

# Net Interchange Schedule Forecasting of Electric Power Exchange for RTO/ISOs

Tom Ferryman, David Haglin, Maria Vlachopoulou, Jian Yin, Chao Shen, Frank Tuffner, Guang Lin, Ning Zhou, and Jianzhong Tong (PJM)

**Abstract**—Neighboring regional transmission organizations (RTO) and independent system operators (ISOs) exchange electric power to enable efficient and reliable operation of the grid. Net interchange (NI) schedule is the sum of the transactions (in MW) between an RTO/ISO and its neighbors. Effective forecasting of the amount of actual NI can improve grid operation efficiency and avoid the volatility of the energy markets due to changes of NI schedules. This paper presents results of a preliminary investigation into various methods of prediction that may result in improved prediction accuracy. The methods studied are linear regression, forward regression, stepwise regression, and support vector machine (SVM) regression. The effectiveness of these methods is compared using the 64 weeks of field measurement data from PJM. The objective is to explore the effectiveness of the prediction methods under different scenarios.

**Index Terms**— power grid operations, net interchange schedule, forecasting, prediction.

## I. INTRODUCTION

To improve the efficiency and reliability of power grid operations, neighboring RTO/ISOs often exchange electric power. Net interchange (NI) schedule is the sum of the transactions (MW) between an RTO/ISO and its neighbors. NI is used in security constrained economic dispatch (SCED) to calculate the economic dispatch of next several hours. Therefore, effective forecasting of the amount of NI can improve the operation efficiency of the RTO/ISO.

To forecast the NI, it is essential to build a forecasting model that describes the relationship between future NI and currently available variables. NI schedule (NIS) can be updated by market participants. The reason for update may depend on location marginal price (LMP), external prices, outages and constraints, forecasted demand, and/or dispatch control signal. To date, the authors of this paper have not identified any literatures that directly relate these variables to future NI.

To build the model to forecast NI, this paper proposes a data-driven methodology. Several forecasting methods (i.e., linear regression, forward regression, stepwise regression, and support vector machine (SVM) regression) are implemented and applied to forecast NI based on the real-world data. A metric is proposed and used as Measure of Effectiveness (MOE) to compare the effectiveness of different forecasting methods and identify recommendable methods.

The remainder of this paper is divided in six sections.

Section II will define the problem of forecasting NIS and introduce a MOE to measure the effectiveness of forecasting. Section III will describe the data used to train and test the algorithms. Section IV will describe the methods used. Section V will present the results for forecasting NIS. Section VI will draw the conclusions and discuss future work.

## II. PROBLEM DEFINITION

The goal is to provide an effective method to predict actual NI of an RTO/ISO. To develop this method, real-world data was assembled (described in section III) from PJM, a cooperating RTO. The original datasets are divided into two subsets, i.e., training dataset and testing dataset. Prediction models are derived based on the training dataset using a number of prediction methods. The derived prediction models are then used to predict the NI data values of the testing dataset to evaluate their effectiveness. This methodology is performed iteratively. At the  $k^{\text{th}}$  iteration, the NI values of week  $k+1$  are predicted based on a model derived from training data spanning weeks 1, ...,  $k$ . The predicted NI values are then compared with the actual NI to evaluate the prediction effectiveness. The training dataset is therefore increasing in size at each iteration by a week's worth of data, approximately 672 (7 days \* 24 hours \* 4 data records/hour) data records/week, assuming no missing data.

