Implementation of a Gas Load Forecaster

At Williams Gas Pipeline

Presented to:

2001 Pipeline Simulation Interest Group (PSIG) Meeting

Salt Lake City, Utah

October 2001

By:

Paul Lamb, Williams Gas Pipeline – Transco

Dan Logue, Energy Solutions International

1 Abstract:

This document explains the use, benefits, and components of a Gas Load Forecasting System (LFS):

- Describes how the LFS is used within the Integrated Gas Management System (IGMS)
- Provides an overview of Artificial Neural Networks
- Provides an overview of gas load forecasting techniques
- Provides practical experiences, findings, and recommendations of the authors

2 Introduction

Williams Gas Pipeline (WGP) is an operating group of The Williams Companies, Inc. It is principally involved in the interstate transportation and storage of natural gas in the United States. WGP operates five pipeline systems, consisting of: Northwest Pipeline; Kern River systems; Texas Gas Transmission; Transcontinental Gas Pipe Line (Transco); and Williams Gas Pipeline–Central (Central). The combined WGP pipeline network has more than 27,300 miles of pipeline and is among the nation's largest-volume transporters of natural gas.

WGP has recently embarked on a series of twenty-one projects to enhance its information technology (IT) infrastructure. One of the IT systems being implemented is the Integrated Gas Measurement Management System (IGMS). IGMS is an integrated suite of software applications that provides operations analysis and decision support functions to WGP's three Operations Control centers. IGMS is being implemented with the assistance of Energy Solutions International, a Houston-based pipeline simulation company, which provides software development and systems integration services on the project

One of the key components of IGMS is the Gas Load Forecasting System (LFS). The LFS is an artificial neural network-based forecasting application that provides short-term (1 to 5 day) gas load forecasts down to the meter level. LFS generates hourly or daily forecasts depending upon the type of historical data available to it. The LFS stores load forecasts in a centralized Oracle database allowing the data to be shared among different applications within IGMS that require the forecasts.

The first implementation of the LFS has just been completed at WGP. The system is currently undergoing testing and will be deployed at two WGP locations by the time this paper is presented. As with many IT system implementations, the process has been a challenge. Nevertheless, the authors have learned much about gas load forecasting and have gained practical experience in making the system actually work and provide meaningful results. This paper starts by describing how the LFS is used within IGMS. Next, an overview of artificial neural networks is provided, followed by a discussion of gas load forecasting techniques. Finally, the practical experiences, findings and recommendations of the authors are presented.

3 The Purpose of the LFS

IGMS is an integrated suite of applications or modules that are integrated within a common IT framework. Figure 1 provides an overview of the system.

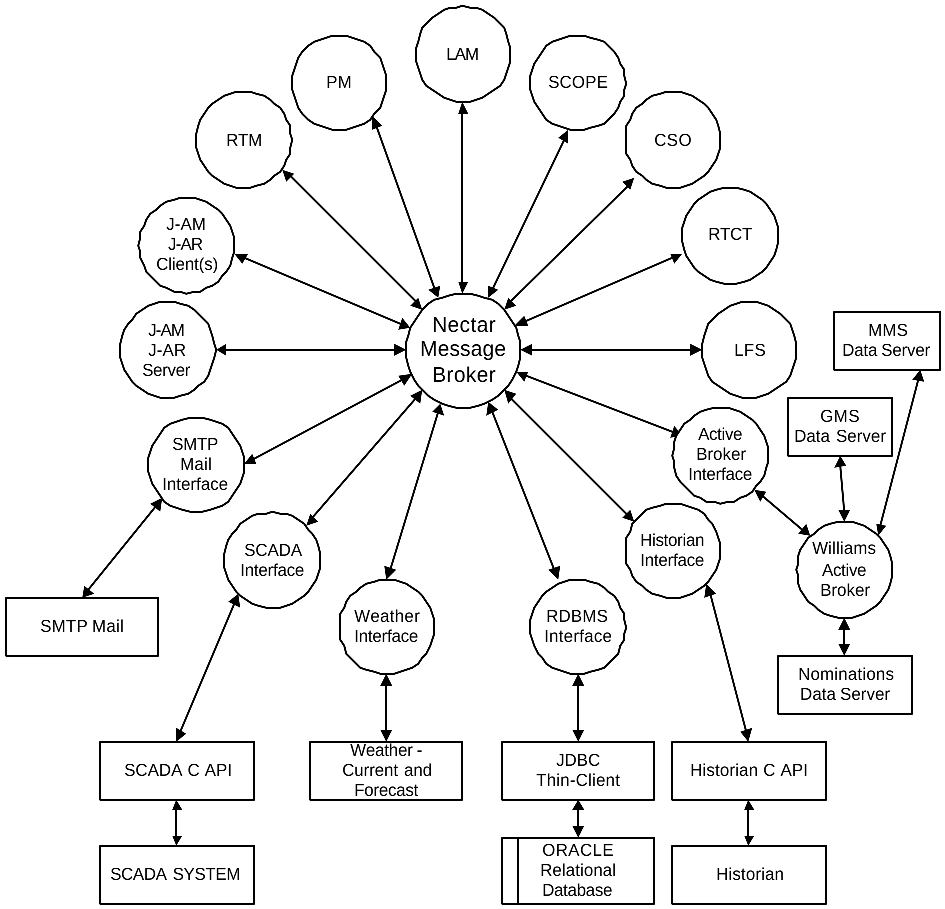

Figure 1: IGMS Overview

Each of the modules provides specific functionality, either for direct consumption by the user or for use by other modules. From a functional perspective, the main IGMS modules are as follows:

- Real Time Model
- Predictive Model (Predictor and Look-Ahead)
- Compressor Station Optimizer
- Load Forecaster
- Data Historian and
- Real Time Compressor Station Tuning

Within IGMS, there are three principal uses for load forecast data. Two of the uses are as input to other modules, basically serving as *data feedstock*. These modules are the Real Time Model (RTM) and the Predictive Model (PM). The third use of load forecast data is for advisories to Operations Control personnel and casual users throughout the company.

