8.5 A Globally Relocatable High-Resolution WRF Realtime Forecast System for Renewable Energy

Aijun Deng*, Younghun Kim, Srivats Shukla, Wander Wadman and Ali Mohammed
Utopus Insights, INC. Valhalla, NY, USA

1. INTRODUCTION

Weather Research and Forecasting (WRF) modeling system has been used in realtime nowcast/forecasting, broadly across various applications, including mesoscale weather prediction (Deng et al. 2012a, Gaudet et al. 2012), aviation planning and optimization (Langelaan et al. 2013), as well as atmospheric transport and dispersion (Deng et al. 2012b, Lauvaux et al. 2012, Barkley et al. 2017) with WRF's chemistry module (WRF-Chem) enabled. WRF is also becoming a useful tool for renewable energy industry to forecast energy generation and demand, as well as operational planning. Utopus Insights, an independent new digital software company spun out of IBM Smart Energy Research, has recently developed a globally relocatable high-resolution WRF-based forecasting system code-named "Nostradamus" to support its time-sensitive, high-resolution atmospheric modeling applications including wind and solar energy forecasting for its clients worldwide. This paper describes Nostradamus and some of its use cases over various locations across the globe, with focus on the instance running for a utility company in Vermont (VT), USA. Overall performance of Nostradamus and the downstream energy forecasting models is presented. Issues related to the surface temperature cold bias are discussed. We also present how Nostradamus is optimized on the HPCC.

2. DESCRIPTION OF NOSTRADAMUS

Nostradamus is one of the Utopus Insights' products that includes Xplore for realtime asset monitoring, Pulse for predicting failure of assets, Weather Insights that provides interactive weather visualization for utility company, and HyperCast for predicting future energy generation, each of which has components for both wind and solar energy. The core of Nostradamus is the WRF modeling system that predicts meteorological variables including wind, wind gust and solar irradiance, etc. These meteorological variables, along with the observed energy production values, are used to drive the downstream energy analytical models, Wind HyperCast, Solar HyperCast, and Demand Forecasting that in turn feed the Peak Load Forecasting application. All the downstream applications use deep learning (neural networks) and/or machine learning techniques at their core to incorporate high-dimensional weather forecasting data for improved energy forecasting accuracy.

*Corresponding Author: Aijun Deng, Utopus Insights, Inc., Valhalla, NY 10595;
aijun.deng@utopusinsights.com

Wind HyperCast uses the Nostradamus-predicted meteorological variables including surface and hub-height winds as well as wind gusts, plus the recent measurement of wind power and wind speed at the turbine site. Further improvement of the wind energy forecasting has been achieved based on the multi-layer feed-forward perceptron, a class of deep learning algorithms. Nostradamus configuration does not yet consider the turbine drag effect on mesoscale wind solution.

Solar HyperCast uses the Nostradamus-predicted meteorological variables including direct and normal radiation components, as well as the diffused component that is an important contributor to solar power production on days of thin clouds. It also uses Nostradamus-predicted snow cover and snow fall amount that, recently added to Solar HyperCast, to improve the forecasting accuracy on the days after a snowfall, when the irradiance might be high, but the snow cover on solar panels drives lower power production.

The physics suite selection for the WRF model used in Nostradamus is similar to those used in RAP and HRRR except for the Noah land surface model. Nostradamus is currently initialized from GFS and runs up to 4 times per day depending on the application, assimilating available meteorological observations within a time window prior to the initialization time, using the multi-scale FDDA technique (Deng et al. 2012a). Nostradamus also has an automated on-the-fly validation system to provide the client with confidence on the forecast accuracy of relevant meteorological features such as wind speed, direction, temperature, etc. Built with high scalability across various computing platforms, Nostradamus can readily be deployed on local HPC clusters, stand-alone machines as well as in the cloud. Utopus Insights has been using Nostradamus successfully around the globe for its clients for various renewable energy applications.