To evaluate the performance of different forecasting methods, a MOE was calculated for each method as follows:

$$\begin{aligned}
 &P(t + \delta) \text{ denotes the predicted NI in MW for time } (t + \delta), \\
 &\quad \text{made at time } t. \\
 &A(t + \delta) \text{ denotes the actual NI in MW at time } (t + \delta). \\
 &T_1(t) = \frac{|P(t + \delta) - A(t + \delta)|}{\max\{|A(t + \delta)|, 10\}} \\
 &T_2(w) = \text{median}(T_1(t) | t \in t_w \text{ and } w > 1) \\
 &MOE = \text{median}(T_2(w), \forall w > 1)
 \end{aligned}$$

A few comments about the MOE are in order:

- The denominator of  $T_1$  might normally be the actual net power exchanged to make  $T_1$  a relative error, and would work fine except for the cases where the actual

net power exchanged was very close to zero. In those cases, the values of  $T_1$  would be grossly inflated, even if the prediction errors were very small. The use of 10 as the minimum alternate value for the denominator is a judgment that reflects domain experts' desire to have error magnitudes much less than 10.

- The use of absolute values in defining  $T_1$  clarifies the desire to consider over estimates and under estimates equally.
- Index  $w$  ranges from 1 to 64 and indicates the numerical label of the week of the data set.
- Set  $t_w$  includes all the time instances of week  $w$ .
- For calculating  $T_2(w)$ , the median of  $T_1$  was made over all the instances in  $w^{\text{th}}$  week. The use of *median* for the calculation of  $T_2$  instead of *mean* or some other measure of central tendency is to prohibit occasional very large prediction errors from causing a prediction method to be rejected, even though it is generally very effective.
- The sampling interval is 15 minutes. The discrete data uses  $t$  as an index and refers to the current time. The forecasting step is  $\delta$ , which indicates a leading time of  $15 \cdot \delta$  minutes. (Index  $\delta$  refers to a lead time for forecasting in the number of increments of 15-minutes.)

According to the discussions with industrial experts at PJM, the desired performance levels thresholds are set up as:

Forecast Lead Time	$T_2$ Limit
Hour 0 to Hour 1	10%
Hour 1 to Hour 2	12%
Hour 2 to Hour 3	15%
Hour 3 to Hour 4	20%

The researchers' efforts are focused on reducing  $T_2$ . This paper presents the results for 15-minute, and 1, 2, 3, and 4-hour forecasts.

Important constraints to this investigation are:

- Computer hardware must make predictions using modest-sized desktop personal computers or smaller.
- Processing must be straight-forward and capable of being used by PJM personnel with typical computer skills.
- Data used for the forecasting process can only be the type indicated by the data provided by PJM. Various kinds of data (e.g., weather) might arguably be useful, but are NOT considered in this proof-of-concept study.
- Processing times must be less than 12 hours to enable overnight processing.

The researchers are aware that this is a time series kind of problem and that an initial step in any time series problems is to account for any long term explainable and predictable trends. That is the goal of this investigation. After this goal is accomplished

to a reasonable level, analysis of residuals using ARIMA and other time series techniques is a natural next step.

### III. DATA AND MODEL

Figure 1 provides an illustration of the data matrix created to support the analysis tools. The far left column, colored blue, is the dependent variable, ActualNetInterchange. This is the quantity that to be predicted at time  $t + \delta$ .

The second sub-matrix indicated in Figure 1, colored light blue, represents the potential explanatory data variables provided by PJM, which may impact on the future NI. The PJM provided real-world data for one time period, 2000/2001. The data provided by PJM including date/time information, actual net interchange, actual load, scheduled net interchange, and 16 other variables that PJM deemed likely to help the forecast.

