Interpreting Human Activity from Electrical Consumption Data through Non-Intrusive Load Monitoring

by

Mark Daniel Gillman

Submitted to the Department of Electrical Engineering & Computer Science

in partial fulfillment of the requirements for the degree of

Masters of Science in Electrical Engineering & Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2014

Signature redacted

Author ....
Department of Electrical Engineering & Computer Science
May 21, 2014

Signature redacted

Certified by ....
Steven B. Leeb
Professor
Thesis Supervisor

Signature redacted

Certified by ....
John S. Donnal
Doctoral Candidate
Thesis Supervisor

Signature redacted

Accepted by ....
Leslie A. Kolodziejski
Chairman, Department Committee on Graduate Theses

Interpreting Human Activity from Electrical Consumption Data through Non-Intrusive Load Monitoring

by

Mark Daniel Gillman

Submitted to the Department of Electrical Engineering & Computer Science
on May 21, 2014, in partial fulfillment of the
requirements for the degree of
Masters of Science in Electrical Engineering & Computer Science

Abstract

Nonintrusive load monitoring (NILM) has three distinct advantages over today's smart meters. First, it offers accountability. Few people know where their kWh's are going. Second, it is a maintenance tool. Signs of wear are detectable through their electrical signal. Third, it provides awareness of human activity within a network. Each device has an electrical fingerprint, and specific devices imply associated human actions. From voltage and current measurements at a single point on the network, nonintrusive load monitoring (NILM) disaggregates appliance-level information. This information is available remotely in bandwidth-constrained environments. Four real-world field tests at military micro grids and commercial buildings demonstrate the utility of the NILM in reducing electrical demand, enabling condition-based maintenance, and inferring human activity from electrical activity.

Thesis Supervisor: Steven B. Leeb
Title: Professor

Thesis Supervisor: John S. Donnal
Title: Doctoral Candidate

Acknowledgments

A heartfelt thank you to my direct and indirect sponsors: US Army Corps of Engineers, Grainger Foundation, Kuwait-MIT Program, Draper Laboratories, PM-Force Sustainment Systems, and RDECOM.

To my peers at MIT, this team meant a lot to me. Thank your for your friendship and mentorship, namely John Donnal, Jim Paris, Will Cotta, Chris Schantz, and Uzoma Orji. I stood on your shoulders. To my mentors here, Dr. Steven Leeb and Dr. Jim Kirtley, I am in your debt. Dr. Leeb, thank you for the patience, persistence, motivation, and direction over the last two years.

To the professional, hard-working, and supportive people at Ft. Devens (Lee O'Donovan, Jesse James, Bill Singleton, Gary McMichael, Will Eicher, Robin Szczuka, Jeff Gladney, and others), a hearty Tango-Mike. I could have accomplished nothing without the unfettered access you granted my team.

To the Sharon Public School District, especially the Cottage Elementary staff, our children are in extremely capable hands. Scotty Schertz, Ken Wertz, thank you for your tireless efforts on behalf of science and your school district.

To CSM Beharie, SGM Santos, LTC Brown, SFC Sartaine, Tom Theaux, and others from RDECOM (including JJ and others from the Institute for Soldier Nanotechnologies at MIT), thank you for giving us access into one of the world's premiere training facilities at Ft. Polk, LA. Robert Wilson, Donald White, MAJ Reed, CW2 Garavito, SSG Roman, you guys make things happen and I sincerely appreciate your support on this project.

Jana, Grant, Sophia, Annie, and unborn baby Gillman, thank you for letting me indulge my passion for learning and for giving me the time and flexibility to finish strong.

THIS PAGE INTENTIONALLY LEFT BLANK

Contents

List of Figures

Chapter 1

Smart Meters Are Not Smart Enough

Figure 1-1: An austere microgrid that would benefit from appliance-level feedback

Commercially available power meters are built for utilities, not consumers. Even as technology shortens the measurement time intervals, low-resolution kilowatt (kW) consumption data does not contain enough information to provide situational awareness or enable intelligent demand response in a power network.

Consumers need appliance-level feedback to reduce electrical consumption and save money. Specifically, they need to know what loads are currently on and how many kW's each load uses. We propose a method to both obtain and display appliance-level information using nonintrusive load monitoring (NILM) technology. Using NilmDB [1] software, NILM Manager analysis tools [2], and custom analysis filters and methods I developed [3], we demonstrate that appliance-level feedback is possible using a single sensor package located at the generator or panel. Field tests at the Cottage Elementary school, Ft. Devens base camp, Ft. Polk Forward Operating Base (FOB), and generator-powered Army field hospital demonstrate NILM's capability and state-of-the-art.

Appliance-level feedback is actionable because it tells the consumer precisely where and how they could reduce every usage. With it, three distinct advantages over today's smart meters become possible:

1. Gain accountability of the entire electrical network; know what's on and how much power each device uses
2. Enable condition-based maintenance; know when equipment is malfunctioning before critical failure occurs and wasteful operation persists
3. Infer human activity from electrical activity; gain awareness about the users on your power system

If you want the "SO WHAT", skip to Chapter 5. It highlights NILM's major capabilities through numerous real-life examples. Read Chapter 2 to learn how the NILM works. Chapter 3 shows the results and lessons learned from testing at Cottage Elementary School. Chapter 4 describes the military microgrid environment (Fig. 1-1) on our three military test sites, ideal test beds for NILM technology. Chapter 6 has a few practical lessons learned that will improve future installs, and Chapter 7 concludes with my vision of what facilities managers need to know about their network and how I propose to present it to them.

I am passionate about this work and believe in its potential. As a Soldier, I foresee this technology saving both blood and treasure if properly implemented. I will forever be available at mark.d.gillman@gmail.com for questions or comments.

Chapter 2

The Nonintrusive Load Monitor

Many current approaches to power monitoring involve unfortunate compromises. Submetering appliances provide highly discretized data, but the requisite hardware and communication requirements are substantial and invasive [4]. Similarly, high data rates can be essential for resolving adequate detail for diagnostics and consumption, but they also severely tax transmission capabilities for remote access or they require local storage that often requires retrieval [5]-[9]. Alternatively, data rates can be lowered to make data handling tractable at the expense of meaning and intelligibility in the data. Resolving these trade-offs require a new approach.

NILM, nonintrusive in the sense that all hardware is upstream from the loads, determines the operating schedule of different loads strictly from measurements made at a utility service entry point. New sensors have extended the NILM concept from its beginnings in electrical monitoring. It is now possible, for example, to nonintrusively monitor the consumption and operation of water at various loads in a building [10]-[12]. New versions of the NILM provide remote access and collation of consumption data and diagnostic information without requiring excessive communication bandwidth. This "small data transmission" approach relies on the availability of inexpensive computation and flexible, custom database management tools for processing data before expensive transmission resources are required, minimizing network demand and distal storage capability.

Nonintrusive Load Monitoring (NILM) has unrealized potential in facility and microgrid applications. Microgrids, such as those powering austere military base camps, are uniquely suited to benefit from NILM. They have finite, often standardized loads (same or similar on every camp), transient occupants (and thus economically indifferent), and rely heavily on outside support for fuel and repairs. Stateside facilities powered off a larger grid share a similar but perhaps lower degree of indifference from the users, tend to have a broader array of electronic loads available to them, and most likely have a dedicated maintenance staff with some continuity. Nevertheless, both austere and enduring facilities can economically benefit from consumption feedback provided by the NILM system.

There are three main subtasks within our monitoring system: data acquisition, data processing, and data management (Fig. 2-1). The output of this system is actionable feedback to the customer.

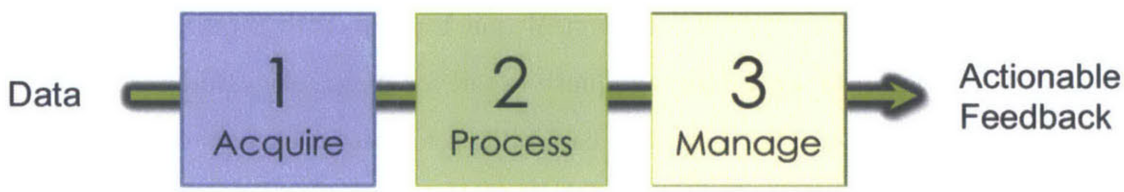

Figure 2-1: Three main subtasks of NILM

2.1 Data Acquisition

The latest MIT version of NILM uses commercial off-the-shelf sensors and a microprocessor to measure voltage and current at a 120-480 V service entry sub panel, generator, or other power source. Fig. 2-2 shows a schematic of the hardware.

We measure voltage and current are using LEM LF 305-S current transducers (CT) and LEM LV 25-P voltage transducers (VT). These sensors scale down the primary voltage and current values through a transformer, where the secondary current is routed through a known resistor across which voltage can be measured. The measured voltage for current is proportional to the primary current by the relation

$$V_M = I_S \times R = I_P \times K \times R \tag{2.1}$$

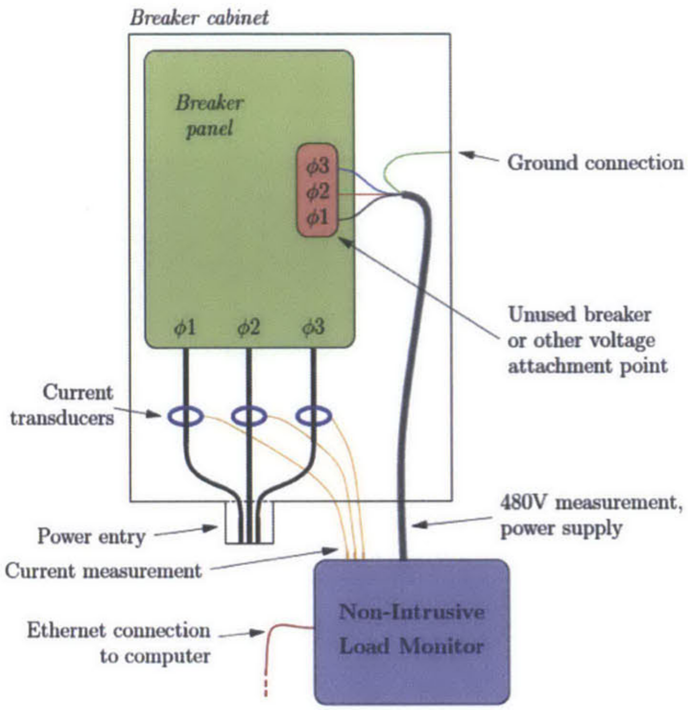

Figure 2-2: NILM Schematic

The symbol K is the transformer ratio. These Hall-effect sensors feature a relatively fast transient response time and bandwidth, more than sufficient to perceive current and voltage changes on the scale of tens of microseconds. A 16-bit ADC data acquisition board (LabJack UE9 Microprocessor) requests samples from the LEM sensors at a programmable rate that is dependent on the desired stream resolution and I/O interface. For all of my field tests, I used 8 kHz over Ethernet cable.

The sample rate can be varied as needed for any given application. The 8 kHz rate is frequently used in our monitoring work because it is sufficiently adequate to resolve utility harmonics with substantial resolution to at least 13th harmonic, useful for transient recognition. This relatively high rate can capture small high frequency signatures like motor principle slot harmonics that are do not present themselves at integer multiples of the utility frequency. The high resolution improves averaging accuracy and serves to reduce the probability of overlapping transients.

The microprocessor, sensors, power supply, and voltage transformer are integrated with a custom PCB (ver 2.1) designed by Dr. John Rodrigeuz and manufactured by

NEMOmetrics. The finished package is bundled into a weather resistant enclosure. See Fig. 2-3.

Figure 2-3: NILM Hardware enclosure

Data is stored and processed on a computer co-located with the measuring devices and connected via Ethernet cable. The processor requirements are somewhere between standard desktop chips and Raspberry Pi A models (256 MB), with the least expensive desktop models of today being more than sufficient. Storing 8 kHz samples has substantial memory requirements. Indeed, one hour of raw voltage and current measurements require more than 0.5 GB of storage per phase. We recommend hard drives built to support heavy use as this drive will be continuously written to. Western Digital Black units have performed well.

Typical installation is indoors, where power is reliable, Internet connectivity is available, and equipment (monitor, computer, and sensors) is unaffected by the weather. The system draws power from the panel through a single-phase auxiliary circuit breaker (15A is more than sufficient). Military specifications require more durable, weatherproof versions as a FOB offers little in the way of shelter near the panels and generators. At our test site in Ft. Devens, the electric panels are outdoors. Moreover, they are separated from the living quarters by more than 20'. We developed a deployable system (DepNILM, pronounced "deep" NILM, Fig. 2-4 that assembled the computer, NILM enclosure, uninterrupted power supply (UPS), ventilation system, and the associated wiring into one Husky-brand Tuff Box. The box sits out in the gravel with a single conduit connecting it to the panel housing three CT cables and five 18 AWG conductors. The conductors connect to the three utility phases, neutral, and ground. A more detailed description is found in Appendix B.

Figure 2-4: Deployable NILM design (DepNILM)

There are several ongoing energy-saving initiatives to monitor and gather data across military installations. Many quality data acquisition meters on the current market are being employed. One popular model used at Ft. Polk, the Base Camp Integration Lab at Ft. Devens, and other locations is the Shark 200. Their sample rate is reportedly between 24-30 kHz, but the data storage is capability is only 4MB. After the preprocessing, the samples are only stored on the order of once every 5 minutes. Such data is useful for general-purpose data logging, but not for high-resolution appliance-level analysis. Any meter can give you a monthly statement, but only a NILM system can itemize it, tell you what might be broken, and offer black-box archives to allow you to infer what user behavior was at the time.

2.2 Data Processing

Once acquired, the next step is to make sense of the data through the process outlined in the right side of Fig. 2-5.

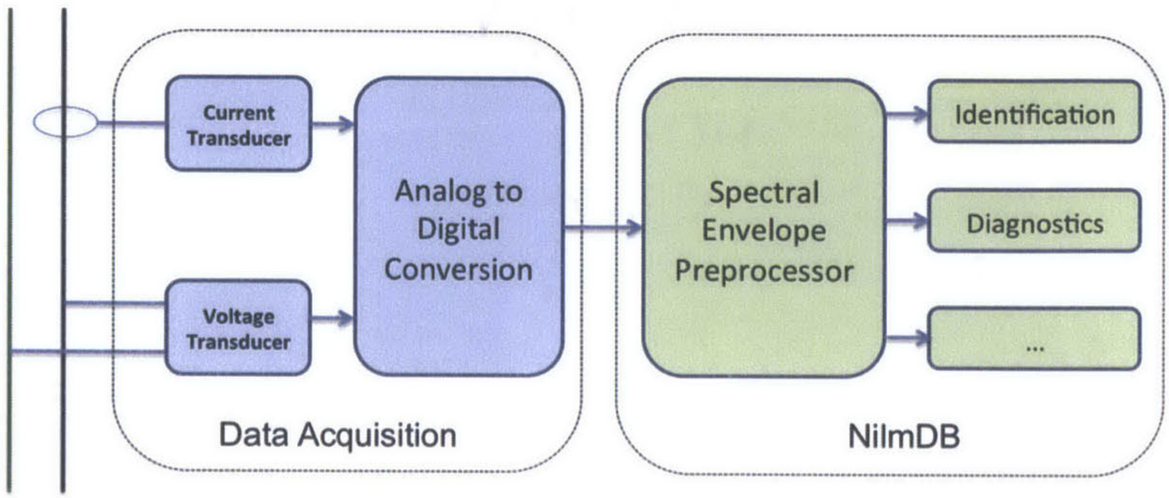

Figure 2-5: NILM Hardware and Software processes

In general, your ability to record, access, and manipulate digital data is limited by memory, processing, and network capabilities. Recall that one hour of raw voltage and current measurements from our NILM at Cottage Elementary School requires more than 0.5 GB of storage. Storing, downloading, and manipulating a single day's worth of data this size requires non-trivial capabilities. Some commercial power monitors, such as the PQube [5], attempt to circumvent large data demands by recording voltage and current data at a much lower sampling rates. The Shark 200 wirelessly communicates data to a central server, through the resolution is low and wireless transmission on military installations can be problematic for security reasons. This data goes onto a memory card that must be manually retrieved, plugged into a computer, and downloaded. Not only is this inconvenient, especially as the number of devices and geographical locations increase, but the resolution of the data samples is orders of magnitude less, thus unsuitable for disaggregating multiple loads from a single point.

When developing diagnostic algorithms, it is important to be able to access the data at many different time scales. Differentiating loads by their startup transients can require sub-millisecond resolution. However, many loads have duty cycles on the order of minutes, hours or even days. Extracting useful information, whether it be from a single-phase line cycle or a year's worth of three-phase transients, requires zoom-able data that can be plotted easily over many orders of magnitude. NilmDB, by Jim Paris, is a database system that provides exactly this functionality. NILM Manager, created by Jim Paris and organized by John Donnal, is the graphical user interface between each NILM and combines multiple systems under one central umbrella.

I applied Jim's software and John's tools to data gathered from three separate field tests. I wrote software to disaggregate loads and classify them. Finally, I developed a methodology for applying DepNILM on a large scale.

2.2.1 NilmDB

While at MIT, Dr. Jim Paris developed a low-level high-speed data management software suite christened NilmDB, short for NILM Database, that preprocesses and decimates the input data efficiently [1]. NilmDB stores data collected at data rates over 8 kHz on multiple data channels. This includes voltage and current measurements on a polyphase utility service entry. NilmDB can process, store, and time-stamp individual data points continuously at these relatively high sample rates. This custom "power system" database management software is also capable of simultaneously tracking and time stamping data from flow and vibration sensors, e.g., for monitoring other utilities like water consumption [12].

Data is grouped into streams, which are analogous to database tables. Streams have a fixed data format specifying the numeric precision and number of data points in each row. Streams also have a flexible metadata dictionary, which can store arbitrary `key:value` pairs. This is useful for storing scaling coefficients, descriptions, and units. A simple example of measuring a three-phase power supply with the NILM illustrates the operation of this database system:

1. An analog to digital converter (ADC) records current and voltage sensor read-

ings for each phase. In this example the system has three phases and therefore 6 sensors are used to measure the power

2. NilmDB assigns timestamps to each set of samples and stores the data as six columns of unsigned integers in a stream called `/acquisition/raw`

3. A daemon process on the machine then extracts this stream from NilmDB, applies the correct scaling factors to recover current and voltage from the raw counts produced by the ADC and calculates the power consumed by each phase

4. This calculation is fed back into the database as a new stream called `/acquisition/power` which has 3 columns of floats representing watts per phase. Both streams can be extracted by other processes for more calculation or for display

For more complex installations, NilmDB might be fed the results of several different ADC's simultaneously resulting in several different `/raw` streams. Additionally, there might be many different analysis daemons each requesting various input streams and producing many output streams. Such analysis daemons might calculate load detection, harmonic content, line frequency, etc.

