Real-Time Itemized Electricity Consumption Intelligence for Military Bases

Omid Jahromi
Belkin International Inc.
Playa Vista, CA, USA
omid.jahromi@belkin.com

Alan Meier
Lawrence Berkeley National Laboratory
Berkeley, CA, USA
akmeier@lbl.gov

Abstract **— This paper outlines the methodology and the interim findings of a NILM demonstration project that Belkin and LBL have been conducting with a grant from the Department of Defense. More detailed findings including accuracy statistics and achieved energy savings will be presented at the workshop.**

Keywords— Energy Conservation, Load Disaggregation, Non-intrusive Load Monitoring, Machine Learning, Three-phase Power

I. INTRODUCTION

The concept of non-intrusive monitoring of energy-using equipment in buildings was first described by George Hart [1]. Detailed energy consumption information obtained through non-intrusive methods can aid consumers and building auditors identify major loads and provide feedback to assist in saving energy [2]. This can be achieved without interfering with internal activities. Furthermore, the cost will be (potentially) much lower. Since then, the capabilities of non-intrusive load monitoring (NILM) systems have improved enormously as a result of parallel advances in hardware, disaggregation algorithms, and growth in computation power. Nevertheless, Zeifman and Roth [3] concluded that no package can reliably identify all of the loads and their associated electricity consumption. Since then, new NILM packages have appeared that take advantage of technical improvements as well as unique characteristics of the devices being monitored [4],[5]. For example, Gupta et al. [5] exploited the interference created by switch mode power supplies (SMPS) as a means to identify individual loads. Advances like this will make NILM practical for energy auditing and analyses in residential buildings.

The next challenge for NILM technology is commercial buildings. Commercial buildings are larger than homes, have more and different appliance inventories, and typically draw 3-phase power. In this paper we describe early efforts to apply NILM technology to a commercial building. We begin by introducing the new NILM technology – the Belkin Echo – and its unique features. We then describe new problems related to installing the technology in a target building, and detecting and assigning loads. Finally, we present early results.

II. NILM FOR DEPARTMENT OF DEFENCE BASES

ESTCP is an environmental technology demonstration and validation program run by the Department of Defense (DoD). The program was established in 1995 to promote the transfer of innovative technologies that have successfully established proof of concept to field or production use. ESTCP demonstrations collect cost and performance data to overcome the barriers to employ an innovative technology because of concerns regarding technical or programmatic risk, the so-called "Valley of Death."

ESTCP Project EW-201335 was awarded to Belkin International Inc. in collaboration with Lawrence Berkeley National Laboratory (LBL). The demonstration objectives are:

1. Provide the DoD with real-time, itemized energy information (by device or category) through a single sensor installed at a building's meter.
2. Validate the performance of Belkin Echo by comparing with baseline data
3. Validate costs by comparing Echo installation costs to those for a comparable sub-metering system.

To achieve the above goals, Belkin has installed two sets of monitoring devices at buildings under test. A first set of monitoring devices are installed at the entry point of electricity into the building (Fig. 1). A second set of monitoring devices including circuit-level sensors and plug-level sensors are installed on nearly all the circuits in the test buildings and as many electrical plugs as practically possible (Fig. 2).

In the rest of this paper, the first set of monitoring devices will be called "Belkin Echo" whereas the second set of monitoring devices will be referred to as "ground truth system" or "baseline data collection system".

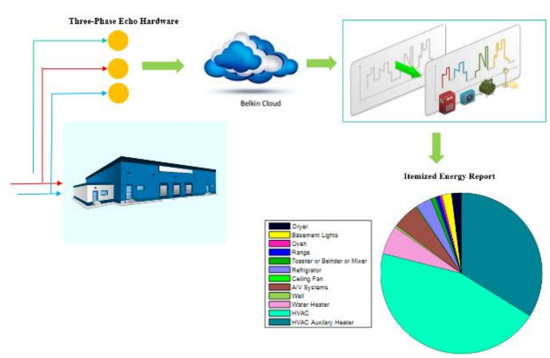

Fig. 1. Belkin Echo system includes hardware, cloud storage and algorithm components.