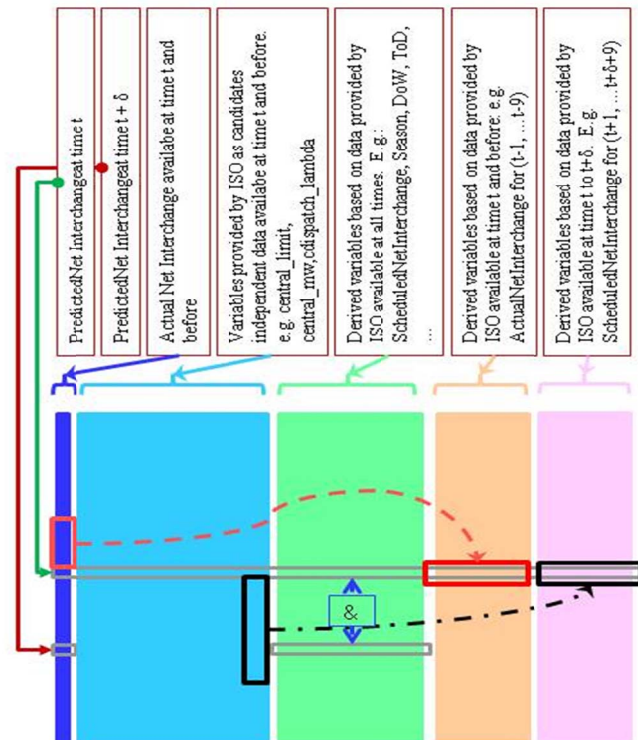


Figure 1. Illustration of Data Matrix.

The data is provided every 15-minutes (with a few exceptions) for 64 weeks between year 2000 and 2001. The data was reviewed to identify data quality issues. Domain knowledge indicated a few of these checks and a data driven test. The data must pass a data check based upon a comparison of the data with itself using the Two-Tiered Chebyshev Method [Amidan 2005]. This data check generates upper and lower limits that indicate values outside of which are considered very likely to indicate the data has a value that should be considered unrealistic.

Overall, only a small amount of data was identified as poor quality. Data that failed any of the above checks were replaced with NA's (a notation for missing data). The third

sub-matrix indicated in Figure 1, colored light green, contains derived variables that can be available to the analysis tools at the time of prediction and before, and after that time. A prediction at time  $t + \delta$  can take advantage of the scheduled net interchange variable, and a few variables that are directly related to the time  $t + \delta$ . The following data were derived from the date/time information and can be used as additional independent data variables:

- **Seasonal Index:** The year was partitioned into 6 seasons (each 61 days long). An indicator variable was created for each season and each date was mapped to the season it was associated with;
- **Day-of-Week Index:** the date was mapped to 7 indicator variables depending on the day of the week at the time  $t$ ;
- **Time-of-Day Index:** the time,  $t$ , was mapped to one of eight bins (00:00 to 03:00, 03:00 to 06:00, ..., 21:00 to 24:00) and indicator variables created to capture that mapping.

The fourth sub-matrix of Figure 1, colored orange, is past actual net interchange values for the time of the prediction (i.e.,  $t$ ) and the preceding 135 minutes, nine reading previous to time  $t$ . For instance, at time  $t = 12:00$ , the actual net interchange values from times 9:30 to 11:45 are included in the fourth sub-matrix.

The fifth sub-matrix of Figure 1, colored pink, is the upcoming scheduled predictions (Scheduled predictions are provided by PJM for several days in advance of “current” time.) for the multiple time periods between time  $t$  and time  $t + \delta + 9$ . This corresponds to using up to 25 additional variables when  $\delta$  is four hours or sixteen 15-minute intervals. For instance, at time  $t = 12:00$  and for a  $\delta = 4$  intervals (or 60 minutes), scheduled predictions from PJM are included from times 12:15 to 15:15 ( $\delta = 4$  intervals + 9 extra time periods of 15-minutes each) in the fifth sub-matrix.

The first sub-matrix consists of the dependent variable to be forecasted. The second, third, fourth and fifth sub-matrices constitute the independent variables that impact on the dependent variables.

Naturally, many variations of the independent variable matrix could have been considered. The version described above is a reasonable choice because it has some sophistication beyond the raw data, but with still a modest degree of increased complications. Use of other variations could result in improved prediction, but are deemed beyond the scope of this investigation.

The same data set was used for all forecasting methods.

