Development and Implementation of a Ductless Heat Pump Metering Plan

Robert Davis and Kevin Geraghty, Ecotope, Inc.

ABSTRACT

The scope and volume of utility conservation programs across the country have led to renewed interest in measurement and verification (M&V). Accelerating energy-efficiency gains and reducing emissions requires that programs actually deliver measurable savings. In the spring of 2009 an extensive M&V project was launched to assess the Northwest Ductless Heat Pump Pilot Project. This project included sub metering of heating, domestic hot water, and total energy. Temperature and other parameters were also used to determine heat-pump performance for the emerging application of offsetting electric space heat as a residential retrofit. This paper describes the development and implementation of the measurement plan used in an ongoing evaluation (95 sites in three states).

The installation procedures and post-install processes described in the paper are designed to:

- Collect data to verify the ductless heat pump (DHP) was operating at or near the manufacturer's performance specifications;
- Monitor occupant interaction with the DHP and the remaining heating system, and determine heating energy savings;
- Check data streams with a custom, add-on automated QC program to identify problems;
- Allow addition/replacement of sensors after initial installation with minimal data loss;
- Summarize other non-space heat usages and the interaction with heating requirements.

The paper will discuss the selection of the final metering package, development of training for installers, interaction between the equipment vendor and technicians to troubleshoot problems, and implementation of data quality control. Illustrations of data collected and data quality issues identified in the QC process will be presented.

Introduction

The scope and volume of utility conservation programs have renewed interest in measurement and verification processes. With so much at stake, it is essential that programs actually do result in measurable savings. Monitoring should also be able to identify issues with the equipment being monitored so that problems can be corrected quickly. Several monitoring systems are available commercially and have received widespread usage. These systems would ideally combine economy, durability and reliability with easy remote access to data streams which are output into malleable formats for analysis. This paper describes aspects of a measurement plan used in an ongoing evaluation of ductless mini-split heat pumps (95 sites). Results from the study should be available in mid-2011.

The study homes are all single-family detached site-built structures, range in size from 1,000 to 3,500 ft^2 and are located in both relatively mild (<4000 base 65° F heating degree day) and harsh (> 9,000 base 65° F heating degree day) climates. Most homes in the study are more than 15 years old, and the typical ratio of heating to cooling usage in the Pacific Northwest (if any mechanical cooling is even installed) is about 10:1. All homes in the study had zonal electric heat as their primary heating source before installation of the ductless heat pump. Sites are spread over three western states (Washington, Oregon, and Idaho) so access via remote connection is very important.

The installation procedures and post-installation processes described in the paper are all designed to: (a) collect data needed to verify the program is accomplishing its goals; (b) minimize front-end data collection errors; (c) check data streams with an add-on automated program to identify problems; and (d) allow addition/replacement of sensors after initial installation with minimal data loss.

Ductless heat pump (DHP) technology, using various sophisticated control and variable-speed operations, has been used extensively in Asian markets for at least 30 years. Although used only sparingly in various American niche markets, it is a potentially very efficient zone system that can supplement or replace existing residential zone heating systems. The fact that these DHP systems can provide both heating and cooling, with a ductless system, makes them an attractive consumer option. Recent changes in the technology – use of inverter-driven compressors – have increased nominal efficiency and there is great interest in the *in situ* performance of DHPs. This is especially true in the Pacific Northwest, where more than 575,000 single-family homes are heated by zonal electric systems (NWPCC 2010). Most of these homes have no mechanical cooling, so this aspect of the technology has other implications.

Electric utilities in the Pacific Northwest are very interested in offering the DHP as an electricity conservation measure. A careful review of field performance, using as many homes as possible, is key to establishing performance of the DHP and establishing its cost-effectiveness.

The primary goals of the field monitoring are to characterize energy usage, ambient temperature, and main living space interior temperature. The main research question is to determine the amount of offset to straight electric resistance heating provided by the DHP technology. Secondary goals are to measure non-heating usage in the home (hot water/base load) and to place heating usage (and cooling usage) into context as a function of house heat loss rate and house type. The data-logging system needs to be able to measure true root mean square (RMS) power and integrate properly to accumulate electricity consumption over at least one year's time.

