Wind Power Forecasting Error Distributions

An International Comparison

Bri-Mathias Hodge, Debra Lew, Michael Milligan
National Renewable Energy Laboratory
Golden, CO, USA

Hannele Holttinen, Samuli Sillanpää
VTT
Espoo, Finland

Emilio Gómez-Lázaro
University of Castilla-La Mancha
Albacete, Spain

Richard Scharff, Lennart Söder
KTH Royal Institute of Technology
Stockholm, Sweden

Xiaoli Guo Larsén, Gregor Giebel
Technical University of Denmark
Roskilde, Denmark

Damian Flynn
University College Dublin
Dublin, Ireland

Jan Dobschinski
Fraunhofer Institute for Wind Energy and Energy System Technology
Kassel, Germany

***Abstract*—Wind power forecasting is essential for greater penetration of wind power into electricity systems. Because no wind forecasting system is perfect, a thorough understanding of the errors that may occur is a critical factor for system operation functions, such as the setting of operating reserve levels. This paper provides an international comparison of the distribution of wind power forecasting errors from operational systems, based on real forecast data. The paper concludes with an assessment of similarities and differences between the errors observed in different locations.**

***Keywords*—wind power forecasting, power system operation, power system reliability, power systems, wind power generation**

I. INTRODUCTION

The amount of wind power being incorporated into power systems worldwide has been increasing dramatically during the past decade. Wind power has no fuel costs and zero emissions, which means that its increased presence in power generation portfolios provides great benefits to society. However, wind power is a variable and uncertain power resource, in contrast to traditional thermal power units. This has led to concerns from utilities and system operators about how increasing amounts of wind power will be handled in system operations [1]. One way to reduce the uncertainty surrounding wind power output is through wind power forecasting systems. Typical systems used in operational forecasting consist of one or more Numerical Weather Prediction (NWP) models that provide forecasts of wind speed on a grid over a geographic area, coupled with statistical techniques that translate the forecasts to local wind plant conditions and convert forecasted wind speed to power [2]. Although these forecasts provide system operators with an expected wind power output level at future times, they are not perfect forecasts. Understanding the magnitude and frequency of wind power forecasting errors can facilitate the integration of wind power through advanced operational techniques, for example, setting dynamic reserve levels [3, 4], using stochastic unit commitment models [5, 6], or through simply increasing operator awareness. Power system operations are already designed to handle a certain degree of variability and uncertainty because load is itself both variable and uncertain [7]. Therefore, we are most concerned with the large and infrequent wind power forecasting errors. Large forecasting events can lead to major economic inefficiencies through non-optimal commitment schedules.

Wind power forecast errors are often a concern in wind integration studies and stochastic unit commitment models. Many of these studies assume that the forecast error distribution follows a normal distribution [3, 8, 9]. However, this is an overly simplistic assumption for most forecasting methods and timescales examined [10, 11]. Other distributions have been examined, including the Weibull [12] and beta [13] distributions; however, in this work we utilized the hyperbolic distribution [10]. We analyzed the forecast error distributions observed in a number of different countries and electrical systems and at two different timescales that are important in the unit commitment and economic dispatch process. Comparisons were made between the different cases, and conclusions on the importance of the differences for power systems operations with higher wind power penetrations were drawn.

II. METHODS AND DATA

In this study, statistical analysis techniques were applied to wind power forecasting data taken from seven countries. Day-ahead wind power forecasts were supplied for seven countries or balancing areas within a country. Hour-ahead wind power forecasts were utilized for two countrywide systems and for one large wind plant in a balancing area. The hyperbolic distribution parameters were fitted to the data using a maximum likelihood method.

Mean, μ, and variance, σ, the first two standardized moments, are frequently used in the characterization of wind forecasting error distributions, providing important information about the distribution. However, considering the third and fourth statistical moments can provide additional information [10, 14, 15]. Skewness, γ, is the third moment and is a measure of the probability distribution's asymmetry. Kurtosis, κ, is the fourth moment and describes the magnitude of the distribution's peak. Conversely, kurtosis can also be thought of as a measure of the thickness of the tails of the distribution. A distribution with a high kurtosis value is leptokurtic; one with a low kurtosis value is platykurtic. Leptokurtic distributions have more pronounced peaks, slimmer shoulders, and longer tails than normal distributions with identical first two moments. In what follows, we refer to excess kurtosis, the kurtosis above that of the normal distribution, simply as *kurtosis*.