Figure 1 shows a few examples of Nostradamus WRF configuration that varies among applications depending on the objective of each application. These applications include Project Hamburg for Wind and Solar HyperCast in Germany, Project Nolan for convection effect on wind energy production, Project ENEL for Wind HyperCast over northwestern India and Project BHARAT for Wind and Solar HyperCast over the entire India. For example, for Project Nolan, the goal was to address the client's typical concern in which insufficient wind power production was observed while the predicted power production based on the weather forecasts by the weather service. It was noted by the client that this typical discrepancy was caused by unsuccessful prediction of convective activity by the

weather service models. Our study indicated that running a WRF model at high resolution (e.g., 1-km) was helpful in reproducing convection for all the test case provided by the client. Case 1-2 April 2017 was a typical case where high wind conditions were predicted for both days based on the weather surface models, while in reality the wind speed was low on the second day as a result of widespread convection that suppressed the wind production. Nostradamus was able to reproduce the observed convection activities. Figure 2 shows Nostradamus-predicted surface wind speed and gust at Nolan, TX, for 1-2 Apr 2017, showing strong wind gusts (>20m/s) on the early morning of 1st but much weaker on the 2nd due to widespread convection.

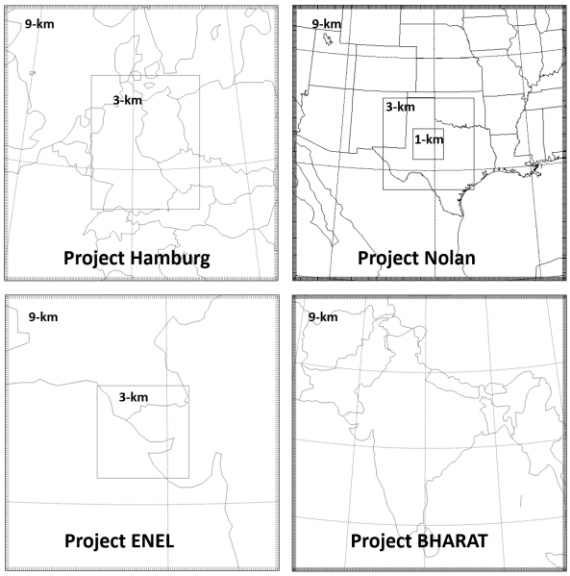

Figure 1. Nostradamus WRF grid configurations for several sample applications across the globe, upper left) Project Hamburg for Wind and Solar HyperCast applications over Germany, upper right) Project Nolan for Wind HyperCast applications in Texas, USA, lower left) Project ENEL for Wind HyperCast over northwestern India, and lower right) Project BHARAT for Wind and Solar HyperCast applications over the entire India.

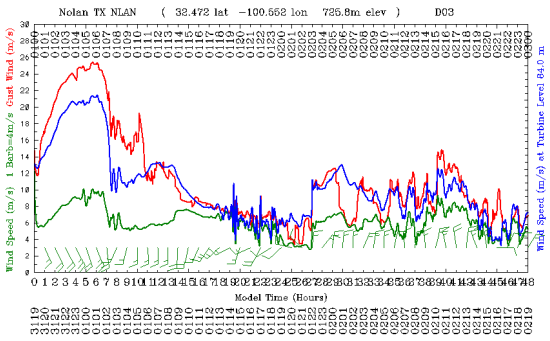

Figure 2. Nostradamus-predicted surface wind speed and gust at Nolan, TX, for 1-2 Apr 2017, showing strong wind gusts (>20m/s) on the early morning of 1st but much weaker on the 2nd due to widespread convection (incorrectly predicted by weather services based on client).

3. NOSTRADAMUS PERFORMANCE FOR VERMONT

3.1 Nostradamus Weather

The Vermont Nostradamus instance consists of three nested grids with 9-, 3-, and 1-km grid spacings, with identical number of grid points (i.e., 300x300x50) for each grid (Fig. 3). The 9-km grid covers a large portion of northeastern U.S, including IN, KY, TN, part of AL and GA, and southeastern Canada. It also covers a large area of Atlantic Ocean so that realistic land-sea thermal gradient is represented in the model. The 3-km grid covers the entire New England states, NY, eastern half of PA, and northern NJ, as well as s portion of Atlantic Ocean.