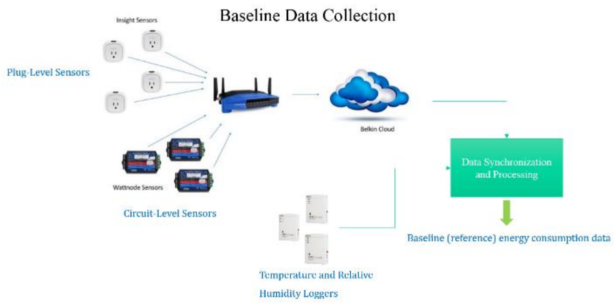

Fig. 2. Baseline data collection system uses circuit-level and plug-level power sensors.

Fig. 3. Belkin Echo hardware being tested in an experimental 3-phase configuration.

III. Hardware Development

Belkin has developed advanced data acquisition hardware that captures voltage, and current data at a very high sampling rate (3.4 kHz) along with a spectrogram of HF noise present on the voltage line (bandwidth exceeding 1 MHz). For the ESTCP project, the hardware was upgraded to handle three-phase installations. A sample of the three-phase data captured by the Echo system is shown in Fig. 13 at the end of the paper.

The Echo sensing hardware features a powerful Digital Signal Processing (DSP) module that performs several signal processing steps including calculating the spectrogram of HF noise, converting current and voltage measurements to frequency domain and representing them in amplitude/phase format (as harmonics of a 60 Hz sinusoidal signal), data compression, and cyclic buffering. The DSP module also compresses the data using both lossy and lossless compression algorithms such that the average bit rate of the data sent to the cloud remains in the order of 50 kb/s.

The Echo hardware further includes multiple communication modules that enable it to send data to the cloud via direct Ethernet link, power line communication (PLC) or cellular (GSM) modem. The entire system is designed such that it produces minimum heat and can be packaged in a sealed enclosure for both indoor and outdoor installations.

The ground truth hardware installed in the test sites include circuit-level power monitors made by eGauge Systems LLC and plug-level power monitors made in-house by Belkin International Inc. In the first test site we managed to install circuit-level power sensor on every circuit in the building. We also installed nearly 80 plug-level power monitors on various outlets to monitor individual plug loads (computers, servers, etc.). In the second site we only installed circuit-level power sensors. Due to space limitations in the circuit barker panels, we could not install sensors on all the circuits, still we managed to install 108 circuit-level sensors based on a priority list that we developed.

This project is funded by the DoD trough ESTCP grant EW-201335.

IV. Installation Sites

Our test sites are located at Joint Base Lewis-McChord near Tacoma, WA. The first site is an office building about 10,000 sq. ft. in size. It has air-conditioning but no windows. The second site is a residential building with 32 suits. It has central heating but not air-conditioning. Both buildings are supplied by 3-phase power. The two installation sites are shown in Figs. 4 and 5.

Fig. 4. Test Site I at Joint Base Lewis-McChord

Fig. 5. Test Site II at Joint Base Lewis-McChord

It is worth mentioning that Joint Base Lewis-McChord has more than 4000 buildings. Finding suitable buildings for the demonstration project required many steps including inspection of the building's electrical layout, considering the type of the structure and its use and also receiving approval from the proper authorities. This part of the project by itself took significant time and effort.

Another significant challenge we faced was that the entry point of electricity to our test sights comprised multiple underground cables bundled together to form a thick conductor. It was not possible to mount a current transformer (CT) over the complete ensemble of cables carrying power so we put our CTs on only one power-carrying cable per each phase. The current measurement was then scaled up to represent true current entering the building. We validated this technique both in the lab and via the ground truth system described subsequently.

V. Algorithm Development

Belkin's machine learning team has developed a three-phase NILM algorithm for this project. The algorithm includes a training sub-algorithm (Fig. 6) and a power estimator sub-algorithm (Fig. 7).

Both the training and the power estimator sub-algorithms rely on discernable power jumps (called 'edges') to identify and separate various loads. At each edge, several features such as real power delta, reactive power delta and the delta values in first and third harmonics of current are calculated and recorded. Features from the HF noise spectrum (see Fig. 13) are also included in the feature vector. A key aspect of the algorithm is that it identifies and subsequently classifies single-phase, two-phase and three-phase edges.