#### IV. ANALYSIS METHODS USED

Several different methods were used, some simple and some more sophisticated. The simple methods are intended to provide a baseline from which one can measure the improvement provided by the more sophisticated methods. To help understand these methods, a few terms are defined as follows:

- $t$  denotes the current time at which the calculations for the predictions began. Only those data that are available at

time  $t$  can be used to make prediction.

- $t + \delta$  denotes the time in the future at which the value of the dependent variable is to be predicted. The symbol  $\delta$  is the lead time of prediction. Multiple lead time  $\delta$ 's were explored: 1, 4, 8, 12 and 16 corresponding to 15, 60, 120, 180, and 240 minutes into the future.
- The data consisted of 64 weeks of data. Each week from week 2 on was predicted using available training data to tune each of the prediction algorithms. Each prediction week entails 672 predictions made by each method (7 days \* 24 hours/day \* 4 prediction periods/hour).

Five primary prediction methods are utilized in this demonstration paper. The methods explored in this initial demonstration are [11], [12]:

- **Scheduled Net Interchange:** PJM currently has a prediction of the NI for each 15-minute period at least 1 day in advance. This information is provided as an independent variable, denoted “scheduled\_net\_interchange.” This value is considered a prediction and used as the “baseline” method.
- **Linear Regression:** This method fits the “actual\_net\_interchange\_at\_time\_tPlusDelta” data using ordinary linear least squares regression using all the independent data variables as regressors as indicated in Figure 1 (all 4 sub-matrices). Based on linear regressions using all data from time 1 to time  $t$ , the method calculates the regression coefficients. Then, predicts the dependent variable value at the time of the prediction,  $t + \delta$ .
- **Forward Regression:** This is an iterative method for selecting regressors from a group of independent variables according to their significance. This method fits the “actual\_net\_interchange\_at\_time\_tPlusDelta” data in a linear way using all the independent data variables as candidate regressors as indicated in Figure 1. The forward regression algorithm starts by using one regressor for the linear fitting, and then proceeds by adding more regressors at each iteration step. The regressors are added in order based on their statistically significance and they are not added if they are not considered statistically significant. The test for statistical significance is performed using the Akaike Information Criterion (AIC) [13],[14].
- **Step-wise Regression:** This method is similar to the forward regression because it still fits the “actual\_net\_interchange\_at\_time\_tPlusDelta” data in a linear way using the all the independent data as candidate regressors, as indicated in Figure 1. It is also an iterative process, which in addition to including the most significant regressors, like the forward regression, can also remove regressors that are deemed insignificant during the regression run. It is possible that the statistical significance of a variable can change with the addition of new variables into the linear regression model. Stepwise regression is able to re-evaluate the decision of which regressors to use at each iteration step. The AIC is also used by this method as the test for statistical significance [<http://www.r-project.org/>].

- **SVM Regression.** The support vector machine (SVM) method is a data mining technique that is commonly applied to solve classification and regression problems. In this case, it performs  $\epsilon$ -regression and a fitting model is derived using a training dataset [15]. The fitting model can then be used to predict the “actual\_net\_interchange\_at\_time\_tPlusDelta” variable in the testing dataset. The SVM method evaluates the fitting function based on a subset of the training data that constitute the support vectors. Linear combinations of the support vectors form hyperplanes that provide a multi-dimensional fitting of the dataset. For example, there are  $N$  training records and each record is a tuple  $(\mathbf{x}_i, y_i)$  with  $i = 1, 2, \dots, N$ , where  $\mathbf{x}_i$  are the regressors and  $y_i$  is the value of the “actual\_net\_interchange\_at\_time\_tPlusDelta” for the  $i_{th}$  record. The hyperplanes  $f(\mathbf{x}) = \langle \mathbf{w}, \mathbf{x} \rangle + b$  should approximate all  $(\mathbf{x}_i, y_i)$  with  $\epsilon$  precision and  $\mathbf{w}, b$  are the model parameters that are derived by the SVM method. The solution of the fitting problem is the solution of the following convex optimization problem.