The 1-km grid covers the entire VT, NH and portion of NY, with the land surface features of Hudson River valley that extends south-north to connect Lake Champlain, the Green Mountain Range that runs primarily south to north and extends approximately 250 miles from the border with Massachusetts to the border with Quebec, Canada. Between the Presidential Range that peaks at Mt. Washington with elevation of about 2 km and the Green Mountain Range, there is Connecticut River valley that runs approximately north-south along the VT and NH border. The 1-km grid also covers portion of Adirondack Mountains in NY with its peak elevation about 1.6 km. These landuse and topographic feature defines the unique climatological features of VT. For example, the south-north-run valleys provide ideal conditions for channeling effect when the dominant background wind is from south or north (Fig. 4). A high wind threat scenario with southerly or northerly flows can deteriorate due to the channeling effect.

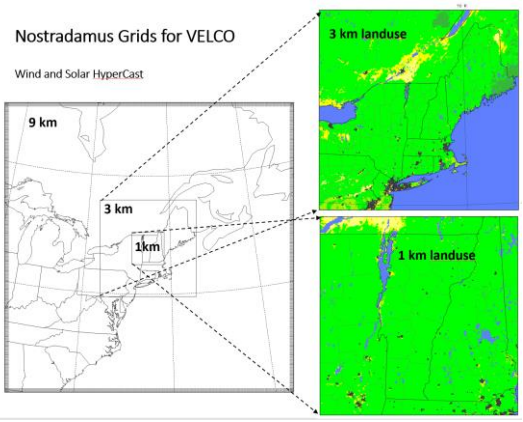

Figure 3. Vermont Nostradamus grid configuration, with 300x300x50 grid points for each of grid. Landuse for the 3- and 1-km grids are also shown.

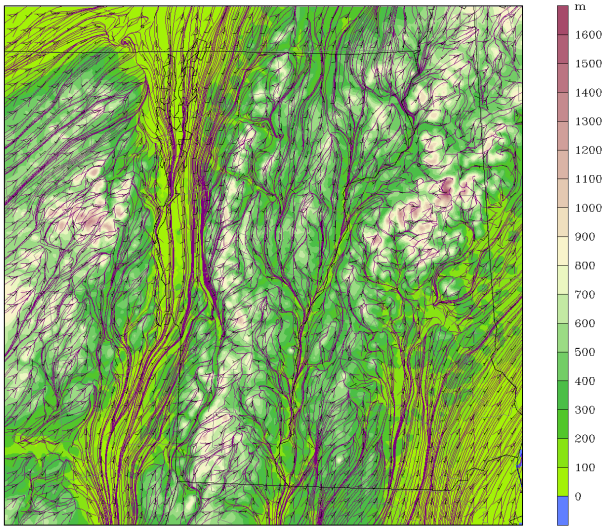

Figure 4. Typical surface southerly flow patterns represented by streamlines produced by WRF on the Vermont 1-km grid. The higher the wind speed, the tighter the streamlines become. Each full wind barb represents 5 m/s wind speed. The color bar denotes the ground elevation.

Utility companies such as those in Vermont are very interested in accurate prediction of special weather events such as damaging wind gust, lightning and icing, in addition to the traditional weather features such as wind speed, wind direction, temperature, and precipitation and etc. that are regularly needed to drive the energy forecasting models. Lightning and icing requires further study and have not been implemented into Nostradamus. For wind gust, a simple algorithm was implemented based on the strength of turbulence (e.g., atmospheric boundary layer depth) and wind shear. Preliminary validation showed that the original implementation intended to overestimate the surface wind gust, thus a scaling factor was introduced which was later tuned to produce reasonable wind gust forecast based on observations.

The case of Oct 29, 2017, was one of many cases with high surface wind weather conditions resulting in power outages. The extremely high wind gust conditions developed ahead of a cold front associated with rapidly intensifying cyclone caused numerous amount of power outage, with many homes having no power for weeks. Figure 5 and 6 shows the Nostradamus-predicted vs observed wind conditions for Rutland and Burlington. For Rutland the model-predicted surface wind gust of 51 mph at 2:00AM EDT, 30 October agreed very well with the observed value of 49 mph at 1:56AM EDT. The model-predicted gust value of 60 mph at Burlington at 3:10 AM EDT 30 October also agreed well with the observed value of 59 mph at 3:54 AM although the model peak is 44 minutes too earlier. The model-predicted daily precipitation amounts and patterns (Fig. 7) also agree well with the observed. These results indicate that Nostradamus has good skill in predicting wind gust and precipitation for use for utilities.