We utilized some standard statistical tools such as histograms, quantile-quantile (Q-Q) plots, and cumulative distribution function plots. It is important to note that the forecast errors were normalized, based on the wind power capacity, for the sake of comparison. Therefore, all of the forecast errors lie on the interval from -1 to 1. The Q-Q plots shown here are normal Q-Q plots that compare the observed distribution to a Gaussian distribution with the same mean and standard deviation as the observed distribution. They include a line that runs through the first and third quantiles of the observed distributions. If the two distributions are identical, the line should pass through all of the points in the observed distribution. The cumulative distribution plots show how likely a random error from the distribution will be less than or equal to the magnitude selected.

III. SYSTEM OPERATIONS WITH FORECASTS

Wind power forecasting plays an important role in reducing the uncertainty of wind generation. Forecasts may be included directly in the unit commitment and economic dispatch scheduling process used to ensure enough generation is available to meet forecast load, or they may simply provide situational awareness for the balancing authority. Day-ahead forecasts are often required for the unit commitment process because the starting of large thermal units can often take 24 hours or more. The forecasted wind power output at this timeframe can be used to optimize the availability of other generation units during the course of the following day. The economic dispatch process sets the final power output for units that are online and is performed closer to the time of realization, often one hour ahead. Variability and forecast errors at smaller timescales are often compensated with reserves held for that purpose. Because wind forecasts can be helpful to system operations in both the unit commitment and economic dispatch phases, we examined the wind power forecasting errors that occur at these two timeframes, in this paper represented by day-ahead and hour-ahead forecasts.

IV. ERROR DISTRIBUTIONS FROM OPERATIONAL SYSTEMS

In this section, we examine wind power forecast error distributions observed in a total of seven countries at the day-ahead and hour-ahead timescales. In this work, we follow the convention that the error is equal to the forecast minus the realized value.

A. Day-Ahead Forecasts

1) United States

Day-ahead forecasts for the United States were taken from the Electric Reliability Council of Texas (ERCOT) interconnection for the year 2010, with an installed wind capacity of approximately 9,000 MW. As shown in Figure 1, the distribution was leptokurtic, with a significant negative skew. The distribution also had a fairly large spread, with minimum and maximum error values above half of the installed capacity. The red line represents a normal distribution with the same mean and standard deviation as the observed errors. Figure 2 shows that the distribution was poorly represented by the normal distribution. The observed error distribution had a more pronounced peak, slimmer shoulders, and fatter tails than the corresponding normal distribution. This is an example of the differences between the observed error and normal distributions, and the other data sets show similar features.

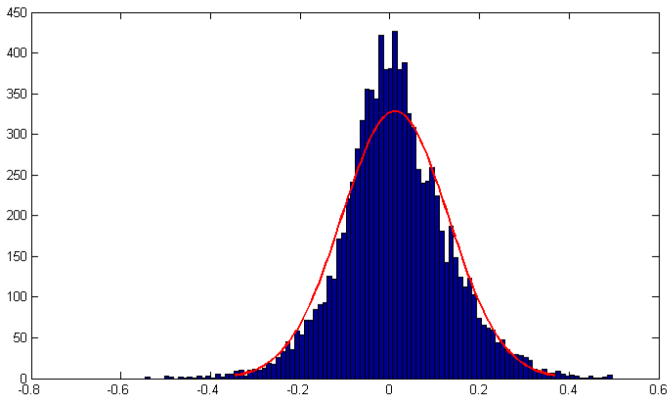

Figure 1. Histogram of the normalized day-ahead forecast errors for the ERCOT system. μ = 0.0117; σ = 0.1187; γ = -0.616; κ = 1.0308

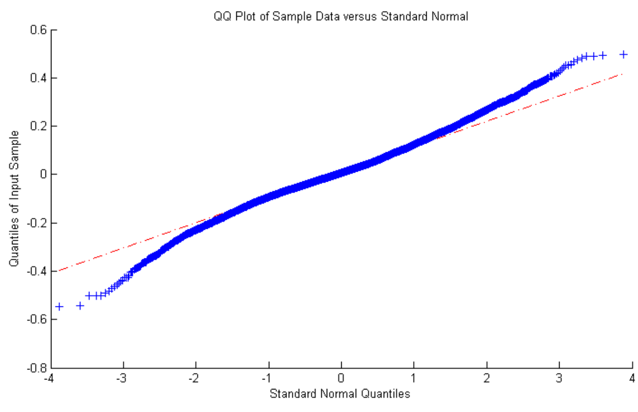

Figure 2. Normal Q-Q plot of the ERCOT day-ahead forecast errors. If the observed errors were well represented by a normal distribution, all observations would lie on the straight line.