$$\begin{aligned} & \underset{\mathbf{w}, b, \xi, \xi^*}{\text{minimize}} && \frac{1}{2} \|\mathbf{w}\|^2 + C \sum_{i=1}^n (\xi_i + \xi_i^*) \\ & \text{subject to} && y_i - \langle \mathbf{w}, \mathbf{x}_i \rangle - b \leq \epsilon + \xi_i \\ & && \langle \mathbf{w}, \mathbf{x}_i \rangle + b - y_i \leq \epsilon + \xi_i^* \\ & && \xi_i, \xi_i^* \geq 0 \end{aligned}$$

The constant  $C > 0$  is a tuning parameter of the optimization process. The problem formulation includes slack variables  $\xi_i, \xi_i^*$ . So far, it was assumed that the data can be fitted by a linear plane. However, it is often the case that the data can only be adequately approximated by a non-linear surface. The SVM method has a clever way of dealing with this issue by projecting the data to a feature space via a mapping  $\Phi$ . In the new feature space the data should be transformed in such a way that a linear fit provides a good approximation. The structure of the optimization equations above remain the same with  $\mathbf{x}$  being substituted by  $\Phi(\mathbf{x})$ . The mapping of the data itself does not have to be implemented by the algorithm, since the algorithm only uses the dot product of the mapped data  $K(\mathbf{u}, \mathbf{v}) = \langle \Phi(\mathbf{u}), \Phi(\mathbf{v}) \rangle$  which is known as the kernel trick. This research utilizes the Radial Basis Function (RBF) kernel. The parameter setting for the SVM method, described in Section IV, are  $C = 50,000$  and  $\epsilon = 2$ . The SVM method is implemented using the LIBSVM software package [16].

The first method uses the scheduled net exchange data currently available from PJM as forecasted values. As such, it was chosen as the benchmark for measuring effectiveness. The other methods were selected as having various degrees of sophistication and anticipated effectiveness. Limited resources did not allow exhaustive investigations of:

- the many other methods that exist,

- use of additional independent data that PJM might be able to provide,
- use of additional derived variables as independent data, and/or
- different variations or tuning parameters of these methods that might provide enhanced capability.
- Any of these areas of investigations could result in major improvements in accuracy.

## V. RESULTS

In this section, the results are summarized. The MOE defined in Section 2 was calculated for each of the 63 test weeks for each method. Table 1 summarizes the results.

Table 1. MOE value over week 2 through 64 for predictions 15, 60, 120, 180, 240 minutes in advance of *current* time.

Methods	Prediction in minutes				
	15	60	120	180	240
<b>Schedule</b>	9.17	9.19	9.23	9.19	9.20
<b>Linear Reg.</b>	7.65	7.82	7.70	7.76	7.77
<b>Stepwise Reg.</b>	7.48	7.52	7.70	7.74	7.77
<b>Forward Reg.</b>	7.53	7.76	7.55	7.58	7.71
<b>SVM</b>	7.91	8.11	8.20	8.14	8.18

Each of the methods demonstrated a notable improvement over the Schedule method. This is not surprising because the schedule values were available to be used by each forecast method.

Figure 2 shows the MOE for each week for the 15-minute forecast for each method. The most notable observation is that the MOE varied over a range from 3 to approximately 20 for all the methods. Also evident in Figure 2 is the fact that the MOEs of the methods are highly correlated, probably due to the volatile weekly behavior of the grid; when some weeks it behaved as predicted and other weeks it did not.

In order to assess the methods’ relative performance, weekly differences between pairs of methods were calculated and presented in the boxplots<sup>1</sup> in Figure 3.

Table 1 and Figure 3 both focused on the 15-minute forecasts. The forecasts for 1, 2, 3 and 4-hour forecasts show very similar patterns. It is notable that the degradation due to time is rather minor. Also, the relative strengths of the methods are similar for all forecast periods tested. The first four boxplots in Figure 3 show that all four forecasting methods do better than the original schedule’s

<sup>1</sup> The boxplots in Figure 3 summary the values of the difference between one forecast method and another forecast method by displaying a horizon line at the median and a box that has upper and lower lines at the approximate 75 and 25<sup>th</sup> percentiles. Small circles are displayed for outliers outside of the box.

