed in section V B.*

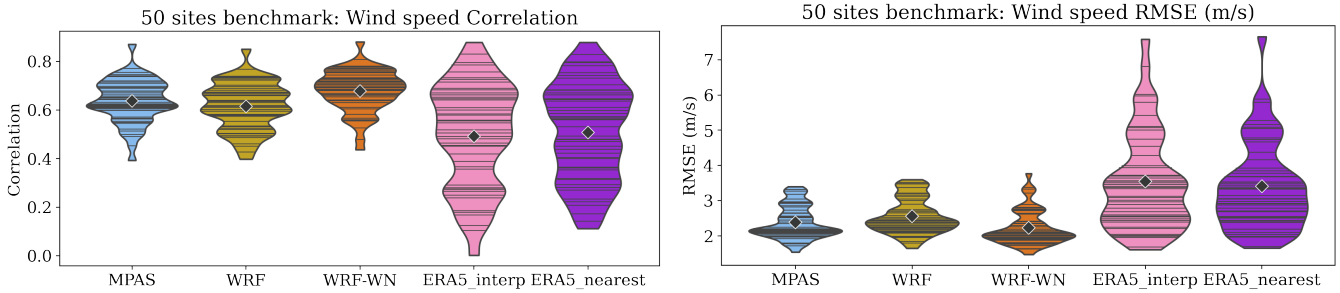

Figure 10: *Results of the validation benchmark for 50 sites, shown using a violinplot (empirical distribution plot) where black lines mark the value at each site. The mean value of the metric across sites is highlighted.*

Firstly, we notice that modelling significantly improves the metrics that would be obtained by the reanalysis alone. The reanalysis greatly underestimates the mean wind speeds and doesn't capture the wind distribution correctly. Therefore, the dynamical downscaling of reanalysis to higher resolutions is seen to be a useful method to characterize the wind speed of a site in greater detail.

Secondly, MPAS simulations are not simply comparable to WRF in accuracy: they show a better performance in every single metric. This is a very powerful indication that the MPAS meshes generated for this purpose are robust, accurate and useful. Note also that not only the mean values across sites are better in MPAS than WRF, but also the spread of the results is consistently lower.

Finally, we see that WRF with analysis nudging (WRF-WN) is much better than standard WRF (at least in terms of daily wind speed simulations). In fact, the metrics improve so much that they are in this case almost always better than the MPAS results. Interestingly, MPAS is the best model in two of the distribution metrics, which are specially relevant for wind resource assessment.

So, the results of this benchmark corroborate the first results seen in the sensitivity test of section IV B.

VI. CONCLUSIONS

This study has put together and validated a robust methodology to run MPAS in contexts such as wind resource assessment. To do so, the first step has been to design equivalent reference WRF simulations that can be used to benchmark the regional MPAS results. Arguably our most significant innovation is the open-source mesh generation strategy that opens the door to generating customized MPAS meshes easily. Finally, the sensitivity test and the benchmark validation of 50 one-year meteorological towers have shown that our MPAS workflow is consistently accurate. The fact that in many occasions wind speeds modeled in MPAS are in greater agreement with real-data measurements than wind speeds modeled in WRF is a strong indication that the smooth transition across scales in MPAS is beneficial.

However, one must acknowledge that the validation performed in this study is not enough to draw strong conclusions on the superiority of MPAS over WRF. MPAS results should be tested further to ensure that the meshes we are using do not create unwanted numerical errors due to their fast resolution transitions. Moreover, in the wind industry the standard downscaling technique is to use WRF with nudging, that in this study appears to generally outperform MPAS. Therefore, we cannot claim that MPAS is better than the state-of-the-art wind resource assessment strategy. In spite of the aforementioned reasons for caution, this study presents a dynamical downscaling technique using MPAS that can already be very useful for the wind industry. After all, running a simulation using two models is an excellent way to obtain an uncertainty estimate.

We aim to emphasize the great potential regional MPAS shows, despite the main focus of the MPAS community not being regional simulations. Some of the topics that may increase the appeal of regional MPAS include:

- Generating MPAS meshes that can resolve higher resolutions using smooth transitions. In the wind industry, grid-spacings of 100 meters are commonly run using WRF.
- Coupling MPAS to a microscale model that can provide higher resolutions and accurately model wind turbulence.
- Adding an implementation of analysis nudging in regional MPAS to evaluate how much the wind time series can be improved by relaxing the simulation towards the reference reanalysis.
- Implementing the physics parametrizations that are available in WRF and widely used in the wind industry, but have not been added to MPAS yet.