Measurement Design

The measurement design had four goals:

1. Deliver heating system energy use once the DHP is installed. This was accomplished by metering the DHP and separately metering all the resistance loads in the zone electric heating system that was displaced (but not removed).
2. Meter the performance and operating patterns of the DHP as it relates to the various determinants of consumption for this type of equipment.
3. Meter the domestic hot water (DHW) usage to help establish regional planning assumption based on sub metering done in the early 1990s but not repeated. This required a submitter on the large resistance load associated with the DHW tank.

4. Meter the total electric energy usage of the home by metering the service drop for the whole house. This had the effect of giving a sum check on the other meters and (with subtraction) allowed a picture of the miscellaneous electric loads in the home. Like the DHW this load was sub metered in the early 1990s and no similar data set had been accumulated since.

The metering approach designed to respond to these goals was called a quad-metering system. The DHP and house service loads were monitored with true power meters. The electric resistance heaters and DHW tank were monitored with simple current transformers (CTs). Three temperatures were measured in the basic metering plan: outdoor ambient, indoor zone where the DHP was installed, and vapor line temperature at the heat pump itself.

Preliminary metering in a small pilot suggested that the cooling signal determination using only indoor temperature was very problematic, and the analyst was left to guess when cooling was occurring in the swing seasons of late spring and early autumn. The controls for the DHP equipment are very interactive and it is possible for simultaneous cooling and heating to occur. Measuring the vapor line temperature allows the analyst to be sure when the unit is cooling and allows a direct accumulation of the total cooling load and the conditions where cooling is supplied while electric resistance heat is also used.

Metering Equipment

To reliably measure whole-house electricity usage, hot water usage, and heating usage (including DHP) over a year's time, the metering equipment needed to be well-designed, durable, and weather-resistant. The hardware selected included industry-standard current transformers, wired thermistors, watt transducers, and pulse counters. Details for two of the instruments used in the project are found in the references (Continental Control Systems 2008, Onset Computer Corporation 2008). The equipment was designed to be installed outdoors, if needed. Data were sampled every five seconds and averaged into five-minute averages. Storage was made into a solid-state data logger equipped with internal Global System for Mobile communications (GSM) type communication technology. Data were uploaded automatically every six hours to a web-based server. From this point, data were screened for anomalous readings through a custom automated process (described in detail below).

Figure 1. Schematic of Data Collection Process

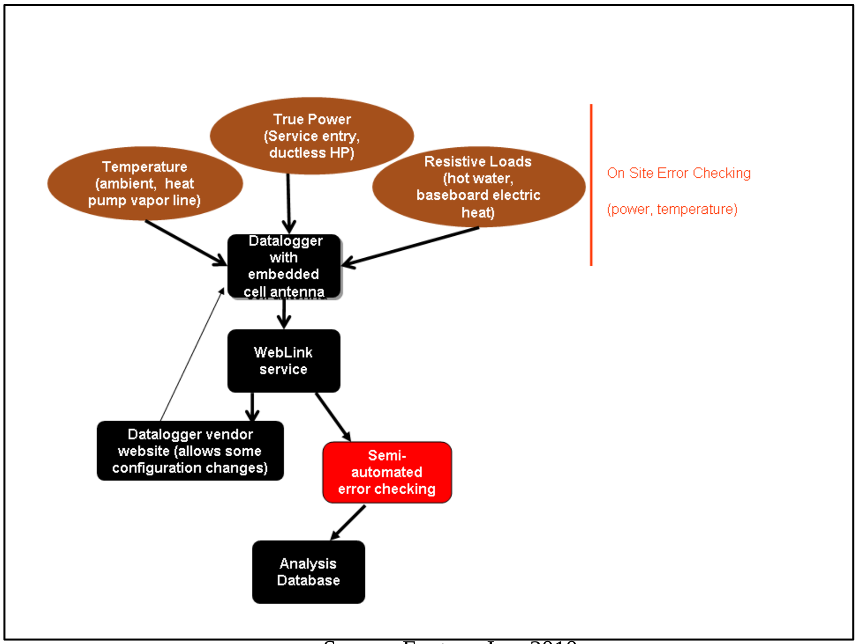

Source: Ecotope Inc. 2010

From the perspective of a year-to-several-years-long data gathering effort such as the one discussed in this paper, the principal advantage of near real-time data retrieval – as opposed to long-term accumulation onsite and one-time retrieval – is to provide an early-warning system for data production or quality problems, so that timely corrections or repairs can be made. With nearly 100 sites in the field producing data at five- or one-minute intervals (roughly an aggregate 300,000 data points per day), this early warning system needed to be highly automated in order not to overwhelm human monitors.