2) *Finland*

The Finnish installed wind capacity was the smallest in the study, with 102 MW of rated power. However, the capacity was spread over 25 sites (77 turbines), with the largest distance between the sites being 630 km. Figure 3 shows the slightly positively skewed and leptokurtic distribution of observed wind power forecasting errors for the Finnish system. The distribution included a number of fairly large positive forecast errors (over-forecasting), with a few exceeding half of the installed capacity. This may have been the result of the smaller number of turbines included in this data set, and possibly erroneous data used in producing the forecasts.

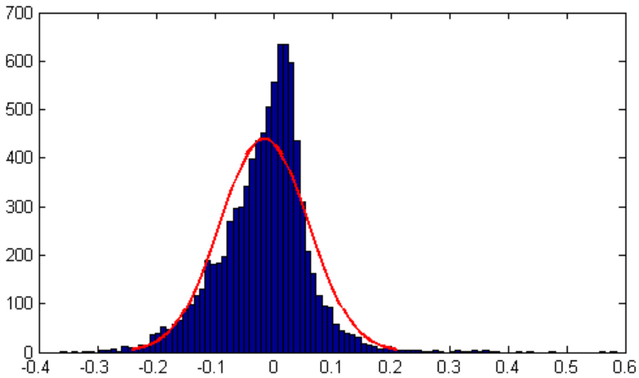

Figure 3. Histogram of the normalized day-ahead forecast errors for the Finnish system. $\mu = -0.0155$; $\sigma = 0.0751$; $\gamma = 0.0720$; $\kappa = 3.1036$

3) *Spain*

The Spanish installed wind power forecasting error histogram is shown in Figure 4. This data was from the year 2010 and included 19,300 MW of wind power capacity. The distribution was leptokurtic and fairly strongly positively skewed. The forecasts also displayed a notable bias, corresponding to more than 15% of installed wind power capacity. The distribution also had distinctly fat tails in both the over- and under-forecasting directions.

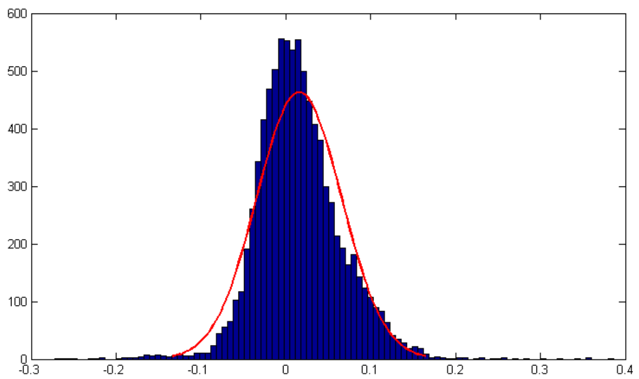

Figure 4. Histogram of the normalized day-ahead forecast errors for the Spanish system. $\mu = 0.1624$; $\sigma = 0.0514$; $\gamma = 0.3855$; $\kappa = 3.0180$

4) *Sweden*

The day-ahead forecasts for the Swedish system (year 2011) included 2,899 MW of installed wind capacity. The forecast errors plotted in Figure 5 showed a slightly leptokurtic negatively skewed distribution. The Swedish errors were interesting for their fairly small spread, with the largest errors being less than 30% of installed wind capacity. This was likely because of the large geographic diversity stemming from the multiple sites covering a large geographic area. It was also interesting to see that the normal distribution would under-represent the negative error tail, but over-represent the positive error tail, because of the skewness. Figure 6 shows the cumulative distribution function of the observed errors, the normal distribution based on those errors, and a hyperbolic distribution fit to the observed errors. It was readily apparent that the hyperbolic distribution provided a superior fit to the data than did the normal distribution, with the hyperbolic line running on top of the observed errors line for much of the cumulative distribution function. The Swedish example was chosen to display the cumulative distribution plot because of the clear example of the improved fit of the hyperbolic distribution. However, other cumulative distribution functions would have similar characteristics.