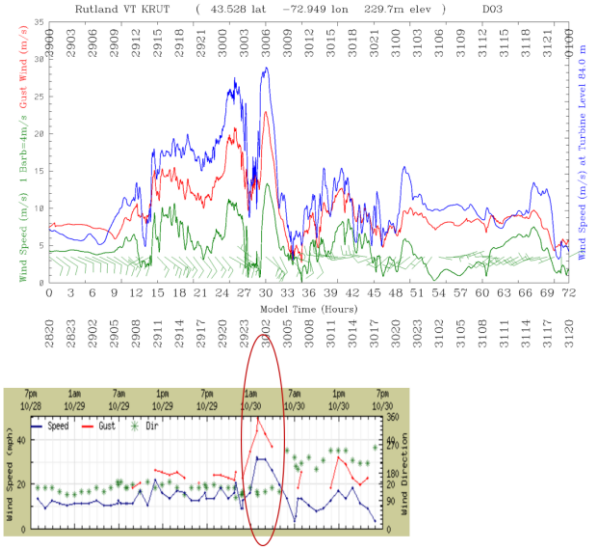

Figure 5. Nostradamus-predicted (upper) and observed (lower) conditions for Rutland, VT. In the upper panel, blue, red and green curves represent hub-height wind speed, surface gust and surface wind speed, respectively. In the lower panel, red, blue and green represent the observed surface wind gust, wind speed and direction.

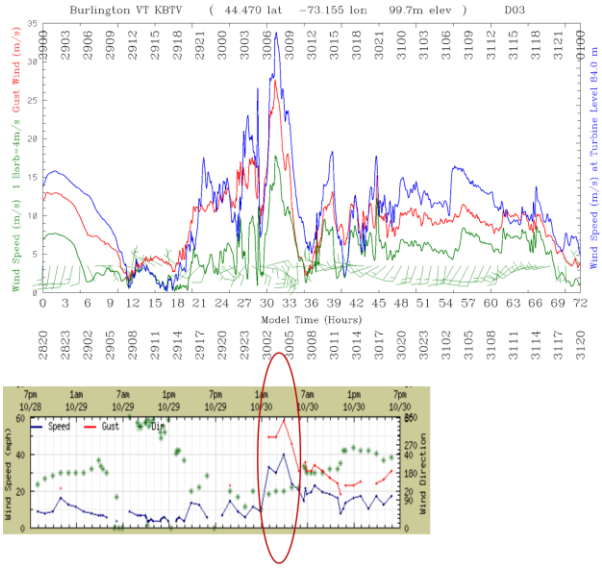

Figure 6. Nostradamus-predicted (upper) and observed (lower) conditions for Burlington, VT. In the upper panel, blue, red and green curves represent hub-height wind speed, surface gust and surface wind speed, respectively. In the lower panel, red, blue and green represent the observed surface wind gust, wind speed and direction.

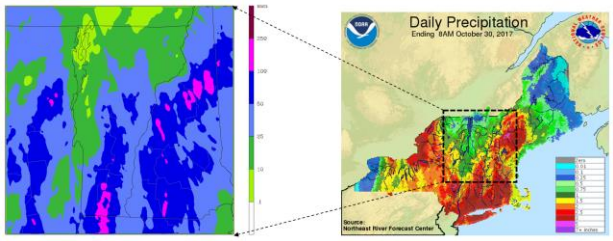

Figure 7. Nostradamus-predicted (left) vs observed (right) 24-h precipitation ending at 8 AM EDT 30 Oct. 2017 on the 1-km grid.

As indicated earilier, Nostradamus outputs such wind speed, wind direction, and temperature, etc. are used as inputs to the downstream energy forecasting models involving machine learninf techniques. It is important for Nostradamus to have good skill in predicting the relavant weather variables. Figure 8 and 9 show typical error scores during a given 3-day Nostrdamus forecast window for the 1-km grid, extracted from the Nostrdamus validation on-the-fly sytem. The validation is currently done using only the 30+ standard METAR surface weather stations located within the 1-km grid. It is shown that mean error (ME), mean absolute error (MAE), and root mean square error (RMSE) for both wind speed and and wind direction (Fig. 8 and 9) are in a reasonable range (e.g., speed MAE ~1 m/s and direction MAE < 30 degrees) given the relatively complex terrain conditions over the region.