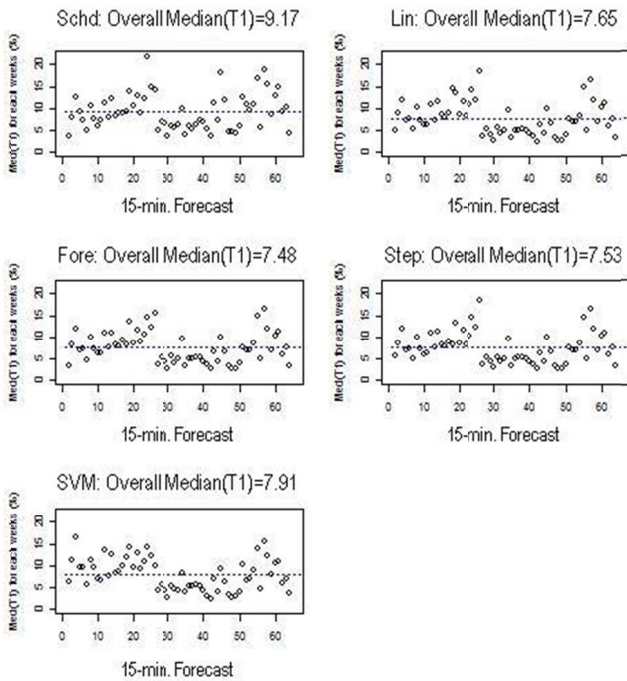


Figure 2. 15-minute Forecast for 5 methods over 63 test weeks.




Figure 3. Boxplots of differences for each pair of methods.

performance. The next 3 boxplots compare linear regression to forward regression, stepwise regression and SVM regression. Linear regression appears to be performing very close to forward and stepwise regression, and a little better than SVM regression. The eighth and ninth boxplots compare forward regression to stepwise and SVM regression. This shows the very close performance of forward and stepwise regression. And it shows forward regression doing slightly better than SVM regression. The tenth boxplot compares stepwise regression with SVM regression. Stepwise scored better.

This is one test. As indicated above, in Section 4, there are many improvements that can be expected if additional regressors were provided, or additional efforts to investigate the use of interactions of the available regressors. It was a bit of a surprise that SVM regression did not perform as well as was expected. This may reflect the fact that SVM has some important parameters for the researchers to set, and is significantly affected by those values. There were some efforts made to pick the best parameter values, but the efforts were basically ad hoc selections by an SVM regression expert and not an exhaustive or structured search.

## VI. CONCLUSIONS AND FUTURE WORK

From a technical perspective, there are still some issues the authors think should be dealt with. It is believed these questions could be addressed and possibly lead to significant improvements in prediction accuracy:

- Are there other independent variables (such as weather or special events) that could be included in the data set that might help the predictions?
- Are there derived variables (transformations or combinations of current or additional variables) that could help improve prediction accuracy?
- Would a second data set yield the same or similar results?
- Would a data set from another utility yield the similar results?
- Could the simultaneous analysis of multiple neighboring utilities result in a significant improvement? (What one utility forecasts it will export to another should sync with what the other utility forecasts will be imported; at least to some extent.)
- Could use of other prediction methods result in significant improvements?
- Could use of time series analysis of the residuals result in significant improvements?
- Could better selection of processing parameters result in significant improvements (especially with SVM)?

These are all questions that hopefully will be addressed in near-term research.

The investigation shows that there appears to be at least one method that improves the accuracy of the current Scheduled Net Interchange estimates, by about 20%. The PJM considers this R&D effort sufficiently promising to endorse it for future efforts and provide additional data for follow-on research. The work thus far continues to encourage the expectation that the methods will provide an accurate NI forecast and help reduce the volatility of the market.