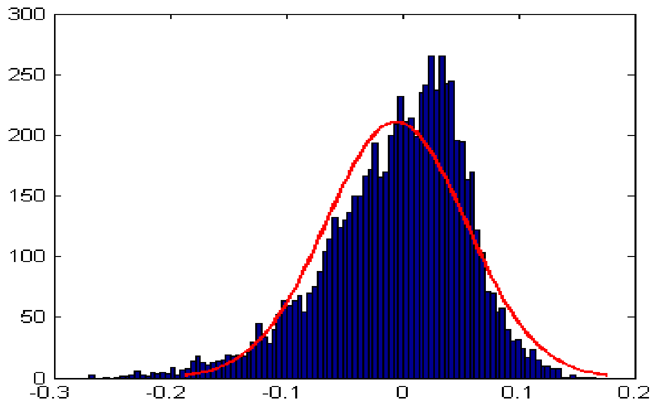

Figure 5. Histogram of the normalized day-ahead forecast errors for the Swedish system. $\mu = -0.0052$; $\sigma = 0.0603$; $\gamma = -0.7252$; $\kappa = 0.7757$

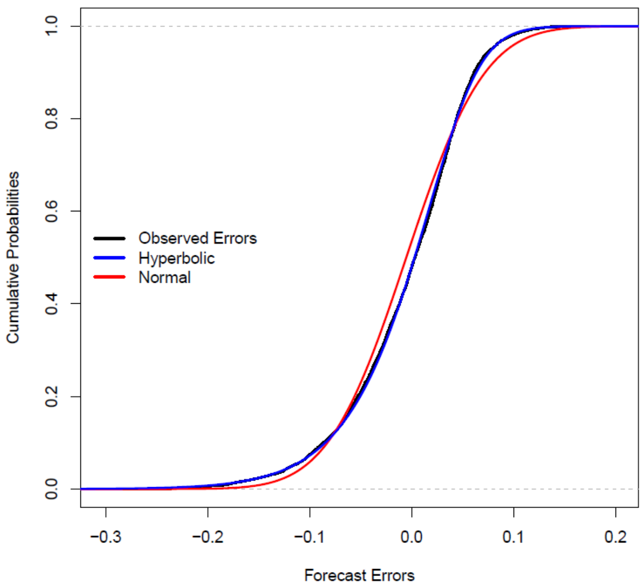

Figure 6. Cumulative distribution plot of the normalized Swedish day-ahead forecast errors

5) *Denmark*

The Danish system data included 3,871 MW of installed wind capacity for the year 2011. The distribution was more strongly leptokurtic than many of the other day-ahead forecast error distributions examined previously, as shown in Figure 7. Also, in contrast to the other countries, the distribution was fairly symmetric, with only a slightly positive skew. The spread of the data was fairly small with relatively few errors greater than 25% of the total installed capacity. This was likely a result of the geographic diversity acquired from the turbines being spread throughout the country. In addition, Denmark has a long history of wind power forecasting and contains relatively easy terrain.

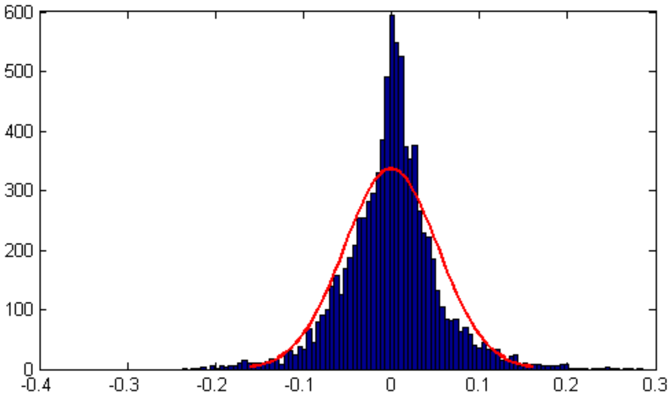

Figure 7. Histogram of the normalized day-ahead forecast errors for the Danish system. μ = -0.0005; σ = 0.0534; γ = 0.1378; κ = 2.3859

6) *Ireland*

The Irish data was from the year 2011 and included 1,557 MW of installed wind capacity. The Irish day-ahead forecasting errors had a small positive skew and were leptokurtic, as shown in Figure 8. There was a fairly large spread to the distribution, with a significant amount of forecast errors approaching 50% of installed wind power capacity. This was likely because of the small geographic area covered by the wind turbines. For reference, the total land area of Ireland is roughly 1/6th the land area of Sweden and 1/5th that of Germany.