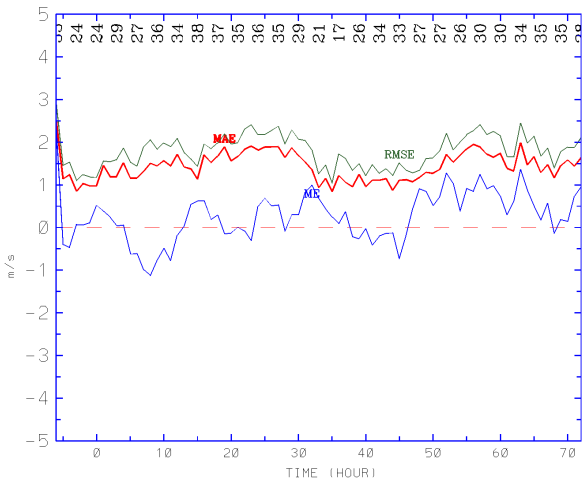

Figure 8. Time series of Nostradamus-predicted surface wind speed errors on the 1-km grid, showing MAE, ME and RMSE, validated against standard METAR observations.

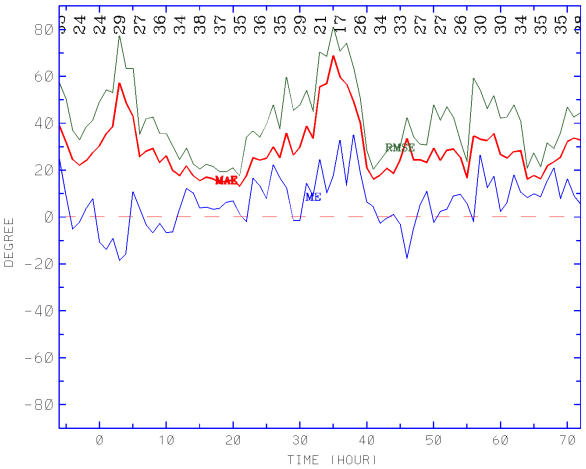

Figure 9. Time series of Nostradamus-predicted surface wind direction errors on the 1-km grid, showing MAE, ME and RMSE, validated against standard METAR observations.

For temperature, Figure 10 show a typical model error for the diagnosed 2-m temperature (T2), and Fig. 11 shows the same but for the temperature predicted at the first model layer (T1) that is located at ~10 m above the ground (AGL) in this case. Notice that there is a systematic cold bias (1-2 K) in both temperatures, with smaller error in T1 than in T2. We noticed that the cold bias is more severe (i.e., 2-3 K) during nighttime into early morning when the atmospheric boundary is stable. We also noticed that T2 occasionally shows unreasonably extreme values on the 1-km grid (not shown) associated with the bad ground temperature values calculated in the Noah land surface scheme (LSM), indicating potential issues with handling the ground temperature in the LSM at high resolutions. Because of this issue and the fact that T1 is in more agreement with the surface temperature observation than T2, we decided to use T1 in place of T2 for the downstream applications as a temporary solution. Perhaps an even better approach is to use the potential temperature and surface pressure to calculate the temperature valid at the surface and use this value in the downstream models. We have implemented this approach into Weather Insights, one of the Utopus Insights' products. An ultimate solution is to improve the Noah LSM for the stable regime and for high-resolution applications.