The RTM requires load forecast data in order to provide boundary conditions at meters that do not have electronic flow measurement (EFM). While becoming less common, several of the WGP pipelines still have a significant number of non-EFM meters. These meters have charts that are read typically once a month. The data is not available in real time, and this is problematic for the RTM. On most pipeline systems, the non-EFM meters are small enough volume and are interspersed among larger EFM meters such that they do not appreciably impact simulation fidelity. Several of the WGP pipelines have entire laterals composed entirely of non-EFM meters. These laterals would normally not be modeled, but the WGP implementation team thought the volumes were significant enough to require modeling.

The other IGMS module requiring load forecasts is the Predictive Models. The Look-Ahead Model (LAM) and Predictor Model (PM) are “fast-forward” simulations, of the pipeline system using the current state of the pipeline as determined by the RTM as a starting state, to determine the response of the system over some short period of time. Typically, Predictive Model runs are at most 24 hours into the future. The distinction between a LAM and PM simulation is the source of the boundary condition imposed on the simulation. Boundary conditions for a LAM simulation are obtained from an operating plan. An operating plan is a chronological list of known pipeline equipment status and set point changes for the pipeline over the duration of the simulation. A PM simulation, on the other hand, is a simulation whereby an engineer or operator manually establishes the boundary conditions. PM simulations are typically run as “what if” cases where the user is interested in evaluating the effect of the user specified boundary conditions on the pipeline.

For a LAM and PM simulation to be accurate, i.e., reflect the true response of the pipeline over time, accurate information must be provided to the simulation regarding meter conditions. Both meter boundary statuses (known pressure or known flow) and specific values for the pressures or flows must be provided to the simulation as a function of time. In WGP’s implementation of the Predictive Models, the LFS will provide the meter flow conditions.

Most seasoned operators have an uncanny ability to intuit gas demand. This comes from years of experience by having repeatedly observed how gas demand fluctuates due to weather, day-of-the-week, time-of-the-day and other factors. This ability, while frequently accurate, lacks several key elements to make it a useful decision support tool. First, given any two controllers, they are more likely than not to provide similar but not identical load forecasts for nearly identical conditions. Therefore, the forecast is dependent upon who is providing it.

Second, the forecasts are not numerically precise, meaning the controllers typically do not provide hard numbers but rather provide it in relative terms or with larger than acceptable confidence bands. For example, the load will be higher than yesterday’s by a little, or it will be around 100 MMCF ± 25 MMCF. Lastly, controllers are not very good at estimating the load forecast profile.

For these reasons, it was decided that the controllers would benefit from being able to view load forecasts directly from an easy-to-use reporting and graphing application interface.

4 Neural Network Fundamentals

There are numerous techniques to forecast gas load. The LFS implementation at WGP utilizes a neural network-based forecast model. Neural networks—or more accurately, Artificial Neural Networks (ANN)—are a non-linear fitting technique. ANNs can approximate any function and are considered *universal approximators*. ANNs are very appealing in application domains where one has little or incomplete understanding of the problems to be solved, but where training data is available (e.g., load forecasting).

4.1 Nodes

An ANN consists of a number of simple processing units (*nodes*) connected by a large number of weighted connections. Nodes perform a simple function; they receive input from neighboring nodes or external sources and compute output signals that are sent to other nodes.

Each node aggregates the signals it receives from other nodes using a *combination function*. Typically, the combination function is additive, i.e., each node provides an additive contribution to the nodes it is connected. The total input to node *j* is simply the weighted sum of the separate outputs from the connected units.

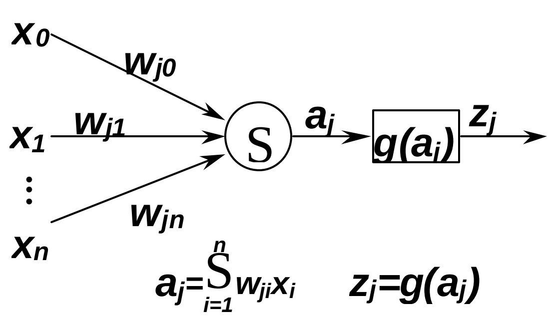

The functions for the nodes (activation functions) are required to introduce non-linearity into the network. Almost any non-linear function will do as long as it is differentiable. It also helps if the function is bounded. For these reasons, the sigmoid function is the most common choice:

$$g(x) = \frac{1}{1 + e^{-x}}$$

4.2 Layers

The nodes are typically arranged into layers of a multi-layer, feed-forward neural network often called a multi-layer perceptron, or MLP. An MLP may have one or more hidden layers. Figure 2 shows a generalized MLP with a single hidden layer.

