LEVERAGING LOAD RESEARCH AND CUSTOMER CONSUMPTION DATA TO IMPROVE CIRCUIT UTILIZATION

Rick Brakken
Nexus Energy Software, Inc. - USA
rbrakken@nexusenergy.com

Taj Ait-Laoussine
Nexus Energy Software, Inc – USA
taitlaoussine@nexusenergy.com

ABSTRACT

This paper discusses a methodology for integrating customer consumption data with GIS and SCADA data to optimize circuit utilization, prevent device overloading, and improve load-flow models, while leveraging sophisticated modelling techniques and software systems that were originally developed for load settlement operations. The paper features a case study of results and benefits achieved by a mid-sized US utility.

INTRODUCTION

Distribution planners and operators have traditionally struggled with the trade-off between reducing distribution costs and improving reliability. One way to reduce costs is to increase system utilization by "right-sizing" devices and deferring system upgrades as long as possible, thereby decreasing the installed capital cost per kWh served. Conversely, a common way to improve reliability is to "gold-plate" the system, increasing capacity to protect against circuit and device overloads, even if subjected to high future load growth. A key contributor to this seemingly irreconcilable conflict is the uncertainty of loading conditions and growth rates at network points downstream of the substation. This uncertainty creates the need to add capacity margin into system design and sizing decisions, thereby increasing cost. Yet, even with this capacity margin, load uncertainty can still result in premature overloads or device failures, resulting in outages and increasing capital costs even more. Acquiring accurate resolution of where load is distributed along a feeder and how this load varies by season, day, or even hour would enable planners and operators to increase system utilization, while at the same time identify and mitigate potential overload or voltage problems.

The question is how does one go about acquiring more accurate load data and incorporating this data into distribution planning and operating processes, and is the benefit worth the effort and expense of doing so? Even with the availability of lower cost metering technologies, the expense and difficulty of directly metering all points within the distribution system and managing the resulting data often makes this approach impractical or unfeasible. The benefits usually do not justify the cost.

Another approach is to combine whatever existing metering data that may be available with customer load data collected for billing purposes, using modelling and aggregation techniques to estimate load at any point in the network. In the past, achieving accurate results with this approach was difficult due to issues in customer data quality and validation, lack of knowledge about where customers were located on the network, difficulty in processing large amounts of data, and a multitude of other issues. However, recent advances in load modelling and forecasting, customer data validation and processing, and utility deployments of GIS have made this approach much more feasible. This paper summarizes the methodology of such an approach, as well as the results and tangible benefits achieved by a utility in the United States. that deployed this approach in a fully integrated system that interfaces the company's GIS, CIS, and distribution planning models. The system has significantly improved decision-making and provided a positive return in the first year that paid for the entire cost of the system.

MAXIMIZING THE VALUE OF EXISTING LOAD INFORMATION

With the emergence of deregulation and competitive energy markets, distribution companies are facing increasing pressure to deliver energy with high reliability and at a low cost. Through mechanisms such as performance-based ratemaking and customer choice, regulatory agencies are also encouraging the operators of distribution networks to seek new and innovative ways to cut costs while maintaining reliability. For competitive distribution companies, this means improving the utilization of their assets to ensure that every single dollar of investment is generating its highest possible yield.

In order to achieve this objective, distribution companies clearly need to acquire a better understanding of loading conditions beyond the substation. As discussed earlier, an obvious approach is simply to collect more information: if every customer or every line transformer is equipped with a meter that captures hourly variations in load, then distribution engineers can readily know the actual distribution of load across their network. This ubiquitous metering approach, however, is not necessarily the ideal solution and for several reasons:

1. It generates a very large amount of information, which

can be more confusing that enlightening – especially if the utility lacks the sufficient processing power and capacity to handle all that data;