NilmDB incorporates robust synchronization to the utility to compute spectral envelopes in the presence of noise and voltage waveform distortion [13]. Through this computation, the harmonic content is separated and grouped into real and reactive components using a series of mathematical algorithms. This step exposes energy consumption and harmonic signatures that are relatively uniquely associated with the physical task performed by a load. Preprocessing raw data eliminates distracting data artifacts, such as the underlying 60 Hz carrier wave of an AC power system, and reveals fingerprint signatures that can be used to identify load operation and diagnostic conditions.

Spectral envelopes are extracted by the preprocessor, which consists of short-term averages of harmonic content present at each of the harmonics of the incoming line frequency. As reviewed by Shaw [14], the in-phase spectral envelopes a_k of an input current signal $x(\tau)$ are

$$a_k = 2/T \int_{t-T}^{t} x(\tau) sin(k\omega\tau)\, \mathrm{d}\tau \tag{2.2}$$

where k is the harmonic index, and the quadrature spectral envelopes $b_k$are

$$b_k = 2/T \int_{t-T}^{t} x(\tau) cos(k\omega\tau)\, \mathrm{d}\tau \tag{2.3}$$

For NILM applications, the time t is referenced such that the term $sin(\omega t)$ is phase-locked to the voltage measurement, which can be achieved using a Kalman filter [15] to determine voltage-waveform zero-crossings. The averaging interval T is one or more periods of the line frequency. Under these conditions, the spectral envelopes are computed as $P_k = a_k$ and $Q_k = -b_k$, and these are the outputs of the preprocessor. The outputs are defined in this way so that the values P_1 and Q_1 correspond to the conventional definitions of real and reactive power. The spectral data is quantized and finally analyzed with software filters known as transient detection algorithms. We no longer use the Kalman filter. Instead, we use the zero crossings only and create a pure sinusoid for use in calculating power. This effectively eliminates noise from the voltage and causes us to rely entirely on current data instead. More can be read about this method, known as Sinefit, in [1, 13].

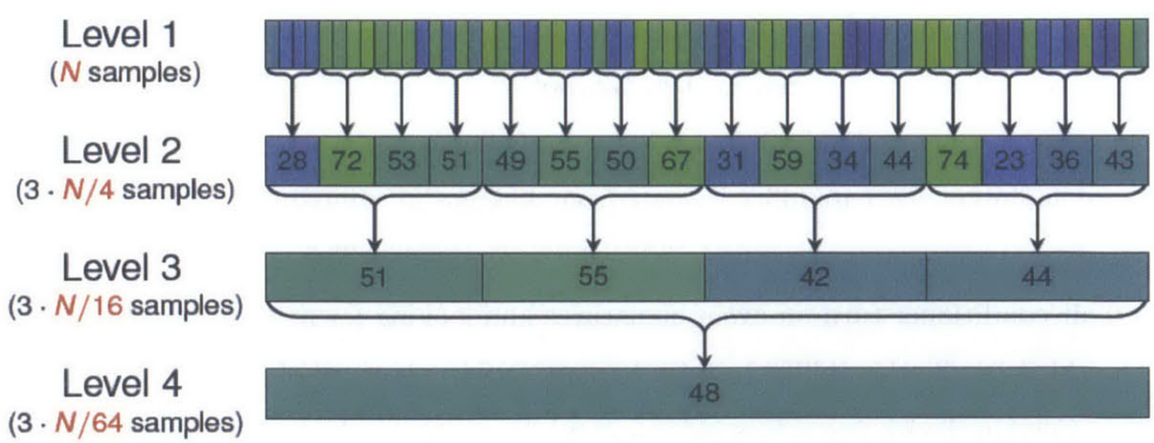

Figure 2-6: Decimation of original signal of N samples compresses and compartmentalizes stored data for targeted retrieval

Equally important, the NilmDB software can store data streams at multiple levels of decimation with a finite, quantifiable bound on data storage requirements. For example, referring to Fig. 2-6, one level of decimation is accomplished by taking every four data points from the raw data and recording the minimum, maximum, and average values. This decimated data is stored in a new file that is approximately 1/4th of the original size. NilmDB automatically decimates recorded data, iteratively storing each copy separately. If N is the original size of the data, the aggregate sum of infinite decimations is $2N$.

2.2.2 Load Discrimination

This section and the subsequent field test results represent my main contribution to the body of knowledge on NILM. We can gather data, and mathematically prepare it for analysis, but then face the problem of finding meaning amongst it. Results are based on appliance-level information. This information is obtained by using edge detection (finding ON and OFF events), load classification (determining what device caused each event), and observations based on these known appliance activities and times. Describing the results along three lines of effort, appliance-level power consumption, time of use, itemized costs, and other accountability information becomes quantifiable when individual device behavior is known. The second result, condition-based maintenance capability, comes from looking at similarities, differences, and patterns amongst the signatures from the same loads. This is akin to comparing 100 air conditioner turn-on event signatures and looking for measurable aberrations. The third result, the ability to infer human activity, is more of an art than the first two. Certain device use relays information about specific human actions. One must understand what "normal" behavior is in the context of the monitored site. I've been a Soldier for 10 years and have lived on FOBs during two deployments giving me a good feel for Soldier patterns of life. I project that understanding on the power grids being observed to extract meaning.

How to pick out individual loads from a central location

In general terms, every electrical device draws apparent (S) power in some combination of real (P) and reactive (Q) component. As described previously, these components are separated through the spectral preprocessing within NilmDB. Real and reactive power waveforms can be further dismantled into their harmonic subcomponents (see Fourier), with the even harmonics canceling because the waveforms are half-wave symmetric. The sum of the real harmonics make up the total real power, though each harmonic has individual characteristics unique to a load. The fundamental component by far has the largest magnitude, and thus demands the most attention. However, when two devices have similar real and reactive draws, their harmonics can often be used to distinguish them.

It is the understanding of these subcomponents of waveforms that enables us to find the individual appliance needles from the proverbial big-data haystack. Through the stream of AC electromagnetic energy measured at the panel, we are looking for subtle energy changes that happen in milliseconds. Because the total power flowing through an electrical panel is simply the sum of everything plugged in and operating, subtle changes in the overall power signature reveal the activity of individual devices.

As observed by Hart [16], when the time scale of observation shrinks to seconds or sub-seconds, distinct shapes in the power waveforms become "readable." I'll illustrate load discrimination with an example. Fig. 2-7 depicts 48 hours of real power data (fundamental component) taken from Ft. Devens over a weekend when troops were present. This 150-man base camp draws power through two 600A 3-phase panels. We monitored each of them separately but plotted both together. With the human eye, little about individual appliances can be ascertained from this graph. However, through our integrated system, we can zoom in and observe a one-hour section with high resolution. This is not possible with standard meters, which often only keep samples every second or more.

Zooming in on a single hour of the weekend data from Fig. 2-7 gives us the more detailed plot in Fig. 2-8. At this scale, appliance-level transients emerge. Each

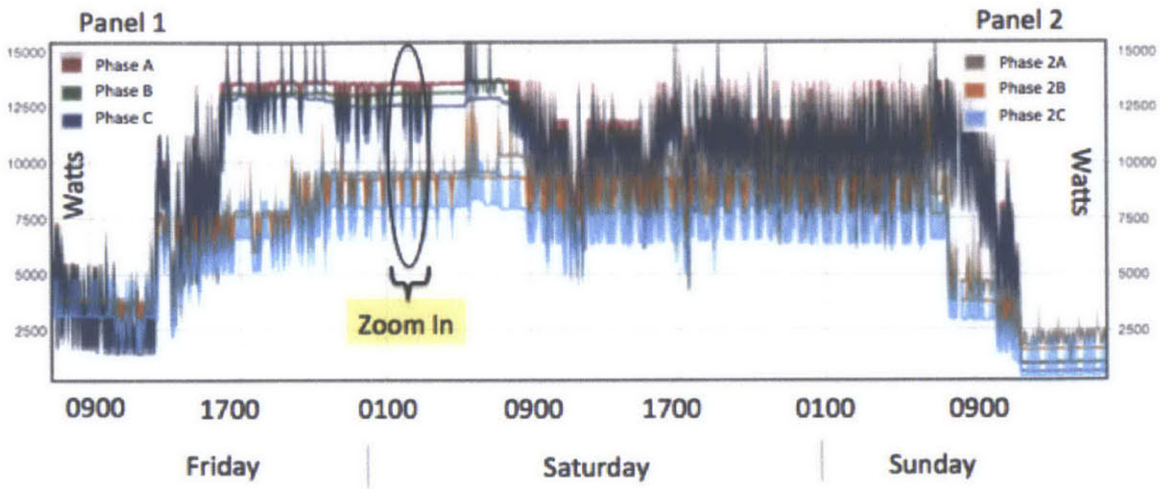

Figure 2-7: 48 hours of real power data displayed on NILM Manager

machine has a unique electrical fingerprint. As human fingerprints have distinctive recurves, deltas, creases and scars, so to do electrical fingerprints have distinctive peaks, steady-state power changes, start-up durations, and other minutia that make them stand out. The heater, for instance, is an almost purely resistive load. On start up, the steady-state change in real power is an almost instantaneous step from Off to On. A refrigerator compressor, on the other hand, is an inductive load. It is characterized by the large inrush current at turn-on required to activate the motor. This particular refrigerator is a single-phase load that cycles automatically to maintain a set temperature. When it turns on, the transient (real power, fundamental component) takes the shape of the blue trace in Fig. 2-9. The peak value and the steady-state change of the transients tend to be highly convergent over thousands of cycles. This statement applies to both real and reactive power, and extends to each load's respective fundamental and harmonic components. Also, the utility phase for a device like this rarely changes because it stays plugged into the same 20A branch line.

Computed or observed features of a device such as peak, steady-state change, or utility phase are what I call discriminating elements. Any machine transient may be characterized by one or many such elements. Once a fingerprint is taken via a training phase, where each machine is cycled independently while the NILM records, a library

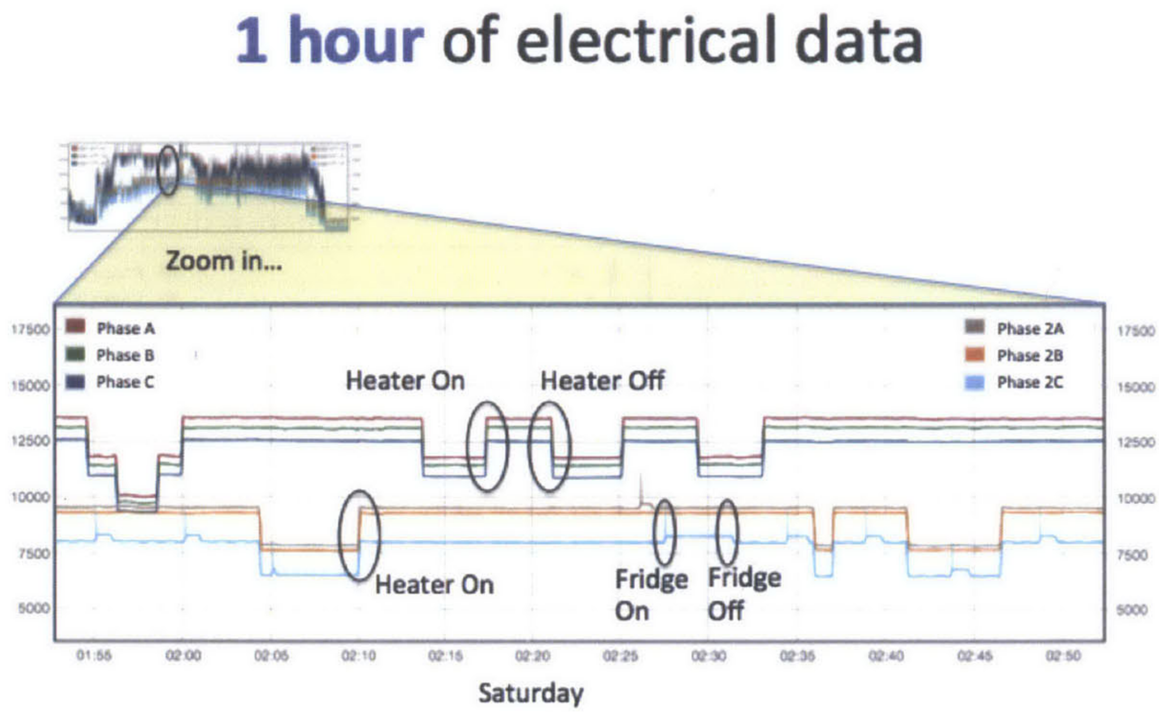

Figure 2-8: One hour of real power data showing distinct transients

of numerical traits about device and its unique characteristics can be created. The benefit of military equipment is that they tend to be highly standardized. Events recorded at one base are readily applicable to other bases. The values each library should maintain are the first differences, steady-state changes, and peaks for each harmonic subcomponent, both real and reactive. The utility phase, while often useful and consistent for larger loads, may be subject to change if a load is moved to another outlet.

Outside of the laboratory, "steady state" is somewhat of a misnomer. The gird is in constant flux and devices are constantly turning on and off. Though steady-state power changes are difficult to quantify, they are arguably the most important. Averaging helps smooth the samples but must be applied selectively. In an electrical system, steady-state change can be reliably measured as the difference between sample averages before and after the transient. Fig. 2-10 is a plot of real power samples of a refrigerator turning on. With Sample 22 as the event reference point, the previous power level is the average of some number of samples before the event. Exactly which

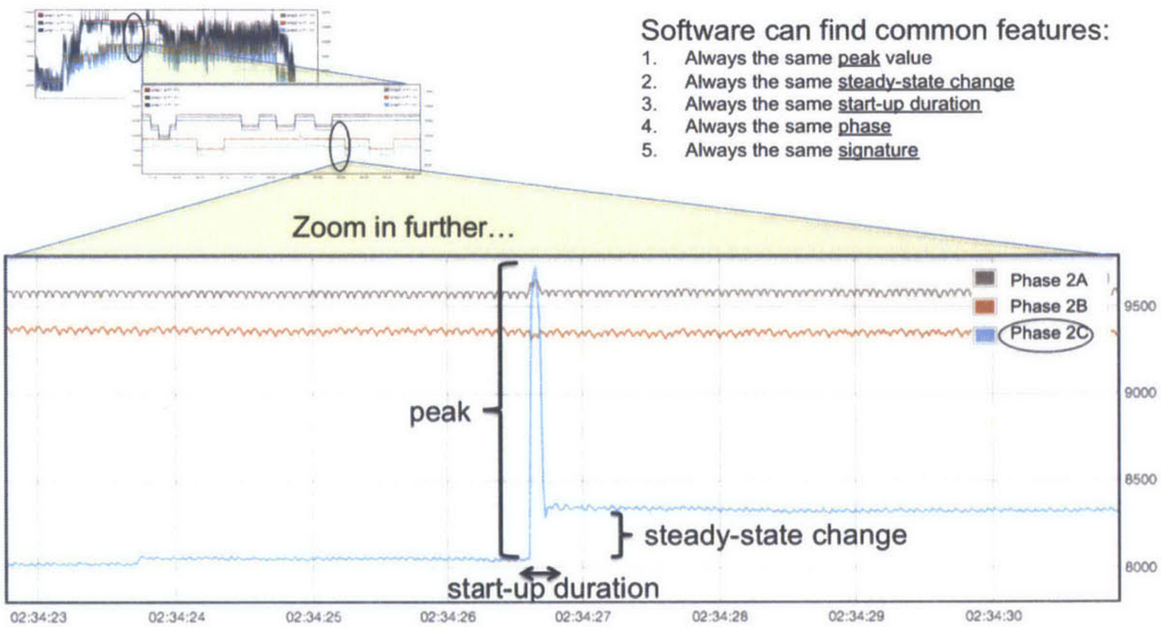

Figure 2-9: 8 seconds of real power showing details of a refrigerator transient

group of samples to use in calculating the post-event power depends on the type of transient (resistive or reactive) and the duration (number of samples) until the event reaches steady state.

In practice, I found that using a short, medium, and long-range averages were all necessary to characterize a given power network of multiple loads. For small events ($<$ 500 W), which includes most electronic loads, the transient is over quickly. Short-range averages seem to work best. For larger loads whose transients end quickly (namely resistive loads), the short or medium-range averages are adequate also. For transients that take longer to settle, the long-range average is most appropriate.

The variables that determine which average scheme to apply are the transient duration, the peak/steady-state disparity, the load type and the load density. In load dense environments where many on/off events are happening rapidly, shrinking the window of time for each transient reduces the likelihood of two transients overlapping and skewing the steady-state change values. Shrinking the number of points used in each average is also advisable. When two different loads have the same long-term steady-state change, their short or medium average may be distinct and thus be used to discriminate them. If the peak and the steady-state change are nearly

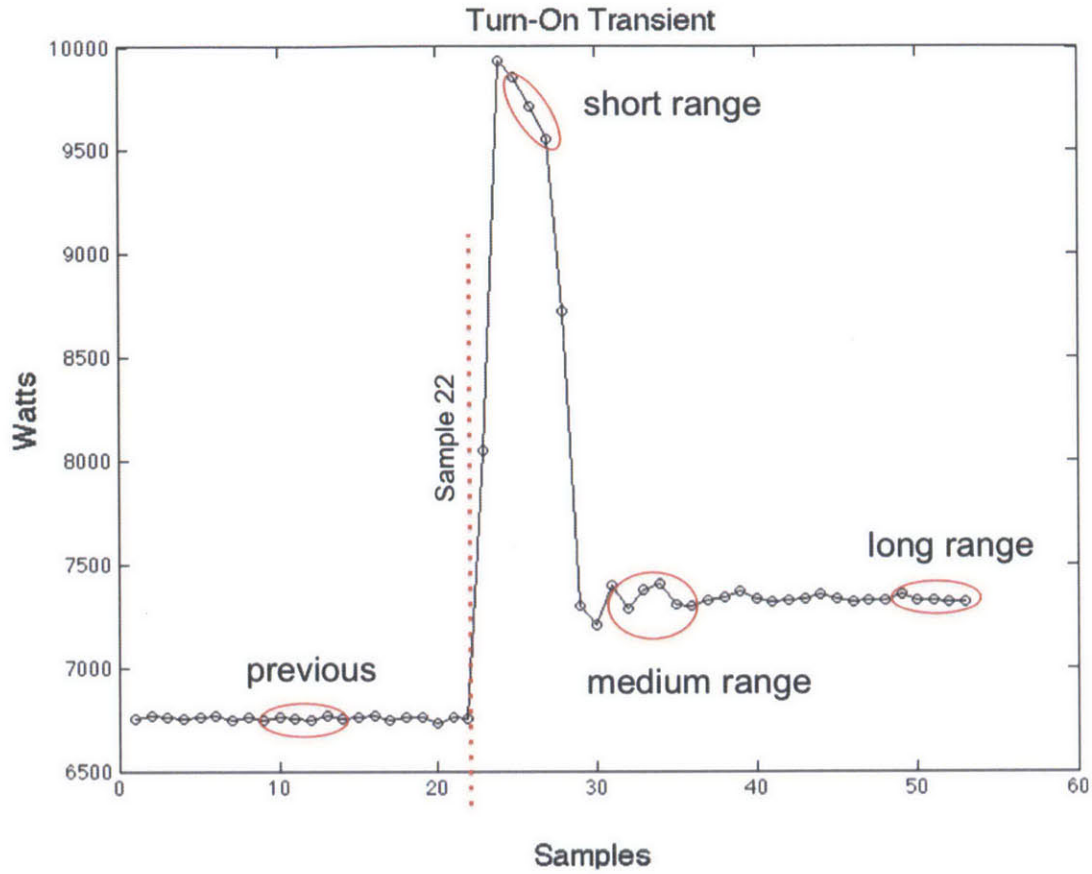

Figure 2-10: Graphical description of edge detection algorithm

equal (common in heaters or other resistive loads where the transients are almost step functions), a short-range average is sufficient. In general, there is an empirical art to calculating the steady-state change that is ultimately unique to each power system. However, similar facilities tend to have similar load types and thus similar patterns.