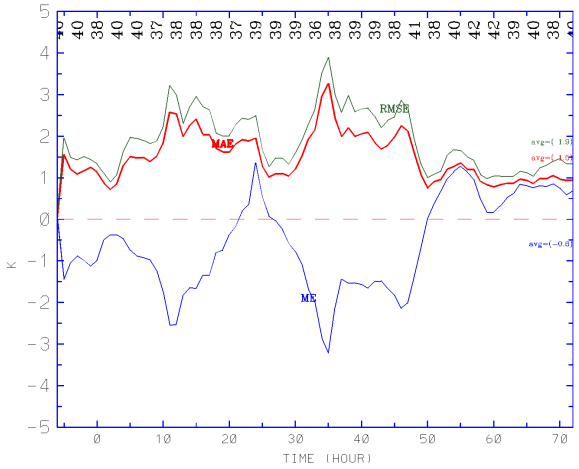

Figure 10. Time series of Nostradamus-predicted 2-m temperature errors on the 1-km grid, showing MAE, ME and RMSE, validated against standard METAR observations.

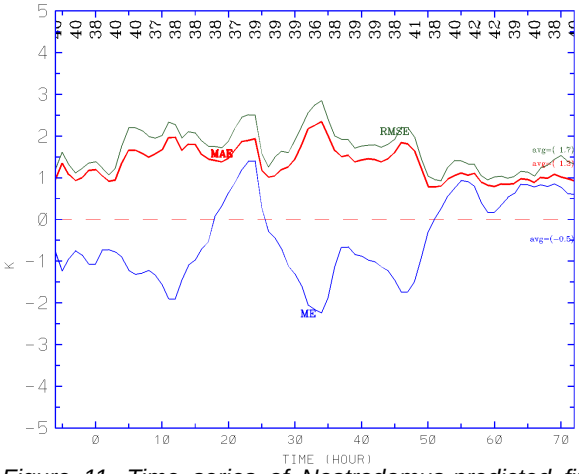

Figure 11. Time series of Nostradamus-predicted first model layer temperature errors on the 1-km grid, showing MAE, ME and RMSE, validated against standard METAR observations.

3.2 Renewable Energy

Wind HyperCast forecast wind power for each wind turbine separately or the entire wind farm depending on the data source and availability, for up to 10 days into the future, at up to 5-minute time resolution. The inputs to Wind HyperCast include meteorological variables predicted by Nostradamus, recent power measurement and recent wind speed measurement at the turbine site at the hub-height level. The output is the power forecast. As shown in Figure 12, the core of Wind HyperCast is a software layer that contains models involving deep learning (neural networks) techniques. Figure 13 shows Wind HyperCast performance for a Vermont wind farm with capacity of 40 MW. It is shown that errors of Wind HyperCast is as low as 5-7% when Nostradamus weather forecasts are used.

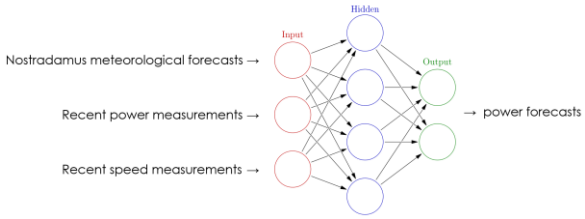

Figure 12. A schematic diagram for Wind HyperCast.

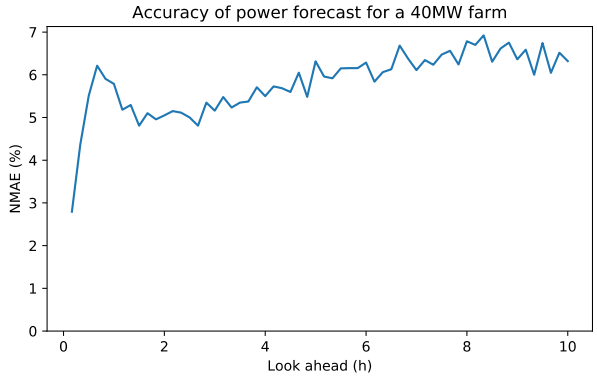

Figure 13. Normalized MAE of Wind HyperCast.

Solar HyperCast is like Wind HyperCast. Its inputs to Solar HyperCast include meteorological variables predicted by Nostradamus and recent solar power measurement, and the output is the solar power forecast. The core of Solar HyperCast is a software layer that contains models involving machine learning. A recent upgrade to Solar HyperCast was to include the snow forecast. Figure 14 shows the Solar HyperCast-forecasted power vs measurement for one of the solar farms in VT. The prediction was from the version of Solar HyperCast that did not consider snow covering the solar panels, thus overestimated the power production for the period of 6-8 January 2018 after a snow storm.