2. It is not always reliable, given that all meters are subject to failure or incorrect reads;

3. It can be very expensive, particularly in larger service territories;

4. By itself (i.e. without strong analytics) it can only represent the past rather than help forecast the future.

Instead of ubiquitous metering, many utilities are electing an alternate route: they are combining existing (or strategically enhanced) metering levels with new software technology that allows them to gain a better insight into the existing meter data they possess. After all, utilities are collecting extensive information about customer behaviour through their metering system. Indeed, the consumption of each and every single residential customer is collected and recorded every month. Larger commercial customers are often equipped with meters that not only record their monthly consumption but also track their maximum monthly demand. And the largest of customers are usually equipped with meters that continuously collect hourly average demand. Utilities therefore already have, as a general rule, access to large amounts of load information. Why not maximize the value of that information by incorporating it into the distribution network planning and operation process? This is where the software technology comes in.

Most distribution companies already collect and store the following information about their customers:

Customer Characteristics. For each customer connected to the distribution network, utilities often store a variety of customer characteristics (in addition to the basic billing demographics) such as rate class, SIC code, weather zone, network connection point, type of end-use, etc.

Monthly Consumption. In order to bill most customers, utilities usually meter their monthly energy consumption.

Monthly Maximum Demand. For a smaller set of customers (usually medium-sized commercial and industrial customers) utilities also often collect a monthly maximum demand value, in addition to the monthly consumption value.

Interval Consumption. For certain large or special customers, utilities meter the consumption on more regular intervals, usually on hourly intervals.

Load Research Data. Finally, utilities often collect load research data, which consists of interval load information for a sample of various customer types. Load research data it is usually collected either for rate-making purposes or for the design and implementation of energy efficiency programs.

The newer software technology capitalizes on these data sources to help utilities accurately monitor the utilization of their distribution assets, ensure their proper and proactive maintenance, and predict where and how to enhance their network to maintain a high level of reliability.

SOFTWARE METHODOLOGY: THE DATA INTEGRATION APPROACH

To be effective, the data integration approach must employ a methodology that comprises five steps: (1) Segment, (2) Profile, (3) Forecast, (4) Aggregate, and (5) Analyze.

Segment

The first step is to segment the customer base. The objective of segmentation is to define various customer groups within which customer consumption patterns are comparable. Some segmentation lines are obvious: residential customers are likely to behave differently than industrial customers; likewise, customers in affluent neighbourhoods tend to consume power differently than customer in more marginalized areas. Other segmentation lines are not so obvious and require a careful statistical analysis to determine their relevance.

The segmentation component grabs all available customer characteristics and load research data to test a variety of different segmentation options. Each tentative customer segment is analyzed to determine its fitness and relevance across several dimensions.

Segmentation is by nature a very complex and computationally intensive process. Some software, however, can streamline the entire approach by allowing users to rapidly test and approve multiple segmentation hypotheses.

Profile

Profiles are essentially complex and highly accurate models for allocating non-interval loads across hourly intervals. Using the available load research data assigned to each segment in the segmentation step, the technology generates models for how the segment is expected to behave under different conditions and at different times. Profiles enable the software to calculate the expected load of each and every customer, and for any hour of the year, without having to know the actually hourly consumption.

Forecasting

One of new software's key strengths is the ability to define weather-sensitive profiles that rely on hourly temperature inputs to calculate hourly loads. The weather-response regression functions generated in this process can accurately predict any customer's demand given a specific set of weather conditions. Thus, the software can actually forecast the load of any customer or customer group given a specific weather forecast. In fact the forecasting capability is not limited to weather, though it tends to be the most often used approach.

Aggregate

The software aggregates profiled and forecasted load according to highly complex aggregation schemes, such as those used in modelling distribution circuits, and can calculate the load at virtually any point on a distribution network, effectively bypassing the need for ubiquitous metering.

Analyze

The last step, analysis, allows the user to view and act on the results. The types of results and analysis made possible by such software are illustrated in the set of figures that follow. Figure 1 presents an example a historical load trends, whereby a user can view results for a given point in the distribution network for a particular timeframe, specific day types and/or weather conditions, and one or more phases. This type of report is ideal for understanding how load varies over time, isolating extremes or uncovering load imbalances. For instance, users can request to view the load at a particular point on a feeder, and obtain the curve of maximum demand, the curve of minimum demand, and the average load curve. Understanding the range of fluctuation of a particular point can be critical to implementing a planned outage, for example.