ACKNOWLEDGMENTS

The measurements and computer power used in this project were provided by *Vortex Factoria de Càlculs*, where I have been working for over three years.

I want to thank Dr. Pedro Peixoto for all the guidance and support as director of this research, and Dr. Bernat Codina for his advise on the writing process. In addition, I must thank my colleagues at Vortex that were always there for discussions and brainstorming. They have taught me a lot about the wind industry, and given me the opportunity to explore the possibilities MPAS has to offer. A special thanks to Pau Casso and Gerard Cavero, who played essential roles in many steps of the results presented in this thesis.

REFERENCES

Bowden, G., P. Barker, V. Shestopal, and J. Twidell, The Weibull distribution function and wind power statistics, *Wind Engineering*, *7*, 85–98, 1983.

Bullock, O. R., H. Foroutan, R. C. Gilliam, and J. A. Herwehe, Adding four-dimensional data assimilation by analysis nudging to the Model for Prediction Across Scales - Atmosphere (version 4.0), *Geoscientific Model Development*, *11*(7), 2897–2922, doi:10.5194/GMD-11-2897-2018, 2018.

Chen, F., Z. Janjić, and K. Mitchell, Impact of Atmospheric Surface-layer Parameterizations in the new Land-surface Scheme of the NCEP Mesoscale Eta Model, *Boundary-Layer Meteorology 1997 85:3*, *85*(3), 391–421, doi:10.1023/A:1000531001463, 1997.

Copernicus Climate Change Service (C3S), ERA5: Fifth generation of ECMWF atmospheric reanalyses of the global climate, 2017.

David R. Stauffer, and Nelson L. Seaman, Use of Four-Dimensional Data Assimilation in a Limited-Area Mesoscale Model. Part I: Experiments with Synoptic-Scale Data, *Monthly Weather Review*, pp. 1250–1277, doi:10.1175/1520-0493(1990)118¡1250:UOFDDA¿2.0.CO;2, 1990.

Engwirda, D., JIGSAW-GEO (1.0): Locally orthogonal staggered unstructured grid generation for general circulation modelling on the sphere, *Geoscientific Model Development*, *10*(6), 2117–2140, doi:10.5194/gmd-10-2117-2017, 2017.

Engwirda, D., Generalised primal-dual grids for unstructured co-volume schemes, *Journal of Computational Physics*, *375*, doi:10.1016/j.jcp.2018.07.025, 2018.

Grell, G. A., and S. R. Freitas, A scale and aerosol aware stochastic convective parameterization for weather and air quality modeling, *Atmos. Chem. Phys*, *14*, 5233–5250, doi:10.5194/acp-14-5233-2014, 2014.

Hersbach, H., et al., The ERA5 global reanalysis, *Quarterly Journal of the Royal Meteorological Society*, *146*(730), 1999–2049, doi:10.1002/QJ.3803, 2020.

Hoch, K. E., M. R. Petersen, S. R. Brus, D. Engwirda, A. F. Roberts, K. L. Rosa, and P. J. Wolfram, MPAS-Ocean

Simulation Quality for Variable-Resolution North American Coastal Meshes, *Journal of Advances in Modeling Earth Systems*, *12*(3), doi:10.1029/2019MS001848, 2020.

Iacono, M. J., J. S. Delamere, E. J. Mlawer, M. W. Shephard, S. A. Clough, and W. D. Collins, Radiative forcing by long-lived greenhouse gases: Calculations with the AER radiative transfer models, *J. Geophys. Res*, *113*, 13,103, doi:10.1029/2008JD009944, 2008.

Imberger, M., X. G. Larsén, and N. Davis, Investigation of Spatial and Temporal Wind-Speed Variability During Open Cellular Convection with the Model for Prediction Across Scales in Comparison with Measurements, *Boundary-Layer Meteorology 2021 179:2*, *179*(2), 291–312, doi:10.1007/S10546-020-00591-0, 2021.

Jiménez, P. A., and J. Dudhia, On the Ability of the WRF Model to Reproduce the Surface Wind Direction over Complex Terrain, *Journal of Applied Meteorology and Climatology*, *52*(7), 1610–1617, doi:10.1175/JAMC-D-12-0266.1, 2013.

Klemp, J. B., A Terrain-Following Coordinate with Smoothed Coordinate Surfaces, *Monthly Weather Review*, *139*(7), 2163–2169, doi:10.1175/MWR-D-10-05046.1, 2011.

Liu, Y., et al., The Operational Mesogamma-Scale Analysis and Forecast System of the U.S. Army Test and Evaluation Command. Part I: Overview of the Modeling System, the Forecast Products, and How the Products Are Used, *Journal of Applied Meteorology and Climatology*, *47*(4), 1077–1092, doi:10.1175/2007JAMC1653.1, 2008.