The training sub-algorithm uses ground truth data to identify the exact time a specific load has been turned ON or OFF. It then finds the corresponding "edge" (a point in time where there is a discernable power delta, either positive or negative) on the data collected by the Echo hardware and labels it. The edges that do not belong to the load under test will be labelled "Other". This way, all the edges in the training data will be labeled as "Load X – ON", Load X –OFF" or "Other". A clustering algorithm is then used to train a model (classifier) that can separate these three classes. The classifier parameters are saved as the 'model' for the load under training.
The above process is repeated for multiple loads (HVAC, Lights, Refractor, etc.) so that a model for each load is obtained. For complex loads such as HVAC, the training algorithm is actually capable of identifying and learning the features of multiple 'modules' inside the load with distinct ON/OFF power characteristics.

The power estimator sub-algorithm uses the training model to identify those edges on the Echo data that best match the features of the trained load. It then 'links' the ON and OFF edges to form a piece-wise constant estimation of that loads power. In many cases, there will be missed edges or falsely detected edges so the linker algorithm is equipped with several logical tests to identify and, as match as possible, correct for a missing edge or false edge.

Our initial results show that the algorithm works well and is highly accurate in estimating major loads such as HVAC (see Figs. 7 and 8). We shall report statistics on the edge detection accuracy and the power estimation accuracy of our algorithm during our oral presentation at the NILM workshop. We shall also describe our ongoing effort to develop a methodology to itemize the constant (baseline) power that is not associated with ON or OFF edges. As can be seen from actual data (see Fig. 13), there is significant amount of power which is not associated with "observed" activity on the power data. We are currently developing a methodology based on inverse-problem theory (Tikhonov regularization) that aims to distribute the baseline power between 3-phase loads and single-phase loads. We will report our progress on solving this problem at the workshop as well.

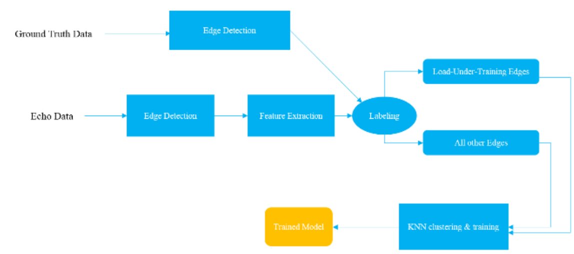

Fig. 6. Training Algorithm

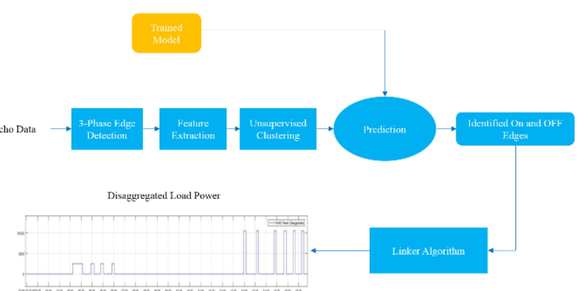

Fig. 7. Power Estimator Algorithm

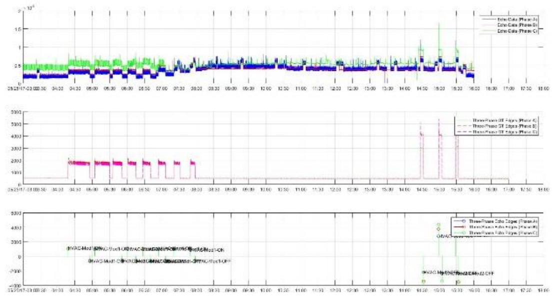

Fig. 8. Sample Edge Detection Results

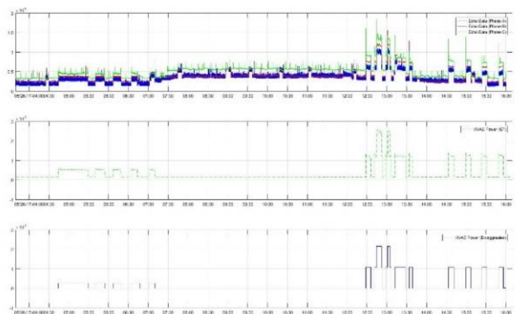

Fig. 9. Sample Power Estimation Results

VI. DISCUSSION

We believe that itemized power consumption information can provide energy managers with significant insight into a building's energy consumption. However, it is not obvious how the end data should be presented to the energy manger. Below, we discuss the merits of several presentation formats.