Figure 1: Historical Load Trend Analysis

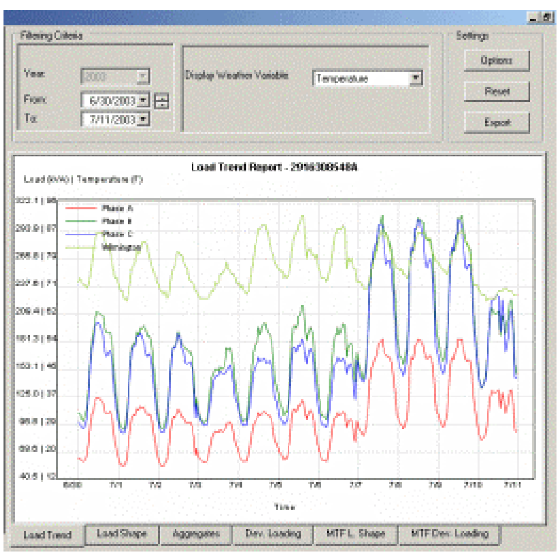

Figure 2 shows a 24-hour load-shape report that displays the maximum, average, and minimum historical load-shape at any point in the distribution system during a user-specified time of year, day type, and weather conditions

Figure 2. 24-Hour Load-Shape Analysis

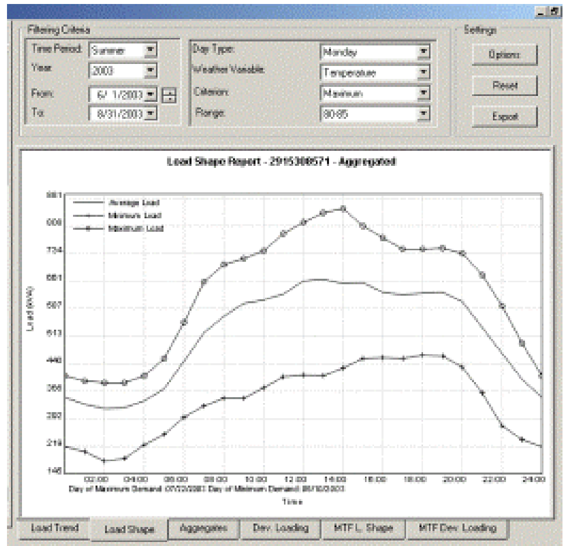

Figure 3 shows a report whereby a user can view lists of selected devices and their characteristics and calculate loading statistics such as average utilization, peak loading, and overload duration. Typically, all results can be sorted and filtered on the fly to enable users to quickly pinpoint equipment that has been problematic or that may be subject to failure in the future. In this example, the report displays line transformers on a particular feeder that have been subject to overload conditions during the previous 6-month period, along with the percent of time that these transformers were in an overload condition during this time. This type of analysis is particularly useful for transformers and other devices where risk of failure and loss of life is a function of not only overload magnitude, but also of overload duration.

Figure 3. Device Loading Analysis

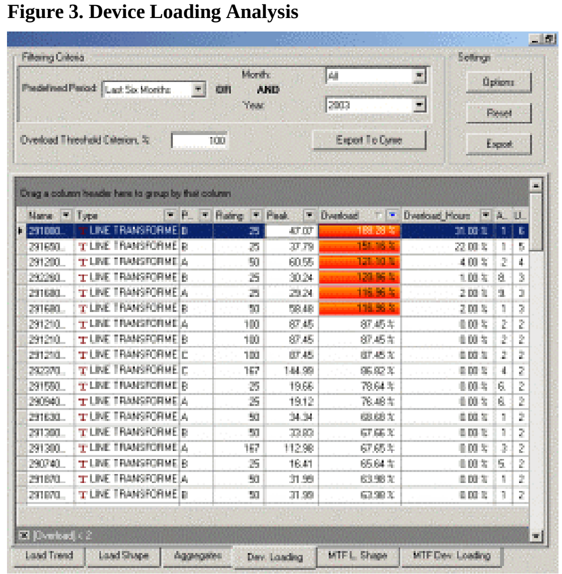

Figure 4 shows a mid-term forecast report. This report provides weather-normalized load forecasts within the next 12 months at any network point based on multiple weather scenarios for consistent risk-based planning.