In general, transient features vary with load class, and distinctive transient features are found in the shape of turn-on transients in real and reactive power and higher harmonic current content, as well as non-line-locked frequency content like motor principle slot harmonics. Turn-off events tend to be more abrupt in time, and are typically identified by clustering the inverse or removal of steady-state characteristics associated with the load.

Edge Detection algorithm

I developed an edge detection algorithm that picks out "events" based on the method shown in Fig. 2-10 which I define as real power changes (fundamental component, specifically) that meet two criteria:

1. fluctuates more than a defined amount within a 2 millisecond interval

2. first difference (between the reference sample and the next sample) is greater than a defined amount

In other words,

$$|\Delta_{P(f)}| > x \tag{2.4}$$

$$|P_{f_{n+1)}} - P_{f_n}| > y \tag{2.5}$$

Where:

1. P_f = Fundamental component of Real Power

2. P_f = nth sample of Power

3. x, y = user-defined thresholds

The reference sample is the 5th sample in a 2-second-wide moving array of data, each row of which is a 60 Hz sample set of time-averaged data. Most start-up transients complete their course and reach steady state within two seconds, though often it is within half a second. The delta value is calculated for each sample by subtracting the 2-point power average before the reference sample from the 2 or 5-point power average after the sample in the following way:

$$\Delta_{P_5} = P_{avg(10-14)} - P_{avg(1,2)}(ONevents) \tag{2.6}$$

$$\Delta_{P_5} = P_{avg(7,8)} - P_{avg(1,2)}(OFFevents \tag{2.7}$$

Using the same reference sample, the first difference at that point is equivalent setting n from equation (2.5) equal to 5. The x and y values are dependent on the noise in the system and the size of the loads you are looking for. In the field tests we conducted, the five largest loads (fans, heaters, compressors, refrigerators, and pumps) were well above the noise layer. Since the cumulative consumption of those loads represents more than 90% of the power consumption in the power systems we studied, they were the focus of conservation and feedback efforts anyway. Even with the variables x and y set as high as 400 W, all events were readily identifiable. An example of the code used for Ft. Devens is in Appendix A under the "EDGE DETECTION ALGORITHM" section.

It is possible pick out events smaller than 400 W. Indeed, loads lower than 100 W are perceptible as demonstrated in the IT room at Cottage Elementary School. The sensor's dynamic range is adaptable to the anticipated current range of the loads. With 16 bits of analog-to-digital conversion, typical building distribution network noise floors observed in the field might provide 13 to 14 usable bits of resolution, or a one part in ten thousand resolution that can be flexibly scaled over the anticipated current range. Our sensors in the field test are calibrated to resolve well in the 300 W to 3 kW range (per phase), targeting the larger machines that would have the greatest impact on consumption. In practice, noise in the system can vary. The presence of electronic loads, variable speed drives (VSD), grid stability, and even the frequency of events. The average "quiet" system is on the order of 50 W (peak to peak), but noisy systems may fluctuate up to 1 kW or more (see Fig. 2-11). Averaging techniques help discriminate smaller load transients, but this study focused primarily on the larger, more costly machinery whose signatures are readily extractable.

It is more cumbersome to distinguish small loads for two reasons: many of their fingerprints have overlapping parameters, and the noise in the system can drown out their shape. As in the larger loads, more criteria must be used to define just what characteristics guarantee an inimitable solution. Lights, cellular phones, laptops, battery chargers, and other electronics tend to require several parameters to be defined in order to uniquely identify them. From my experience in facilities whose consumption

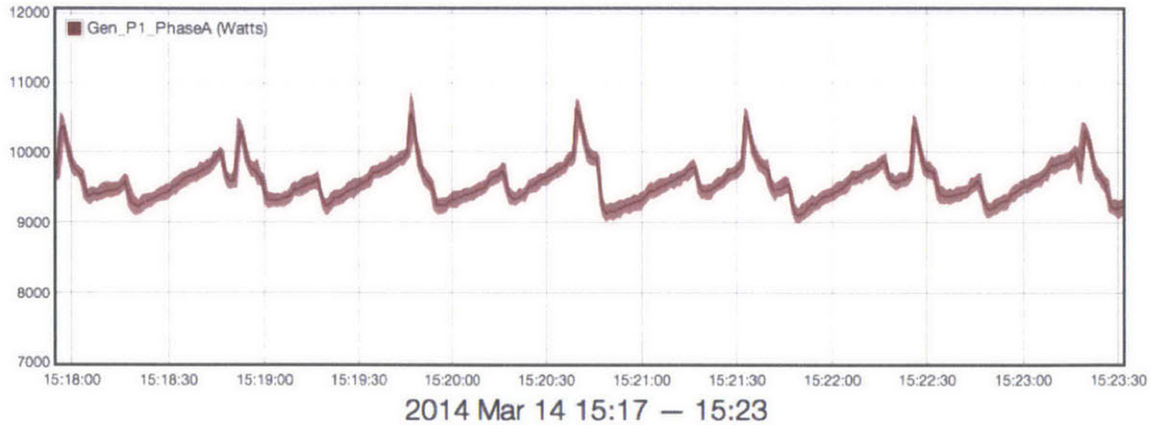

Figure 2-11: Power fluctuation due to variable speed loads

volume average less than 500 kWh, the bang is not worth the buck as far as cost savings goes, and the focus should be on the larger loads. This changes when interpreting human behavior is the goal as battery-charging events become more important than the overall consumption.

Load Classification algorithm

The classifier filter matches each event to a corresponding device. When events are found and discriminating elements recorded, these elements are compared to exemplars in the library. A comparison is made between each edge detected to a list of known values. This implies that a library of load characteristics for each device in the system must first be created during a training phase. Exemplar data can be generated in a variety of ways, including manual activation of individual loads during a training period, or automated machine learning on a data set. To ensure reliable identifications, rich streams of relevant data should be used to create exemplars, including real and reactive transient power shapes, higher harmonic content, and state models of load behavior.

Most often, the values we compare include derivatives and steady-state changes of real power, reactive power, and some harmonic component. More dimensions are possible and sometimes necessary to discriminate individual loads on busier systems. Using a python script, I created a filter bank of known load values using IF statements.

For instance, the heater transient of a 60 kW ECU can best be characterized by the change in steady-state wattage, between 2.9 kW and 3.6 kW. A single-dimension comparison is sufficient to uniquely identify this load because no other load comes close to this value.

Other loads like evaporator fans often require multi-dimensional criteria. A particular fan has a steady-state change of 500-700 W, a range that also includes the refrigerator and water pumps. Here, reactive power change, peak values, or other harmonic component changes are also required to make the classification unique. In practice, as many discriminating elements as necessary are used to classify a load uniquely.

The need for a training phase is obvious, though over time the eye gets accustomed to what to look for. HVAC units, pumps, compressors, fluorescent lights, and others all tend to be ubiquitous devices in any facility. Rational inferences as to what loads would be present coupled with an intermediate level of understanding of "common" signatures make virtual training possible for a human. For instance, most compressors have a large power spike at turn on. Most heater events look like step functions. Most laptop computers and TV screens "look the same" graphically. While every brand may differ, a trained operator can sift through the events using NILM Manager and make highly accurate assumptions about what is going on without knowing anything about the site. It is entirely possible that this same capability can be programmed by a skilled computer scientist so as to remove humans from the loop entirely, though this eventuality will be left to future research.

For a given electrical appliance, the discriminating elements tend to fluctuate over a measurably small range. Thus, empirical limits can be established based on observations of only a few events under varying conditions. For the refrigerator, varying conditions consist of varying ambient temperature and internal temperatures. Since the power draw primarily emanates from the compressor motor, a refrigerator in a hot room that has just been loaded with warm food will theoretically draw the maximum current. Likewise, the refrigerator that has been filled with ice and is located in a cool space will draw the least. The power fluctuation between these extremes tends to be within reasonable bounds, about 12% standard deviation. Generally speaking, standard compressors controlled by thermostats (like many air conditioning units and refrigerators) consume steady amounts of power while running. Instead, it is the run-time durations that will fluctuate under varying temperature conditions as described above. In other words, rather than modulating its speed based on temperature feedback, motors simply stay on for longer or shorted periods but run at the same speed. An example of the classifier filter used used for Ft. Devens is in Appendix A under the "LOAD CLASSIFICATION ALGORITHM" section.

Many load identification methods have been researched, a brief summary of which can be found in [17]. In simpler systems like residences with perhaps 10 loads, most of which are for short-term use, load classification is a straightforward comparison of real and reactive power magnitudes. In this study of a commercial three-phase system at the Cottage school, there were over 60 loads that, for the most part, operated automatically. Too many signatures had similar kW and kVAR characteristics at the fundamental frequency to use only fundamental features. To these features, we added harmonics, first differences, and utility phase.

2.3 Data Management using NILM Manager

Since NilmDB must be collocated with the sensors in order to process the high bandwidth data and to efficiently calculate secondary streams based on this data. It is inconvenient and impractical to continuously visit each installed NilmDB, so we have designed a centralized management infrastructure.

The primary difficulty in monitoring power consumption is extracting useful information from the very large amounts of data collected. In order to develop algorithms for this purpose, data should be accessible over many different time scales. Loads can be differentiated by their startup transients, which can require sub-millisecond resolution. However, many loads have duty cycles on the order of hours or even days. To effectively process this data, a view of the usage by sub-second, hour, or month must be available and easily plotted. NilmDB [1] is a database system that provides exactly this functionality. Dynamic access, graphical tools, and data manipulation is made possible by NILM Manager [2].

Former Army officer John Donnal developed tools to analyze and manage the decimated data in a compact and user-friendly web accessible platform. He calls is NILM Manager, through which he accesses the decimated files. Specifically, the application provides access to remote NilmDB's allowing any user with an Internet connection to view data graphically and manipulate the streams (zoom in and out). The block diagram in figure 3 shows the interaction between NILM Manager, NilmDB, and the user. Most of the graphs throughout this thesis are courtesy of NILM Manager.

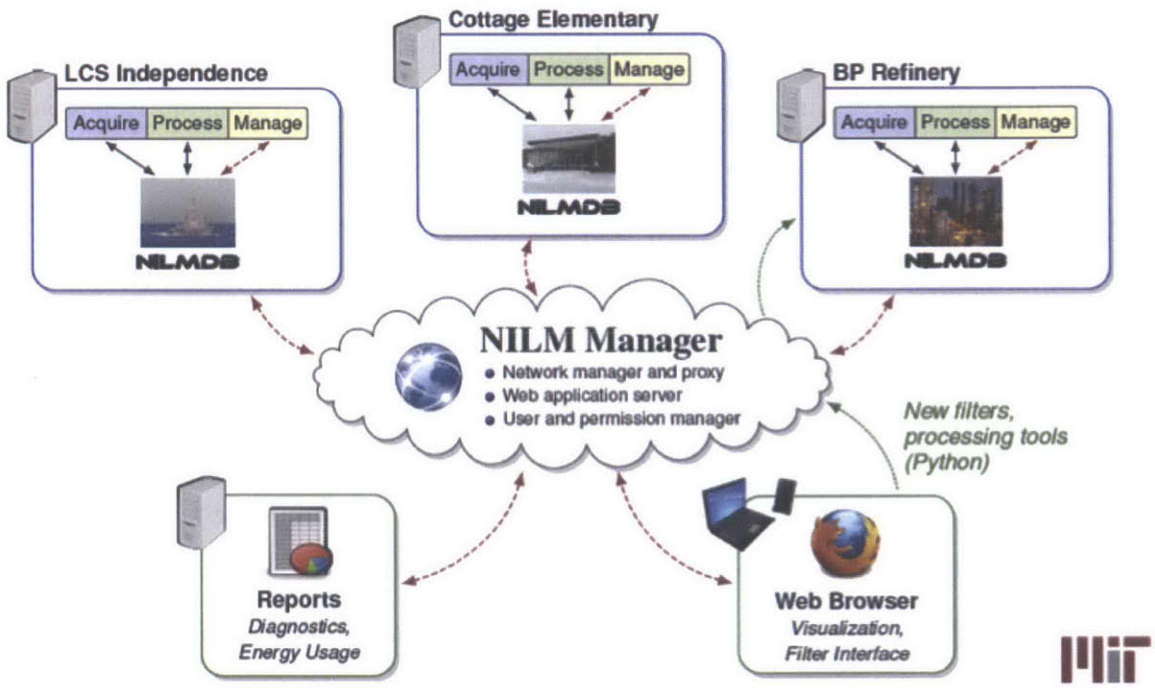

Figure 2-12: Nilm Manager interfaces between NilmDB and the user

NILM Manager is a secure central portal serving NILM computers, communicating over the web. It supports remote access to information from NilmDB installations anywhere in the WWW-accessible world. NILM Manager minimizes network traffic with essentially no limits on data analysis by offering two key transmission-bandwidth minimizing features. NILM Manager can access useful information at any time scale, from fractions of milliseconds to years, by transceiving only the dedicated, decimated packets of data appropriate for a particular time scale, always limiting requests to a finite packet size, e.g., what can comfortably fit on a computer screen for viewing. This reduces data transmission significantly, on the order of bytes rather than gigabytes.

In the "reverse" direction, NILM Manager can transmit small packets of interpretable Python code to a distal NilmDB installation, enhancing any particular NilmDB with new data analysis capabilities flexibly programmed in a MATLAB-style environment. New analysis is conducted on the distal NilmDB computer, without the need for high bandwidth data transmission. New NilmDB database streams can be constructed from these new signal-processing requests, and interrogated with small packet transmissions back to the NILM Manager. NILM Manager can coordinate the operation of dozens to hundreds of NilmDB installations.

Security of data is important. With NilmDB and NILM Manager, bulk data is stored locally and all data transfers are encrypted via VPN. NILM Manager brings every installed NILM under one central viewing umbrella as shown in Fig. 2-12. The Manager tool acts as a central server handling authentication and authorization. Such data protection is on par with most commercially available security systems. NILM Manager also greatly enhances analysis capabilities with the ability to interactively zoom over any recorded period of time. Analysis previously limited to one hour at a time due to transferring, processing, and graphing constraints, becomes faster and easier without sacrificing resolution [2].

Wrapping the NilmDB installations behind a manager allows the database to focus on processing data rather than authentication mechanisms, user interfaces, and the numerous other aspects of a modern web platform. This architecture also allows placement of NilmDB's behind network address translation (NAT) and Firewalls. Most residential networks share a single public IP between all the connected systems through NAT. Devices behind NAT cannot be contacted by the public Internet without special configuration on the router. Setting up such a configuration may not be possible or desirable. We avoid this issue by having the NilmDB instances initiate SSH tunnels to the NILM Manager. Each NilmDB knows the location of the NILM Manager (which has a public address) and uses preloaded authentication credentials to establish a tunnel. Because NilmDB initiates the connection with the NILM Manager rather than vice-versa, NAT and Firewall traversal is not an issue. Log on to see NILM Manager data and capabilities at `http://wattsworth.net`.

2.4 Methodology for Quantification of Results

2.4.1 Calibration

NILM Manager plots are in analog-to-digital converter units. Each sample n must be calibrated to reflect true voltage, current, and power according to the linear equation (2.8).

$$y_n = mx + b \tag{2.8}$$

NILM Manager uses linear equations because the voltage and current transducers are defined by the transformer equations, which are also linear. Voltage and current are also offset, oscillating between 0 and 65536, the discrete range of sample magnitudes defined by the 16-bit ADC. In the background of the website, these scaling factors are independently settable for each data stream.

An inexpensive FlowPro space heater (model 01038) is useful for converting arbitrary units to real ones. The heater is a good reference for current and power because it draws more than 10A of current and has a very good power factor (better than 0.99). Using a Fluke 321 clamp multi-meter measurement of grid voltage and space heater current at Ft. Devens, the space heater's real power consumption (P) is Eqn. (2.9) and the scaling factor for power is (2.10).

$$P_{known} = V_{knownRMS} \times I_{knownRMS} \times pf \tag{2.9}$$

$$P_{known} = P_{unknown} \times x_{Power} \tag{2.10}$$

The voltage and current offsets are found by taking the average of positive and negative peaks from the plot (within the same line cycle) during steady state, shown in (2.11) and (2.12).

$$V_{offset} = mean(V_{peak} + \mid -V_{peak} \mid) \tag{2.11}$$

$$I_{offset} = mean(I_{peak} + \mid -I_{peak} \mid) \tag{2.12}$$

The scaling factor for voltage is found by assuming the grid voltage to be purely sinusoidal (though it can be verified, this tends to be the case) and determine the peak voltage by the Eqn. (2.13).

$$V_{knownpeak} = V_{knownRMS} \times \sqrt{2} \tag{2.13}$$

Calibration from Latex comparison On the data plot, determine the peak in arbitrary units and finally resolve scaling factor $x_{Voltage}$ using Eqn. (2.14)

$$x_{Voltage} = \frac{V_{knownpeak}}{V_{arbitrarypeak}} \tag{2.14}$$

In finding the scaling factor for current, no assumptions are made about the quality of the sinusoid. While recording with the NILM, cycling the space heater produces Fig. 2-13. Isolating a few (at least one) line cycles of the space heater before and after turn-on, compute the RMS values of current.

