

Advanced Load Profiling Techniques for Energy Management and Procurement

*Lindsay Audin, CEM, CEP
President, Energywiz, Inc.*

ABSTRACT

Understanding how energy is used by a facility, both retrospectively and in real time, is assisted through analysis of interval data showing demand in time periods much shorter than may be seen on a monthly utility bill. Metering and software to simplify that task are rapidly advancing, and several Excel-based techniques are now available to make that task even easier.

INTRODUCTION

Data covering energy usage in short time periods (typically 1/4 hour to an hour) are called "interval data." While providing much greater detail than the usual monthly usage and peak demand data found on most commercial and industrial utility bills, handling of such information (which may involve over 35,000 data points per year) can be challenging. Using a modern personal computer (i.e., 700 MHz Pentium III or faster) equipped with Excel 2000 (or more recent versions) can, however, quickly digest such files and convert them into two- and three-dimensional (3D) charts, such as seen in Figure 1.

To disseminate some of these techniques, AEE (www.aeecenter.org) offers an online course, and (in early 2004) the New York State Energy Research and Development Authority (NYSERDA) sponsored a series of live seminars on them. Many large end users and energy practitioners are now applying the techniques for their own use. Further information on the basic 3D load profiling methods may be found in the December 2001 and September 2003 "Tips of the Month" at www.energybuyer.org.




Figure 1

These methods have been advanced to help end users better understand how their energy usage is impacted by such variables as outdoor temperature, operating conditions, and equipment choices. Such understanding is essential for successful participation in demand response programs, load shedding, power procurement, and other potentially profitable ventures. To that end, this article reviews several new procedures for charting energy-related information:

- 3D profiling of hourly temperature.
- 2D plotting of power demand and temperature together on one chart to show their relationship.
- 3D profiling of hourly real-time electricity pricing in a deregulated market.
- 3D profiling of hourly cost (i.e., usage multiplied by price on an hourly basis).

- date-embedded scatter plotting using 3D and pivot charting methods.

All techniques and macros covered in this article were developed by the author through his firm, Energywiz Inc., of Croton, New York. Note that older versions of Excel (i.e., prior to Excel 2000) may not have all the capabilities discussed herein.

3D Temperature Charting

While the general relationship between temperature and cooling/heating loads is well understood, knowledge regarding how temperature varies from hour to hour, day to day, and across a full year is often little more than anecdotal. When an electric bill showing the monthly peak demand is received, some thought may be given to how hot it was that month, but few energy managers check to see if that same (or higher) temperature occurred on other days that saw the same (or lower) demand. The ability to see the patterns of hourly temperature in 3D across a month, season, or year at one glance, however, makes comparison to 3D plots of hourly demand relatively easy.

Hourly load data may be obtained through utility metering and/or an energy management system (EMS). Hourly weather data are available from the National Climatic Data Center (NCDC) at the National Oceanic and Atmospheric Administration (NOAA), a federal agency, or from an EMS equipped with an on-site weather station. Data from federally-operated weather stations (scattered across the nation) may be purchased via www.ncdc.noaa.gov, or from a private firm that re-formats such data into more convenient arrays (e.g., www.weatherbank.com).

A 3D temperature profile may be created using Excel's 3D area or surface charting techniques in the same manner as done for electric demand. Figures 1 and 2 are 3D area charts for (respectively) load and temperature for one month. The large bump in Figure 1 occurred when an extra chiller and its auxiliaries were turned on and allowed to run unnecessarily for a full 8-hour work shift.

When these 3D area charts are placed as objects into their respective Excel data table worksheets, touching the cursor to anomalies (such as the bump) reveals the day, date, time, and demand (or temperature) at the point that is touched. The row containing the interval data for that day is also automatically outlined, thus allowing (as discussed in

the next section) easy and quick creation of a 2D chart showing simultaneous demand and temperature.




Figure 2

2D Profiling of Hourly Temperature and Load

Excel allows the use of two Y axes to show the variation of two variables against a third (such as demand and temperature both graphed against the same time scale).

To create a 2D chart with two Y axes, use Excel's "secondary value axis" capability. Group the rows of data for temperature and demand in the same data table and chart them both against time. In the resulting chart, the temperature points will initially be so close to the X-axis as to not be visible. Place your cursor on one of the temperature points and right click on it to see a list that includes "format data series." Click on the "axis" tab and choose the "secondary axis" radio button. A new axis (on the right side) of the chart will be created showing temperature on a scale that makes the temperature curve readily visible.