On-Site Error-Checking Procedure

Error checking is not delayed until after data have been recovered. At each field site, approximately ten sensing elements (current transformers, temperature sensors, etc.) were deployed to measure temperatures and electricity consumption. A field installation guide was developed in the early stages of field installation, and it covered most elements of field installation. Onsite installation managers were required to fill out a detailed site protocol, including types of sensors and individual sensor serial numbers (since these are the primary identifiers of sensors once data returns from the data logging vendor). Since the work was geographically dispersed, five different field installers (and six different electricians) were involved in installations. This meant that it was likely problems would occur despite careful attention to the installation protocol.

No assumption was made, at least by the project manager, that the error checking offered by the data logger manufacturer (as part of configuring sensors) or the automated data checking, was the best way to minimize field installation errors. Also, since many sites are located at least two hours' drive from the installers' office, it is critical to minimize installation errors.

The most complicated part of the installation is the group of equipment which measures power usage by the house and ductless heat pump. After this hardware was installed and the data logger initialized, the field installers were instructed to check apparent power readings against handheld measurements to determine reasonable equivalence. “Reasonable” means that, depending on which load was measured, the handheld measurement would not necessarily be exactly that of the data logger, given the fact that the power factor was not always measured by the handheld device. Key to making this process less confusing was selection of the proper logging interval; with the combination of current transformers (50 amp for the DHP or 100 amp for the whole-house service) and a 30-second logging interval, the checkout math became very straightforward and helped minimize setup problems since the site installer could see quickly if data logger measurements agreed with the handheld measurements.

Temperature measurements were not monitored as closely. Partly this is because one of the measurements (heat pump vapor line temperature) is used as an indicator of mode of operation (heating or cooling); the actual temperature measured is less important than the divergence in temperature from ambient temperature. (That is, if the vapor line temperature is considerably above ambient temperature, the unit is in the heating mode; if the reverse is true – vapor line is very cold – the unit is in cooling mode.) This indicator is important as a delineator of heating/cooling energy usage.

This situation is different in about one-third of the sites. In these sites, the accuracy of the temperature measurements is critical, since it is directly tied to unit thermal output and accuracy. Unfortunately, no systematic checkout procedure was used at the time of installation (more attention being paid to the much more intricate device used to measure airflow), and we learned after reviewing data that one of the thermostats displayed nonlinear response (about 15% of sites). These sensors were replaced, but in some cases there had been months of delay. These are non-accumulation sites (meaning the total usage over a desired time interval is not the objective of the measurement), but focus on unit performance at different outdoor temperature bins; nevertheless, it would have been much better to have performed a careful checkout of sensors at installation. More discussion of this issue follows the main error checking discussion.

Overview of Data Management

The data logging vendor offered two interfaces for clients to gather and interact with site data remotely once it had been delivered to the web-served data warehouse: first, a website interface, and, second, a "web services" interface where our computers could directly retrieve data from the data warehouse using the Simple Object Access Protocol (SOAP) internet web services protocol (Onset Computer Corporation 2009).

We invested in the latter method – automatic SOAP calls using in-house client routines – because it was the most automated method of delivering site data to our local repository. For timely data-monitoring purposes we did not, and do not, believe that the website point-and-click interface scales adequately beyond more than a handful of sites.

The system we established automatically retrieves all new site data from the warehouse once a day via command-driven batch files, and subjects it to range and sum checks. Because one of our site monitoring channels is total service power consumption, we are able to compare service consumption against the sum of sub metered power consumption channels (usually electric resistance, domestic hot water, and ductless heat pump). The difference between the service load and the sum of these sub metered loads, constituting lighting, kitchen appliances, and plug loads, should of course never be negative. In practice this summing constraint has proven to be one of the most useful ways of detecting data quality problems.

