

On of the Performance of the WRF Numerical Model over Complex Terrain on a High Performance Computing Cluster

Nicholas Christakis

Department of Physics and IACM, FORTH-GR
University of Crete
Heraklion, Greece
nchristakis@tem.uoc.gr

Theodoros Katsaounis

Department of Mathematics and Applied Mathematics and
IACM, FORTH-GR
University of Crete
Heraklion, Greece
thodoros@tem.uoc.gr

George Kossioris

Department of Mathematics and Applied Mathematics and
IACM, FORTH-GR
University of Crete
Heraklion, Greece
kossioris@math.uoc.gr

Michael Plexousakis

Department of Mathematics and Applied Mathematics and
IACM, FORTH-GR
University of Crete
Heraklion, Greece
plex@tem.uoc.gr

Abstract— In this paper the performance of the Weather Research and Forecasting (WRF) numerical model over complex terrain on a High Performance Computing (HPC) cluster is discussed. Simulations were performed with two different nested domain configurations and the patterns and variability of near-surface winds over two selected areas of the Greek territory were compared with data from surface meteorological stations. Various verification statistics were utilized to gauge the ability of the model for accurate predictions and its performance for the different configurations was assessed. Finally, suggestions were made concerning performance optimization for high resolution simulations.

Keywords— *HPC cluster; WRF; nested domains; high resolution simulations; model performance; surface wind patterns*

I. INTRODUCTION

In atmospheric physics and meteorology, it is critical to obtain accurate forecasts of meso- to micro-scale weather phenomena within reasonable time scales and sufficient lead time (at least 48 hours). In order to achieve that, high spatial and temporal resolutions are required and hence, sufficient computational power is needed. In recent years, significant effort has been put in analyzing the performance of Numerical Weather Prediction (NWP) models and in particular WRF with different high performance computational setups in

single-domain simulations with uniform spatial resolution [1]-[5]. However, not much work has been carried out on high resolution nested domains simulations, which are essential for high fidelity forecasts. Recently [6], the performance of an older version of WRF was assessed in different nested domain configurations. In the present paper, we discuss the performance of WRF Version 3 on an HPC cluster (72 cores in total) and the cost to produce in two different nested domain configurations accurate forecasts when the highest spatial resolution is 3-km and 1-km, respectively in regions of complex topography. For this study, the near-surface wind patterns and variability over two selected areas of Greece were chosen and the sensitivity of the WRF model for the two different nested domain configurations is measured against two wind-measuring stations consisting of 30 m masts equipped with vane anemometers. The presented work is based on preliminary results of on-going work within the auspices of the AKAIPRO research project [7]. We chose to simulate a typical summer period, 11-17 July 2013, characterized by moderate to strong north or north-easterly winds. This particular period was chosen because of the fairly stable atmospheric conditions in the region.

The period between 11 and 17 July 2013, was typically characterized by Etesian winds. They are northern sector

winds blowing over the Aegean Sea during summer and early autumn (see e.g. [8]).

The remaining of the paper is organized as follows: Section II describes the numerical and experimental setup. Section III contains an analysis of the mesoscale model performance with the different nested domain options. Our conclusions are summarized in Section IV.

II. NUMERICAL AND EXPERIMENTAL SETUP

The WRF model [9], utilized in the AKAIPRO research project, is a state-of-the-art regional to global-scale numerical weather prediction model that is used for operational forecasting as well as academic research all over the world. The model can be built in serial, parallel (MPI) and mixed-mode (OpenMP and MPI) forms and has been specifically designed to perform well on massively parallel computers [6]. For this study, all simulations were performed with WRF Version 3.4.1.

A. The TYPHOON HPC Cluster

The TYPHOON HPC system is a Fujitsu blade cluster with 12 Intel Xeon E5-2430 @ 2.2 GHz cpu's. Each blade hosts a pair of cpu's and contains 16GB of RAM and a 146GB SAS hard drive. It has been deployed by the AKAIPRO research project [7] to pursue work in weather forecasting, development of climate models and other critical research. The jobs submitted to the cluster have included both small WRF runs on separate nodes or large runs on the entire TYPHOON cluster. The GNU compilers were tested and worked but using the Intel C/C++/Fortran compilers resulted in substantially better performance. Best results were achieved using a distributed-memory parallelism (dmpar) build, which enabled MPI. However, due to the small size of the cluster, a hybrid build (dmpar+smpar) was used in all runs reported below. It has to be noted that hyper-threading was turned off, since it had only a marginal or negligible improvement in computing performance.