A 2D two-axis chart (see Figure 3) was used to highlight a problem with the training of a chiller plant operator who erred by running an unnecessary chiller. The original usage anomaly (found while reviewing the 3D profiles of hourly weather and load data in Figures 1 and 2) was highlighted by the unusually high 1/4-hourly consumption seen during an 8-hour period of one day.

The large hump seen in Figures 1 and 3 revealed a sudden rise (and fall) of about 500 kW coincident with the shift changes at the chiller plant. When queried regarding this result, the operator indicated that it was an unusually hot day (quite plausible for June), so all chillers were turned on. As may be seen in Figure 1, however, no other day in the same month saw operation of the extra chiller. On all days, room temperatures were maintained at the same level.




Figure 3

By charting hourly temperature against time (the thin graph lines in Figure 3), it was possible to see quickly a later June day that was somewhat hotter (both at peak and across the day). By charting both that hottest weekday in June and the subject weekday, along with the loads seen on both days (the thicker graph lines), it was obvious that the extra chiller did not need to run that month.

To avoid possible occupant cooling complaints, some chiller operators have developed a bad habit of simply turning on all plant equipment early in the day and letting it run, regardless of actual outdoor temperatures. In this case, that habit cost (at New York City rates) an extra \$20,000 in unnecessary demand charges on that one day.

3D Power Price Profiling

In deregulated power markets (and where some very large power customers are permitted to access wholesale markets), power pricing is developed every hour. Customers may purchase power on a real-time and/or day-ahead basis using such hourly pricing. Doing so sharpens sensitivity to market-based pricing while creating profitable options for peak shaving when power prices are unusually high.

A 3D profile showing how such pricing varies over days and hours during a year may be created using the same Excel charting method as used to create a 3D temperature or load profile. Hourly market pricing may be available through the local ISO, some power marketers, or (perhaps) your utility.

Depending on the mix of generation feeding into a power pool, hourly pricing may be quite sensitive to temperature and dispatch of generation and transmission systems. It is not unusual, for example, to see prices jump by 50 percent to 100 percent (or more) when temperatures rise and are sustained for several hours, and/or when a low-cost generator (e.g., coal or nuclear) or transmission line is suddenly shut down, forcing use of more expensive peaking power plants using natural gas or diesel fuel. Where a significant portion of power is produced using natural gas, a sensitivity to very low temperatures (when natural gas use for space heating competes with winter power generation, driving gas prices upward) may also occur.

Assessing the economic viability of demand control and/or on-site generation under hourly pricing regimes requires some knowledge of the price sensitivity of markets to temperature and other variables. While expensive computer models are used by power marketing firms and utilities for this purpose, any Excel user may get a good sense of the vulnerability of the pricing of his local market by comparing hourly pricing to hourly temperature.

Figure 4 shows how day-ahead hourly pricing in the NY City area varied during calendar 2003. Note the two sharp peaks (in early March and mid-August). The first is due to very low temperatures at a time of very high natural gas pricing (thus forcing up the price of gas-fired generation) even though the peak system-wide power demand was well below peak power supply. The second peak is due to high temperatures, causing a high system-wide electric demand for cooling, thus forcing use of the most expensive power sources. Since (under typical ISO rules) all generators called upon to run are paid the same hourly price

for their power, the market clearing price for power (at which supply meets demand) is quite high compared to the rest of the year.

