A Mass-Conserving Wind Field Model Run Directly from Tower Observations

R. J. Conzemius[1], L. M. Sheridan[1]
[1]WindLogics, Inc. Grand Rapids, Minnesota
Corresponding author address: bobc@windlogics.com

Introduction

For new wind farm projects over the past several years, two different methods of assessing the wind resource have become rather popular. Interestingly, these two methods require vastly different amounts of computing power. On the side of ease and simplicity, linearized models such as WAsP allow wind farm developers to make a fast and reasonably accurate estimate of wind power climatology at a site. On the other side, methods that fully involve numerical weather prediction models can provide more spatial and temporal detail about the wind flow at a site, but they require greater care and take more time to run.

For WAsP, the user interface allows a relatively novice user to navigate the wind resource assessment process. The program is also relatively simple in its input data requirements; only one set of anemometer/wind vane measurements is needed. Because the wind resource grid computations can be completed in a matter of seconds, it allows the user to quickly estimate the wind speed climatology and its variation across a particular site, making it the most popular wind resource assessment package. However, WAsP does not provide information such as spatial variation of wind direction, diurnal and seasonal power variation (although these can be calculated from tower data), and it generally uses only one tower as input, making it more difficult to apply at large development sites, where multiple towers are present. The tower data are pre-processed into climatology prior to running WAsP and are not input hour-by-hour, which removes the temporal detail.

On the other hand, modeling techniques that utilize the full set of Navier-Stokes equations can provide a more accurate and detailed depiction of flow such as variation of wind direction, diurnal and seasonal variation in wind speed, and more sophisticated treatment of flow over complex terrain. These techniques occasionally have biases caused by inaccurate or incomplete description of physical processes, such as land-atmosphere energy exchange, or boundary layer turbulence. The input data usually comes in the form of 3-D gridded data from sources such as NCEP/NCAR Reanalysis data or Rapid Update Cycle (RUC) model analyses. Tower data might not be used directly as input, but they can be used near the end of the resource assessment process to correct any biases in the model output. The modeling proceeds serially in time, retaining temporal detail in the output product. One disadvantage of this method may be the long run times required to complete the modeling—often on the order of weeks. In general, there is a large disparity in computing time between these two methods.

We propose here a third method, which, like WAsP, uses tower data as input to the numerical computation but also performs physical calculations to compute the 3-D flow. The new method allows the direct use of hourly data from multiple towers in the flow field computation, removing both the climatological pre-processing requirement and single tower limitation of WAsP. The method also has the advantage of providing the temporal and spatial detail characteristic of the more complex methods but the required run times are significantly shorter because mesoscale numerical weather prediction models are not used. We present here a basic overview of the methods involved and the results of testing at a dozen wind energy development locations.

General Methods

The working engine of this method is the nSwift model (Geai 1987), developed by ARIA Technologies. The process works by taking input from multiple towers and interpolating the observations onto a model grid. The interpolated and gridded data are then adjusted using a mass conservation constraint to provide a physically consistent description of the flow. The resulting gridded flow field is corrected back to the tower observations if the mass adjustment introduces any bias in wind speeds. The tower data are input on an hourly basis, and a gridded, three-dimensional flow field is output at each time step of tower data. The 3-D model output can then be averaged in time to provide wind climatology at the desired hub height. For purposes of discussion, we will designate this method as Towers to nSwift (TTS).

The first step is to gather hourly tower data into an input meteorological file. Since there no climatological filtering is applied, all available towers are used simultaneously. Likewise, if a given tower has multiple levels of observation, all these levels are used. To account for stability effects, temperature data are taken into account as well, provided they are available. Because the flow field computation extends above the uppermost tower levels, the tower observations are overlain with gridded RUC wind vector and temperature data, horizontally interpolated to the tower locations. The result is a column of observations that includes all tower levels and extends through the top of the model domain. These column observations serve as the input data to the process.

The nSwift model is then run over the entire dataset of hourly observations. Typically, a full year (8760 hours) of observations is used. To speed the model computations, the runs are spread, by month, over 12 processors. At the end of the runs, the monthly files are compiled, and the wind climatology is created by averaging in time at the chosen hub height level.