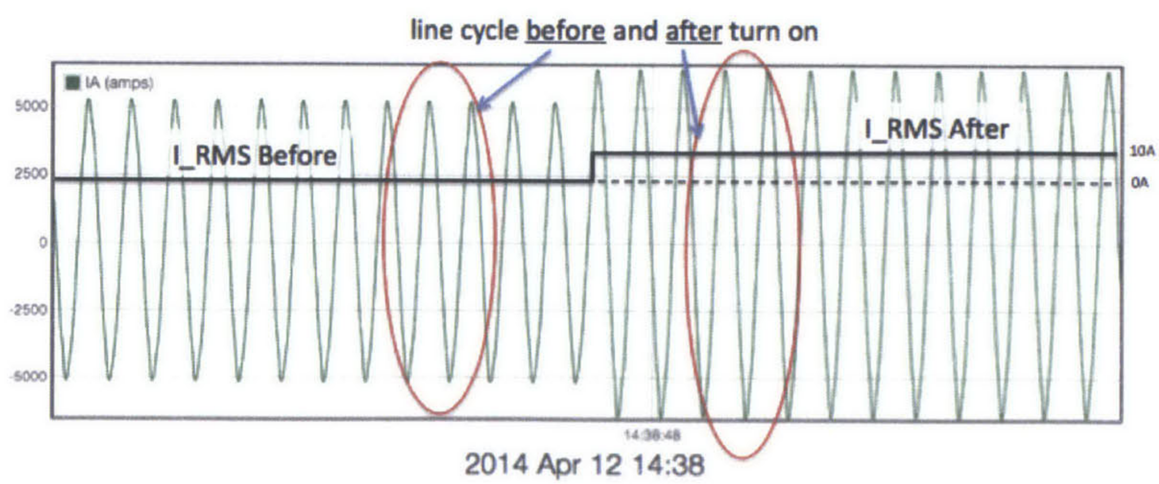

Figure 2-13: Current plot of a space heater turning on

Using the RMS current value from the Fluke meter, a comparison using Eqn.

(2.15) is possible.

$$I_{knownRMS} = I_{RMSafter} - I_{RMSbefore} \times x_{Current} \tag{2.15}$$

Solving for $x_{Current}$, the scale factor for current is Eqn. (2.16).

$$x_{Current} = \frac{I_{RMSafter} - I_{RMSbefore}}{I_{knownRMS}} \tag{2.16}$$

I'll insert some notes about the various tests site calibrations also. Fort Devens has sub meters at every tent permitting a true comparison between power and scaled power plots in NILM Manager. Individual appliances were cycled in isolation, averaged by like item, and an appropriate scale factor for power on each panel was calculated. The FOB on Ft. Polk did not have the same true data to compare against. The generators at the hospital had digital control panels permitting instantaneous balancing between the current and voltage data and the generator output.

2.4.2 Determining Power Consumption

To determine the power draw for an individual device, three pieces of information are required: the ON-time, OFF-time, and the average power draw when on.

The ON-time and OFF times are determined by transient detection algorithms and software filters as explained previously. The average power draw for each device is determined by taking the mean power drawn as determined by averaging several on/off events. Despite what happens during the initial moments of start-up, the power consumption of most devices can be modeled using step functions. Fig. 2-14 shows the signature of a commercial boiler draft fan turning on. Because it runs for several minutes at a time, we can assume for now that the area between the data curve and dashed step function during the first second becomes increasingly negligible the longer the machine operates. While fans, pumps, and other motorized devices do not draw constant power, the model suitably approximates the total power draw. We will demonstrate this later in the field exercises.

With the on and off times stored, the magnitude of the step function multiplied

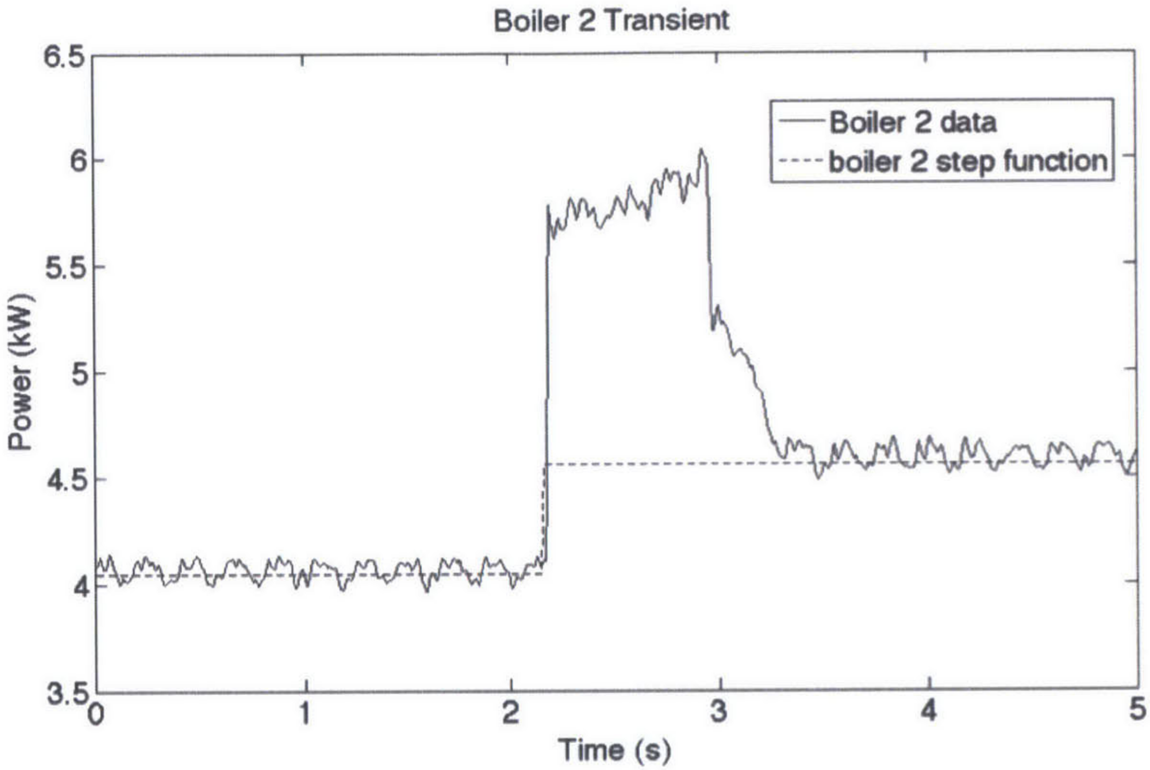

Figure 2-14: Boiler transient modeled with a step function

by the time duration (converted to kilowatt hours) becomes the power consumption of each device. It is simply the area of a rectangle. At any one time, the total power is the sum of the draws from each individual device. In Fig. 2-15, 4 heaters turn on, run for a while, and turn off. They are each three-phase loads, so the total power becomes the summation of rectangular areas of each respective phase of respective device. If a single phase of a heater consumed 3 kW when running, and a heater ran from 11:08 to 11:18, the power is

$$P = 3[kW] \times .167[hrs] \times 3[phases] = 1.5[kWh] \tag{2.17}$$

Using this same math, the total power for the four heaters in Fig. 2-15 is approximately 3.9 kWh.

Certain load types defy this logic. This includes variable speed or variable frequency drives (VSD, VFD), battery chargers, and other modulating or exponentially decreasing devices. For the purposes of energy scorekeeping, these types of loads are

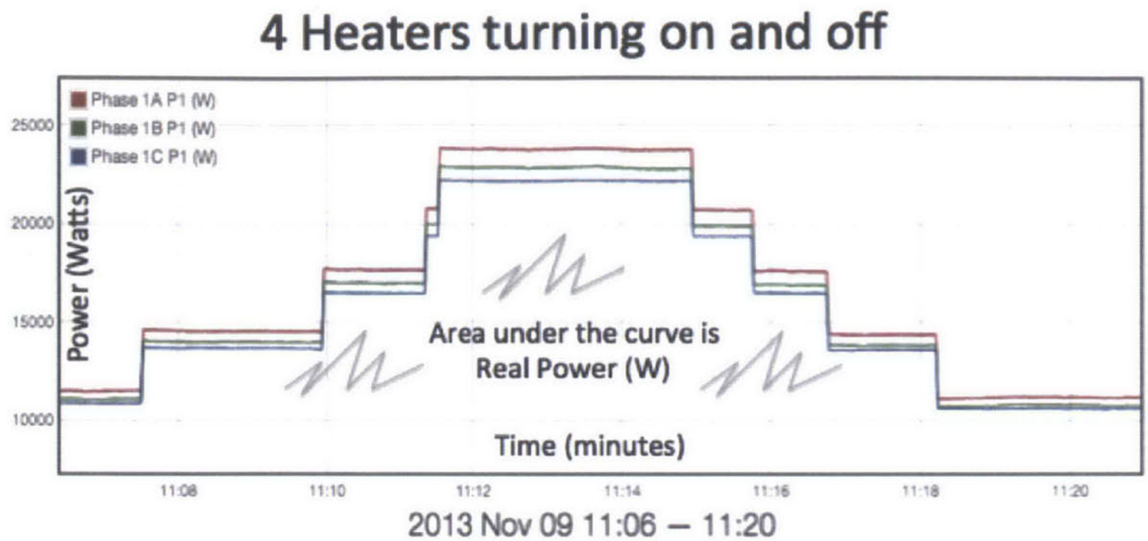

Figure 2-15: Power consumed is the area under the curve

only very loosely approximated in my work. For gathering behavioral information, however, such as what time a pump turned on, the shape of the power waveform after the transient my be irrelevant.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

Field Test at Cottage School

The Cottage Elementary School in the Sharon School District in Massachusetts has served as a fascinating and representative test bed to demonstrate the NILM approach for monitoring. The school is actively used by hundreds of students and teachers. The load sizes, types, and levels of automation seen here are uncommon to residences, where many field tests have been conducted previously [18]-[23]. Many of the devices are systems of loads, an extension of multistage loads. The boiler, for instance, has a draft fan, blend pump, actuators, burner controls, and a transformer igniter. Each component has a unique signature and a prescribed sequence of operation in non-pathological operation.

Cottage school's electrical network suffers from two common problems that are not atypical to public buildings. First, the building is over 50 years old and has seen major renovations since opening. Second, upon investigation, the design of the electrical system (the blue prints) and the actual wiring of the panel do not match completely. Further, what is provided in the plans (including the scribbles on the electrical panel) was incomplete. Major equipment has been replaced and rewired. Certainly the maintenance team has changed over multiple times, leaving several unknowns as to what load is connected to what breaker. Through or initial testing, we were able to ascertain the majority of the unknowns.

An electrician installed the NILM system on a 3-phase sub panel known as the emergency panel (EBPP). Fig. 3-2 shows the connection scheme. In the interests

Figure 3-1: Cottage Elementary School in Sharon, MA

of student education, a website was also established to graphically display the pre-processed power consumption data from that panel (`http://wattsworth.net`). The EBPP is the critical electrical node servicing the school's communications, heating system, kitchen appliances, septic system, and other important loads. In the event of a power outage, the backup generator supplies power to this panel enabling the school to provide shelter, heat, food, and communication capabilities to the surrounding community. There are more than 30 sub panels at Cottage, but the EBPP accounts for about 1/4th of the school's total electrical power consumption during winter months.

3.1 Electrical System Background

In cold weather, the largest power draw on this panel is from the machinery involved in creating and distributing heat. Cottage's heat system is a closed-loop reverse-return hot water system regulated by an integrated building control system (see Fig. 3-3).

Operation of the heat system depends on several user-established inputs. If the

Figure 3-2: Installation of NILM at Cottage Boiler Room panel

outside air temperature is below a prescribed level, the boilers will operate according to water temperature settings in the loop. If the return-loop temperature is below a lower threshold, the boiler will operate until the water reaches an upper threshold. To prevent cracking inside the boiler, a blend pump mixes return water with supply water. Cottage's boilers heat water using natural gas, but the electrical signatures of the draft fan and blend pump are detectable during operation. The Variable Frequency Drive (VFD) circulation pumps pressurize the supply loop and move the water through the piping system to the school. The VFD operational speed depends on system pressure. Head pressure, like voltage, maintains the desired flow inside the system. Upper and lower limits are set and measured by the pressure differential. When the pressure is too low or too high, the pumps will speed up or slow down by increasing or reducing voltage frequency.

Cottage's emergency panel has many other loads unrelated to the heat system, including the IT Room, hot water pumps, large kitchen appliances, etc. These are listed by circuit-breaker number in Fig. 3-4. The minimum and maximum kW values correspond to the range of their loaded and unloaded power draw. Some devices, including lights, are frequently on or off, while others continuously operate. Others quietly consume power keeping their internal systems running on standby, even when not in full use. Also of note is the number of 3-phase loads, indicated by multiple breakers with the same label, such as the Make-Up Air fan in the kitchen. The total

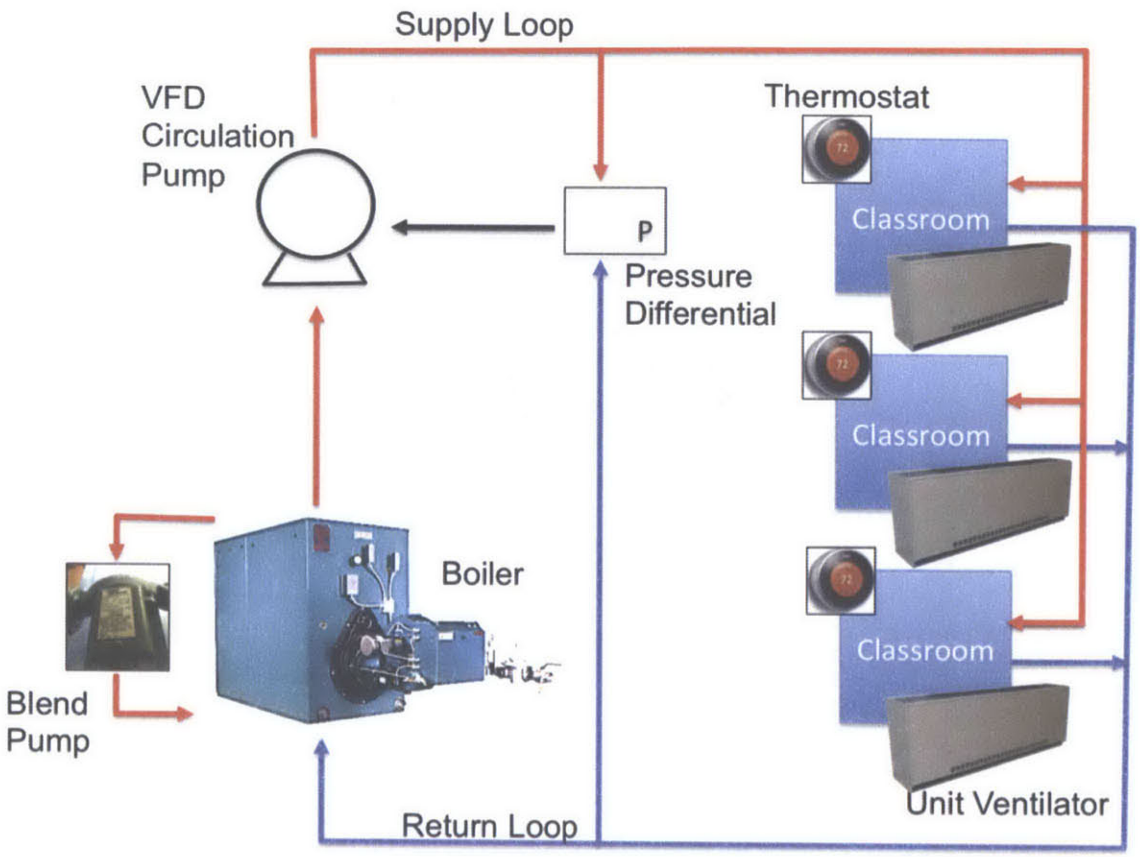

Figure 3-3: Heat system at Cottage

draw from such loads is the sum of the power drawn on each phase. For instance, the circulation pumps are 3-phase VFD motors drawing a maximum of 1.5 kW per phase, or 4.5 kW total. Other loads in the building create a base load present on the panel electrical phases. Loads not of interest for tracking the heat system, appliances, and the IT room were identified as "unknown" base load not of interest in the current survey.

3.2 Load Disaggregation at Cottage

Power signals at a central point are simply the sum of each individual load's power draw. As an example, Fig. 3-5 depicts one hour of the collective power signal on phase A on Monday, March 26, 2013, from 12:00-1:00 PM. The following figures will demonstrate how the filters decompose this signal into its individual loads. With the DC

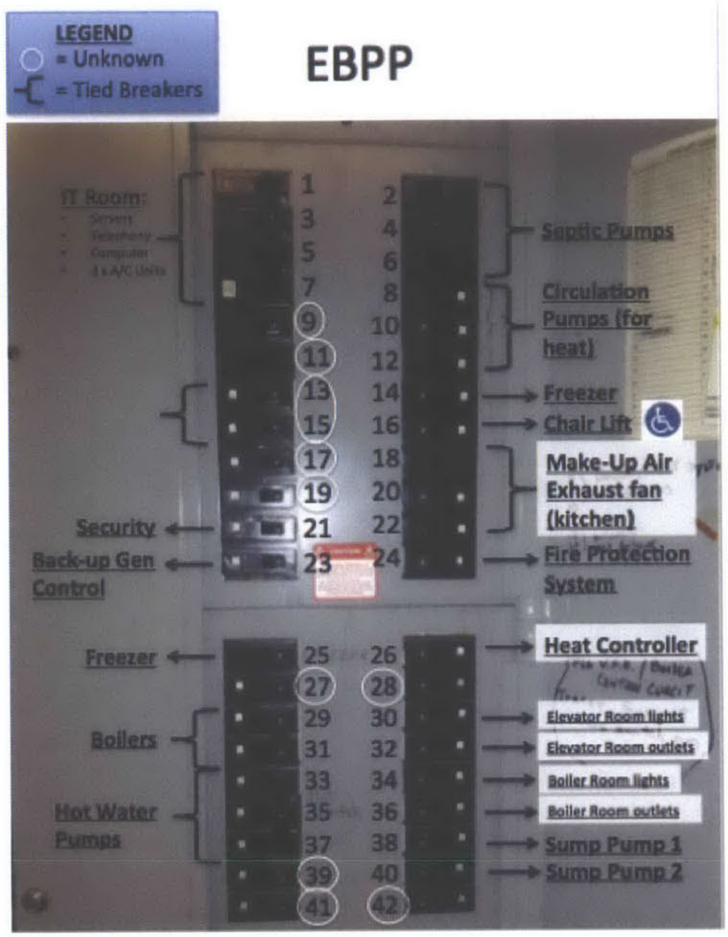

Figure 3-4: Load Schedule of Emergency Panel labeled during NILM training phase

offset removed, each transient is then modeled with a step function, the superposition of which effectively reconstructs the original signal. The filtering software detected 23 transients during this hour. Two key features are apparent from the graph: the transients and the baseline. One refrigerator transient is circled. The baseline, about 3.2 kW, is the power draw of all machines that remained on for the entire hour. Graphically, it is the low point on the plot. From Fig. 3-8, we know that other phase-A loads that make up the baseline are the Freezer, Make-Up Air Unit, circulation pumps, and several smaller loads (control equipment, communications equipment, etc.).