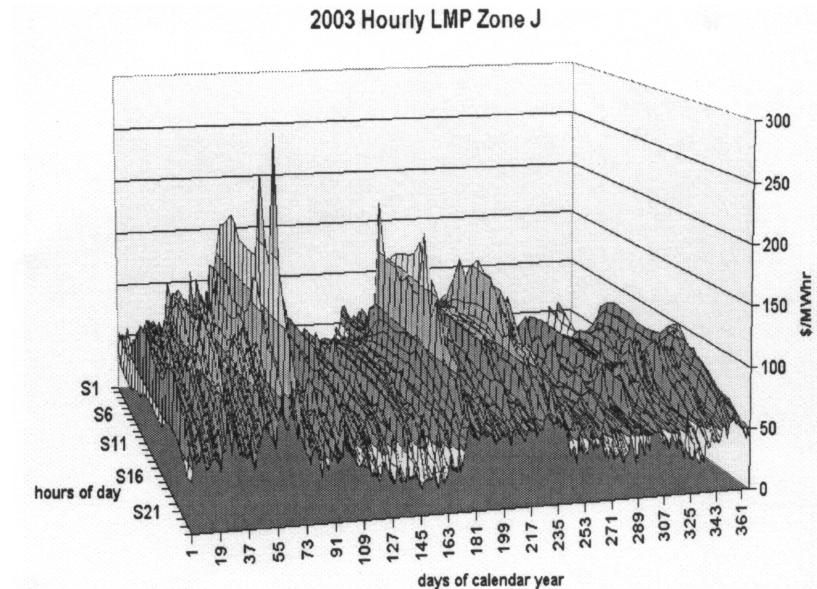


Figure 4

Note, however, that power prices for the remainder of 2003 remained significantly lower. A customer considering installation of a peak-shaving generator must consider the possibility that his on-site power plant may only run for a few hours on a few days in such a year, thereby limiting his ability to cut his costs.

If such a market pricing pattern were seen to be consistent over several years, it might make sense to buy weather insurance that pays off only when temperatures occur outside a defined band. The insurance claim could then be used to cover the high cost of power during such brief periods instead of investing in a rarely-used (and potentially expensive) on-site generator.

With such an hourly pricing chart in hand, a customer considering a *fixed annual* power price may then quickly see the premium he is paying for that zero-risk option. By right-clicking the Z-axis and choosing "format axis," one may then (under the "scale" tab) set the value at

which “category (X) axis crosses at” to equal the fixed commodity price. The degree and frequency of times at which hourly pricing is lower (or higher) is then readily obvious. For an accurate comparison, any fixed ancillary costs charged by the ISO should be deducted from the fixed marketer price so that only the difference due to market fluctuations is visible. Such costs may be available through the marketer, your utility, or may be posted at the ISO’s web site.

3D Profiling of Hourly Power Cost

With hourly power pricing and usage both in hand, it is then possible to assess the simultaneous impacts of price and usage. By creating a formula that multiplies each cell of price (in \$/MW-hr) by each cell of use (in MW-hr), we then see how *cost* (not price) varies by the hour. Since high hourly usage is typically coincident with high hourly price, the combined impact of these two variables may result in extremely sharp spikes in hourly cost.

In Figure 5, we see the impact for a building with electric chillers. Since a comparable profile under a fixed price regime would be shaped just like Figure 4, one may see the magnifying effect of hourly pricing on some days and its constricting effect during most other days (when price is much lower).

Once again, such charts may assist in understanding the options for power supply and demand control. Note, for example, the twin




Figure 5

peaks seen on one high load day in the middle of the year. The gap between them is the result of running an on-site peak shaving generator on that rather high-priced day.

Date-embedded Scatter Plotting

Determining the sensitivity of a building's heating fuel usage to outdoor temperature is one of the mainstays of basic energy analysis. When only monthly energy data are available, however, accuracy and detail may be quite limited. Lacking any better data, many practitioners must content themselves using monthly degree-days or average monthly temperature. Using Excel's XY scatter plotting process allows easy charting of these variables and may be supplemented with a trendline and an approximate equation describing the general relationship (see Figure 6).




Figure 6

When daily fuel usage is available from gas meters equipped to provide hourly (or daily) consumption, that technique may be greatly improved. Variations due to weekends, holidays, production/schedule changes, temperature controls (e.g., automatic setbacks), and usage anomalies then become apparent.

In a typical scatter plot of temperature vs. energy use, points are

usually distributed in a linear grouping since energy use (for heating or cooling) is generally proportional to the temperature difference between indoors and outdoors (with a constant added to cover non-climatic factors such as domestic hot water usage and line losses). This method demonstrates how consistently a building's heating system (or, if plotting electricity, its cooling system) responds to outdoor temperature. Other factors (such as humidity, infiltration, wind speed, night/week-end setback systems, and system failures) may tend to scatter the points away from simple proportionality.

Excel can calculate a "best fit" linear equation for the points in a scatter chart and display a trendline fitting that equation. Points far above or below the trendline may indicate an operational problem (e.g., failure of an HVAC control). Placing a cursor on a point in a standard two-dimensional (2D) Excel scatter plot causes a boxed note to appear containing temperature and energy use (or whichever two variables were used to create the scatter plot).

While useful, this form of analysis has limited value since it is difficult (especially when examining data covering many months) to determine *when* a particular anomaly occurred, or if a point that looks correct is really in the wrong place. Without such time-based knowledge, it may be difficult to determine causes for (and pursue corrections of) anomalies.