We check each batch of new data for the expected time gap between successive observations (five minutes or one minute, depending on the site). We also take the opportunity to check the timeliness of the most recent data obtained in our retrieval request. Given that the site loggers call to transfer accumulated data to the warehouse every six hours, a "most recent time" significantly in excess of six hours indicates trouble. The daily retrieval and data-checking process currently takes about two hours to run each night.

In its current tuned state the system works well as an early-warning system which alerts us only to problems important enough to pay attention to, but stays silent on negligible data problems.

But the system was not easy to develop. Issues we grappled with include: (a) the inherent complexity of SOAP; (b) incomplete documentation (for example, undocumented limits on the size of individual data requests to the vendor's web services server, and undocumented delivery of observation times all in UTC rather than local time, irrespective of any user setting); (c) developing a system for mapping the retrieved data to data-stream characteristics known from onsite installation; (d) a needlessly convoluted and deeply-nested XML data schema; (e) occasional data delivery gaps from the vendor's server in response to web services requests – sometimes remedied by later identical requests, sometimes retrievable only by website point-and-click methods, sometimes irretrievable.

Error Checking Details

As stated above, the automated error checking focused on both temperature and energy readings. A program was written that could search each site data file quickly and identify sites with anomalous readings.

Temperature error-checking was straightforward. Reasonable temperature ranges were assigned to the ambient (outdoor) and heat pump vapor line sensor channels; it was possible to look through what would eventually become several dozen sites to notice if problems were occurring with these channels. It would be more accurate to say, however, that really only very high or very low temperatures would be identified as issues since the range of possible ambient temperatures was expected to be between -15ºF and 110ºF, and the range of possible vapor line temperatures was expected to be between 25ºF (cooling operation or defrost) and 150ºF (heating operation).

Checking electricity usage was more involved but the basic concept is simple: compare all sub metered usage with the total service entry usage over each logged period to make sure the sum of sub metered usage does not exceed the service entry usage. It would be impossible for this to really happen unless there is a measurement or data collection problem.

The key to the daily and detailed error checking is the use of custom programming that is part of exploratory data-analysis software. This approach allows construction of targeted programs which can quickly comb through hundreds of thousands of data points per day and produce automated, compact text files which are e-mailed to the field monitoring program manager each morning. The contents of the text files indicate the condition of each of the 95 sites and flag problems.

The only compelling reason for frequent automated data retrieval is timely data-quality monitoring. It is therefore ironic that the selected vendor's data-retrieval system has proven to be the source of some data quality problems. We cannot say how many of these problems would be remedied by a different implementation of this technology – to date we have only used one data logging system – but there is an inherent complexity to the process which may make the data-retrieval system more fragile and failure-prone than a simpler alternative of long-term onsite accumulation. Although an occasional failed upload call is not a problem given the loggers´ data storage capacity, site loggers intermittently unable to call in may deteriorate into states in which they neither respond to remote commands nor record data; in such a situation a site visit offers the only hope of putting the logger back to work. In addition there has been at least one episode of widespread synchronous call-in failure in which the GSM vendor was apparently the source of the problem.

But far more of our data problems originate with the data loggers themselves rather than the data transfer mechanism. Data loggers are temperamental and not straightforward to install and configure correctly. Timely data retrieval and scrutiny is essential in detecting and attending to subtle configuration problems, and logger-originated data quality issues, such as data corrupted by electromagnetic interference, cannot be mitigated in all cases.

Empirically, a fairly high percentage of critical problems at a typical site surface, and are resolved, within the first week or so. This occurs despite the care taken onsite to check that sensors are configured and recording properly. Other factors in addition to installation mistakes can cause problems. A site can have problematic GSM phone communication despite an apparently strong signal onsite; sensor serial numbers (typically nine digits) can be recorded incorrectly so that incoming data streams cannot be automatically matched to known characteristics data. Data-averaging intervals can be set incorrectly so that what is supposedly an average or amperage value over five minutes is in fact the last recorded instantaneous value in a five-minute period. Timely data retrieval and scrutiny is critical in getting through this initial trouble-prone state to achieve a relatively stable state of acceptable functioning.