Step functions with a magnitude equal to the average delta kW values for each device were used to model the changes in steady state power level as loads activated. Actual turn-on transients for most machines are not clean step functions but in fact vary according to the physical task it performs [24]. In Cottage, most loads were

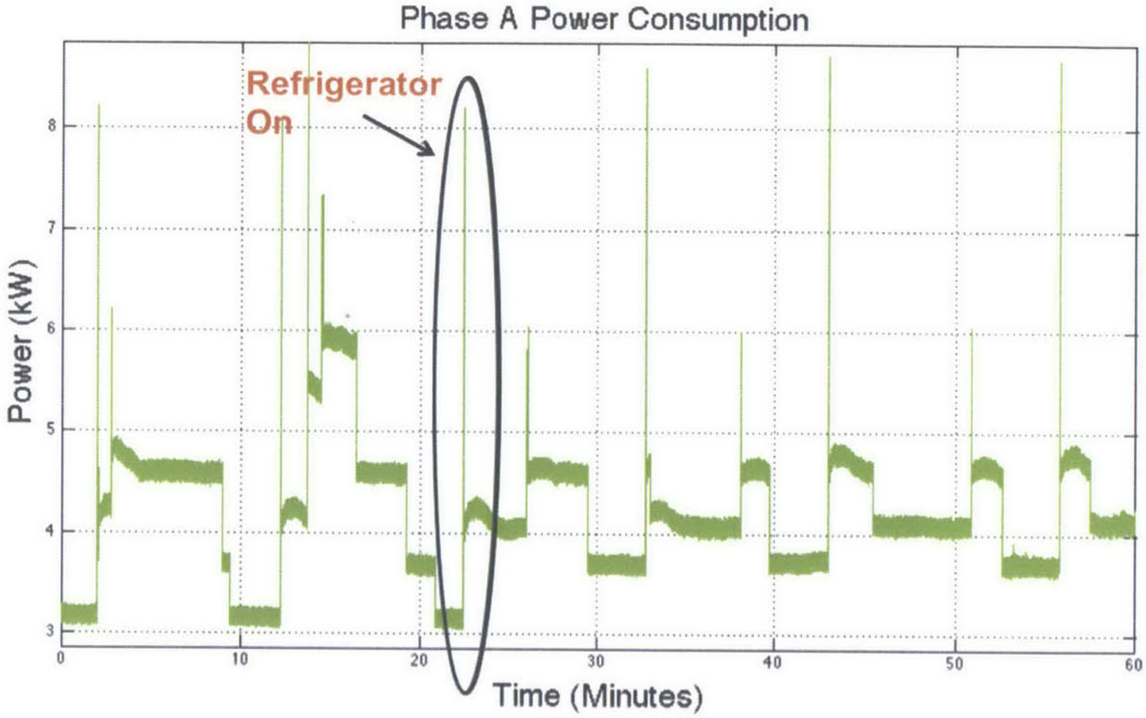

Figure 3-5: One hour of power data shows individual device transients

distinguishable using relatively simple characterizations of the "load transient," i.e., just the change in steady power consumption.

For example, see Fig. 3-6. The Boiler transient in the top left shows power over time. Because the fan must physically push air that is initially static, it requires more force at first to overcome inertia. In the plot, we see power rising between 2 and 3 seconds. As more laminar flow is reached, however, the power requirements on the pump quickly reduce and approach steady state levels by 4 seconds. Note that inrush currents peak at about 2 kW for fractions of a second. Power levels off at a level in the range of 0.46 -0.58 kW at steady state. The other transients follow similar patterns for moving refrigerant and sewage. Given that these transients each reach a quasi-steady state within a few seconds and given that the operating durations are on the order of minutes, the step function is a good approximation (less than 5% error) to use to determine kWh consumed. Recalling that real power consumption is the area under the power curve, the turn-on and turn-off transients disclose the duration of each machine's operation. Using this logic, the boiler 1 pump can be modeled using a step function of +/-0.51 kW, the refrigerator +/-0.89 kW, and the septic pump +/-1.38 kW. The baseline, or the DC offset, is about 3.2 kW. Each machine's operation over this one-hour period, graphed separately, is shown in the bottom right corner of Fig. 3-6.

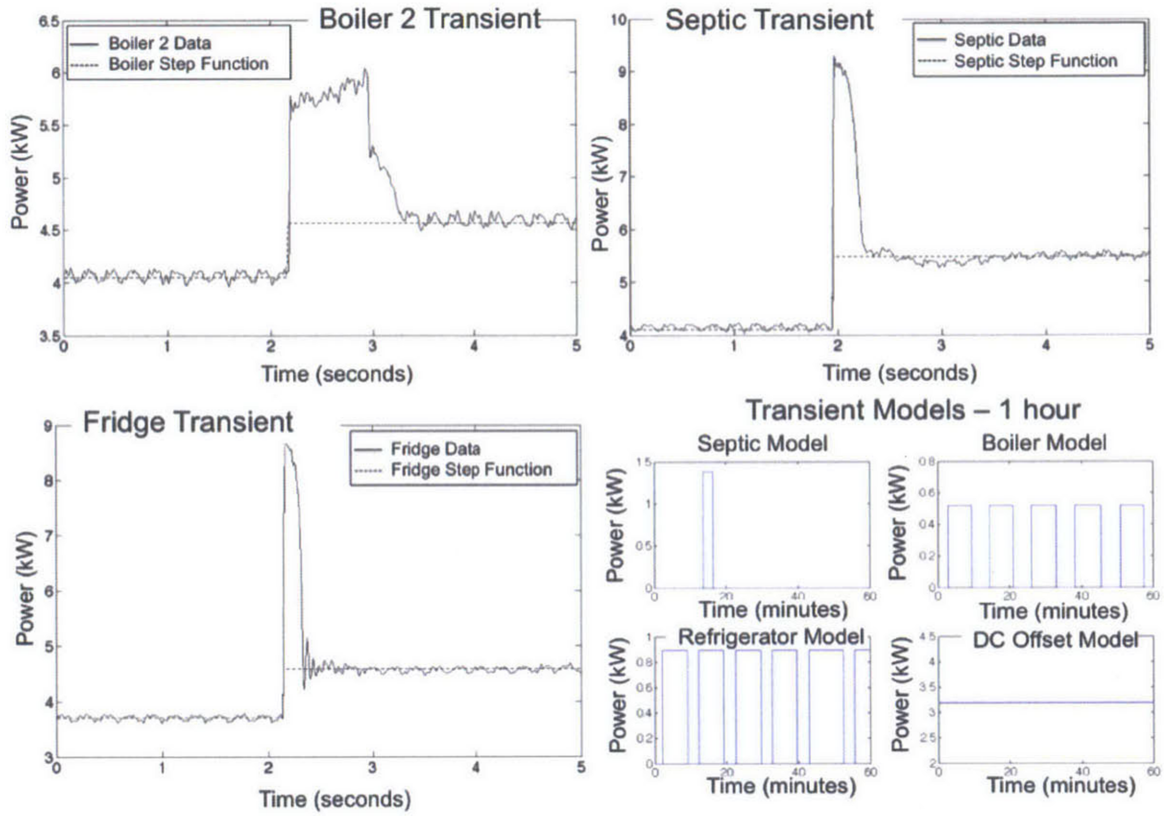

Figure 3-6: Three device transients modeled by step functions

The superposition of these three individual transient models closely approximates the original power signal (Fig. 3-7), validating the efficacy of this hasty modeling method. Once the edges can be detected, named, and kWh can be approximated, we can then keep score of each machine's activity and cost.

3.3 Measuring Appliance-Level Power Consumption

We trained the NILM at Cottage manually by cycling each load. Because the school was operational, this was not done in isolation. The list of unique loads with their

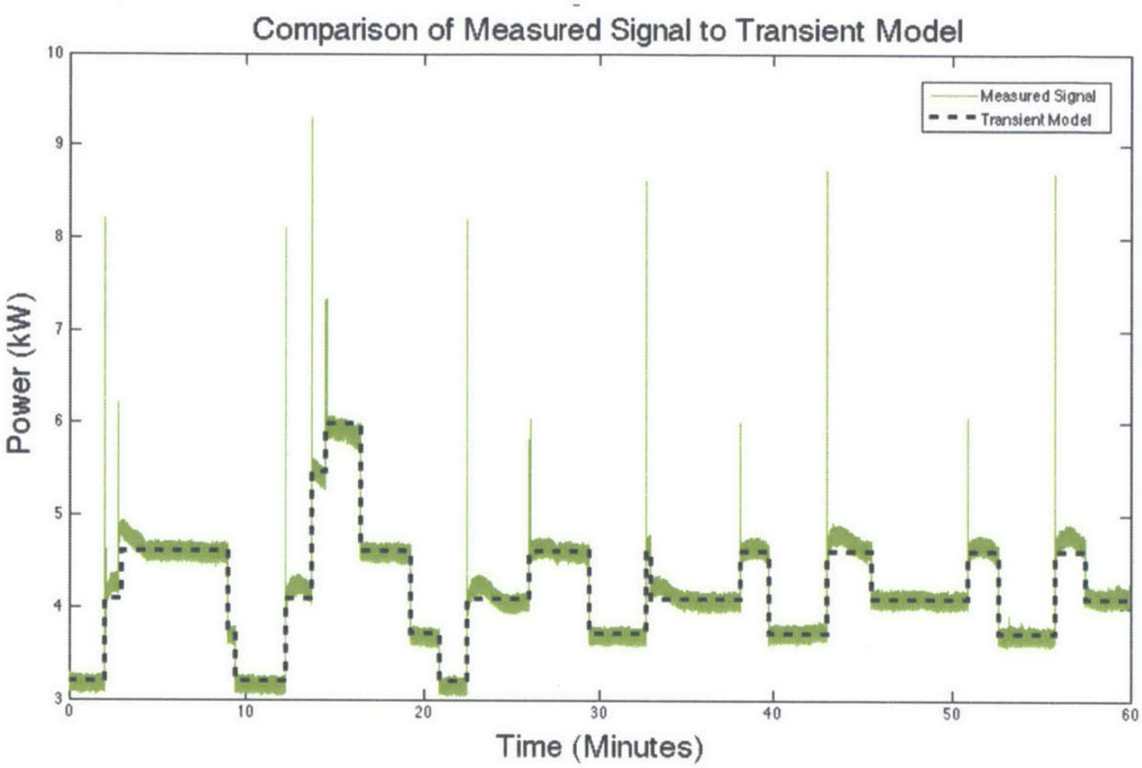

Figure 3-7: Comparison of true power data and transient models

respective power draws is presented in Fig. 3-8.

The NILM system detected some 5100 events over a period of 6 winter days in 2013 (March 11-13 and March 24-26). Peak hours featured more than 50 events, while the minimum number in a one-hour period was 18. Software corroborated the results. When events were classified, flags were raised if the same machine turned on twice without turning off in between. All such errors were checked graphically. In total, more that 98% of the events were classified without error. These are corrected by hand. The software sometimes misses and event, but the eye can generally distinguish what happened. Most often, the issue is simultaneous events. A few events were also missed because they occurred too close to the hour (within a few samples). This issue has been remedied for future experiments by eliminating the one-hour file sizes, opting instead to concatenate stored file segments into one large file. With the errors visually corrected, daily run times, cycle durations, and power consumption costs were tallied based off of the NILM output. The results are shown in Fig. 3-9. From the monthly

Breaker	Load	Min (kW)	Max (kW)
1	IT Room:		
	Extreme	0.14	
	UPS	0.06	
3	IT Room:		
	Phones	0.09	
	SMC	0.07	
	SMC	0.07	
	Sonic Wall Video	0.1	
	AC Pump	0	0.06
	UPS	0.045	
5	IT Room:		
	cable amplifier	0.025	
	PA/clocks	0.16	
	desktop CPU	0.07	
	apple CPU	0.025	
	server	0.07	
	UPS	0.21	
	UPS	0.13	
7	IT Room		
9	unknown off/on a lot	0	1.37
13	unknown 208V	0.9	0.9
15	unknown 208V	0.25	0.25
21	security	0	
23	generator controls	0	
25	freezer	1.2	1.2
27	unknown ~17 min run time	0.05	0.6
29	Boiler System:		
	Boiler 1 Draft Fan	0	0.86
	Boiler 2 Draft Fan	0	0.74
	Transformer 1		
	Low-flame solenoids 1		
	Transformer 2		
	Low-flame solenoids 2		
	High-flame solenoid 1		
	High-flame solenoid 2		
	control	0.27	
31	Boiler:		
	Boiler 1 Blend Pump	0	0.35
	Boiler 2 Blend Pump	0	0.54
	boiler control	0.13	0.13
33	Y-pump Kitchen	0.11	0.11
35	R-pump	0	0.07
37	Y-pump School	0.11	0.11

Breaker	Load	Min (kW)	Max (kW)
2	septic pumps	0	1.3
4	septic pumps	0	1.3
6	septic pumps	0.02	1.32
8	VFD circ pumps	0	1.5
10	VFD circ pumps	0	1.5
12	VFD circ pumps	0	1.5
14	freezer	0	
16	chair lift	0	
18	Make Up Air Fan	0	0.28
20	Make Up Air Fan	0	0.28
22	Make Up Air Fan	0	0.28
24	fire protection	0	
26	heat control	0.28	
28		0	
30	elev rm lights	0	
32	elev rm outlets	0.07	
34	boiler rm lights	0	
36	boiler rm outlets	0	
38	sump pump	0	
40	sump pump	0.05	
42	unknown	0.09	0.56

Legend

Always On Baseline Load

Figure 3-8: Individual loads by circuit breaker number power bill, Cottage paid just over 9 cents per kWh to the utility company.

	Boiler 1 Pump	Boiler 1 Fan	Boiler 2 Pump	Boiler 2 Fan	Circulation Pumps
Average Number of times used per day	15.3	15.3	63.5	63.5	1.0
Average duration ON each time (min)	9.7	9.7	7.2	7.2	1440.0
Average total duration ON per day (min)	289.7	289.7	544.2	544.2	1440.0
Average Daily Cost	**$0.15**	**$0.37**	**$0.44**	**$0.63**	**$8.71**

	Septic	Head End Room	Freezer	Refrig-erator	Unknown	Make-Up Air
Average Number of times used per day	2.5	1.0	120.0	69.7	19.2	0.8
Average duration ON each time (min)	2.7	1440.0	8.0	16.2	21.9	420.0
Average total duration ON per day (min)	6.7	1440.0	957.3	1075.8	412.2	478.7
Average Daily Cost	**$0.04**	**$2.61**	**$1.23**	**$2.11**	**$0.31**	**$0.62**

Figure 3-9: Itemized costs of Cottage Loads

The heat system represents the highest cost on the EBPP, more than $10 per day. It is made up of the 3-phase circulation pumps and the boilers. Broken down into its subsystems the largest single loads are the circulation pumps. One VFD pump is always on while the heating system is on, though the speed and thus power draw fluctuates. From recorded data, these pumps consume, as a rough average, 1.3 kW per phase costing over $8 per day. Combined, the creation and transmission of heat represented almost 11% of the monthly bill in March. Note that this does not include the contribution of the univents in all of the classrooms that distribute the heat to the tenants.

The reconstructed model (Fig. 3-7) accurately models the original signal, validating NILM's disaggregation method. While the model is visually similar in its basic shape, there are elements of the original that are clearly not in the reconstructed model. First, the power peaks, including their peak amplitudes, are not shown as explained in Section II. Second, the slow, smooth fluctuations, such as the subtle changes in the variable speed drive, are not accounted for. In general, these represent room for improvement but do not invalidate the approach. While important, the precision of the kWh measurements is secondary to the accuracy of cataloging the individual device patterns from an aggregate feed.

Understanding the details of electrical systems empowers decision makers to make changes without service interruptions or sacrificing environmental comfort levels. Systems like Cottage that employ integrated control systems are commissioned when first emplaced. Over time, as equipment or conditions change, these settings require updates to keep the system optimal. Department of Energy calls this "continuous commissioning," or updating system controls over time as conditions change [25]. NILM is able to provide early warning that conditions have changed.

Some limitations became obvious from this experiment. The higher the load count, the higher the likelihood of ambiguous results. Two (or more) loads may turn on, off, or one-on/one-off at the exact same time. Higher sampling would improve resolution, but this would bring the added requirement of more memory and, in this case, more bits of resolution on the ADC. Previous research has advocated collecting all questionable identifications after filtering in order to run "anomaly" algorithms. These make successive comparisons of the anomaly delta kW against both combinations of known transient delta kWs and known machine states (on or off) [24]. Also, only changes are visible with the NILM. If loads rarely (or never) cycle, i.e. they are always on, then they are not uniquely distinguishable. The sum of continuous loads comprises the baseline load, which can be discretely determined only by shutting everything off and then back on one at a time.

3.4 Condition-Based Maintenance

The boiler system is a complex and critical device in the heat loop. Understanding the operation details allows us to look for problems. The burner, a Gordon-Piatt R8-3-G-15, runs off natural gas. Its electrical components include the draft fan (the major load), solenoids along the gas line, a transformer, a modulation motor to control the air-gas mixture, and various controls. On call for heat, the draft fan turns on and purges the system by pushing out any remaining combustible vapors. The fan is powered across 240V, thus the signature on two phases (Fig. 3-10). After a programmed 90 seconds passes, the ignition sequence commences. Once properly lit, the flame sends heat into the boiler chamber until the system requirements are met, upon which the gas turns off and the fan shuts down. This process is repeated every time several times an hour for as long as the system continues to call for heat.