For example, load forecasts could be created for up to five different weather scenarios, such as extreme heat (e.g. one-in-25-year scenario), semi-extreme heat (e.g. one-in-10-year scenario), average, semi-extreme cold (e.g. one-in-10-year scenario), and extreme cold (e.g. one-in-25-year scenario), giving planners and operators an accurate view of system conditions that might be encountered under extreme, but realistic weather conditions, even if recent historic conditions have only been moderate or mild.

Figure 4. Mid-Term Forecast Analysis

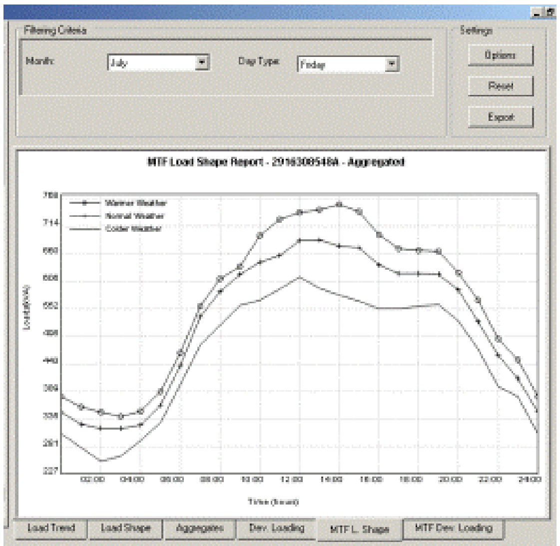

Figure 5 shows a load decomposition report that would enable users to identify the source of loading at any point within the network, from different customer segments or rate classes, down to individual customers. With such information, distribution planners and operators could identify the source of unfavourable or "spike" load patterns that may be driving the need for system capacity upgrades so that more cost-effective alternatives may be identified. These alternatives may include re-allocating or balancing load on the feeder, or perhaps offering customer incentives to reduce demand spikes such as demand response pricing or load control programs (particularly if the unfavourable load pattern is driven by a few select customers).

Figure 5. Load Decomposition Analysis

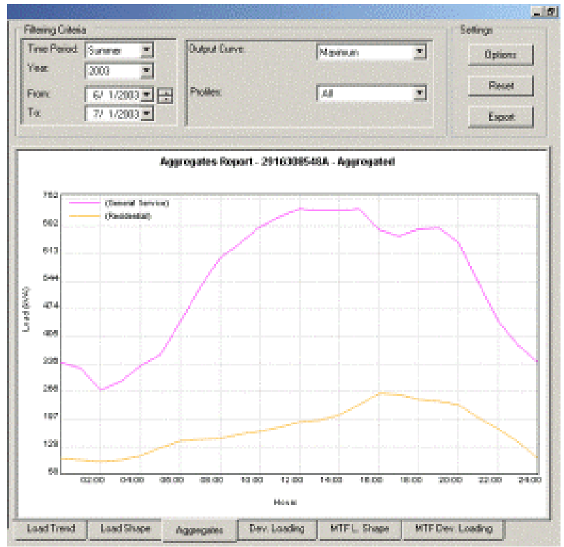

METHODOLOGY IN ACTION: CASE RESULTS

In the winter of 2003, a mid-size US utility serving about 400,000 customers in the Northwest United States, was facing the ubiquitous challenge of improving the reliability of an aging distribution infrastructure amid increased regulatory controls and financial pressures. They had been making incremental improvements of circuit-level utilization over the past decade, but had reached the limits of what planners thought they could safely extract from the system. They therefore decided to embark on a pilot implementation of a software solution based on the methodology described above.