The data loggers have a web-based remote management interface which permits the resolution of some problems without repeated site visits, e.g., mistaken data averaging intervals. At regularly scheduled data upload intervals there is a window of opportunity to send configuration instructions to the data logger, and with reasonable luck these are in fact executed. In addition, problems which are essentially ones of interpretation (incorrect sensor serial number, incorrect pulse count multiplier applied to power consumption data) can also be corrected remotely. But there remain certain problems that can only be addressed with site visits. The following table summarizes important site interventions:

Table 1. Site Interventions

Total sites	95
Data logger replaced	8
Other critical interventions requiring a site visit	3
Important configuration issues resolved using remote interface	7
Other critical first-week data quality issues cleared up without site visit	6
Site visits to fix signal interference problems (desirable, not essential (in most cases)	23

A rough guess is that without the feedback provided by timely data monitoring, about 20 of our 95 sites would not have produced usable data. Final data quality in the remaining sites also will benefit from the data-monitoring effort. The wide scope of signal interference problems, for example, was only evident to us because of automated data monitoring, and this in turn permitted us, in concert with the vendor, to develop a strategy to reduce signal interference in ongoing data collection. We think this adds up to a compelling case for a strategy of investing carefully in near real-time data monitoring, with its attendant expenses. Given all the costs of recruitment, equipment purchases, and installation, a site failure rate of over 20% is simply too high in most field monitoring situations.

There is a second important reason to develop such near-real-time monitoring machinery, and that is it enables much faster learning and response to problems on the part of all participants, including both Ecotope and the data logging vendor. It cannot be stressed too much that successfully executing a long-term data logging program should be thought of as a process of adaptive learning. Ecotope's installation procedures and rates of problem site occurrence improved over the course of the project. Thanks to systematic data monitoring we were able to bring a number of equipment problems to the vendor's attention rapidly and forcefully, which was useful to the vendor and for us. There is no ready-made template for large-scale projects of this sort. All participants need to adapt, and to learn, and a systematic and well-thought-out system for timely data monitoring and analysis allows that to occur.

Coefficient of Performance (COP) Measurements

As mentioned above, about one-third of the sites were used to estimate *in situ* system efficiency (coefficient of performance, or COP). The COP is the ratio of heating (or cooling) output from the DHP to the power needed to run the compressor and indoor and outdoor fan. (Output is converted from Btu/hr to kW so the numerator and denominator are in the same units). Another way of expressing the COP is in efficiency percentage, with a COP of 1 meaning 100% efficiency. The COP measurement is very useful for comparison to AHRI-rated performance, and also to inform a parallel analysis of utility bills (also used to determine savings from application of the ductless technology).

Two temperature sensors were added (to measure change in temperature across the indoor unit) and a small vane anemometer was installed to provide a proxy measurement for airflow. (This device accumulated pulses in a manner similar to that for the electric energy current transformers/watt transducers. Different pulse rates could be compared with a one-time calibration to determine cubic feet per minute [CFM] of airflow.) The product of temperature split and airflow is thermal output in heating or cooling. Since energy usage/power of the ductless heat pump and outdoor temperature are also unknown, system COP can be calculated as a function of outdoor temperature bins.

Graphical analysis is a powerful tool for quick evaluation of a site. The following graphic shows these elements for a site where everything is working correctly. This site is located in Idaho Falls, ID, a cold site (about 9,500 base 65° F heating degree days); the monitoring period shown includes relatively mild outdoor temperatures (between about 20°F and 50°F).

This is a one-ton ductless heat pump which typically moves about 200 CFM (which corresponds to a "medium" flow setting generally associated with the way most homeowners run their DHP with the "AUTO" setting; CFM are indicated by the red dots). When the system is in running in a lower power mode (indicated by smaller blue bars near the bottom of the graphic), the *in situ* COP is between 3 and 4. When the outdoor temperature is colder (20°F-30°F) and the system is under more load (using about 500-650 Watts), COPs are less favorable but still typically between 2 and 3. This performance is within expected ranges and indicates a system that is producing good data. The graphic also allows a quick evaluation of the return and supply temperatures at the DHP (green and orange lines, respectively). The return temperature is usually in the mid-70s°F (expected values since the DHP is mounted high on the wall of the main living space) and the supply temperature (in heating) is typically in the mid- to high 90s°F (also expected).