Using advanced Excel techniques, that limitation has now been overcome.

Enter the 3D Scatter Plot

By taking advantage of Excel's 3D and pivot chart capabilities, time-based data (e.g., day of the week and date) may be embedded into each point of a scatter plot. When touching such a point with the cursor, the text block that appears now contains not only temperature and use data, but also the day of the week and date when that use occurred (see Figure 7). It is then possible to relate actions taken on that day to possible deviations or problems (e.g., holiday closure, EMS programming alterations, failure of the outside air damper to close).

This "date-embedded scatter plot" technique was developed by Energywiz Inc. in 2002, and has been successful in pinpointing and correcting a variety of problems related to heating and cooling of commercial and industrial facilities. A macro for this technique may be obtained by contacting the author at energywiz@aol.com.




Figure 7

While appearing to be a 2D scatter plot, the chart is actually a top view of a 3D chart in which one is looking down on vertical cylinders whose tops appear as colored points. The height of each cylinder is the day and date at which a given temperature and usage occur (see Figure 8).

The ability to express day and date as a cylinder height is based on Excel's use of a day-numbered calendar which begins on January 1, 1900. Each day thereafter is assigned a sequential number. Those numbered days actually create cylinder height: April 1, 2004, for example, is 38,078 days "high." When formatted by Excel, such numbers are converted back into the standard calendar days and dates that appear in the text boxes discussed above.

How Has This Technique Been Used?

In a building used as a testing laboratory, it revealed problems with the programming of a night/weekend setback heating system that were causing excess consumption of natural gas for heating. In this facility, occasional jumps in natural gas use were generally ignored because they were thought to be related to production of large quantities of gas-heated hot water occasionally produced for testing consumer appliances.

In the example seen in Figure 9, most of the *weekend* days (when the lab is shut down and interior temperatures are allowed to float

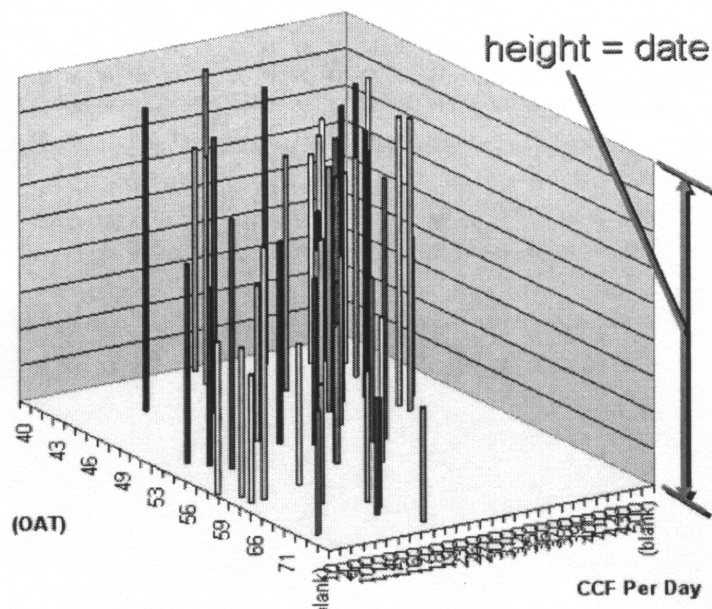


Figure 8

downward) are encompassed within the dotted outline, but several (in dashed circles) are weekend days using significantly more fuel than other weekend days at the same average outside air temperature (OAT). See, for example, two points occurring at 53 degrees OAT. The one in the circle used 19,000 cubic feet of gas (i.e., 190 CCF) on a Saturday, while another used only half as much (90 CCF) on a Sunday having the same average temperature. The lab was shut down on both days.

In this particular case, an error in EMS programming caused outside air dampers to fully open between midnight and 5 AM on some *weekend* days, causing the gas-fired heating system to run much more while the building was unoccupied.

Because that same error also re-closed the dampers before any building personnel arrived to witness the problem, the excess usage continued for some time before being discovered. When one looked only at *monthly* gas consumption numbers, usage was not significantly higher than expected. Likewise, daily usage when the problem occurred was comparable to that of normal *weekdays*, so the 2D scatter plot points for those problem days actually looked normal because they were near the trendline.