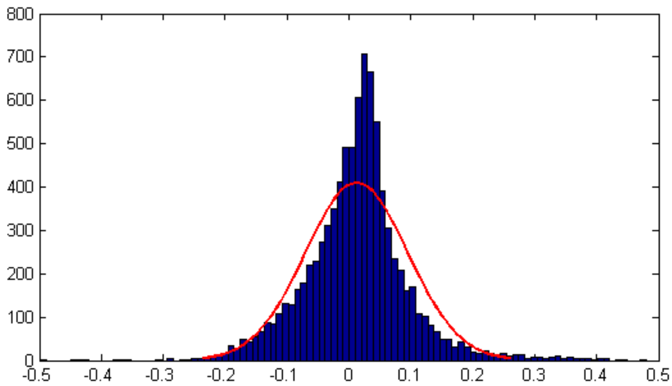

Figure 8. Histogram of the normalized day-ahead forecast errors for the Irish system. μ = -0.0123; σ = 0.0827; γ = 0.3063; κ = 3.0311

7) *Germany*

The German data was from the year 2010 and covered the total installed wind capacity in Germany ranging from 25.18 GW in January 2010 to 26.39 GW in December 2010. The power measurement was based on an up-scaling algorithm based on spatially distributed reference wind farms that included about 25% of the total capacity. The forecasts were used and published by the German transmissions system operators and were based on combinations of power forecasts from different providers and on different NWP models. The day-ahead forecasting errors had a slightly negative skew, and were leptokurtic, as shown in Figure 9. The spread of the data was relatively small, with all errors less than 30% of installed wind capacity. This was because of the large number of turbines included in the analysis, as well as the geographic spread of the locations used.

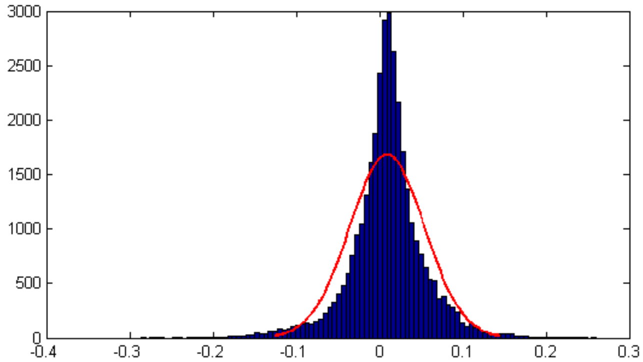

Figure 9. Histogram of the normalized day-ahead forecast errors for the German system. μ = 0.0092; σ = 0.0450; γ = -0.2891; κ = 3.5896

B. *Hour-Ahead Forecasts*

In power system scheduling, short-term wind power forecasts are necessary to set the generating unit output levels in the dispatch process, which often coincides with intraday market timing. These shorter term forecasts are used to reduce the uncertainty from day-ahead forecasts; consequently, only these forecast errors must be balanced by reserve power [16]. Although the dispatch interval may vary between systems, we used a one-hour interval as a representative example.

1) *United States*

The U.S. hour-ahead forecast error distributions came from a single wind plant in the Xcel Colorado service territory with approximately 300 MW of capacity. Because this data came from a single plant, the benefits of geographic diversity were not apparent. This was clear when looking at the extreme values shown in Figure 10. The maximum errors for the single plant were approximately 80% of the total capacity. It must be noted that these large values were likely because of the manual curtailment of wind plant output. These hour-ahead forecasts also had a much greater kurtosis value than the day-ahead forecasts observed in the previous section. The practical implication of this is that the forecasts were often more accurate, but had occasional instances when they were very inaccurate.