Figure 2: Well-Behaved COP Site

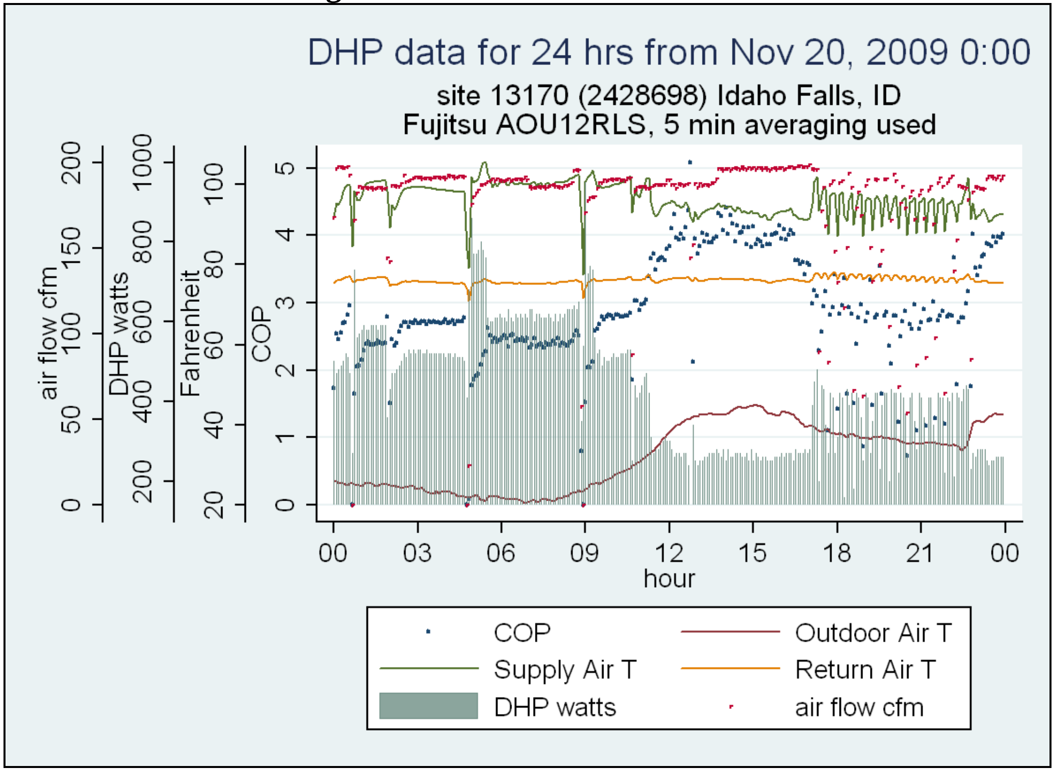

Source: Ecotope Inc. 2010

The next graphic shows the same measured fields, but for a site with a problem in temperature and airflow measurement. This is a site where, over the course of several months, the vane anemometer became dirty and the anemometer wheel stopped rotating. Note the airflow readings (blue dots) are very sparse and essentially give zero readings. The supply air temperature sensor also had a problem (which was not noticed during initial installation), but this is not shown in this graphic. Temperatures were not double-checked at most COP sites, which was a serious oversight. It has turned out that a small number of sites do have this problem; however, most of the affected sensors have been replaced at this point.