The nSwift model (as well as its predecessor, Minerve) is also heavily used in WindLogics' resource assessments. Typically, WindLogics performs a multiply nested run of the MM5 numerical weather prediction model (Grell et al. 1994) over a 12 month period. The gridded hourly output from the innermost MM5 grid is then used is input to the nSwift model, which then computes the finer resolution flow over the site. The TTS foregoes the MM5 run, which saves more than half the computing time.

Comparisons with Other Resource Assessment Methods

We tested this new method at twelve wind development sites that were instrumented with multiple towers. At each of these sites, we also conducted the more traditional WindLogics wind resource assessment technique. The traditional WindLogics technique is to use RUC analyses as input to MM5, the MM5 model is run for a year, and the MM5 output is utilized as input to the Minerve model (predecessor to nSwift). The Minerve output is then averaged, and monthly average wind speeds are adjusted back to tower-observed monthly average winds speeds, for all available tower locations, to provide a 2-D resource grid at hub height. Because the primary stages are (1) MM5 run, (2) Minerve run, and (3) adjustment, we abbreviate this process as MMA.

We also ran WAsP at each of the sites. Where available, we used roughness data in addition to terrain elevation data, and the same model grid was used for both WAsP, MMA, and TTS. We used the same period of tower data for all three methods.

The analysis was performed at the height of the tallest available tower anemometer data. No extrapolation to typical hub heights was performed as we did not want to introduce any shear-related uncertainty into the comparisons. In this way, the output annual average speeds can be compared directly with the tower data. In order to make these true comparisons between model and observations, it was necessary to withhold data from at least one tower from the modeling process, then use the data from the withheld tower to calculate errors (modeled minus observation) in the modeling process. All available towers were used as input, except for WAsP, where only one tower could be used (we use the highest level for each tower), and all other towers were used to test the model. To provide a larger sample of error estimates, we repeated the testing procedure, swapping the input and test towers, so that each tower was used as both an input and a testing tower. For both processes, the number of rounds at a particular site was equal to the number of towers.

The test sites were chosen to represent a wide variety of terrain and wind resource conditions. Nine of these sites were located over the central part of North America (from the Midwest to the southern Plains of the United States), two in eastern North America, and one in the Pacific Northwest. For each site, we chose one year of contiguous meteorological tower data and only used towers that were collecting during the selected measurement year. Because we wanted to limit our tests specifically to the wind field modeling part of the resource assessment process, no post-processing or long-term normalization was applied. The tests were made by comparing tower wind speed observations directly against model-predicted wind speeds only at the highest levels where those speeds were measured. To ensure equivalent sample sizes for annual average computations, we only used hours when wind speeds from both test tower and model output were available. We averaged these speeds to compute the annual average wind speed and reported errors as root mean square error (RMSE) in annual average wind speed.

Results

Table 1 presents the results of comparing the performance of WAsP, MMA, and TTS. The results of the testing reveal that, in general, the accuracy of the particular resource assessment method chosen is proportional to the complexity and amount of computer time required. The lowest average root mean square error (RMSE) among the methods was achieved by MMA. On a site-specific comparison, MMA provides the lowest RMSE for 7 out of the 12 cases. Following MMA in both complexity and average RMSE, the TTS methodology produces the lowest RMSE for 4 out of the 12 test sites. Least complex and time-consuming, WAsP provides the lowest RMSE for only one of the 12 locations. In a head-to-head comparison with WAsP, TTS has the smaller RMSE in 9 out of the 12 cases, and MMA also has a smaller RMSE in 9 of the 12 cases. Figure 1 shows a graphical depiction of the results.

We have found that the accuracy of the MMA process comes about largely from the adjustment of the model-output monthly average wind speeds back to the tower observation values. The model output fields, on their own, occasionally have a bias that would result in larger errors, but the adjustment reduces those errors significantly. It appears that the nonlinear modeling process provides a more physically consistent depiction of the flow over topography, even when there is a wind speed bias present. The adjustment essentially eliminates the bias while still maintaining some physical consistency in the flow.