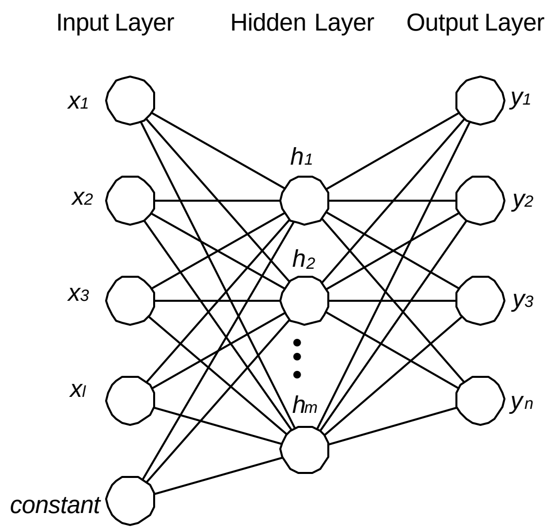

Figure 2: A Generalized MLP

Nodes are categorized as to which layer they exist. Nodes in the first layer, or input layer, receive input from external sources. Nodes in the middle layer, or hidden layer, receive input from the input layer nodes and send output to the output layer. Output layer nodes receive input from the hidden layer nodes and send data out of the network.

The next layer is the hidden layer. Research has shown that neural networks with two hidden layers can approximate any function. Practically speaking, the vast majority of problems can be fit reasonably well with just one hidden layer. Since ANNs with two or more hidden layers often have solution stability problems and produce more local minima, the LFS MLP has only one hidden layer.

The number of nodes in the hidden layer is important. Using too few nodes can result in the network failing to discern the relevant features in a complicated data set. This is called *underfitting.* Conversely, having too many nodes can potentially result in *overfitting.* Since ANNs can approximate essentially any function, they can also fit noise perfectly. Overfitting occurs when the ANN has fit the data well but fails to generalize the data because it is fitting noise. Overfitting produces a low training error but high generalization error. Overfit ANNs fail to predict accurately.

There are a number of “rules of thumb” for selecting the number of hidden layer nodes, but often the best method to determine the optimal number of hidden layer nodes is trial and error.

4.3 Training

The process of adjusting the weights so that the ANN learns the relationship between inputs and desired outputs is called *training*. To train a network, some measure of how well it performs is needed. This measure is called the objective function. The input data set is divided into two sets: a training set and a prediction set. A minimization calculation is performed to determine the lowest possible error of the prediction set. There are many methods available for performing the minimization, but by far the most common method is called backpropagation.

Backpropagation is the most common method used to train multi-layer, feed-forward networks. It can be applied to any feed-forward network with a differentiable activation function. Backpropagation works by minimizing a suitable error function with respect to the weights. If the activation function is differentiable, the activations of the output units become differentiable functions of the input variables, weights and bias. It is then possible to evaluate the derivative of the error with respect to weights, and these derivatives can then be used to find the weights that minimize the error function.

5 Implementation of the Neural-Network Based LFS

This section describes the following: the architecture of the LFS neural network, the available data, the nature of the problem, fitting techniques, and the configuration of input data.

5.1 Architecture

The LFS neural network is a three layer MLP as shown in Figure 3.

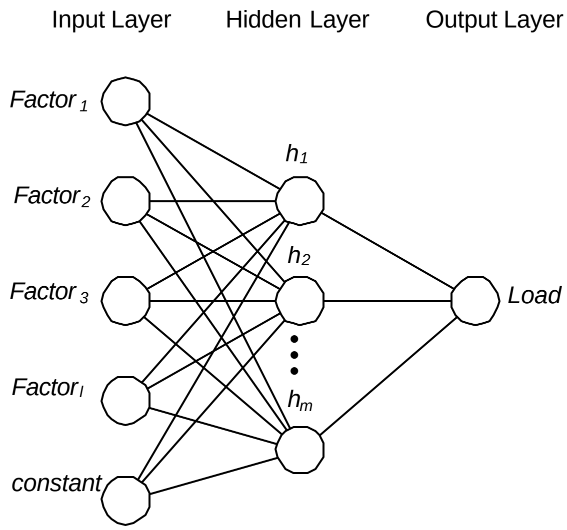

Figure 3: Load Forecaster MLP

The number of nodes in the input layer is based on the number of factors included in the model and how the factors are encoded. The determination of factors and encoding is discussed in the next section.

5.2 Available Data

The most difficult part of creating a good model is selecting and encoding the data for training and validation data sets. Selecting data involves:

- Identifying relevant data
- Identifying data sources
- Determining the quality of the data

There are a number of factors that affect load. The importance and use of information varies on a meter by meter basis, but in general, WGP has determined the following factors are important:

- Weather
- Past Load
- Calendar
- Nominations
- Economics

It is important to note that to use an ANN as a forecasting tool requires all the inputs into the ANN for training need also be supplied to the ANN for forecasting. In selecting factors for training, one must consider whether forecasts of those factors are readily available and are accurate.

5.2.1 Weather

There is a strong observed correlation between weather and gas load, especially temperature and to a lesser extent wind speed. In most cases, as temperature goes down load goes up and vice-versa.

5.2.1.1 Factors

In order to obtain weather data to support the LFS, WGP contracted with a commercial weather service vendor to obtain electronic access to its historical and forecast weather database. The following are a sample of the available weather effects:

- Temperature
- Dew Point
- Relative Humidity
- Wind Speed
- Wind Direction
- Cloud Cover
- Precipitation
- Wind Chill
- Heat Index

5.2.1.2 Locations

Selection of weather locations is based upon the relationship of load demands versus pipeline geography. The weather in an unrelated geographical region may have an impact on another region if economic ties or factors are present linking one region to another.

5.2.1.3 Forecasted Weather

A selection of the weather forecast period and the frequency and criteria for weather updates must be made. The forecast period will determine the upper limit that is possible for the gas load forecast, and the frequency will determine whether improvements in the weather forecast as the current day proceeds are reflected in the load demand forecast.

5.2.1.4 Historical Actual Weather

Historical weather information for a period of 2-3 years for the desired weather effects at the desired cities must be obtained to begin creating forecast models. Generally, the weather vendor provides this information soon after a weather vendor contract has been awarded.