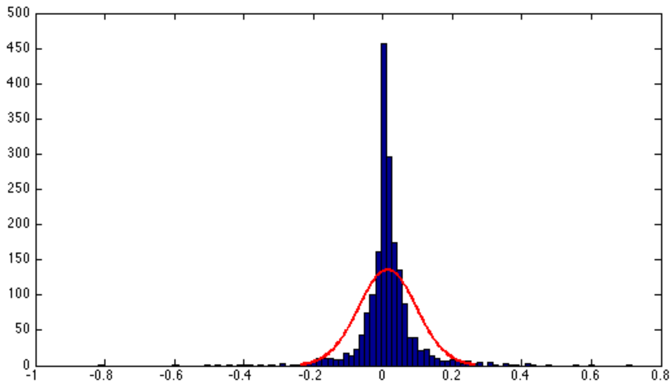

Figure 10. Histogram of the normalized day-ahead forecast errors for the Xcel Colorado wind plant. μ = -0.01; σ = 0.08; γ = - 0.01; κ =17.62

2) *Spain*

The hour-ahead forecasts for the Spanish system included 20,091 MW of wind power capacity, the second largest amount in this study. Figure 11 shows the histogram of the forecast errors. One important aspect to note is the smaller range of values shown in the Spanish data, with forecast errors above 10% of capacity being very rare. Part of the explanation is that the forecasts were hour-ahead instead of day-ahead data; the smaller forecasting interval reduced the uncertainty in the forecast considerably. The fat tails shown in Figure 11 resulted in a poor fit to the normal distribution. Further verification of this finding was provided by the dramatic deviations in the tails in the normal Q-Q plot shown in Figure 12.

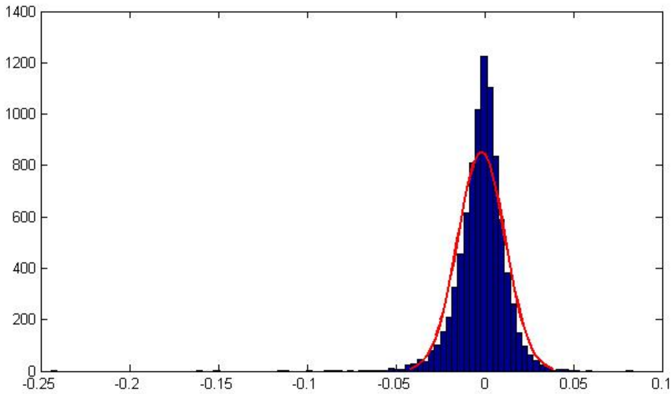

Figure 11. Histogram of the normalized hour-ahead forecast errors for the Spanish system. $\mu = -0.0018$; $\sigma = 0.0133$; $\gamma = -1.6585$; $\kappa = 20.2385$

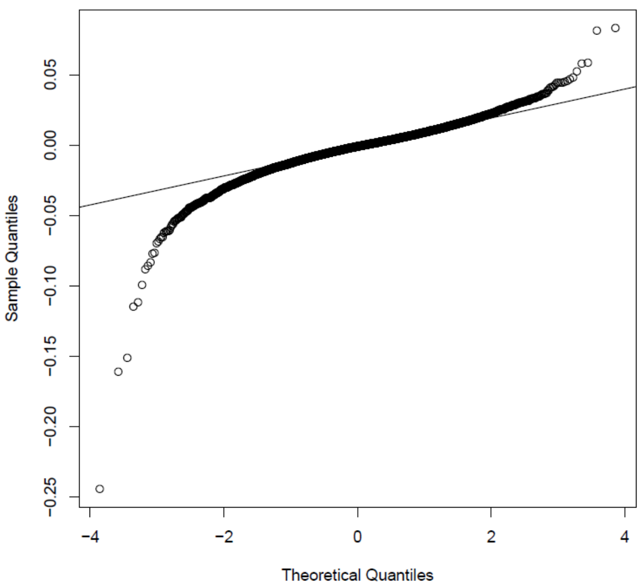

Figure 12. Normal Q-Q plot of the Spanish hour-ahead forecast errors. If the observed errors were well represented by a normal distribution, all observations would lie on the straight line.

3) *Germany*

The hour-ahead forecasts for the German system were for the same set of wind plants described in the day-ahead section. The histogram of the hour-ahead forecast errors is shown in Figure 13. The spread of the errors was very small, with no errors above 10% of installed capacity. As mentioned for the Spanish system, the large amount of wind turbines considered (~25 GW), with the resulting geographic diversity, was an important factor in the smaller spread of the error distribution, as was the usage of online power measurements that underlay a high-quality data check. The distribution was leptokurtic and slightly negatively skewed.