The TTS process does not provide as small an RMSE as MMA does. We hypothesize the larger RMSE is due to several factors. First, the full set of nonlinear equations used in MM5 provides a slightly more physically consistent depiction of the flow for input to Minerve during the MMA process. For tower inputs alone, there are no such physical constraints applied. Second, the MM5 input to the flow field model (Minerve or nSwift) is gridded, allowing for greater representativeness and spatial coverage of the input data, whereas the tower observations input to nSwift are not as regularly spaced and thereby under-sample portions of the flow. The third reason is related to the second in the fact that tower locations at wind development sites are usually biased toward high terrain, where the wind speeds are stronger. The interpolation of these stronger wind speeds over the remaining area, where the terrain is generally lower, results in an area-wide overestimate of the flow. This error has been reduced by adding algorithms in nSwift to correct interpolation of the 3-D wind fields according to the topography and to correct mass-adjusted fields back to tower observations. Nevertheless, we suspect the RMSE would improve if towers were more evenly spread in high, medium, and low terrain locations.

Table 1. RMSE for Resource Assessment Methods at the Twelve Test Sites

Site	WAsP	MMA	TTS
1	0.29	**0.10**	0.18
2	0.10	**0.03**	0.07
3	0.36	**0.15**	0.30
4	0.67	**0.41**	0.54
5	**0.04**	0.06	0.26
6	0.17	**0.04**	0.45
7	0.11	0.28	**0.04**
8	0.90	**0.14**	0.22
9	0.77	0.68	**0.44**
10	0.06	0.09	**0.05**
11	0.36	**0.28**	0.52
12	0.54	0.41	**0.23**
Average	**0.36**	**0.22**	**0.28**

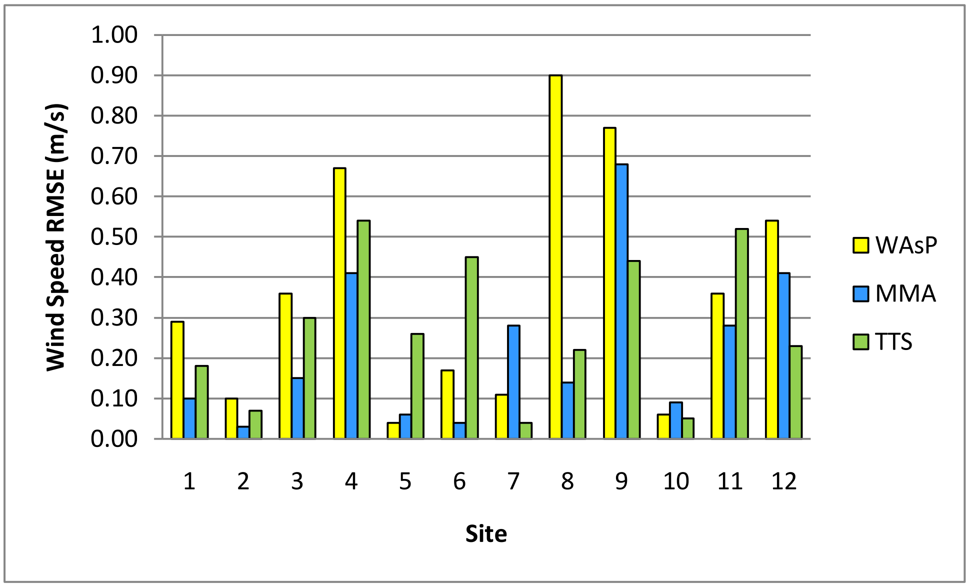

Figure 1. Graphical depiction of RMSE at 12 sites for each of the compared methods.

Overall, the MMA methodology provides an almost 40% smaller RMSE when compared with WAsP. Moderate in both complexity and time consumption, the average RMSE produced by the TTS methodology falls evenly between those of WAsP and MMA, being approximately 22% lower than WAsP and almost 22% higher than MMA.

References

1. Grell, G. A., J. Dudhia, and D. R. Stauffer, 1994: A description of the fifth-generation Penn State/NCAR mesoscale model (MM5). *NCAR Technical Note*, NCAR/TN-398+STR, 117 pp.
2. Geai. P. 1987. *Méthode d'interpolation et de reconstitution tridimensionnelle d'un champ de vent : le code d'analyse objective SWIFT*. Paris: EDF-DER (Ref. HE 34-87.03).