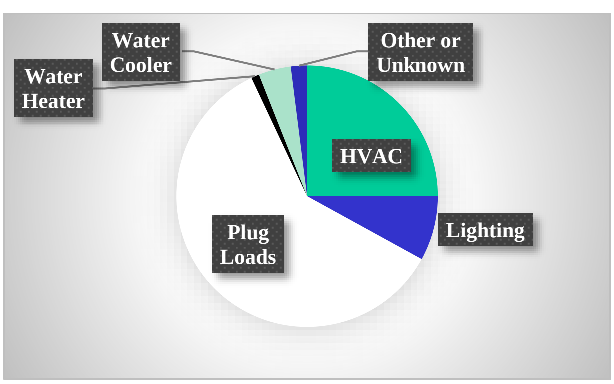

Fig. 10. Itemized Energy Consumption Results Presented in the Monthly Format; This charts shows data for Site I during October 2016.

A fist method to present the results is a long-term summary of energy consumption broken into different categories of devices (Fig. 10). This presentation format can give envelopes for potential savings that can be obtained if certain energy saving measures (retrofits) are implemented. However, without additional context, it does not guide the energy manager towards any specific actions. For example, knowing that 35% of the energy is spent on HVAC by itself does not tell the manager if the building's air conditioning system is efficient or inefficient.

A second method of presenting itemized energy data is to plot them in a daily bar chart as shown in Fig. 11. A daily chart like this can show the change in energy consumption pattern during weekdays and weekends. From this, some insight can be gained as to how much of the energy consumed is related to occupancy.

If properly plotted, a chart like Fig. 11 can also show which loads are relatively constant and which loads are variable in a day-to-day scale. We have learned that it is better to identify the constant loads and position them at the bottom of the bar chart so that they can be visually identified as a "color band" with constant height in the chart.

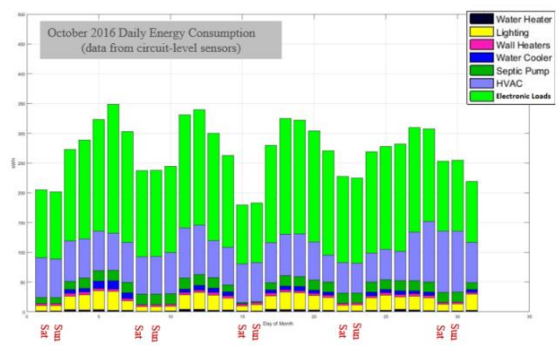

Fig. 11. Itemized Energy Consumption Results Presented in the Daily Format; This chart shows data for Site I during October 2016

A third way to present itemized energy consumption information is presenting the power consumption of a single-load at a once-per-minute or finer time scale (Fig. 12). This mode of presentation can provide useful insights into the behavior of appliances such as HVAC, refrigerator, sump pump or area lights. The plot in Fig. 11 shows a fine-time-scale presentation of HVAC power in test Site I. This plot also shows temperature and relative humidity (RH) inside this building.

A careful examination of this plot reveals that the HVAC first operated in the heating mode during the early morning hours (5-7 AM) and then in the cooling mode in the afternoon (12 -5 PM) of the same day. This, we believe, indicates an inefficient use of the HVAC system. This building could use better insolation so that it will not get too cold during the night and too hot during the day. This plot also suggests that a better control rule for the HVAC could be devised to make its operation more efficient.

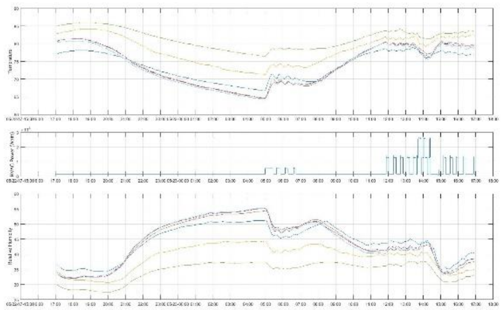

Fig. 12. HVAC power consumption during 24 hours (middle plot) shown together with temperature at six internal locations (top plot) and relative humidity at same locations (bottom plot) in Test Site I.