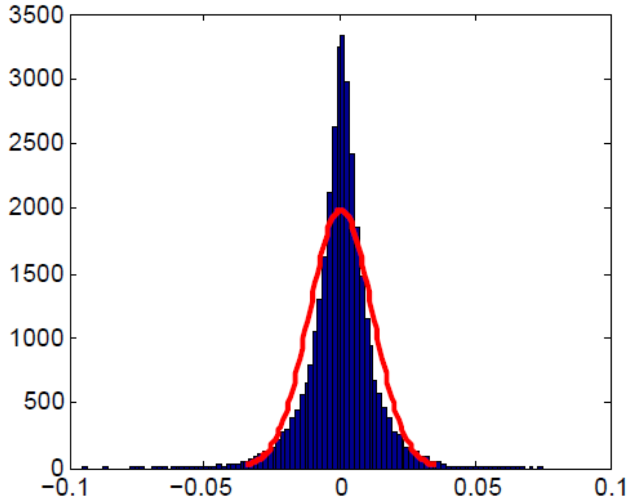

Figure 13. Histogram of the normalized hour-ahead forecast errors for the German system. $\mu = 0.0004$; $\sigma = 0.0116$ $\gamma = -0.2194$; $\kappa = 3.7389$

V. COMPARISON

The wind power forecasting errors shown in this study followed at least one common theme, regardless of country, forecasting period, or installed wind capacity considered: they were all leptokurtic distributions that were poorly represented by normal distribution. However, the distributions shown varied considerably based on each of the aforementioned criteria. As might be expected, the hour-ahead forecasts had much higher kurtosis values than those made at the day-ahead timescale. This would be expected from the reduction in uncertainty that occurs between making a forecast in the day-ahead time frame versus a single hour ahead. These distributions had many more very small forecast errors but still had large forecast errors in extreme cases with high power ramps, as represented by the relatively fat tails of the distributions. Generally speaking, the larger the installed wind power capacity, the smaller the spread of the distribution. This was related to the geographic diversity of having more turbines experiencing different weather conditions at the same time, though one exception of this is the ERCOT day-ahead data set. Most of the wind capacity installed in Texas is found in a narrow corridor in the northwest panhandle of the state. Additionally, wind turbines in the United States tend to be built in clustered plants, with a high density of turbines in a small area. In some of the European countries considered, the turbines are built in smaller groups, with less dense clusters of wind power. This geographic distance means that the forecasting errors between individual turbines are not as well correlated.

TABLE I. DAY-AHEAD FORECAST SUMMARY

	U.S.	Finland	Spain	Sweden	Denmark	Ireland	Germany
Installed Capacity (MW)	9,000	102	19,300	2,899	3,871	1,557	26,000
Data Set Length (Hours)	9,504	8,760	8,760	7,370	8,760	8,760	8,760
Forecast Horizon (Hours Ahead)	8–32	12–36	1–48	16–40	12–36	6–144	12–48

VI. CONCLUSION

This study examined the day-ahead and hour-ahead wind power forecasting errors seen in operating practice in seven countries. The distribution of forecasting errors was shown to be poorly represented by the normal distribution often assumed in wind integration studies. The distributions were found to be more leptokurtic, with an important distinction being the heavier tails seen in the operational forecast error distributions. Extreme errors that are not represented by normal distribution can have a large economic impact on integration planning studies and system operations. We recommend that future integration studies use representative wind power forecasting error distributions to guide the process instead of making the normal distribution assumption. In this study, the hyperbolic distribution was found to better represent the entire wind power forecasting error distribution. Further investigation is planned on the significance of the differences found in the country-to-country variations of wind power forecasting error distributions. Likely causes of such differences—such as country-specific geographic features, forecasting methods, model input parameters, and long-term wind resource quality—will be analyzed. The use of this information can be important in system operations, impacting operational and planning policies. An examination of how these country-specific error distributions could impact issues such as wind power curtailment policies and thermal generator flexibility is planned. Additional work is also planned on disaggregating forecast error distributions based on time of day and prevailing weather patterns to extract more useful information for system operations.