B. Experimental Setup

From a number of observational sites, two in northeastern Greece were chosen for this investigation, with wind speed and direction measurements recorded at 20m, 28m and 30m above ground level (agl). The locations of the observation (surface) stations and their codes are shown in Fig. 1. The two sites (Kavala-kav and Loutro-ltr) were chosen because of the convoluted topography of the region, with the mountainous bulks at the Greece-Bulgaria border lying at their north, the Aegean sea at their south, the Evros river plain and the Dardanelle Straits at their east and the peninsula of mount Athos at their west.

Wind speed and wind direction were recorded by vane anemometers, placed on 30 m masts located at the observation stations mentioned above, for the week between 11 and 17 July 2013. The masts were deployed as part of the AKAIPRO project. The 20 min average wind speed and wind direction at


28 m above ground level were used for comparison with the WRF simulations. In turn, the WRF-simulated wind speed and wind direction were estimated at 28m by interpolation.



Fig.1 Locations of the observational sites together with their codes.

C. Numerical Setup

For the numerical experiments, the initial and boundary conditions were derived from the National Centers for Environmental Prediction (NCEP) model analysis by WRF preprocessing. An improved representation of topography with $60\text{ m} \times 60\text{ m}$ spatial resolution that has been specifically designed for the AKAIPRO project and an updated land-use dataset with 24 land-use categories from the U.S. Geological Survey (USGS) were used, in order to ensure more accurate surface conditions. These simulations were run without nudging [10]. The two different nested domain configurations are shown in Fig. 2.



(a)




Fig.2 The two different nested domain configurations (a) 3-domain configuration C1, with horizontal grid resolutions d01: 27 km, d02: 9 km, d03: 3 km, (b) 4-domain configuration C2, with horizontal grid resolutions d01: 27 km, d02: 9 km, d03: 3 km, d04: 1 km.

In both configurations, domains d01 and d02 were exactly the same. In C1, domain d03 covers an area of 1038 km × 1065 km, while in C2 domain d03 covers an area of 669 km × 453 km and domain d04 an area of 247 km × 223 km. It should be noted that the area covered by d03 in C1 is larger than the area covered collectively by d03 and d04 in C2 with the number of grid points in d03 of C1 being 122,830 whereas the number of grid points in d03 and d04 of C2 are 33,673 and 55,081, respectively. The ratio of the number of grid points of d03 and d04 of C2 to the number of grid points of d03 of C1 is approximately 3/4. However, for reasons of numerical stability, the timestep employed in the C2 configuration for d01 (the outer domain) is 27 s, half of the one employed for d01 in the C1 configuration (54 s). Small timesteps had to be employed in order for the numerical stability criterion to be satisfied [9]. In both configurations, the timestep of each nested domain was one third of the timestep of its parent domain. Only in C2 was the timestep of d04 taken equal to the timestep of d03, since this did not violate the stability criterion and helped speed up the calculations.

Two-way nesting was chosen, since it has been found to perform in general better than one-way nesting [11]. All modelling domains had 50 layers in the vertical dimension, with the model top being set at 50 hPa. The Noah land surface model was used for surface layer parametrization [12]. The Kain-Fritsch scheme was used for parametrizing cumulus convection [13]. Shortwave radiation processes were handled using a cloud radiation scheme [14] and the Rapid Radiative Transfer Model (RRTM) scheme [15] was applied for longwave radiation processes. A total of 9 different combinations were tested on the three-domain nested configuration C1 (see Fig. 2) for the microphysics and planetary boundary layer (PBL) parametrizations, as these are among the processes considered to have a stronger influence on wind fields close to the surface [16]. After careful analysis of the results (not shown here), it was concluded that the combination that performed better was the one utilizing the Thompson scheme for the microphysics [17] and the MYNN2.5 scheme for the PBL [18]. Hence, it was decided to

perform the analysis on the four-domain nested configuration C2 (see Fig. 2) utilizing this combination of physical parametrizations. For all simulations, 12-hour model spin-up time has been allowed.