We are currently investigating the best methods of summarizing and presenting itemized energy data to energy managers such that they are guided towards a set of effective energy saving measures.

Upon completion of the demonstrating project, Belkin and LBL will provide a detailed report on Edge Detection Accuracy as well as Energy Estimation Accuracy of the Echo technology to the DoD. In parallel, LBL will use the itemized energy information to make energy saving recommendations to both Site I and Site II energy mangers.

VII. Conclusion

The overall goal of the project is to demonstrate a new technique of non-intrusively metering individual appliances and equipment in DoD buildings. Instead of building managers being given a single, overall consumption for a building, the Belkin technology would give "itemized" consumptions of made devices inside the building. The technology will save energy by delivering detailed, actionable information to building managers about equipment failures, operational problems, and opportunities for retrofits.

We are continuously updating our algorithms and analytical tools based on the data that we have captured from the two test buildings. At the workshop we shall present our latest results including statistics on accuracy of the system and the actual energy savings in test sites I and II.

VIII. Acknowlegement

This project is funded by the Department of Defense trough ESTCP grant EW-201335. Lawrence Berkeley National Lab's work was supported in part by the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies, of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

References

[1] Hart, George William. 1992. "Nonintrusive Appliance Load Monitoring." Proceedings of the IEEE 80 (12): 1870–91.

[2] Berges, Mario E., Ethan Goldman, H. Scott Matthews, and Lucio Soibelman. 2010. "Enhancing Electricity Audits in Residential Buildings with Nonintrusive Load Monitoring." Journal of Industrial Ecology 14 (5): 844–58. https://doi.org/10.1111/j.1530-9290.2010.00280.x.

[3] Zeifman, M., and K. Roth. 2011. "Nonintrusive Appliance Load Monitoring: Review and Outlook." IEEE Transactions on Consumer Electronics 57 (1): 76–84. https://doi.org/10.1109/TCE.2011.5735484.

[4] Froehlich, J., E. Larson, S. Gupta, G. Cohn, M. Reynolds, and S. Patel. 2011. "Disaggregated End-Use Energy Sensing for the Smart Grid." IEEE Pervasive Computing 10 (1): 28–39. https://doi.org/10.1109/MPRV.2010.74.

[5] Gupta, Sidhant, Matthew S. Reynolds, and Shwetak N. Patel. 2010. "ElectriSense: Single-Point Sensing Using EMI for Electrical Event Detection and Classification in the Home." In Proceedings of the 12th ACM International Conference on Ubiquitous Computing, 139–148. UbiComp '10. New York, NY, USA: ACM. https://doi.org/10.1145/1864349.1864375.

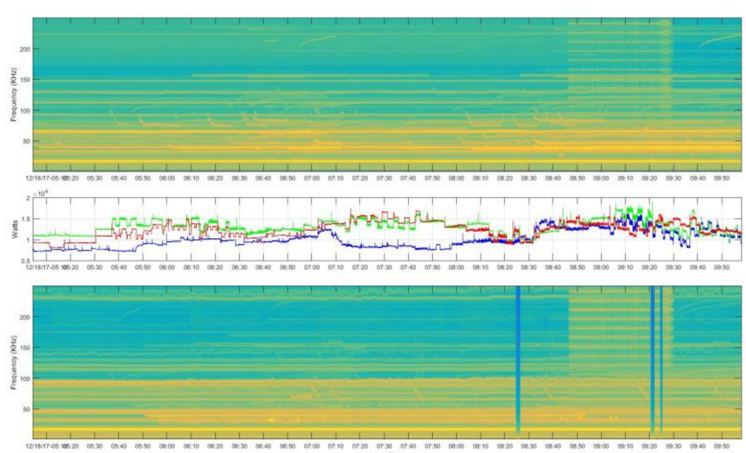

Fig. 13. Sample plots showing data captured by Belkin Echo sensors at Site II. The top and bottom plots show spectrograms of HF noise observed across the voltage lines (in Delta configuration). The middle plot shows real power for each phase (Y configuration). The horizontal axis is local time.