5.2.1.5 Selection of Weather Vendors

A selection of the weather vendor must be made. This involves many factors including the following:

- Price
- Forecast capabilities – frequency, and accuracy
- Time in business

Some large energy companies have existing weather services. Sometimes, the demand forecast results in a reassessment of the existing weather service used by the company. In some instances, a company may purchase information from multiple weather services, and choose between the data supplied for performing the forecast.

5.2.2 Past Load

There are advantages to obtaining historical load from the corporate database versus directly from the SCADA database. The SCADA databases have a tendency to contain outliers, missing data attributed to instrument outages, and excessive noise. Company historical data tends to be *scrubbed* and more suitable to load forecast training.

5.2.3 Calendar

The patterns of gas usage definitely depend upon simple calendar information. Date or day-of-the-year alone is insufficient for providing the information in a manner usable by correlating techniques. Instead, effects must be divided into relevant components such as the following:

- Hour of the day
- Weekend
- Day of the week
- Month of the year
- Daylight
- Holiday
- Daylight-savings

5.2.4 Nominations

Nominations represent the nominated or expected gas demand for a meter. This would seem to be a valuable input that would correlate well with the future gas demand. In fact, for meters at fixed flow rates, or industrial customers, nominations may be the only reliable source of information for the gas demand forecast.

However, for meters driven heavily by factors such as weather and calendar inputs, nominations are not reliable inputs for the demand forecast. Nominations in these areas are typically daily nominal values that may or may not attempt to compensate for minor variations of the weather, and in a majority of cases, the nominations are meaningless.

Therefore, nominations can be used for some meters in lieu of forecasts generated by data-fitting techniques, but they are not used as inputs to a data-fit of gas demand.

In order to obtain hourly values for demand forecast, the daily nominations are optionally distributed using the prior days' hourly load pattern or a previously determined daily load pattern.

5.2.5 Economics

A fundamental problem with the use of economics as forecast inputs are that the future economic values, either for the transmission company or competitors, are rarely available. Retrieval of historical price information for training is just the beginning of the work required to reasonably use economics as a training input. An analysis must be performed of customer zones, competitive and exclusive markets and possible contractual obligations.

The use of economic inputs for training is available but currently not used due to the following factors: difficulty, cost of data retrieval and organization, expected maintenance difficulties and low potential contribution to demand forecast.

5.2.5.1 Prices

Historical and current price information for the various customer zones can typically be retrieved from corporate databases.

5.2.5.2 Competitor Prices

Daily competitor price information can be obtained from on-line sources such as the Gas Daily, but company changes can make this an extremely high maintenance training input. These company changes can be any of the following: name changes, mergers, and splits. In addition, the determination must be made of the markets in which each company is active and in direct or indirect competition with the transmission company.

5.3 Nature of the Problem

5.3.1 Base Load Changes

Changes in interruptible transportation packages either on the transportation pipeline or on competing pipelines result in significant base load changes that must be accounted for in the forecasting technique.

5.3.2 Weather Forecast Inaccuracy

The weather is typically the largest contributor to the forecast of load. The load forecast accuracy is only as accurate as the weather data being used, and the weather forecast is only good for a one to three day period. Even with a reasonable forecast of highs and lows for the day, it is difficult for weather services to accurately predict hourly weather several days in advance.

5.3.3 Meter Group Distribution

In some situations, forecasting gas load must be accomplished at a group level instead of at the individual meter level. For instance, meter flows for several meters all feeding an LDC may add up to a predictable flow given weather and other inputs, while individual meter flows are strictly determined for operational reasons by the LDC.

Distribution of the gas forecast to the meter level can be done in several ways:

- Distribution ratios of the total predicted flow based upon the historical flows over a user-specified period
- Fixed percentages
- Equal percentages

The distribution ratio is more difficult than the other methods but performs better because the flow distribution is based upon real-time data read from the meter flows.

5.3.4 Multiple Weather Locations

Meters must be able to use more than one weather location to account for connections to large LDC networks or interconnects to other pipelines.

5.3.5 Operational "Noise"

Historical load information over a 2-3 year period generally contains outliers, outages, or anomalous behavior and is not pristine data. Sorting through logbooks in operations could be used to eliminate some ranges of data, but this is quite tedious and will not eliminate all bad data. Thus, any data-fitting technique must be robust and not adversely skew the forecast when the data set has reasonable operational variations present. Data that is drastically incorrect such as zero flows for periods where flow is present should not be selected. Periods of time used to fit a forecasting model should be representative of typical pipeline operations.

5.3.6 New Meters

When new meters are added, some mechanism needs to be in place to forecast the meter's load until enough history is gathered to forecast the meter load. A simple and effective method is to temporarily use the forecast from a nearby and similar meter, and to scale the forecast using the relative meter average flows.

5.3.7 Type of Meter

The meter service dramatically affects the load pattern. The following are the primary meter services:

- Residential
- Industrial/power generation
- Pipeline Interconnect

Forecasting gas demand is easiest during the winter for meters that continually flow gas to residential areas, because the residential area loads are largely used for heating of homes and have cyclic daily gas usage during the winter.

Conversely, it is most difficult to forecast gas demand for meters delivering to industrial areas or power generation plants, because these meters have intermittent and irregular flow patterns with little or no correlation to the weather, calendar, or any available input.

5.3.8 Receipt Versus Delivery

Forecast of receipt meters is generally more difficult   Meter flow relationships to the related input variables may change due to seasonal differences in the weather-to-demand relationship and even pipeline operational changes may affect the relationship.