Figure 3: COP Site with CFM and Supply Air Temperature Problems

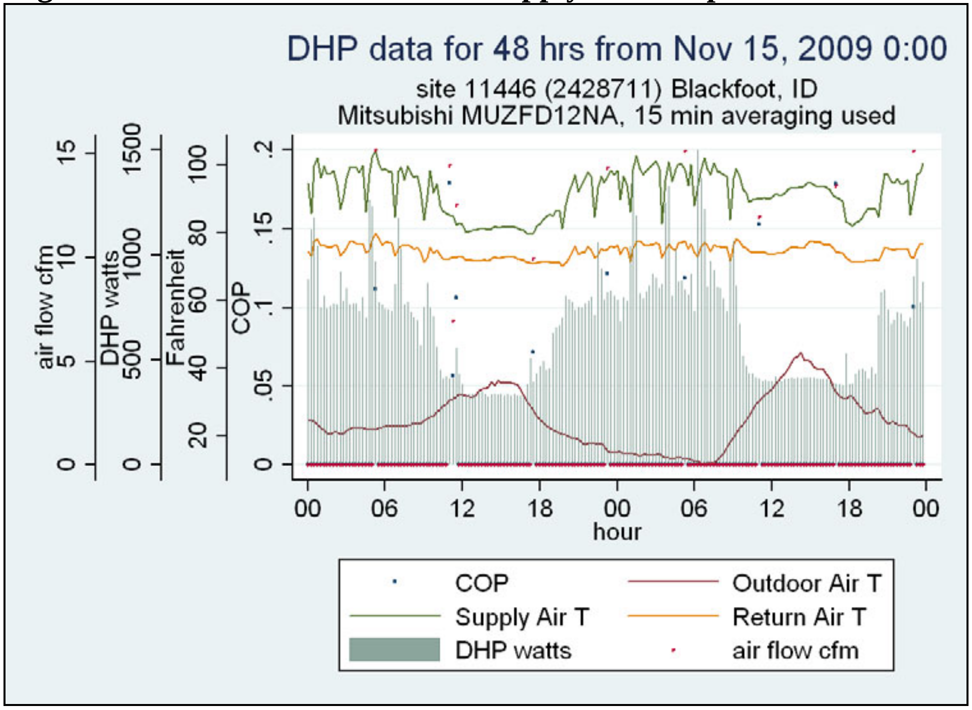

Source: Ecotope Inc. 2010

The consequences of errant measurements at the COP sites are not as dire as for the year-long accumulation sites, since the performance is described in relation to outdoor temperature bins rather than accumulated over the entire year. However, if a problem is not detected relatively quickly, it is quite possible that a unique part of a heating (or cooling) season could be lost.

An automated procedure was not established for reviewing these data streams; they have been examined through a more labor-intensive manual process. At this point (most of the way through the heating season), about three-quarters of the COP sites have produced enough detailed data for productive use. Without an automated review process, more data problems would have occurred and it is certain that more data loss would have resulted. Preparation of an automated system is underway; this will be most helpful in extending the useful life of systems during the summer of 2010 (cooling COP measurement).

Conclusion

Monitoring systems were installed in a total of 95 residences in three Northwest states over the course of eight months. Installations were completed by several installers and electricians, resulting in varying installation quality. Mistakes in configuration and data collection occurred and continue to occur over the remaining months, but a standardized review process has minimized data loss.

Critical elements of minimizing data loss are:

1) The ability to keep track of the data streams as they are uploaded via automated process from the data logging company to local machines

2) An understanding of reasonable data values in the data streams via error checking and graphical analysis
3) The ability to quickly make simple changes in analysis problems (to process data that appear to be bad but that are merely a result of scaling factor problems, etc.)
4) A quick response by field personnel to fix persistent field problems with sensors/data loggers

It is critical to review the error file daily and flag problem sites for quick investigation. The site might be having problems, but apparent errors could be the result of internal software glitches, or could be short-term quasi-problems that are solved during the next data upload. On each end of the project, the personnel involved need to understand the equipment and be able to fix problems reliably and quickly.

What is an acceptable error level? The accuracy of one-time measurements depends on the technician's experience, understanding of what is to be measured, and equipment calibration. For longer-term projects, with more moving parts and intermediate steps, some amount of data pollution and loss must be assumed. But having a way to ascertain there is a problem, and do it regularly, is crucial to minimizing data loss.

References

Continental Control Systems. 2008. **WattNode Advanced Pulse Installation and Operation Manual**. Revision 2.34, July 21, 2008.

Northwest Power and Conservation Council (NWPPC). 2010. **6th Northwest Power Plan** (supporting data spreadsheet). Portland, OR. Pre-publication edition, February, 2010.

Onset Computer Corporation. 2008. **Hobo U30 Station User's Guide**. Document D-12269-D.

Onset Computer Corporation. 2009. **Hobolink Web Services Developer's Guide**. Document 12789-A.