Ma, Y., Y. Yang, X. Mai, C. Qiu, X. Long, and C. Wang, Comparison of Analysis and Spectral Nudging Techniques for Dynamical Downscaling with the WRF Model over China, *Advances in Meteorology*, *2016*, doi:10.1155/2016/4761513, 2016.

Murthy, K. S., and O. P. Rahi, A comprehensive review of wind resource assessment, *Renewable and Sustainable Energy Reviews*, *72*, 1320–1342, doi:10.1016/J.RSER.2016.10.038, 2017.

NCAR, MPAS-Atmosphere Meshes, 2019.

Olson, J. B., J. S. Kenyon, W. A. Angevine, J. M. Brown, M. Pagowski, and K. Sušelj, A Description of the MYNN-EDMF Scheme and the Coupling to Other Components in WRF–ARW, *NOAA Technical Memorandum OAR GSD*, p. 61, doi:10.25923/N9WM-BE49, 2019.

Pinard, J. D. P., R. Benoit, and W. Yu, A west wind climate simulation of the mountainous Yukon, *http://dx.doi.org/10.3137/ao.430306*, *43*(3), 259–281, doi:10.3137/AO.430306, 2010.

Powers, J. G., et al., The weather research and forecasting model: Overview, system efforts, and future directions, *Bulletin of the American Meteorological Society*, *98*(8), 1717–1737, doi:10.1175/BAMS-D-15-00308.1, 2017.

Ringler, T., L. Ju, and M. Gunzburger, A multiresolution method for climate system modeling: Application of spherical centroidal Voronoi tessellations, *Ocean Dynamics*, *58*(5-6), 475–498, doi:10.1007/S10236-008-0157-2, 2008.

Ringler, T. D., J. Thuburn, J. B. Klemp, and W. C. Skamarock, A unified approach to energy conservation and potential vorticity dynamics for arbitrarily-structured C-grids, *Journal of Computational Physics*, *229*(9), 3065–3090, doi:10.1016/J.JCP.2009.12.007, 2010.

Ringler, T. D., D. Jacobsen, M. Gunzburger, L. Ju, M. Duda, and W. Skamarock, Exploring a Multiresolution Modeling Approach within the Shallow-Water Equations, *Monthly Weather Review*, *139*(11), 3348–3368, doi:10.1175/MWR-D-10-05049.1, 2011.

Sakaguchi, K., et al., Exploring a Multiresolution Approach Using AMIP Simulations, *Journal of Climate*, *28*(14), 5549–5574, doi:10.1175/JCLI-D-14-00729.1, 2015.

Santos, L. F., and P. S. Peixoto, Topography-based local spherical Voronoi grid refinement on classical and moist shallow-water finite-volume models, *Geoscientific Model Development*, *14*(11), doi:10.5194/gmd-14-6919-2021, 2021.

Skamarock, W. C., J. B. Klemp, M. G. Duda, L. D. Fowler, S.-H. Park, and T. D. Ringler, A Multiscale Nonhydrostatic Atmospheric Model Using Centroidal Voronoi Tesselations and C-Grid Staggering, *Monthly Weather Review*, *140*, 3090–3105, doi:10.1175/MWR-D-11-00215.1, 2012.

Skamarock, W. C., M. G. Duda, S. Ha, and S. H. Park, Limited-Area Atmospheric Modeling Using an Unstructured Mesh, *Monthly Weather Review*, *146*(10), 3445–3460, doi:10.1175/MWR-D-18-0155.1, 2018.

Tammelin, B., et al., Production of the Finnish Wind Atlas, *Wind Energy*, *16*(1), 19–35, doi:10.1002/WE.517, 2013.

Thuburn, J., T. D. Ringler, W. C. Skamarock, and J. B. Klemp, Numerical representation of geostrophic modes on arbitrarily structured C-grids, *Journal of Computational Physics*, *228*(22), 8321–8335, doi:10.1016/J.JCP.2009.08.006, 2009.

Xu, K.-M., and D. A. Randall, A Semiempirical Cloudiness Parameterization for Use in Climate Models, *Journal of the Atmospheric Sciences*, *53*(21), 3084–3102, doi:10.1175/1520-0469(1996)053¡3084:ASCPFU¿2.0.CO;2, 1996.

Yang, H., M. Gunzburger, and L. Ju, Fast spherical centroidal Voronoi mesh generation: A Lloyd-preconditioned LBFGS method in parallel, *Journal of Computational Physics*, *367*, doi:10.1016/j.jcp.2018.04.034, 2018.