5.3.9 Gas Storage

Meters connected to storage are difficult to forecast because the flows are generally set to satisfy short-term to long-term storage management needs. Also, the flows are often at fixed rates and have intermittent periods of zero flow.

5.4 Data Fitting and Forecasting Techniques

5.4.1 Neural Networks

Neural networks are a generalized non-linear fitting technique as described above in Section 3.

In some cases, data fit can be improved with the use of multiple neural networks combined to forecast an entire year. For example, a seasonal neural network could be configured to use four neural networks to fit the following periods:

1. November - March
2. April - May
3. June - August
4. September - October

5.4.2 Regression Techniques

Prior year load and a linear addition due to temperature difference is a simple forecasting technique where linear regression can be applied. Another example of a correlation for forecast calculation presented in the AGA Gas Engineering Operating Practice Series (page 142, 1992) is as follows:

$$\text{Sendout} = S_0 - 20.9(T) - 7.5(T_r - T) - 2.7(T_y - T) - 1.2(T_d - T) + 12.4(W)$$
$$-0.2(W \times T) - 0.1(R) - 61.9(\sin d) - 0.9(\sin d \times T)$$

where: S_0 = sendout at 0° F (1736.7 MMcf/day $^5/_8$ 0° F)

T = average temperature at official site

T_r = average temperature at remote site

T_y = average temperature yesterday

T_d = average dew point temperature

W = average windspeed

R = solar radiation in langleys per day

$\sin d$ = seasonal effects

Using regression techniques requires prior knowledge of the structure of the input to output relationship. Considering the number of inputs used, the number of primary, secondary and cross terms results in an extremely large number of inputs to the correlation. If such a complex function is not used, then data fit and forecast will suffer.

5.5 Configuration of Input Information

5.5.1 Selection of Data Periods

Data periods used for neural net training should be selected based on seasonal variance and pipeline operation techniques, which may vary throughout a calendar year. This is done primarily for two reasons. First, the characteristic between load and weather/lag data may differ from one calendar season to the next. Secondly, the relationship between load and demand may change when a company changes operational techniques throughout the season. Training data sets should be selected with these factors in mind so that the forecasted load values will not deviate from the actual load values due to varying operational techniques or seasonal variances. These considerations will often explain unexpected results. For example, the forecast accuracy for one set of meters in the summer was better when the neural net is trained on yearly data, while the winter months performed better when trained on seasonal data. In this example, the pipeline operation technique varied from the summer and winter months, while the relationship between weather and load did not drastically change. The summer load values are heavily dependent on the relationship between weather and load demand, while the winter months were affected by the change in pipeline operations.

5.5.2 Delay Effects

The selection of delay effects may either increase or decrease the accuracy of the load forecasts. Using a load lag has the benefit of accounting for baseline load changes into the forecasted values. If a meter's baseline load changes during a year (a power plant may increase its capacity) the load forecasts could be thrown off without using a load lag parameter in the neural network. However, if a meter on a certain day does not follow a characteristic flow pattern, the forecasted flow corresponding to the lag value used during this time may be skewed or inaccurate. The benefits of including a long-term baseline change into the forecast often outweigh the rare occurrences of the flow forecast being skewed due to an anomaly in a load.

The inclusion of a weather lag is often quite beneficial to the forecast values. Correct neural net training could not be managed with the inclusion of a load lag without the corresponding weather lag. Using the weather lag alone may often increase the accuracy of the forecasts. A short-term weather lag delay of one hour is often enough to notice an improvement in forecast accuracy, as the neural net will include the characteristics related to the weather value and the weather value one hour ago. In this case, the neural network will inherently account for load demand relationships based on temperature changes within the lag time. For example, the inclusion of a weather lag will correctly predict the load increase at a power station when the current temperature is 95 °F and the previous hour's temperature was 80 °F. In this case, the load will be slightly higher than if the current temperature is 95 °F and the previous hour's temperature was 92 °F.

5.5.3 Training Periods Versus Prediction Periods

A training data set should represent and contain the relationship between the neural network input data and the meter loads to be predicted. The proper inclusion of prediction periods offers tangible benefits to load forecast training for a variety of reasons. The proper inclusion of load forecast periods in training a neural network means that the forecast period must not be included in a training set of data. Otherwise, an irrelevant forecast accuracy will be attained, as the neural network would have included the prediction set patterns into the training set. A valid forecast accuracy may only be attained when the prediction set is not included within a training set.

The inclusion of a prediction period within the training set allows the neural net to give the user an indication of the forecast accuracy. More importantly, the prediction period allows the user to get an indication of how the inputs to the neural networks affect the load forecasts for a meter or meter group. By generating multiple training sets with different inputs, the user may compare the forecasted loads to the actual loads during the prediction periods. This offers the immediate advantage of being able to determine the effect that the various inputs to the neural network have on the accuracy and patterns of the load forecast.

5.5.4 Structure of Data Inputs

The quality of the data fit for the forecast depends, to some degree, on the way the input data set is structured. Two examples are as follows.

Example 1. Temperature and wind speed versus temperature and wind chill.

Example 2. Day-of-week represented as a number from 1-7 versus a continuous sine wave.

5.5.5 Data Scaling

The neural network weights in the input layer essentially serve as input scaling values. The learning rate of a neural network is generally increased when the inputs are pre-scaled between minus 1.0 and 1.0 before going to the neural network. Different scaling values may be selected, as the neural network performance may be “tuned” to an optimum value.