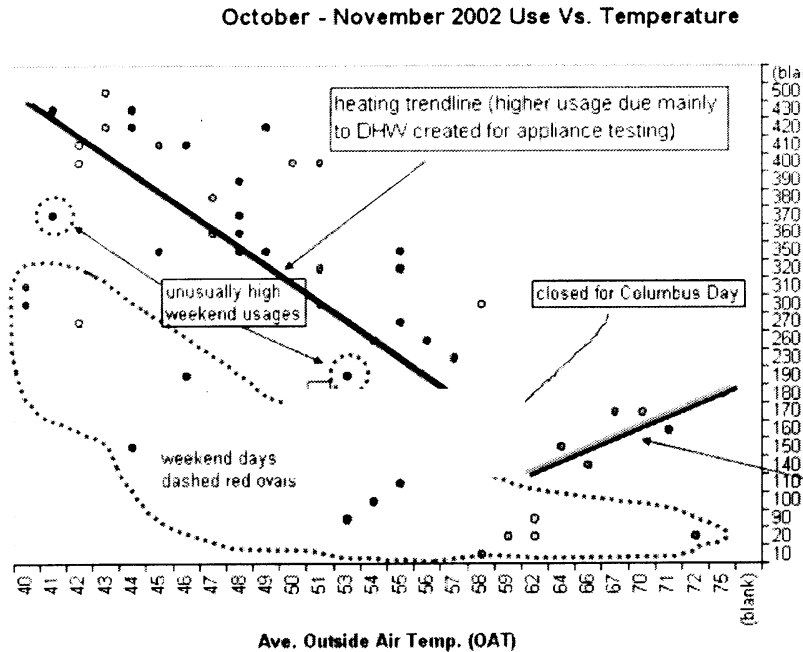


Figure 9

During a routine check of dates when "normal" plot points occurred (possible only with the date-embedded scatter plot system), *week-end* days with high usage were found that should have been grouped with other weekend days at lower usage levels. When building personnel were asked to check those dates for any unusual activities, night operation of the gas-fired heating system (logged by the EMS) pointed to unusually long run times. Examination of EMS programming found the startup/shutdown time error, which was corrected by building personnel at no cost.

Upsides and Downsides

Some care is needed when analyzing date-embedded scatter plots. In Figure 9, note that neither the horizontal nor vertical axes are fully linear across their entire ranges. One does not, for example, see any points with temperatures of 59, 60, or 61 degrees. In a standard scatter plot, one would see all temperatures represented on the X-axis, even when no data points existed for those temperatures. In Figure 9, the slightly non-linear

axes are due to the small (two-month) data sample used.

This problem is usually overcome simply by using a larger data set (120 points may be sufficient) that represents every temperature, and enough usage levels appear to create fully linear axes. While the data in Figure 9 are still accurate, and differences in energy usage derived from comparing text box data seen at different points are correct, one must avoid depending solely on the use of axes' scales to quantify anomalies.

Unlike the 2D scatter plot, the date-embedded scatter plot technique does not offer either an automatic trendline or R-squared options. Since a trendline that is based on potentially erroneous data is, however, of dubious value, its absence is not a problem. A graphic line may be added using Excel's drawing commands that, to the eye of the user, shows a linear trend among the points.

Once the date-embedded scatter plot process is used to "scrub" a given data set clean of any problems, the resulting data may then be transferred to a standard 2D plot to produce a trendline and R-squared value that reflects proper operation of a facility. The trendline equation may then be safely used as a predictor of daily usage using actual (or forecast) OAT data. Such a formula may be quite handy for developing and adjusting monthly gas purchasing volumes.

ABOUT THE AUTHOR

Lindsay Audin (energywiz@aol.com) is the president of Energywiz, Inc. (www.energywiz.com), an energy consulting firm serving a variety of large end users, government agencies, energy suppliers, and other consultants. His 28 years of experience in the energy services industry include eight years as energy manager for Columbia University and 12 years with private engineering and energy consulting firms in New York City prior to founding Energywiz in 1996. In 1993, Audin was named AEE's International Energy Manager of the Year, and in 1996 it inducted him into its Energy Manager's Hall of Fame. In 1992, he founded the New York Energy Buyers Forum, a non-profit organization of large energy users that focuses on regulatory and technical ways to cut energy costs. Audin writes the "Energy Wiz" column in *Engineered Systems*, is a contributing editor to *Building Operating Management* magazine, and delivers energy-related seminars on a variety of subjects for AEE and other organizations.