The performance of each different nested domain configuration was performed by computing various statistics for both the wind speed and direction. For the wind speeds the standard BIAS and RMSE measures and their relative counterparts are used, while for the wind direction the relative BIAS and the DACC statistics are used (see [16]). Let m_i and o_i denote the modelled and observed values of wind speed, respectively. The first metric we use is the Mean Error or BIAS, which measures the overall overestimation or underestimation of modelled wind speed values and is defined by:

$$BIAS = \frac{1}{N} \sum_{i=1}^N (m_i - o_i) \quad \text{and}$$

$$\text{Relative BIAS} = \frac{BIAS}{O_{avg}} \times 100$$

where $O_{avg} = \frac{1}{N} \sum_{i=1}^N o_i$ with N the number of records. The

second metric we use is the Root Mean Square Error (RMSE) between the modelled and observed values and measures the amount of scatter of the wind speed errors:

$$RMSE^2 = \frac{1}{N} \sum_{i=1}^N (m_i - o_i)^2 \quad \text{and}$$

$$\text{Relative RMSE} = \frac{RMSE}{O_{avg}} \times 100$$

To measure the directional accuracy of the wind, besides the relative BIAS metric, we also adopted the *DACC* metric. For any two angles α and β the circular distance is defined as $\Delta\theta(\alpha, \beta) = \min|\alpha - \beta, 360^\circ - (\alpha - \beta)|$. Then, the *DACC* metric measures the percentage of times in which the circular distance between modelled and observed wind directions is lower than a threshold, chosen as 30°:

$$DACC = \frac{\sum_{i=1}^N \begin{cases} 1, & \text{if } 0 \leq \Delta\theta_i \leq 30 \\ 0, & \text{otherwise} \end{cases}}{N} \times 100$$

At this point, it should be noted that configuration C1 was tested for two different timesteps for the outer domain d01, namely 54 sec and 27 sec (same as C2 d01 timestep), in order to ensure reproducibility of the results. No significant differences in wind speed and direction were found in the results (with the mean relative RMSE not exceeding 8.7%) and it was observed (see Fig. 3) that WRF scales linearly for this configuration of nested domains; configuration C1_27 (27 sec timestep) at all instances took almost double the time of configuration C1_54 (54 sec timestep) to complete the performed tasks. Linear scaling of WRF has already been confirmed for single domain simulations [19]. Hence, for the rest of this paper, only configuration C1_54 will be considered and will be referred to simply as C1.




Fig. 3 Performance of WRF for C1 configuration. CPU time, required for one hour simulation, was recorded. C1_27 refers to C1 configuration with 27 sec domain d01 timestep and C1_54 refers to C1 configuration with 54 sec domain d01 timestep. The CPU-time ratio between the two different cases is approximately 2 (it varies between 1.84 and 2.26).

III. RESULTS

In this section, all comparisons and statistical analyses have been performed with WRF output from domain d03 of configuration C1 and domain d04 of configuration C2. Time series of the simulated wind speed at 28 m agl (20-minute intervals) and wind roses for energy and direction were compared against the observed near-surface wind speed and direction for the two configurations. A typical time series plot and wind roses are shown in Fig. 4. The model appeared capable of reproducing diurnal variations of wind speed. The main cause of errors in statistics appears to be differences in the time series caused by sudden peaks or drops of winds at a 3-6 hour time-scale of either the predicted or the observed wind speeds. Such differences might be the consequence of subgrid scale effects of unresolved topography by the model [20].





Fig. 4 A typical time series plot (top) and wind roses (bottom) for C1 configuration for the station of Loutro. The 28


m wind speeds and directions, both observed (red) and predicted (blue), are presented for a 7-day period.

Fig. 5 shows the BIAS, relative BIAS, RMSE and relative RMSE of the wind speed measured at 28 m agl, for the two stations and each different configuration. For both stations, the relevant metrics of the two configurations were comparable (less than 13% relative difference), thus indicating that both configurations predict similar wind speeds. However, it is interesting to observe that configuration C1 outperforms slightly the higher resolution configuration C2. This might be attributed to the fact that configuration C2 utilizes a smaller timestep, hence performing more integration steps for the numerical solution of the differential equations involved. This leads to higher accumulation of numerical error, which is depicted in the deterioration in the statistics of C2. Another possible source of error is the fact that the two stations were positioned close to the boundaries of domain d04 of configuration C2.