## VII. REFERENCES

- [1] Abur Ali, and Antonio Gomez-Exposito, *Power System State Estimation*, Taylor & Francis, ISBN:0824755707, March 2004.
- [2] Amidan BG and TA Ferryman. 2005. "Atypical event and typical pattern detection within complex systems," In *Proceedings of 2005*

*IEEE Aerospace Conference*, pp. 3620-3631, DOI: 10.1109/AERO.2005.1559667

- [3] A. Monticelli, "Electric power system state estimation," *Proc. IEEE*, vol. 88, pp. 262-282, Feb. 2000.
- [4] Smith, J.C., and B. Parsons, "What does 20% look like?," *IEEE Power and Energy Magazine*, vol. 5, no. 6, pp. 22-33, 2007.
- [5] DeMeo, E.A.; Grant, W.; Milligan, M.R.; Schuerger, M.J., "Wind plant integration," *IEEE Power & Energy Magazine*, vol. 3, no. 6, pp. 38-46, Nov.-Dec. 2005.
- [6] Jain, A.; Shivakumar, N.R., "Power system tracking and dynamic state estimation," *2009 IEEE/PES Power Systems Conference and Exposition (PSCE 2009)*, pp. 8, Seattle, WA, March 15-18, 2009.
- [7] Sinha, A.K., Mondal, J.K., "Dynamic state estimator using ANN based bus load prediction," *IEEE Transactions on Power Systems*, vol. 14, no. 4, pp. 1219-25, Nov. 1999.
- [8] Mallieu, D.; Rousseaux, P.; Van Cutsem, Th.; Ribbens-Pavella, M., "A new hierarchical approach for dynamic state prediction and filtering in electric power systems," *Power Systems and Power Plant Control. Proceedings of the IFAC Symposium*, pp. 405-10, 1987.
- [9] Rousseaux, P.; Van Cutsem, Th.; Mallieu, D.; Ribbens-Pavella, M., "Dynamic state prediction and hierarchical filtering for power system estimation," *Large Scale Systems: Theory and Applications 1986. Selected Papers from the 4th IFAC/IFORS Symposium*, pp. 721-6 vol. 2, 1987.
- [10] Martin, K. E., "Phasor measurement systems in the WECC," in *Proc. IEEE/Power Eng. Soc. General Meeting Panel Session PMU Prospective Applications*, Montreal, QC, Canada, Jun. 18-22, 2006.
- [11] Haykin, S., "*Adaptive Filter Theory*", Prentice Hall, Upper Saddle River, NJ, 4<sup>th</sup> edition, 2002.
- [12] Draper, N.R., and H. Smith, "*Applied Regression Analysis*", John Wiley & Sons, New York, NY, 1966.
- [13] Akaike, H., "*A new look at the statistical model identification*", *IEEE Transactions on Automatic Control* 19 (6): pg. 716-723, 1974.
- [14] Akaike, H., "*Likelihood and the Bayes procedure*", *Bayesian Statistics*, pg. 141-166, 1980.
- [15] Smola, A. J. and Scholkopf, B., "*A tutorial on support vector regression*", *Statistics and Computing* 14: 199-222, 2004.
- [16] ] Chang, C-C., Lin, C-J., "*LIBSVM : a library for support vector machines*", *ACM Transactions on Intelligent Systems and Technology*, 2:27:1--27:27, 2011.

## VIII. BIOGRAPHIES



**Thomas Ferryman** has over 30 years of experience in a variety of R&D projects, including as Chief Scientist of the Statistical and Sensor Analytics Group at Pacific Northwest National Laboratory, as Chief Systems engineer at Lockheed Aircraft Services Company in charge of the R&D for a \$140 million modification to the AC-130 Pave Specter Aircraft, and as Branch Head of the U.S. Navy's Gun Weapon System Analysis group.