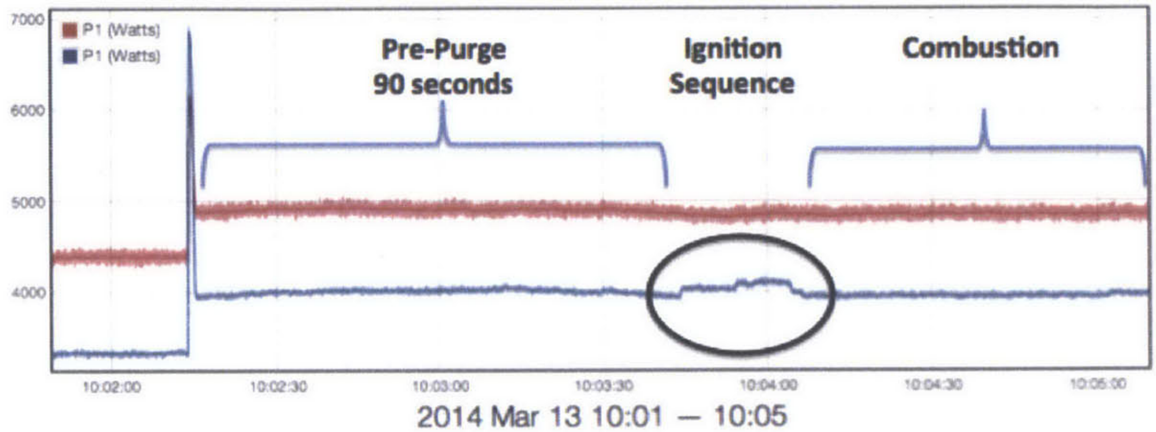

Figure 3-10: Cottage Boiler power consumption on two phases

Focusing in on the ignition sequence, each stage is visible through the high-resolution electrical data we collect. Even amidst the noise floor, the subtle fluctuations of less than 100W are clear with each of the "steps" in power up and down (Fig. 3-11)

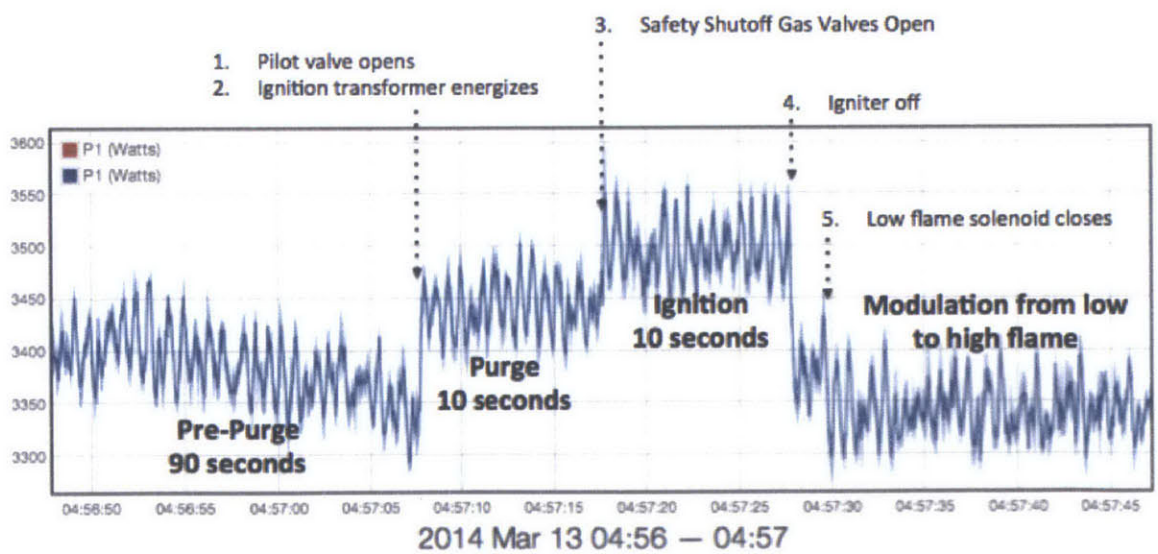

Figure 3-11: Cottage Boiler power consumption on two phases

Variations in this stage may indicate mechanical or electrical problems. Having such data available offers the proverbial airplane "black box" to maintainers. A scenario might transpire as follows. Teachers in the morning notice the room is getting colder and turn their thermostats to the maximum setting. When that fails to work, they call maintenance. Someone comes to the room. Verifying the problem, he or she will start tracing the issue from the thermostat to the uni-vent in the classroom. If the damper is closed and the fan is working, they'll conclude that problem lies beyond the classroom. From their building control panel, they will look at the temperature readings for that classroom and perhaps verify that everything is working properly except the heat. They look at the boiler and see that it is no longer running. By now, several teachers have made the same complaint, and so maintenance moves to the next level and enters the boiler room to investigate. Despite several attempts, it refuses to start beyond the purge stage. Now they have a technical that will require an expert. This is where the black box comes in.

A technician arrives knowing nothing about what happened but knowing a lot about boilers. However, with the data in the right hands, he (or his company engineer team) can rewind to the last five start-ups. Knowing the sequence of operations, they can compare what right looks like with what the current situation is and isolate the problem quickly. I assume a lot here, but in an age of highly standardized equipment such capabilities are not beyond reason. Further, programming detection software to raise alarms when the waveforms deviate beyond some threshold is also practical.

Another result made possible by the NILM is a comparison between the burners. Each has two main electrical components, a draft fan motor and a blend pump. The make and model of the two draft fans are dissimilar. The blend pumps also differ in model. It is common practice to set unoccupied times on building such as this school. It allows the school to maintain a colder temperature during off hours.

There is a balance between how low the temperature dips and how much energy is required to warm it back up the next morning. This ramp-up period was monitored closely so that a comparison between the boilers could be made. Boiler use is frequently alternated between Boiler 1 and 2 for maintenance purposes. On March 11th, Boiler 1 operated alone from midnight to 8AM. On March 12th and 13th, Boiler 2 ran alone during the same time frame. The temperature profiles for those days being similar (lows of 36, 38, and 32 degrees, respectively), we determined that while Boiler 1's blend pump uses 20% less power, the draft fan uses 50% more power when running. The duration times of operation varied drastically, with Boiler 1 staying on nearly 2 hours longer to create (presumably) the same amount of heat. Their operation profile differed as well. Boiler 1 ran 15 times with an average duration of about 22 minutes compared to Boiler 2, which ran 26 and 29 times on consecutive days, respectively. Fig. 10 contains a summary of their head-to-head statistics, revealing that Boiler 1 is about 22% more expensive to operate than Boiler 2 and also puts more about 28% more hours on the machinery for comparable work.

Boiler Cost Comparison (8-hr warm-up)	**Boiler 1** 11 March	**Boiler 2** 12 March	**Boiler 2** 13 March
Draft Fan (Make, Model)	**General Elec** Model 5KC49N	**Marathon Elec** Model EPL 56B34D202BEP	
hp	1.5	1.5	
FLA (115V)	9.2	6.7	
kW	0.83	0.71	
ON duration (min)	337	262	257
cost	**$0.42**	**$0.28**	**$0.28**
Blend Pump (Make, Model)	**Taco** Model 0012-F4-1	**Taco** Model 0012-F4-1	
hp	1/8	1/8	
FLA (115V)	18.4	13.4	
kW	0.33	0.51	
ON duration (min)	337	262	257
cost/day	**$0.17**	**$0.20**	**$0.20**
Total Cost/Day	$0.59	$0.48	$0.48

Figure 3-12: Two different boiler fans and pumps compared head to head

3.5 Interpreting Human Activity from Electrical Activity

NILM also measured the effect of a major change to the system. On March 25, the weather turned warmer. This led to complaints from the teachers about the heat in the rooms. In response, the maintenance technicians throttled all heat valves remotely from their central control station. A corresponding three-phase power reduction of 2.7 kW was observed instantly (Fig. 3-13). Because the VFD pumps are pressure controlled, a sudden decrease in demand caused an increase in pressure, and the active pump responded by slowing down significantly. This decrease was observed for the remainder of the school day (about 6 hours), only to increase again during the evening hours when the weather cooled off and demand again increased. In total, this saved about $1.50. Knowing the actual savings, rather than relying on assumptions or rumors, empowers the customer with actionable feedback for future decision-making.

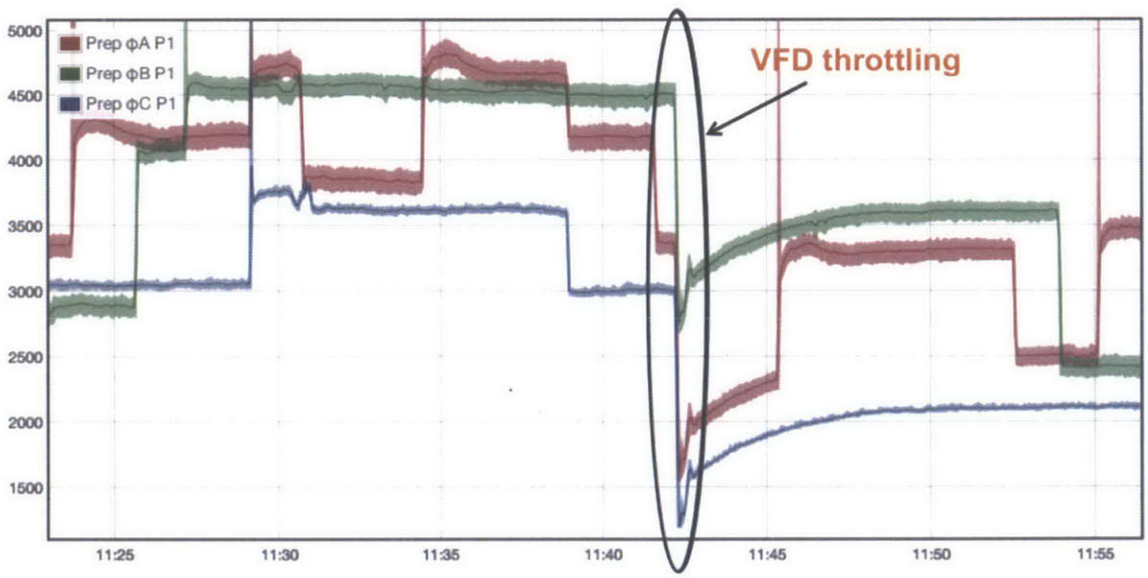

Figure 3-13: Cottage VFD reducing power rapidly as demand for heat is turned off centrally

3.6 Other Results

In this experiment, NILM showed promise as a plausible sensor for natural gas sub-metering. Since the burner specifications [26] and boiler hours-of-operation are known, then the amount of natural gas consumed by the boilers is estimable. Note that natural gas is not sub-metered at Cottage. Using some hasty calculations, the Gas utility billed the school for 6554 ccf during the month monitored. Using data from the six-day period highlighted above, the combined (both boilers) average run-time is 8 minutes and 25 seconds per cycle. This is the duration that the draft fan is operating. From the burner manual (Gordon-Piatt R-8 Model), the first 90 seconds (on a timer) of fan time purges the system. No gas flows into the boiler. For the next 10 seconds afterwards, low-flow gas is injected into the burner to facilitate ignition. Considering only the high gas consumption time, we arrive at 6.75 minutes per cycle. From the results, the boilers run an average of 90 cycles per day, which equates to 10.1 hours of high-gas operation time per day. The firing rate of the burner is 2136 MBH according to the data plate, which represents the maximum numbers of BTUs per hour through the burner. Thus, the total number of MBTUs per month is

$$2136[MBTU/hr] \times 10.1[hr] \times 30[days] = 647.208[MBTU/month] \tag{3.1}$$

From the utility statement, the gas conversion rate is 1 cf =1.02 MBTU. Converting the MBTUs to cf, we estimate the monthly gas consumption of the boilers during this month to be 6345 ccf, which closely resembles the 6554 ccf utility bill. There are other gas appliances, including 3 roof-top HVACs (RTU) and an oven/stove in the school kitchen, but it is interesting to note that the math is in the ballpark and merits further study. Summing all of the gas-burning equipment firing rates together, the boiler alone is 80% of the total. It is known that the burner gas input valve modulates between 50-100% while operating, but during cold seasons it is assumed to run full open. We have future plans to install a camera on the meter and compare the boiler run times to actual gas consumption for verification.

The power study uncovered useful information during the training phase as well. First, there are at least 24 loads that are always drawing power, 14 of which are in the Head-End (IT) room housing all of the network switches and other communications equipment. Including the Uninterrupted Power Supply (UPS), they draw a collective 1.2 kW at rest (while school is not in session). The UPS was in permanent bypass mode because it was not operating correctly, which the network administrator knew.

It was already scheduled for replacement. What was not known was that, even in bypass, the UPS continued to draw about 0.4 kW at a monthly cost of about \$26 just to cool itself and maintain standby posture. Measurements of the total load connected to the UPS also led to the recommendation to reduce the size of the replacement UPS from 10kVA to between 5-8 kVA as their maximum load was less than 2 kW.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 4

Real-World Army Test Sites

Figure 4-1: Ft. Devens Base Camp Integration Lab

The military Forward Operating Base (FOB) is an ideal laboratory for demand response initiatives. Like any office building, FOB occupants are transient, work-oriented, and thus tend to be energy indifferent in regards to cost or conservation. Also, base camp electrical loads include both basic life support and work-related devices that are common among commercial and workplace environments. Finally, these military facilities are typically compact, islanded from the power grid, and easily measurable from a few central locations.

Most military bases contain an array of permanent, semi-permanent, and temporary structures with a wide variety of electrical loads, planning, and implementation. The occupants of these facilities are the ones who use the energy, through they are transient and generally disinterested in the energy bill. The power to these facilities may be grid power or the facility may be islanded with their own power source, such as a base camp running off a generator. These types of facilities are ideal test beds for demand response technology because they stand to benefit the most. Enduring structures capitalize on more efficient, centralized appliances for heating, ventilation and air conditioning (HVAC). Hasty structures, like that pictured below, utilize decentralized ad hoc systems that require more control to prevent waste. At more austere locations, Forward Operating Bases (FOB) pay an extremely high price for their inefficiencies when considering the Fully Burdened Cost of Fuel (FBCF) to keep electrical microgrids operating on the front lines [27].

The United States Department of Defense is actively working to reduce its overall energy demand [28]. This is in accordance with section 431 of the Energy Independence and Security Act of 2007 (H.R. 6/P.L. 110-140 of December 19, 2007) requiring federal building energy use to be reduced 30% by FY2015. As the world's largest consumer of energy [29], electrical costs represent 12% of the DoD's overall bill and 33% of their operational energy costs. The recently founded Office for Operational Energy Plans and Programs (OEPP)) has listed "Establishing a consumption baseline" as a central priority, indicating the insufficiency of the measurement tools currently being used to identify margins for improvement [30]. New metering systems are being implemented at installations across the United States and abroad in attempts to provide this baseline. A large database to collect and manage the massive amounts of data from these metering programs is being developed. Despite all of these efforts, without appliance-level information, the effect of future conservation efforts is going to be unquantifiable.

Curtailing electrical usage presents a unique problem because it is harder to understand and quantify. An aviation unit, for example, can use less fuel by reducing their flight hours. A naval ship can conserve by reducing the speed of their vessel. In these military branches, clear ratios between operations and energy consumption allow prudent cuts to be made. On the contrary, at ground facilities, where hundreds of unique electrical loads exist, exactly how to conserve energy is less evident. With no way to distinguish between critical vs. non-critical loads, no quantification of waste, and no ground-level capability to manage loads effectively, commanders will continue to tell Soldiers to conserve without telling them how. This is akin to doing the same thing while expecting different results.

Ongoing research on FOB energy conservation, including battery-bank enabled peak shaving and generator microgrids , does not take advantage of the "most neglected energy resource – efficiency" [31]. Appliance-level feedback through NILM harnesses the potential of demand response through consumer awareness.

The following is a detailed explanation of the three military test sites which provides the background for the experimental results that follow in Chapter 5 and 6.

4.1 Ft. Devens

Soldiers need shelter and basic amenities to sustain them. Units must be ready to deploy at a moment's notice, often to austere locations. To accommodate their expeditionary nature, the Army developed preconfigured deployable base camps known as Force Provider (FP) packages. Ft. Devens, MA, has a fully functional base camp known as the Base Camp Integration Lab (BCIL). It includes tents, latrines, showers, laundry, and a kitchen. The entire camp, including their electrical sustainment loads, is a rapidly deployable standardized package known as FP-150, meant to accommodate that many troops. The espoused goal of the BCIL is to "ease the logistical burden on war fighters, and they can more clearly focus on the mission" [32]

Army units use the facility typically during weekend training events. While the base can be configured as a generator-powered microgrid, it most often runs off shore power through two 600A electrical panels. Two DepNILMs installed at these panels have been collecting 8kHz 3-phase voltage and current measurements of all electrical activity within the dashed box in Fig. 4-1 since September 2013, surviving the New England winter. As of this publication, they are still active and their live data is viewable through the website `http://wattsworth.net`. Power distribution is reconfigurable by Soldiers using the mobile LEX-brand distribution boxes. One example of the layout is shown in Fig. 4-2.

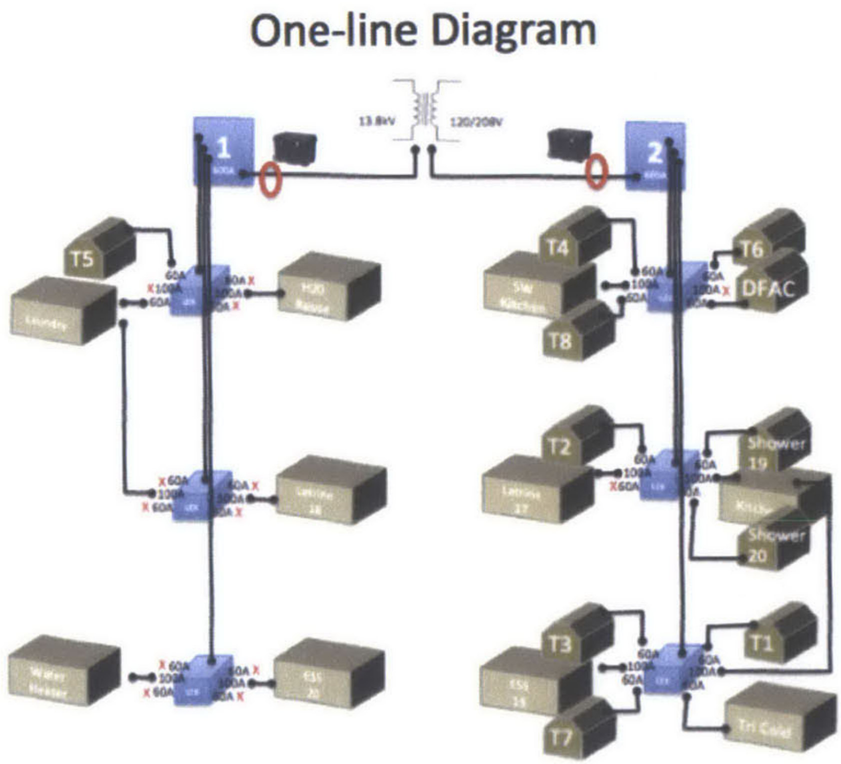

Figure 4-2: One line diagram of Ft. Devens

Atypical to the standard FP-150 package, the BCIL has some electrical upgrades. First, it is wired to operate off either shore power or generator power (see the one-line diagram in Fig. 4-2). Note that all structures are electrically reconfigurable. Tent 5, plugged into a distribution box powered off panel 1, can easily be moved to panel 2 by plugging it into a different distribution box. Second, the six 60 kW generators can be operated independently or synchronized to operate as a microgrid. Finally, dozens of commercially available monitors (Shark 200 and others) have been installed within each FP-150 camp to provide detailed feedback about power and water consumption at every tent. These monitors wirelessly transmit sampled data to a central server where data is stored for analysis.

We arranged with the Project Manager for Force Sustainment Systems (PM-FSS) to install our monitor on one side of the camp. We replicated their monitoring system with just two NILM systems, one on each 600A panel (Fig. 4-3). We also added vastly improved capabilities over the off-the-shelf smart meter: appliance level consumption, maintenance feedback, human patterns of life data, and an accurate water consumption estimate.