In Fig. 6 the BIAS, relative BIAS, RMSE, relative RMSE and DACC of the wind direction at 28 m agl are presented for both stations and each different configuration. Similar behaviour, as in the case of wind speeds, is exhibited with C1 outperforming C2 in the prediction of the wind direction. It should be noted that for both stations the BIAS and relative BIAS are negative.




(a)




(b)

Fig. 5 BIAS, relative BIAS, RMSE and relative RMSE of the wind speed measured at 28 m agl, for each different configuration C1 and C2 at the stations of (a) Kavala and (b) Loutro.



(b)

Fig. 6 BIAS, relative BIAS, RMSE relative RMSE and DACC of the wind speed measured at 28 m agl, for each different configuration C1 and C2 at the stations of (a) Kavala and (b) Loutro.



(a)

We also compare the performance of the two different configurations and the results are presented in Fig. 7. It can be seen that despite the fact that the total number of grid points for C2 is less than for C1 (as discussed in Section II), the speedup in calculations for C1 is approximately 2.4 times greater (the CPU-time ratio varies between 2.24 and 2.6). We believe that this is attributed to the known constraint of the WRF framework that it integrates each different domain sequentially, even though there might be no interaction among the different domains (see e.g. [6]). In both configurations, the timesteps employed were the maximum ones allowed by the numerical stability criterion [9].




Fig. 7 Performance of WRF for C1 and C2 nested domain configurations. CPU time, required for one hour simulation, was recorded.

The scope is to optimize model performance by obtaining accurate results in higher resolution and minimizing the computational cost. In order to do so, the following results are utilized:

a) The ratio of the number of grid points of the two innermost domains of C2 to the number of grid points of the innermost domain of C1 is 3/4.

b) The CPU-time ratio of configuration C1 to configuration C2 is approximately 2.4

c) WRF scales linearly with time

Hence, by keeping the timesteps of configurations C1 and C2 constant (54 sec and 27 sec, respectively for the outer domain), in order to make the CPU-time ratio of C1 to C2 equal to 1, then the ratio of the number of grid points should approximately be 3/10. Different C2 configurations should be tested, where the ratio of the number of grid points of their two innermost domains to the number of grid points of domain d03 of configuration C1 is equal to 3/10. This might as well be a step towards alleviating the problem of worse statistical performance of configuration C2.

IV. CONCLUSIONS

In this paper, preliminary results from the AKAIPRO research project have been utilized in order to assess the performance of WRF numerical model in nested domains simulation on an HPC cluster. Two different configurations have been tested (three-domain C1 and four-domain C2 nested configurations) and various verification statistics were computed, in order to gauge the model performance. It was found that configuration C1 performed 2.5 times faster than configuration C2 and it gave slightly better statistics. However, since it is obvious that higher resolution depicts more accurately regional topographical characteristics, a modification is proposed for configuration C2, based on the results analysis, in order to improve run times and statistics. Further work is already under way in order to test this proposition, as well as different domain configurations.

ACKNOWLEDGMENT

G. Kossioris, Th. Katsaounis and M. Plexousakis have been partially supported by the European Union Seventh Framework Programme (FP7-REGPOT-2009-1) under the Grant Agreement no. 245749 (Archimedes Center for Modeling, Analysis and Computation, University of Crete, Greece).

REFERENCES

[1] D.J. Kerbyson, K.J. Barker and K. Davis, "Analysis of the Weather Research and Forecasting (WRF) Model on Large-Scale Systems", in

Parallel Computing: Architectures, Algorithms and Applications, C. Bischof et al., Eds., John Von Neumann Institute for Computing, Julich, NIC Series, 2007, pp. 89-98.

[2] G. Shainer et al., "Weather Research and Forecasting (WRF) Model Performance and Profiling Analysis on Advanced Multi-Core HPC Clusters", in Proceedings of 10th LCI International Conference on High Performance Clustered Computing, March 2009, NCAR, Boulder, CO.