6 Evaluation of Results

6.1 Accuracy Criteria

Accuracy of the load forecast is calculated as follows:

$$\text{Accuracy} = 100.0 - \text{Error}$$

Error can be calculated in many different ways; the typical definition of error is as follows:

$$Error = 100 \times {}^{1}/_{N} \times \Sigma_i \left[|Predicted_i - Actual_i| / Actual_i \right]$$

This accuracy definition is problematic for low flow meters or meters that will occasionally go to zero. Even if points of zero flow are eliminated, errors at low flow rates mask the errors at higher, more significant flow rates. Thus, a more robust and meaningful accuracy statistic for flow load forecasting is a flow-weighted average error calculated as follows:

$$Error = 100 \times {}^{1}/_{N} \times \Sigma_i \left[(|Predicted_i - Actual_i| / Actual_i) \times (Actual_i / \Sigma_i Actual_i) \right]$$

After rearrangement, this calculation becomes the following:

$$Error = 100 \times \Sigma_i (|Predicted_i - Actual_i|) / \Sigma_i Actual_i$$

This accuracy definition has three useful characteristics. First, the total sum of the actual flows for a training set does not approach zero in any meaningful data set. Second, the errors on the higher flows are given more weight and when the flows are high, the weights are relatively equal. Finally, the calculation is easy to implement.

6.2 Importance of Prediction Accuracy

If the data in the training set does not fit well, it is unlikely that the data forecast will be any better. However, although the data fits very well during training, if proper precautions regarding prediction accuracy are not used, then the accuracy of the forecast may be quite poor.

Poor prediction accuracy with good training accuracy is a clear symptom of overfitting data.

6.2.1 Season of Year

The relationship between the demand for flow and neural network inputs may change slightly with the seasonal variances in the flow demand to the neural network inputs. This behavior is readily apparent when the flow forecast is inaccurate for a specific season.

6.2.2 Training Accuracy

The training accuracy is affected by the characteristics of the neural network. The user may improve or degrade the neural net accuracy by altering the selection of input variables or training parameters. The training parameters include the available calendar effects that are relevant to the forecast parameters (holidays, day-of-week effect, hour-of-day effect, weekend effect, month-of-year effect, holiday effect, daylight effect, etc.) Additionally, the choice of lag effects (weather lag and load lag) may also greatly affect the accuracy of the load forecast.

6.3 Sensitivity Analysis

6.3.1 Significance of Inputs

A great benefit of applying neural networks to load forecasting is that the relationship between load and the relevant input variables to the neural network does not need to be known ahead of time. The neural network may potentially become a black box of mathematical manipulation to the user. In order for the user to get an indication of feedback from the neural network of the relative relevance that the input variables have on the forecasted loads, the neural network must have the ability to perform a sensitivity analysis and provide the results to the user in a format that makes sense. By performing a first order sensitivity analysis, it is possible to determine the amount of effect that each input has on the forecast. The sensitivity should be normalized so that the summation is 100% to make the data meaningful. This is a great benefit from neural networks because this is a means of determining the inputs that are beneficial to the forecast. Any input that has a low sensitivity to the forecast value may not need to be included in the inputs to the neural network. This knowledge saves time and makes computations in training and performing the forecast more efficient.

The weights of input connections to a node within a neural network determine the significance that each input has on the output value of the node. Essentially, the greater the weight, the more significant each input is to the output. The results of this type of sensitivity analysis for an LFS training set are shown in Figure 4 below.

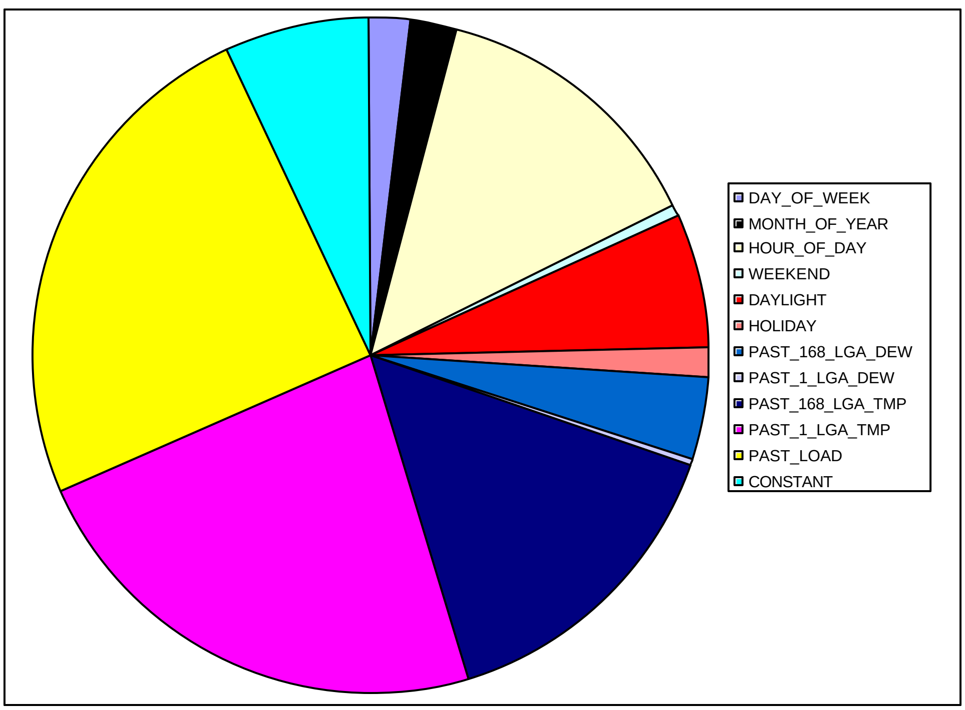

Figure 4: Sensitivity Analysis Pie Chart

6.4 *Overrides*

The inclusion of meter overrides is necessary in any professional installment of a load forecaster. A user may apply a constant bias to a meter to account for offsets in predicted flow or a percent bias to account for a variation in flow peaks for a meter. The use of overrides is necessary to account for known operational changes at a meter station.