In the pilot implementation, five circuits were analyzed. The utility company provided load research data for the circuits, which the software segmented according to rate class, load factor, winter ratio and circuit—a total of 22 different segments were created in all. From there, user analysts developed 21 four-season, three-day profiles (weekend, Saturday, and Sunday/Holiday). All profiles, except for one—irrigation—were weather sensitive, based on weather data from local weather stations. The software analyzed customer billing data provided by the utility, assigned each customer to a segment and profile, and then calculated the hourly loads for each individual customer. The utility provided a GIS-based connectivity model for the pilot circuits so that customer load could be aggregated up to the circuit level.

The results of the analysis provided valuable, and in some cases, surprising, information on network load locations and phase balancing. The utility input the calibrated load data from the software model into their load flow modelling application, so that it could compare its conventional load allocation process (which uses the connected kVA amounts to allocate load) with the new approach. The analysis showed significant geographic variations in the allocation of load over time, as well as substantial changes in estimated voltage levels. The results helped the utility understand how load moved across its circuits (especially longer circuits) during peak and off-peak days, and ultimately provided a better representation of asset utilization on different circuit segments. In some cases, areas that were thought to be overloaded were shown to be more lightly loaded, while other areas exhibited the opposite pattern

Figure 6 illustrates the load reallocations for a portion of one of the pilot circuits.

Figure 6. Load Reallocation—Partial Circuit

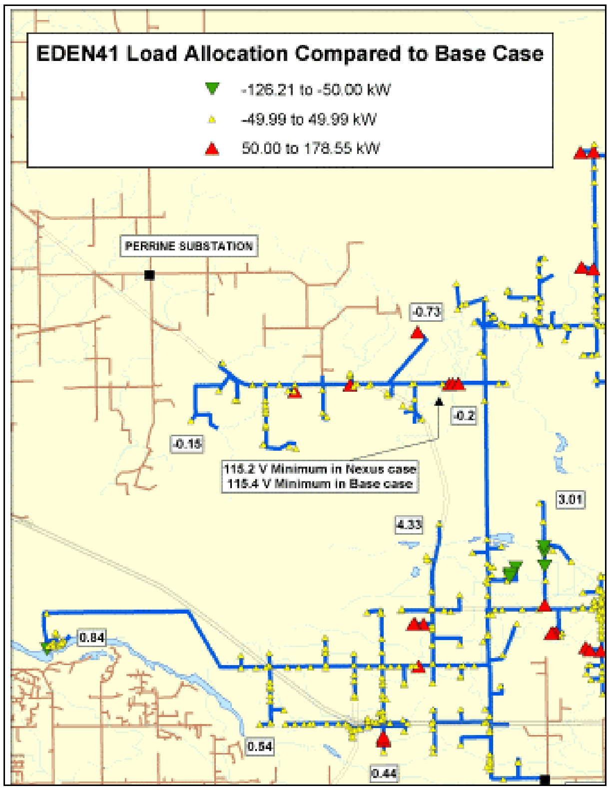

These results had a significant impact on this utility's allocated capital budget for the five circuits. Of ten projects scheduled, four were deferred—one by as much as five years—and one project was accelerated, representing an overall deferment of 46% of the 2005 capital budget for the pilot area. The changes result in an overall net present value savings of 12.4% of the 2005 budget for these circuits, when analyzed over time.

Based on the success of this pilot implementation, the utility has decided to implement this software for nearly all of its distribution circuits, representing about 25,000 miles of feeder lines. The implementation will first target those circuits slated for significant capital investment within the near term, and will incorporate input and output interfaces from the software to the utility's load flow modelling software, SynerGEE, to facilitate use the data to model load flow on circuits prior to future capital allocation processes. The full implementation will also include development of interfaces to all other relevant data systems (CIS, GIS, SCADA, etc.) to enable automatic updates. To improve the accuracy of the profiling models, additional levels of customer segmentation will be analyzed.