[3] A.R. Porter and M. Ashworth, "Configuring and Optimizing the Weather Research and Forecasting Model on the Cray XT", in Proceedings of Cray User Group, 2010, 12 pp.

[4] F. Lu, J. Song, X. Cao and X. Zhu, "CPU/GPU computing for long-wave radiation physics on large GPU clusters", Computers and Geosciences, doi:10.1016/j.cageo.2011.08.007, 2011.

[5] T. Langkamp and J. Bohner, "Influence of the compiler on multi-CPU performance of WRFv3", Geosci. Model Dev., 4, pp.611-623, 2011.

[6] P. Malakar et al., "Performance Evaluation and Optimization of Nested High Resolution Weather Simulations", in Euro-Par 2012, C. Kaklamanis et al. Eds, Springer-Verlag Berlin, 2012, pp. 805-817.

[7] AKAIPRO, Study of extreme weather phenomena in local regions and their impact to civil protection. Research Program "SYNERGASIA", Greek Secretariat of Research and Technology, 2011-2013. <http://idarkstar.tem.uoc.gr/~plex/newakapro>.

[8] V. Kotroni, K. Lagouvardos and D. Lalas, "The effect of the island of Crete on the Etesian winds over the Aegean Sea" Q J R Meteorol Soc, 127, pp. 1917-1937, 2001.

[9] W.C. Skamarock et al., "A description of the Advanced Research WRF Version 3", NCAR Tech Notes-475+STR, 113pp., 2005.

[10] T.L. Otte, C.G. Nolte, M.J. Otte and J.H. Bowden, "Does nudging squelch the extremes in Regional Climate Modeling?", J. Climate, 25, pp. 7046-7066, 2012.

[11] L.M. Harris and D.R. Durran, "An idealized comparison of one-way and two-way grid nesting", Mon. Wea. Rev., 138, pp. 2174-2187, 2010.

[12] F. Chen and J. Dudhia, "Coupling an advanced land-surface/hydrology model with the Penn State/NCAR MM5 modeling system. Part I: Model description and implementation", Mon. Wea. Rev., 129, pp. 569-585, 2001.

[13] J.S. Kain J. S., "The Kain-Fritsch Convective Parameterization: An Update", Journal of Applied Meteorology, 43, pp. 170-181, 2004.

[14] J. Dudhia, "Numerical study of convection observed during the winter monsoon experiment using a mesoscale two-dimensional model", J. Atmos. Sci., 46, pp. 3077-3107, 1989.

[15] E.J. Mlawer, S. J. Taubman, P. D. Brown, M. J. Iacono and S. A. Clough, "Radiative transfer for inhomogeneous atmosphere: RRTM, a validated correlated-k model for the longwave", J. Geophys. Res., 102, pp. 16663-16682, 1997.

[16] F.J. Santos-Alamillos, D. Pozo-Vazquez, J.A. Ruiz-Arias, V. Lara-Fanego and J. Tovar-Pescador, "Analysis of WRF model wind estimate sensitivity to physics parametrization choice and terrain representation in Andalusia (Southern Spain)", J. Applied Meteorology and Climatology, 52, pp. 1592-1609, 2013.

[17] G. Thompson, P.R. Field, R. M. Rasmussen and W.D. Hall, "Explicit forecasts of winter precipitation using an improved bulk microphysics scheme. Part II: Implementation of a new snow parameterization", Mon. Wea. Rev., 136, pp. 5095-5115, 2008.

[18] N. Nakanishi and H. Niino, "An improved Mellor-Yamada Level-3 model: Its numerical stability and Application to a regional prediction of Advection fog", Boundary-layer Met., 119, pp.397-407, 2006.

[19] C. Kruse, "WRF Performance and Scaling Assessment", Computational and Information Systems Laboratory, UCAR, 2013, <https://www2.cisl.ucar.edu/siparcs/calendar/2013-08-02/christopher-kruse/wrf-performance-and-scaling-assessment>.

[20] P.A. Jiménez and J. Dudhia, "Improving the representation of resolved and unresolved topographic effects on surface wind in the WRF model", J. Applied Meteorology & Climatology, 51, pp. 300-316, 2012.