ACKNOWLEDGMENT

B.H., D.L., and M.M. thank David Maggio at ERCOT and Keith Parks at Xcel Energy for supplying data used in this study. R.S. and L.S. thank Mikael Kröckel and Oskar Sämfors at Svenska Kraftnät for supplying data on Sweden. X.L. and G.G would like to thank the Nordic Energy Research project IceWind for funding under TFI PK-int01.

This work was also supported by the U.S. Department of Energy under Contract No. DE-AC36-08-GO28308 with the National Renewable Energy Laboratory.

REFERENCES

[1] H. Holttinen et al., "Impacts of large amounts of wind power on design and operation of power systems: Results of IEA collaboration," Wind Energy, vol. 14, pp. 179–192, 2011.

[2] G. Giebel, R. Brownsword, G. Kariniotakis, M. Denhard, and C. Draxl, "The State-Of-The-Art in Short-Term Prediction of Wind Power: A Literature Overview", 2nd Edition, Project report for the Anemos.plus and SafeWind projects. 110 pp. Risø, Roskilde, Denmark, 2011.

[3] R. Doherty and M. O'Malley, "A new approach to quantify reserve demand in systems with significant installed wind capacity," IEEE Transactions on Power Systems, vol. 20, pp. 587–595, 2005.

[4] E. Estanqueiro et al., "Energy storage for wind integration: Hydropower and other contributions," IEEE Power and Energy Society General Meeting Proceedings, San Diego, California, July 22–26, 2012.

[5] A. Tuohy, P. Meibom, E. Denny, and M. O'Malley, "Unit commitment for systems with significant wind penetration," IEEE Transactions on Power Systems, vol. 24, pp. 592–601, May 2009.

[6] P. Ruiz, C. R. Philbrick, and P. Sauer, "Wind power day-ahead uncertainty management through stochastic unit commitment," IEEE Power and Energy Society Power Systems Conference and Exposition Proceedings, Seattle, Washington, March 15–18, 2009.

[7] H. Holttinen et al., "Variability of load and net load in case of large-scale distributed wind power," 10th International Workshop on Large-Scale Integration of Wind Power into Power Systems as well as on Transmission Networks for Offshore Wind Farms Proceedings, Aarhus, Denmark, October 25–26, 2011.

[8] V. Pappala, I. Erlich, K. Rohrig, and J. Dobschinski, "A stochastic model for the optimal operation of a wind-thermal power system," IEEE Transactions on Power Systems, vol. 24, pp. 940–950, May 2009.

[9] M. Ortega-Vazquez and D. Kirschen, "Estimating spinning reserve requirements in systems with significant wind power generation penetration," IEEE Transactions on Power Systems, vol. 24, pp. 114–124, 2009.

[10] B.-M. Hodge, E. Ela, and M. Milligan, "The distribution of wind power forecast errors from operational systems," 10th International Workshop on Large-Scale Integration of Wind Power Into Power Systems as well as on Transmission Networks for Offshore Wind Farms Proceedings, Aarhus, Denmark, October 25–26, 2011.

[11] P. Pinson, "Estimation of the uncertainty in wind power forecasting," Ph.D., Energetique, L'Ecole des Mines de Paris, Paris, France, 2006.

[12] K. Dietrich, J. Latorre, L. Olmos, A. Ramos, and I. Perez-Arriaga, "Stochastic unit commitment considering uncertain wind production in an isolated system," 4th Conference on Energy Economics and Technology, Dresden, Germany, 2009.

[13] H. Bludzuweit, J. A. Dominguez-Navarro, and A. Llombart, "Statistical analysis of wind power forecast error," IEEE Transactions on Power Systems, vol. 23, pp. 983–991, August 2008.

[14] B.-M. Hodge and M. Milligan, "Wind power forecasting error distributions over multiple timescales," in IEEE Power and Energy Society General Meeting Proceedings, Detroit, Michigan, July 24–29, 2011.

[15] C. Lowery and M. O'Malley, "Impact of wind forecast error statistics upon unit commitment," IEEE Transactions on Sustainable Energy, in press.

[16] J. Dobschinski, B. Lange, K. Rohrig, Y. Saint-Drenan, A. Wessel, and L. von Bremen, "The potential of advanced shortest-term forecasts and dynamic prediction intervals for reducing the wind power–induced reserve requirements," European Wind Energy Conference and Exhibition Proceedings, Warsaw, Poland, April 20–23.