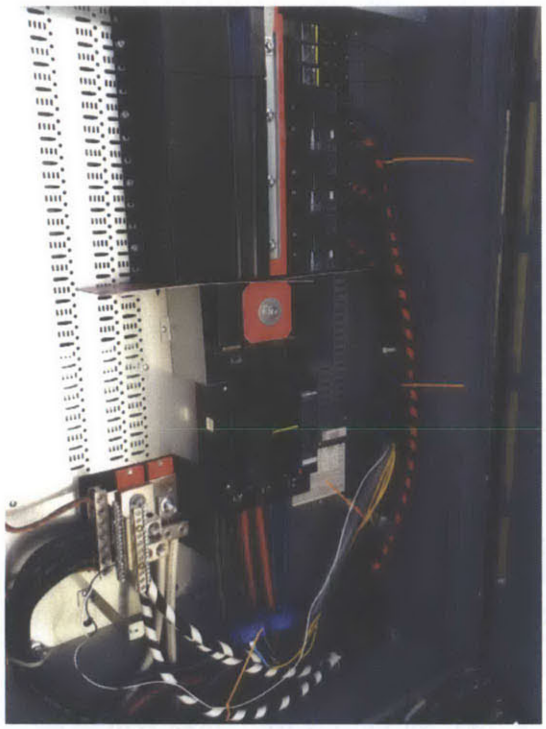

Figure 4-3: NILM Installation at Ft. Devens on 600A Panel

The primary power-consuming loads at the BCIL are listed in Fig. 4-4. Some loads like the space heaters are single-phase, while others are 3-phase. The "Units" column indicated the number of unique devices on the camp by type. The "Watts" column is the average number of total power used by each load during steady-state operation. This data was obtained during the training phase, where each device was cycled in isolation. Where there were several loads of the same type, an average value was determined from several cycles of each load. Except for the ECU compressor, steady-state operation is relatively stiff. As demonstrated in [3], Turn-On and

Turn-Off evens can be accurately modeled as step functions. As further evidence, the power consumption on the BCIL was measured both by commercial power meters and through NILM analysis during two separate 48-hour periods. The NILM system detects every On and Off event, determines the run-times for each major appliance, and calculates the total power consumption as the sum of power used by each individual load. Using the numbers in Fig. 4-4, the NILM estimate was within 5% of the true measurement of the commercial meter.

Load	Description	Major Sub-load	Phase	Watts	Units
60K ECU	HDT F-100 60K BTU ECU		3		11
		Digital Scroll Compressor	3	516 - 4951	
		Supply Fan	3	1603	
		Two-speed Condenser Fan	3		
		Heater	3	9563	
Water Pump	Liberty LSG203H		3	3600	4
Reefer	Tricold Thermo King 600cu ft		3	3100	1
Window HVAC	Freidrich EQ08M11		1	1353	6
Space heater	Marley C1512		1	1488	8
Kitchen refrigerator	Traulson TU072HT-X0017		1	524	1
2 vent/light set			1	338	4

Figure 4-4: Top power-consuming loads at Ft. Devens

Among the top-consuming loads, the environmental control unit (ECU) is the most complex. Each ECU has four main sub loads: a supply fan, condenser fan, heating coils, and a compressor. There is one ECU per tent, and there are 11 tents (8 for sleeping/multipurpose use, 2 for showers, and 1 for the dining facility (DFAC)). Each ECU is controlled independently and has four user-selected modes: Heat, Cool, Fan, and Off. There is also a thermostat dial permitting fine temperature adjustments. Vent mode activates the supply fan, which circulates inside air through the ECU. Except in off mode, this fan is always running. Heat mode activates the heating coils. Heat is transferred to the room as the circulating air blows across the hot coils [33]. In cool mode, the compressor starts up followed by the condenser fan. Operation follows the traditional refrigeration cycle. Two things are of interest here. First, this compressor has "soft start" controls that use power regulation to limit current peaks during turn-on. Second, the unit features a digital scroll compressor that cycles frequently between loaded and unloaded states, an energy saving technique. The duty cycle for the loaded state is recalculated every 15 seconds. The compressor may stay loaded 100% of the time (all 15 seconds) if the room is very hot, 10% of the time (1.5 seconds) if the room has met the target temperature, or any 10% interval in between. Fig. 4-5 shows a graph of compressor power consumption during start up (from Off mode).

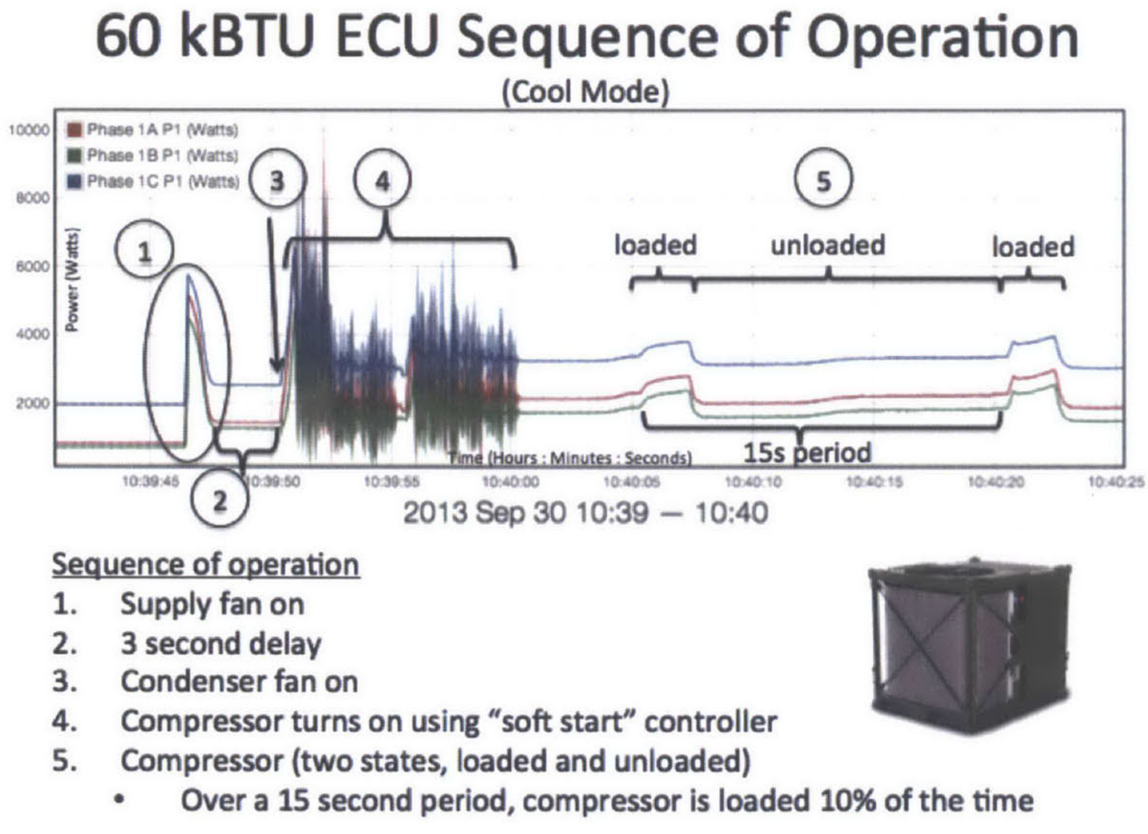

Figure 4-5: Power signature of HDT F100 60 kBTU ECU

Other organic loads on the base camp are the fluorescent lights (white and black), pipe warming cables (to prevent freezing), kitchen appliances, laundry machines, outdoor halogen lights, an ice machine, vents in the kitchen and laundry room, smaller pumps for the sink, washing machine, and water purification units, and various commercial kitchen appliances. While some of these loads draw significant power, the infrequency of use make them of lesser importance in overall conservation efforts. Their contribution to the overall bill on this training base is negligible during the several training rotations analyzed.

Lighting merits further discussion. The cost of lighting the entire 150-person camp for an hour is approximately 4.5 kW (110 bulbs x 40 W/bulb). This is about half the cost of heating just one tent for an hour on a cold day. While small in comparison, there is still value in determining light bulb energy consumption, both for cost itemization as well as human behavior indicators. However, in practice there are significant variations in the off transients of these lights (100-200 W per light switch) because they are easily reconfigurable. At such low power levels (about 150 W per strand of 4), the magnitude of the noise floor as well as the presence of many smaller electronic devices exhibiting similar signatures complicate consistent and accurate detection of lights. For the purposes of this study, lights were ignored.

Tactical Operations Centers (TOC), where headquarters elements use a lot of electronics equipment, have not been studied extensively at Ft. Devens. This is due to the fact that many units rotating through are there for only a weekend. They bring only what the need, which tends to be a personal device and battery chargers. For those that do set up a HQ, they typically use a vacant tent that does not draw power from our monitored panels. More on TOC power consumption can be found in the section on Ft. Polk, however.

4.2 Ft. Polk Forward Operating Base

Ft. Polk, LA is an active duty Army post and home to one of three Army Combat Training Centers. Military units ready to deploy rotate through Ft. Polk to conduct large-scale military exercises. Within the vast wilderness of the Ft. Polk training area, several FOB-style base camps have been built for training use that emulate conditions in current theaters of operation (Fig. 4-6). Compared to Ft. Devens, the FOB used in this experiment is much larger, utilizes single split-phase power, and does not have the laundry, kitchen, shower, or latrine amenities all on the same panels. Thus, only a portion of the activity of the base was electrically observable through the two monitored panels. DepNILM collected data from these panels between February

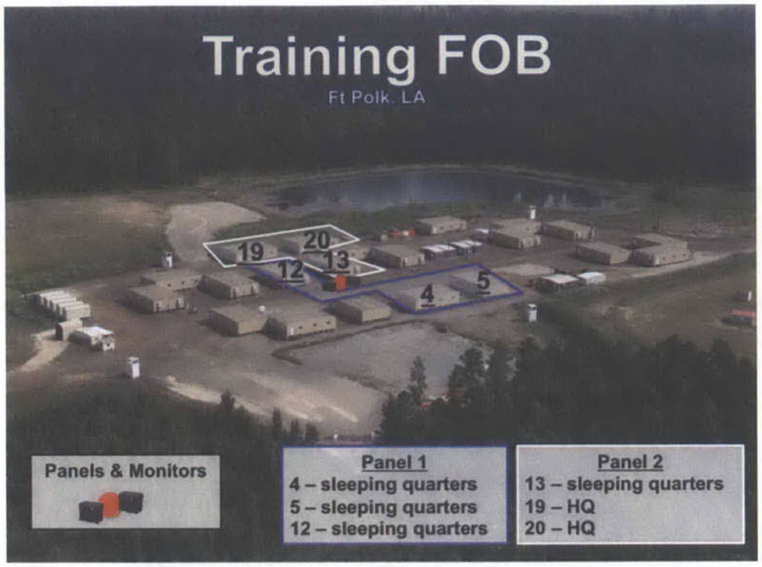

Figure 4-6: Unnamed training area on Ft. Polk used as a FOB

and April 2014 (Fig. 4-7).

Panel 1 is connected to three buildings all used as sleeping quarters. Two HQ buildings as well as one sleeping quarters draw power through Panel 2. The only difference between the HQ and sleep buildings is the plywood partitions in the HQ buildings for offices, workspaces, conference rooms, et cetera. Otherwise, they are equivalent dimensionally, each an open bay 36' x 60' room. The organic loads in each building are also identical, consisting only of lights and environmental control. Each has two thermostat-controlled Bard W484 (10 kW, 34 kBTU) Wall Mounted Package Air Conditioners (capable of providing heat and air conditioning [34]. Users control temperature and modes, which include single-stage heat, cool, and fan (Fig. 4-8). The operation is similar in most instances to the Ft. Devens ECU. One difference is that the supply fan cycles off when the heater or compressor does (with a programed 1 minute delay). Another distinction is that the compressor does not modulate. This makes power consumption much more stable in steady state, albeit at a higher kW value. Both heat and compressor signatures from the Bard ECU are displayed in Fig. 4-9.

Internally, there are three rows of two-lamp T8 florescent lighting, 4 fixtures per row, for a total of 24 bulbs per building (Fig. 4-10). Outside, there are two small

Figure 4-7: DepNILM installation at Ft. Polk FOB

floodlights, one over each door. Each row of lights has its own switch, as does each outdoor light. Occupants have over 40 duplex outlets to connect training and personal items to. In sleep quarters, these may include coffee makers, personal devices, or various battery chargers. In the HQ buildings, loads typically include TV screens, projectors, computers, external monitors, printers, battery chargers, radio equipment, switchgear, electronic clocks, and personal devices.

Troops at the FOB are transient. During the two-month period we monitored, a Combat Support Hospital (CSH) unit occupied part of the footprint for about 2 weeks. They set up a company-sized HQ in building 19, a battalion-sized HQ in building 20, and used the sleeping quarters in in buildings 12 (18 personnel) and 13 (30 personnel). An Engineer unit occupied building #5 (47 personnel) for one week.

Figure 4-8: Bard Wall-Mounted ECU

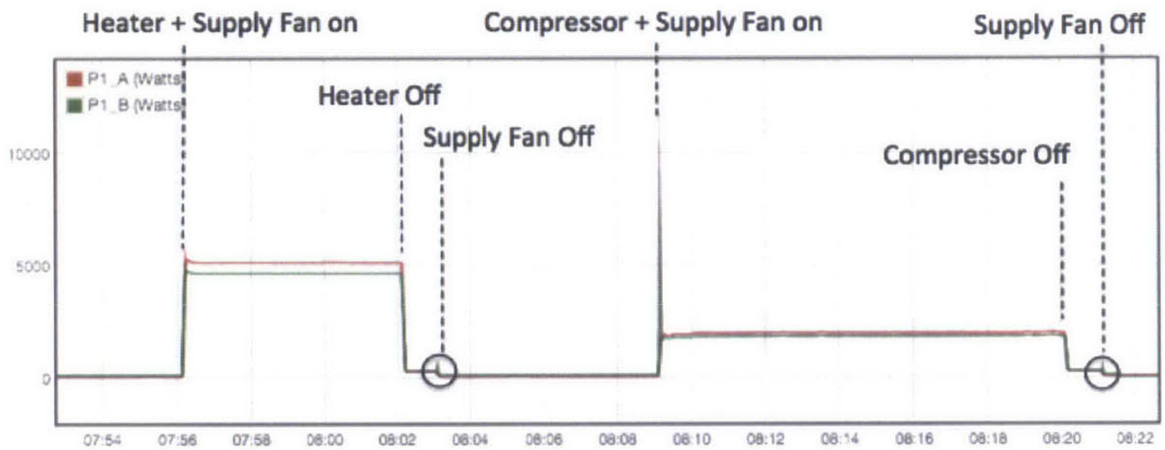

Figure 4-9: Power signature of Bard Wall-mounted ECU

4.3 Ft. Polk Field Hospital

Not all training units operate out of buildings at Ft. Polk. A Combat Support Hospital (CSH) unit can build a field hospital using modular tents in their inventory (Fig 4-11). They erect the field infirmary in any environment and perform most of the functions of a normal hospital. Mobile generators supply their power. Data from this site was collected during a two-week exercise in March 2014. (Fig. 4-12).

Hospital tolerance for power interruptions is extremely low, so they incorporate redundancy by powering key facilities off of two separate generators or simply paralleling their sources. Fig 4-11 and Fig. 4-13 shows two generators each powering half the tent in key facilities, where each tent has two sub panels. Loads connected to the

Figure 4-10: 36' x 60' bays at Ft. Polk FOB

Figure 4-11: Example of field hospital layout

generator were not constant. While we tested, the generator malfunctioned several times. This prompted the Soldiers to replace the generator with another, which meant moving the NILM to a second generator. When the second generator (and then a third) had the same malfunction, they elected instead to reconfigure loads to other generators. Finally, the DepNILM was moved a third time to monitor a different load set.

The loads in the hospital are numerous and technical. The primary consumer is the ECU. Its operation is almost identical to the Ft. Devens version with the exception that the Ft. Polk ECU compressor modulates to save power using a variable frequency drive (VFD) controller [35]. Another large load that normally consumes

Figure 4-12: DepNILM installation on a CSH generator at Ft. Polk

very little power is the Pressurized Oxygen Gas System (POGS). They cycle only in short bursts while in standby mode. DepNILM monitored the intensive care units (ICU), operating room (OR), Pre-Op, emergency room (ER) (Fig. 4-15), medical supply tents (CMS), specialty ward, laboratory, and the hospital operations center (Ops). Some of these rooms have running water available.

In general, starting from the flight line (Fig. 4-14), the flow of casualties is triage, ER, and either Pre-Op for surgeries, ICU for prolonged treatment, or the Specialty wing. Because this was a field exercise, the hospital staff did not actually use many of their electrical devices as they might in a real-world situation. For example, it unnecessary to use the portable x-ray, rapid blood infuser, suction tools, or intravenous line pump on healthy individuals. Other loads of interest are the lights, heating (Bair Hugger) blanket systems, numerous laptops, coffee makers, vital sign machines, ventilators, suction machines, and other various surgical devices.

Only actual electrical activity is measurable. The distinction between training and combat operations becomes important here. In training, electrical equipment is often used only notionally. For example, following triage, a casualty may be taken to the Emergency Medical section, surrounded with mobile lights, tested for vital signs, perhaps rapidly infused with blood, administered fluids through a controlled IV pump, X-rayed, given oxygen, ventilated, and then taken to either pre-op or intensive

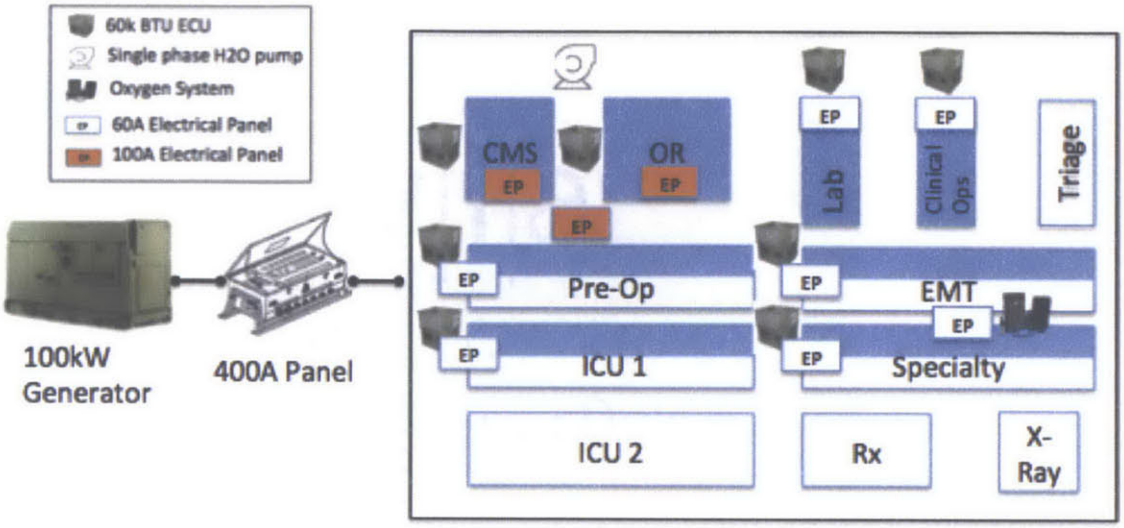

Figure 4-13: Wiring diagram of one section of hospital

Figure 4-14: Notional casualties being air-evacuated to the CSH for treatment

care. Each of these medical measures, consume electrical power, but only during real-life situations. In training, most equipment is used only notionally concealing most medical activities from the electrical monitor.