By training the models to include snowfall rate and accumulative snow amount, HyperCast-predicted solar power is much closer to the measurement. Figure 15 shows the comparison among HyperCast with snow, without snow and measurement for the entire month of January 2018. It is shown that improvement due to including snow features are evident after 10 January 2018 when the snow features were implemented. The overall average error in forecasted power production has reduced by ~ 0.5%-1% (~15% for snow days) for most of the farms. Although the forecasts on the day of snow have shown significant improvements, there was a degradation (under prediction) for clear-sky days, indicating a need for further study, perhaps by refining the classification during model training.

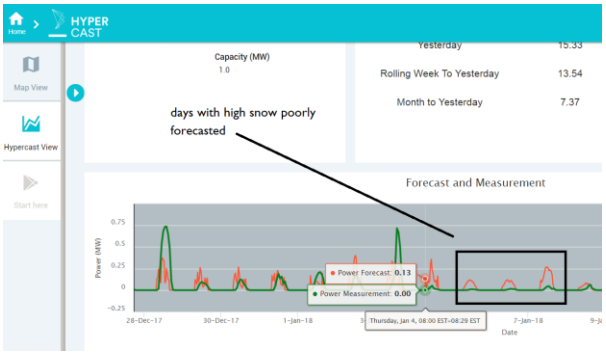

Figure 14. Predicted vs. observed power curves from Solar HyperCast for one of the solar farms in VT, between 28 December 2017 and 9 January 2018. Green curve is the measurement and read is the model-predicted.

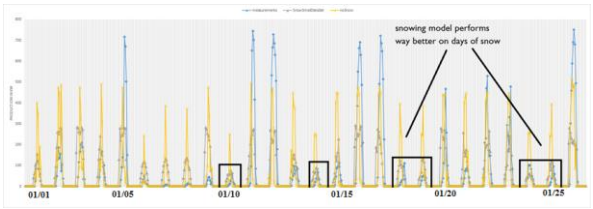

Figure 15. Predicted vs. observed power curves from Solar HyperCast for one of the solar farms in VT, between 1 2017 and 26 January 2018. Blue curve: measurement, yellow: HyperCast without snow, grey: HyperCast with snow

4. NOSTRADAMUS OPTIMIZATION

Given the Nostradamus configuration for Vermont (i.e. 300x300x50 grid points for each of the three grids), a series of timing tests were conducted on the HPCC that has a large number of computing cores and 100GB/s InfiniBand to find the 'sweet spot' or number of CPU cores where optimal run speed is achieved so that our utility client receives the forecast products at the earliest possible time. Figure 16 demonstrates the relationship between the time taken to run the Vermont WRF configuration for 24 hours and the number of CPU cores assigned to the task, for tests without using I/O quilt technique (blue curve) and those with using optimal I/O quilt (orange curve). It is shown that the 'sweet spot' is found at 17-18 nodes (or 476-504 CPU cores). It is also shown that near the 'sweet spot', with using I/O quilt the run time can be further reduced by ~20% (i.e., 56 min. without I/O quilt vs 45 min. with I/O quit). The 11 min./day reduction can be translated to ~2 hour/10-day which is quite significant for the client that needs a 10-day forecast.

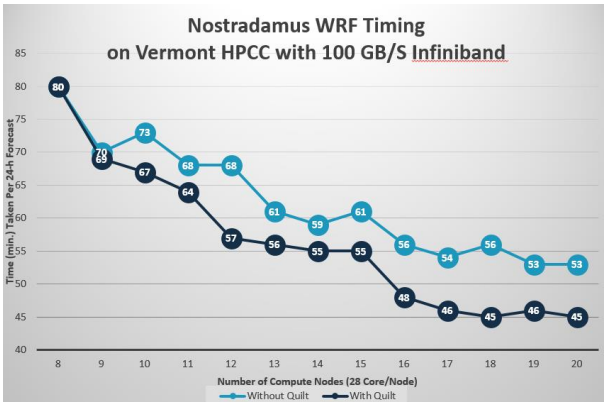

Figure 16. Nostradamus WRF timing as a function of the number of assigned CPU cores.