6.4.1 Operational Changes

Expected anomalies in operation may affect the load forecast accuracy, either for particular days or for extended time ranges. It is often necessary to retrain the neural network for performing the forecasts on a particular meter with recent data if anomalies are detected for an extended amount of time. This is often indicative of a change in the neural-network-inputs-to-meter-flow relationship or even to operational changes of the pipeline transmission company.

6.4.2 Data Fit

Adjustments should be made to account for observed trends in deviations of the forecast from actual flows. The user may wish to adjust either the selection of neural network inputs, or to make changes in the neural network parameters to account for observed variances in the forecast.

7 Considerations

The following section discusses important aspects of the application implementation and deployment.

7.1 Automated Forecast

Regardless of the data fitting technique selected, some mechanism must be created to automatically process information, and produce a load forecast.

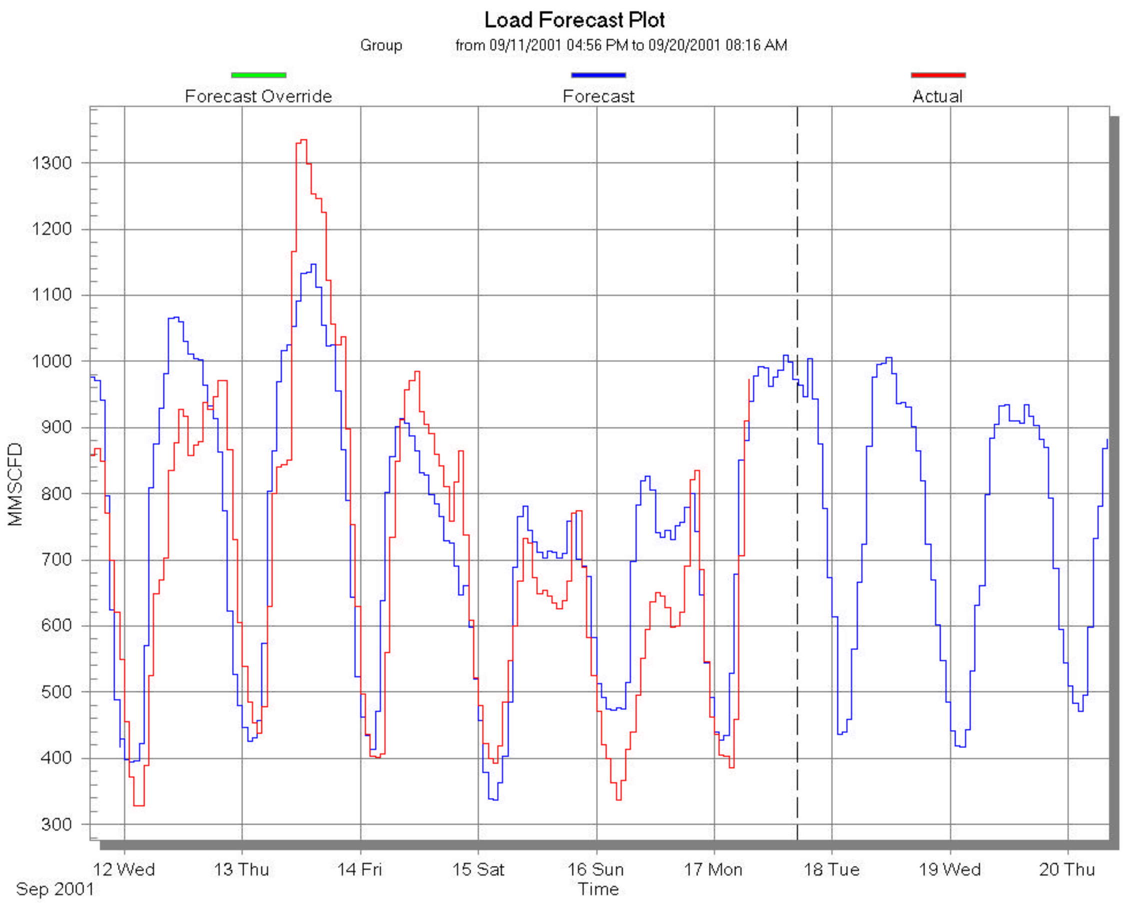

7.2 Data Setup

The most time consuming process in training a neural network is acquiring the relevant input data for training. The input data used for training must be complete for the entire range of the training period for each input variable selected as part of the neural network. Once this data has been attained, it should be stored in a format in which a neural network may read and use the data as needed. Relational databases are well suited for this type of data storage and use.

Once inside a relational database, the input data may be used and manipulated by the neural network in any user-selectable combination that is relevant to a load forecaster. In addition, a means of obtaining current and forecasted weather data and storing them for use by the load forecaster is essential. Automating this portion of load forecasting may save hundreds of man-hours of data entry each year. Several data sources are commercially available for acquiring current and forecasted weather data.

7.3 Neural Network Selection

The neural network selection has a direct impact on the accuracy of load forecasts and may be automated by a professional installation of a neural network. The automation of forecasts is made applicable when the neural net software automatically generates forecasts for several meters. To do this efficiently, the load forecasted must use the neural network parameters obtained from a training set to regenerate the network for application in forecasting. Thus, it is essential to store the neural network parameters in a meaningful and useful format. Again, relational databases are well suited for this purpose.

Due to the wide applicability of neural networks, it is necessary for the user to manually relate a certain neural network to a meter or group of meters. The use of an interface to the neural networks and meters greatly eases this task.