He has had in-depth involvement in developing analysis tools for a large variety of projects; including NASA Ames' Aviation Safety program including the development of algorithms for analysis of digital flight data, unstructured text data and fusion of multiple aviation data types (digital flight data, digital ATC data, unstructured text data, and survey data). (The NASA work won the 2005 R&D 100 Award for The Morning Report. See [www.pnl.gov/statistics](http://www.pnl.gov/statistics))

Dr Ferryman has developed analytical tools for a variety of domains including national security, energy, and environment.



**David Haglin** received his MS and PhD from the University of Minnesota in 1987 and 1989 respectively. He held the position of professor of Computer Science at Minnesota State University, Mankato and has recently moved to the Computational Sciences and Mathematics Division at the Pacific Northwest National Laboratory, where he is a senior research scientist. His research interests include data mining and graph algorithms, particularly in a high-performance computing context.

**Maria Vlachopoulou** received her B.S. degree in Electronic and Electrical Engineering from the University of Sheffield, U.K. in 2006. In 2010 she



received her M.S. degree in Electrical and Computer Engineering and her M.S. degree in Industrial Engineering (Operations Research) from Purdue University. She is currently with the Pacific Northwest National Laboratory as an engineer. Her research interests include the application of algorithms and statistical methods to the area of energy and power systems and energy markets.



**Jian Yin** is a senior scientist at the Pacific Northwest National Laboratory. He received his Ph.D. from the University of Texas at Austin. His research interests include data intensive computing, distributed systems, cloud computing, performance, and high performance computing. He holds five patents. Several of his papers have over a hundred citations. He has been awarded the best paper award from WWW'2001 and MASCOTS'2010.

**Francis Tuffner** (M'08) received his B.S. and M.S. degrees in electrical engineering from the University of Wyoming in 2002 and 2004. In 2008, he finished his Ph.D. in electrical engineering, also at the University of Wyoming. He is currently with the Pacific Northwest National Laboratory as a power system engineer. His research interests include signal processing applied to power systems, embedded control devices, and digital signal processing. (No photo available.)

**Chao Shen** received his BS and MS degrees from the Department of Computer Science and Engineering in Fudan University, China in 2006 and 2010 respectively. He is currently working on a Ph.D. in Computer Science at the School of Computing and Information Sciences, Florida International University. (No photo available.)



**Ning Zhou** (S'01- M'05- SM'08) received his B.S. and M.S. degrees in automatic control from the Beijing Institute of Technology in 1992 and 1995, respectively. In 2005, he received his Ph.D. in EE with a minor in statistics from the University of Wyoming. From 1995-2000 Ning worked as an assistant professor in the Automatic Control Department at Beijing Institute of Technology. He is currently with the Pacific Northwest National Laboratory as a power system engineer. Ning is a member of the IEEE Power Engineering Society (PES). His research interests include power system dynamics and statistical signal processing.



**Guang Lin** Received his B.S. degree in Mechanics with a minor in Electric Engineering from Zhejiang University, China in 1997, M.S. degree in Mechanics and Engineering Science from Peking University, China in 2000 and M.S. degree in Applied Mathematics from Brown University in 2004. In

2007, he received his Ph.D. in Applied Mathematics from Brown University. He is currently with the Pacific Northwest National Laboratory as a research scientist. He is currently served in the editor board of "International Journal of Uncertainty Quantification". He has had in-depth involvement in developing uncertainty quantification tools for a large variety of domains including energy, and environment. His research interests include diverse topics in computational science both on algorithms and applications. A main current thrust is uncertainty quantification, large-scale data analysis and multiscale modeling in a large variety of domain.

**Jianzhong Tong (PJM)** is a Senior Strategist at PJM Interconnection. He has authored and co-authored more than 60 publications and holds a US patent for Method of Voltage Stability Analysis of Power Systems. He earned his BS, MS and Ph.D. in Electrical Engineering. (No photo available.)