5. SUMMARY AND CONCLUSIONS

Utopus Insights has developed a globally relocatable high-resolution WRF-based numerical weather prediction system – Nostradamus, for its realtime renewable engery hypercast applications worldwide. Based on our systematic validation for meteorology for a period of more than 8 months, the current Nostradamus configuration predicts wind speed and direction with reasonable overall model mean absolute errors (typically < 1 m s-1 for wind speed and <30 degrees for wind direction). However, the surface air temperature prediction shows an overall cold bias (~1-2 K). In addition, compared with the temperature at the first model level, the diagnostic 2-m temperature field tends to have larger error (~2-3 K). Validation based on renewable energy production indicated improved Solar HyperCast forecasting with the added snow feature (e.g., 15% error reduction on snowy days for a group of solar farms in the North-Eastern US). Similarly, errors of Wind HyperCasts decreased from 14-16% to as low as 5-7% by also exploiting Nostradamus forecasts of temperature, air density and wind gusts in addition to those of wind speed.

On a HPCC with a large number of computing cores and 100GB/s InfiniBand, Nostradamus timing scales approximately linearly with number of compute nodes up to 18 nodes (or 504 cores), 'sweet spot' used for the operational Nostradamus. We also found that near the 'sweet spot', with using I/O quilt the run time can be further reduced by ~20%.

Future work includes 1) improving Solar HyperCast for clear sky days, 2) exploring turbine drag effect on mesoscale wind prediction, 3) assimilating METMASS from wind farms into Nostradamus, 4) exploring the readiness of MPAS for use in hypercast applications.

6. REFERENCES

Barkley, Z. R., Lauvaux, T., Davis, K. J., Deng, A., Miles, N. L., Richardson, S. J., Cao, Y., Sweeney, C., Karion, A., Smith, M., Kort, E. A., Schwietzke, S., Murphy, T., Cervone, G., Martins, D., and Maasakkers, J. D.: Quantifying methane emissions from natural gas production in north-eastern Pennsylvania, Atmos. Chem. Phys., 17, 13941-13966, https://doi.org/10.5194/acp-17-13941-2017, 2017.

Deng, A., D.R. Stauffer, B.J. Gaudet, and G.K., Hunter, 2012a: A Rapidly Relocatable High-Resolution WRF System for Military-Defense, Aviation and Wind Energy. 13th Annual WRF Users' Workshop, Boulder, CO, June 25-29, 11pp.

Deng A., T. Lauvaux, K. Davis, N. Miles, S. Richardson and D. Stauffer, 2012b: A WRF-Chem realtime modeling system for monitoring CO2 emissions. 13th Annual WRF Users' Workshop, Boulder, CO, June 25-29, 6pp

Deng, A., D.R. Stauffer, B.J. Gaudet, J. Dudhia, J. Hacker, C. Bruyere, W. Wu, F. Vandenberghe, Y. Liu and A. Bourgeois, 2009: Update on WRF-ARW end-to-end multi-scale FDDA system, 10th Annual WRF Users' Workshop, Boulder, CO, June 23, 14 pp (Available on the web at: http://www.mmm.ucar.edu/wrf/users/workshops/WS2009/abstracts/1-09.pdf).

Gaudet, B.J., A. Deng, D.R. Stauffer and N.L. Seaman, 2012: Eddy seeding for improved WRF-LES simulations using realistic lateral boundary conditions, Preprints, 13th WRF Users' Workshop, Boulder, CO, 25-29 Jun., 7 pp.

Lauvaux T., N. L. Miles, S. J. Richardson, A. Deng, D. Stauufer, K. J. Davis, G. Jacobson, C. Rella, G.-P. Calonder, and P. L. DeCola, 2013a: Urban emissions of CO2 from Davos, Switzerland: the first real-time monitoring system using an atmospheric inversion technique. J. Appl Meteor. and Climatol., 52, 2654-2668.

Langelaan, J. W., A. Chakrabarty, A. Deng, K. Miles, V. Plevnik, J. Tomazic, T. Tomazic, and G. Veble, 2013: The Green Flight Challenge: Aircraft Design and Flight Planning for Extreme Fuel Efficiency. J. of Aircraft, 50, No. 3, 832-846.