7.4 Forecast Calculation

The values from the load forecast application are the results of applying the neural network to predicted and current data. It is necessary to store these values in a format that is easily manipulated and that may be retrieved easily by a variety of applications. While the load forecasts alone may provide immediate value and benefit to an energy company, integrating the forecast values into additional software such as predictive modeling provides greater value and incentive for integrating a load forecaster into a company's daily operation.

7.5 Historical Update

Assessment of the continued accuracy of the trained neural networks requires the load forecast to be recalculated when the actual weather becomes available. This provides the quality of the data fit if the weather is 100% correct.

7.6 Data Handling and Storage

All of the information used or produced by the load forecast application is stored in an Oracle relational database. The advantages, flexibility, and general access capabilities provided by the relational database far outweigh the additional cost and maintenance.

7.7 Processing

CPU and memory requirements must be considered when establishing a platform for a load forecaster. The neural network training involves solving a highly CPU-intensive mathematical algorithm that is well suited for today's fast CPUs. The necessity of automating the data retrieval and storage of the neural network parameters, training data, and results often requires increased CPU and memory requirements in addition to that required by the load forecaster. Typical installations of the load forecaster involved a dual 500 MHz Pentium III server with one gigabyte or more of memory.

7.8 Presentation of Results

The load forecasts must also make the forecasted values easily accessible and presentable in a meaningful manner to the user. By graphically representing and comparing the actual flows to the forecasted loads, the user has a means of efficiently comparing and using the load forecaster. The automation of presenting the load forecasted values is of the greatest benefit to energy companies, as the quick and efficient retrieval of this information may be used to make business and management decisions that keep a company running and profits up. The increased need for companies to make faster and more accurate decisions, which may be attributed to growing deregulation in the energy industry, makes the load forecaster an essential and immediately beneficial tool for business operations.

7.9 GUI Tool

The amount of time required to setup, train, and attain forecasted results from a neural network is greatly reduced with the use of a user-friendly GUI. A user should not be allowed to establish a neural network configuration that is not feasible. In addition, a GUI should present the neural network data and computed values in a manner that is easy to use, manipulate, and comprehend to an ordinary user. These aspects are essential in providing a utility that a company may use in its daily operations to make faster and more accurate business decisions.

7.10 Web Pages

The availability and wide spread accessibility of the world wide web makes it a suitable place to interface with load forecasting for large companies with economic load interests in more than one region or building. A load forecaster with the integrated ability of storing results in a manner easily viewed on the internet has extended utility and value for larger companies, allowing management to make decisions based on immediate and current data utilizing one of the most efficient and economical means available.

7.11 Aggregation and Filtering of Forecasted Loads

In regions where a company supplies multiple deliveries with similar flow characteristics, applying a neural network to a group of meters instead of training each meter individually may save time. In addition, in regions where meters have varying flow patterns, the load forecaster application should have the ability to aggregate the totals for each meter. This will allow the user to easily view the significance of load in particular regions efficiently and in a manner that holds the greatest amount of physical relevance to the user. This gives the user confidence in having the complete knowledge necessary to make decisions based on data presented in one location. This saves the user time in using data, and also allows the forecasted results to be presented clearly and meaningfully.

7.12 Maintenance

The automation of neural network software and data gathering allows for a minimal of maintenance activity. In addition, automated maintenance procedures behind the scenes may be easily provided when a commercial relational database is used for data storage and manipulation. In addition to being readily available and understood, various tools are available for diagnostic information, such as memory/data usage and processing requirements.

7.13 Monitoring

The load forecaster should include automated accuracy analysis and notification. When the forecast accuracy for a meter falls below a predetermined level, a list of users may be notified via email that this event has occurred. A user may then make the necessary changes to the neural network, retrain the network, and achieve the desired accuracy level. This type of notification is essential so that critical decisions are not made based on inaccurate data.

7.14 Retraining

The need to include additional data and retrain a neural network is often required by a user. The ability to do this in a friendly manner with a minimum of effort is an essential part in the usability of any load forecast application. The most often-used features of forecasting are made as efficiently and as friendly as possible to provide an integrated mechanism and tool to provide the most amount of value to a company in the easiest manner.

7.15 Integration

A load forecasting tool obviously provides immediate benefit, value, and utility to an energy company. In addition to having the data make faster and more accurate decisions, management acquires a tool that may provide accurate data for future planning.

The benefits of handling the load forecast data with a relational database provides increased value and benefit when used with the combination of other modeling tools. Use of data by predictive modeling is an obvious and tangible benefit. Using the load forecaster with predictive modeling may allow an energy company to have accurate access to additional data such as available capacity and pipeline linepack up to several days in the future. In addition, the integration of load forecast values into survivability studies might make critical differences in management decisions.

8 About the Authors

Paul Lamb, B.Sc., is currently Project Manager for the Integrated Gas Management System (IGMS) at Williams Gas Pipeline in Houston, Texas. Paul has worked in the pipeline simulation field for various companies, including Southwest Gas Corporation, Stoner Associates, and Williams Gas Pipeline-Transco. Paul received an Engineering Science (B.Sc.) degree from the New Mexico Institute of Mining and Technology in 1986.

Dan A. Logue, M.Sc., B.Sc., is currently an officer of and Technology Advisor for Energy Solutions International, Inc. In 1995, Dan founded Wright, Logue & Associates, which later merged with LICEnergy in April 2000 to form Energy Solutions. From 1986 to 1995, Dan developed and implemented optimization and Multivariable Process Control (MPC) at Setpoint, Inc. Dan received his Chemical Engineering degree (M.Sc.) from Louisiana State University in 1986 and his Bachelor's in Chemical Engineering (B.Sc.) from the University of Houston in 1984.