Figure 4-15: Hospital staff treating notional casualties

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 5

What Actionable Feedback Looks Like

Gathering data is the easy part. Many devices can tell you how many kW you are using, but NILM breaks it down to the appliance level. This provides a level of understanding that leaders can base decisions on. First, they can know what is on and how much power each device is using. Second, they can identify maintenance problems before the equipment breaks down and before undetectable waste mounts up. Third, they can have a certain level of awareness about the occupants of their facility. All of this is available remotely. The experiments at Cottage were a trial run. The following results are based on the three in-depth military microgrid field tests at Ft. Devens and Ft. Polk. NILM provides decision makers the technical information they need to direct change. Experiences from each of the locations previously described demonstrate this.

5.1 Accounting for every kWh

Accountability of the costs of any electrical network is possible through appliance-level feedback. Saving electricity should be a calculated decision and not an arbitrary cut. Knowing where each kWh is going is more useful than knowing the monthly, hourly, or even instantaneous kW totals. When the cost and behavior of each individual appliance is understood, demand response can be quantified. DepNILM permits users to distinguish between mission critical loads, quality-of-life loads, expendable loads, and even waste. Several examples follow.

Fig. 2-7 shows the overall kWh consumption data for a cold weekend in November 2013 at Ft. Devens. The following timeline provides some context (Fig. 5-1). An Infantry unit of 90 personnel came to train over a weekend. At noon on Friday, the advanced party consisting of two personnel arrived, signed for the tents, and turned on the heaters in preparation for the main body. At approximately 1600 hrs, the other 88 personnel arrived. After making camp, the Soldiers conducted pre-marksmanship instruction in preparation for the next day's weapons range. At approximately 2100 hrs, they completed the class, conducted hygiene, and went to sleep around 2200 hrs.

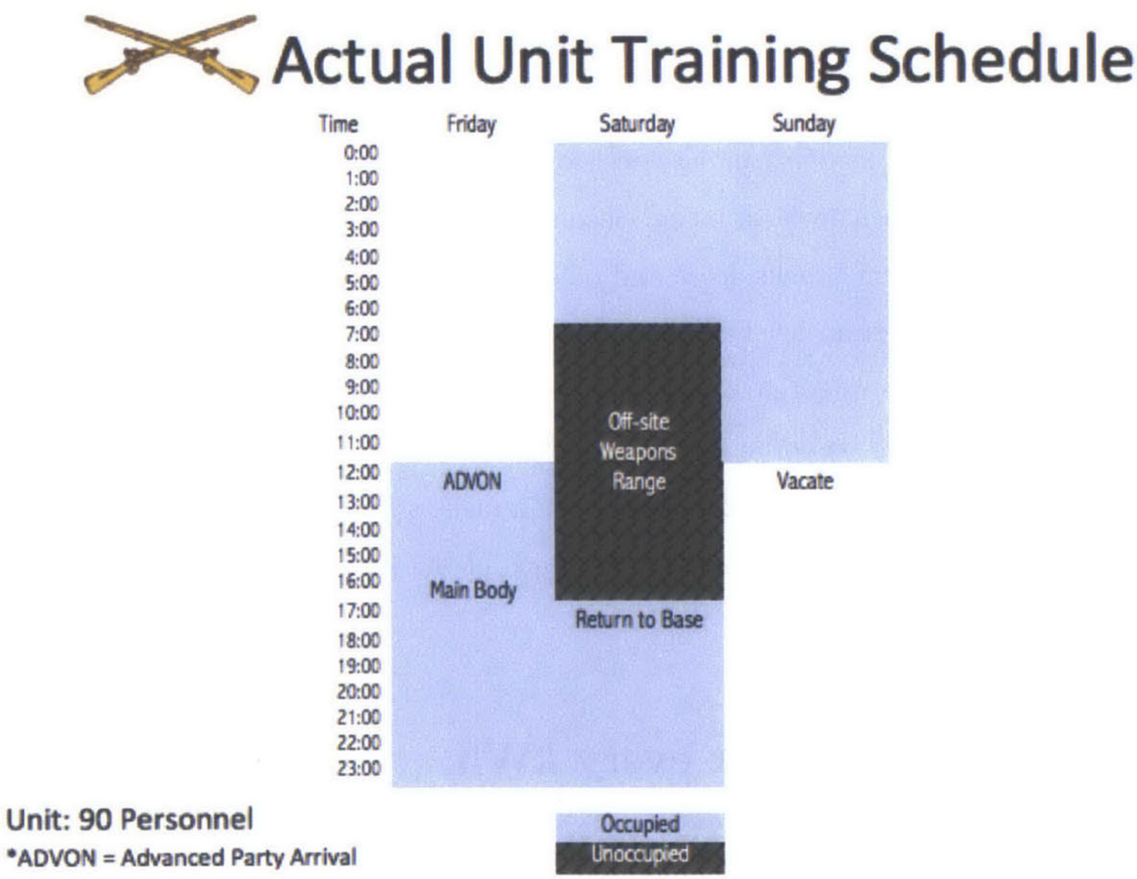

Figure 5-1: Training unit timeline showing when the base is occupied vs. unoccupied

Saturday morning, the unit rose at 0500 hrs, performed morning hygiene, consumed an MRE (Meal Ready to Eat), and then departed for the weapons range around 0700 hrs. Two personnel stayed behind to watch over their gear. At approximately 1600 hrs, they returned to the base camp, ate, cleaned weapons, and conducted hygiene before retiring around 2200 hrs. On Sunday, they arose at 0600 hrs, packed their things, and the majority of personnel vacated by 0900 hrs. All personnel were gone by 1200 hrs, and they turned off the majority of their loads before they left.

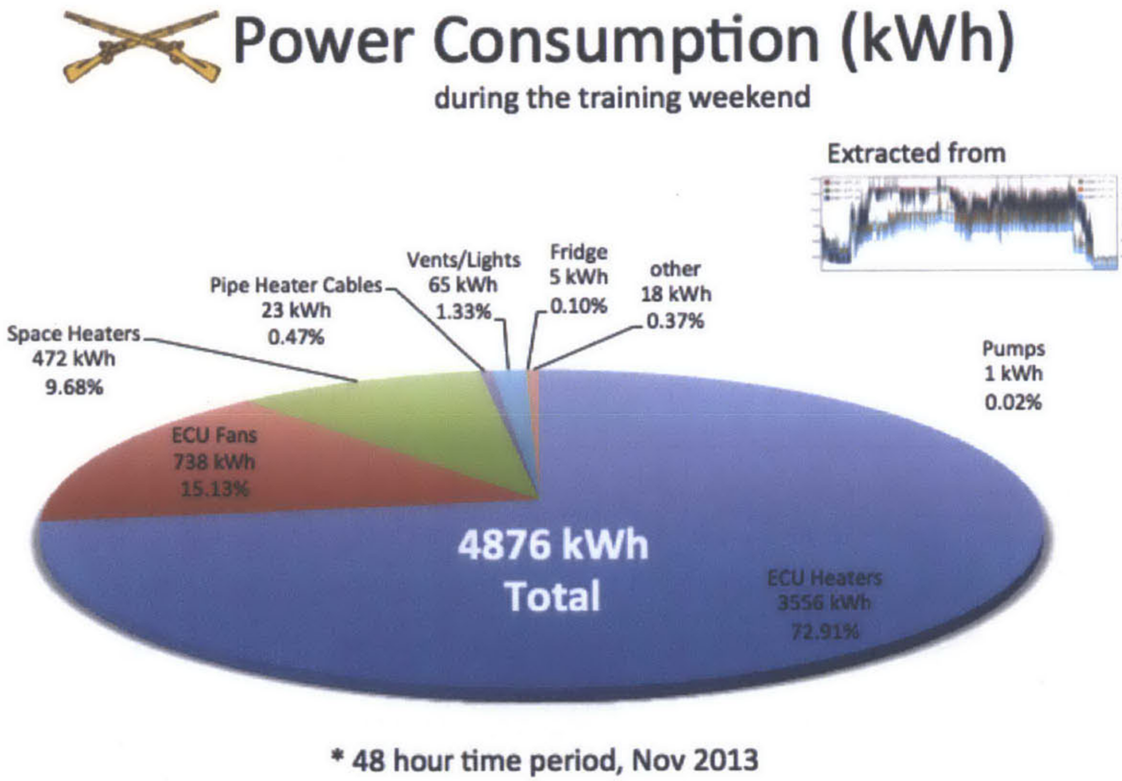

Figure 5-2: Largest loads at Ft. Devens itemized by kWh consumption

DepNILM used this highly repeatable load behavior to itemize the consumption of the largest loads on Ft. Devens, primarily those named in Fig. 4-4. Fig. 5-2 displays each load type in relation the others. 73% of the energy over this weekend went towards ECU heating coils. If the supply fan, which circulates the inside air across the heating coils, is considered part of the heat system, 88% of the cost over the weekend is attributed to 11 ECU machines. Adding in the smaller space heaters (which include the window units A/Cs), used in the showers, latrines, and kitchen, 98% of the total cost came keeping the rooms warm. The fridge, vents, pipe heater cables, pumps, lights, and all other loads not accounted for consume negligible power in comparison to the heaters.

Appliance level detail enables important comparisons. As discussed in the training schedule above, the base camp was essentially vacant during a 9-hour period on Saturday while the unit left for the weapons range. Permanent party staff at the BCIL maintain that to prevent the pipes from freezing, some heaters are left running. Altogether, the maintenance personnel describe the unoccupied winter loads as 100m of pipe heating cables, heat in the latrines (4 space heaters), kitchen (1 wall-unit heater), shower tents (two 60k BTU ECUs), and heat in 1 living space for miscellaneous use.

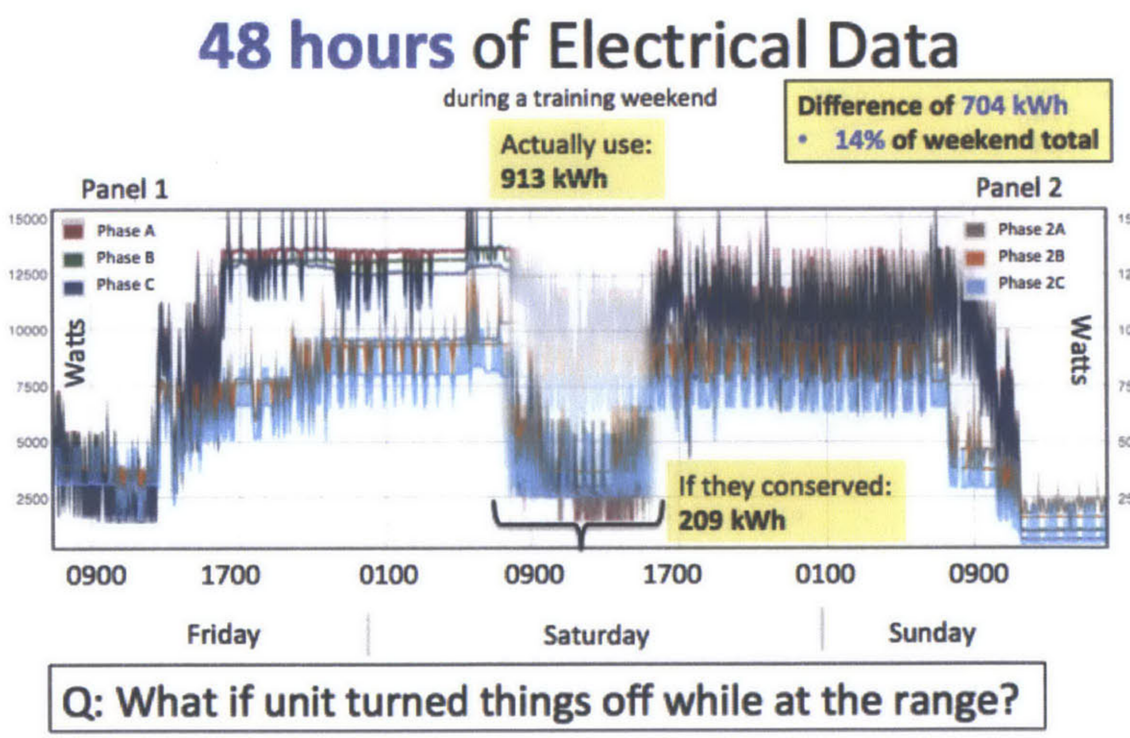

Figure 5-3: NILM measured the projected savings of conservation as 14% of the weekend total

Rather than turning things off, all heaters remained on while the unit left for the range. Fig. 5-3 shows what effect conservation efforts would have had during that 9-hour period. On 5 November, days prior to the unit's arrival, Ft. Devens was in "unoccupied" status running the previously described loads. That power consumption data superimposed over what actually took place graphically illustrates what effect turning off the ECUs would have had. During that period, actual use measured 913 kWh. Had the Soldiers not heated their empty sleeping quarters while away, the number would have been around 209 kWh, an amount measured during the same time period on a different day while the site was actually unoccupied. When the base is unoccupied, only the rooms with water are heated. Over the 48-hr weekend, this translates to 14% of the total energy cost going towards heating unoccupied living quarters. This shows that with absolutely no loss of comfort or mission readiness, the weekend power bill could have been 14% less.

The conversion from kWh to gallons of diesel fuel is Eqn. (5.1). From [36], the energy density of diesel fuel (ρ_{diesel}) coupled with the efficiency (η) of a 60-kW generator operating at full load is 11.76 kWh per gallon. Using the energy savings of 704 kWh for ΔE, the cost of waste in this example equates to 60 gallons of fuel for this 48-hour period.

$$\Delta E \times \frac{1}{\rho_{diesel}} \times \frac{1}{\mu} = Gallons_{diesel} \tag{5.1}$$

Viewing several weekends together, patterns of baseline consumption emerge. For example, if a Commander in charge of several similarly sized FOBs has the type of feedback shown in Fig. 5-4, then each of the base camps could be compared. A similarly sized unit with a much higher consumption could be questioned about its higher consumption with some basis of comparison. Determining what "typical" consumption looks like allows quantifiable comparisons of outliers, which is the basis of accountability.

A second example comes from Ft. Polk. During a two-week exercise, a CSH unit and an engineer unit occupied the buildings at the Training FOB. Because of the high tempo of training during field exercises, trainees pay understandably little attention to electricity costs. 50 Soldiers occupied building #5 during part of this two-week period. When they vacated the building, the two Bard ECUs were left

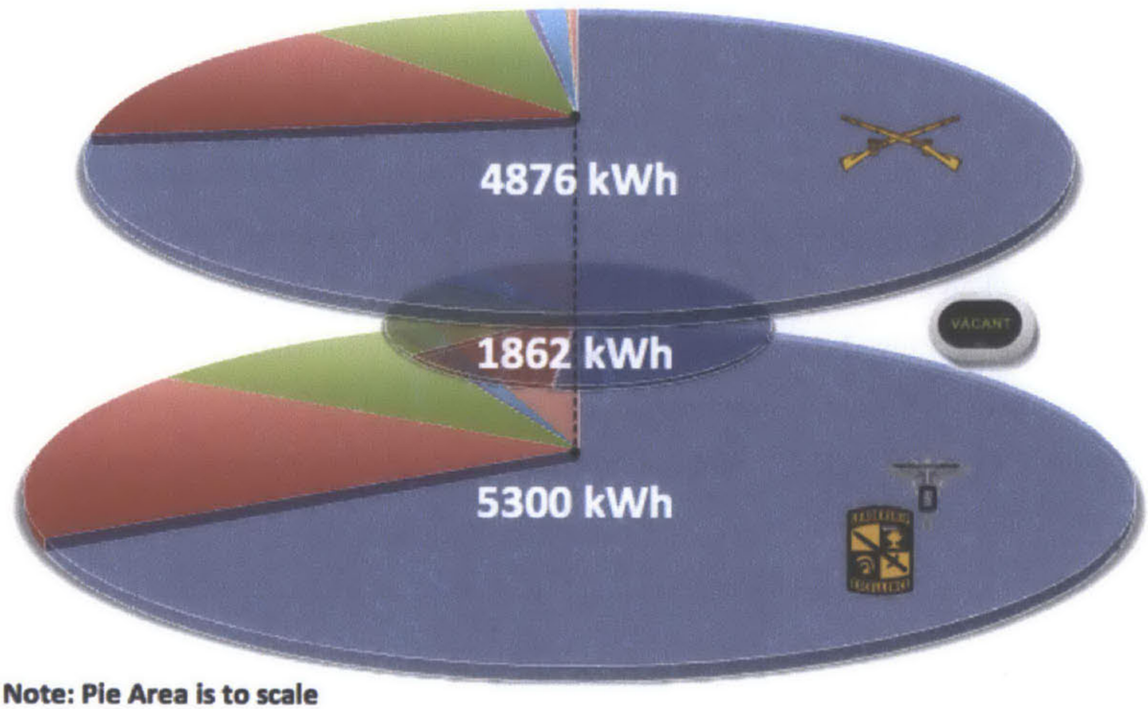

Figure 5-4: Power usage comparison across three different 48 hr periods: Infantry unit (Top), Unoccupied (Middle), and ROTC/Dental Unit (Bottom)

running. Inexplicably, within the same room, the two ECUs were set in opposition to each other, one on "heat" and the other on "cool". Predictably, during cool evening hours, they dueled over the thermostat, alternating operation every 30 minutes or so depending on the outside temperature. Fig. 5-5 shows their electrical signatures over an 8-hour period while the temperature was in the 60s. This data represents the total power consumption of buildings 4, 5, and 12. During this period, building #5 was the only consumer of the three.

The energy wasted simultaneously heating and cooling an empty room for several days is not insignificant. Every hour during this otherwise idle period cost about 4 kWh total, the heater expending about 11 kW for 16 minutes and the compressor about 3.3 kW for 24 minutes. By the same logic of Eqn. (5.1), this translates to over 8 gallons of pure waste and unnecessary wear on the equipment per day. Temporary occupants often operate equipment in ignorance of the costs. Maintenance staff may not have time to make daily checks at every site. However, you cannot hide from the current. DepNILM data is available through the Wattsworth website and routine

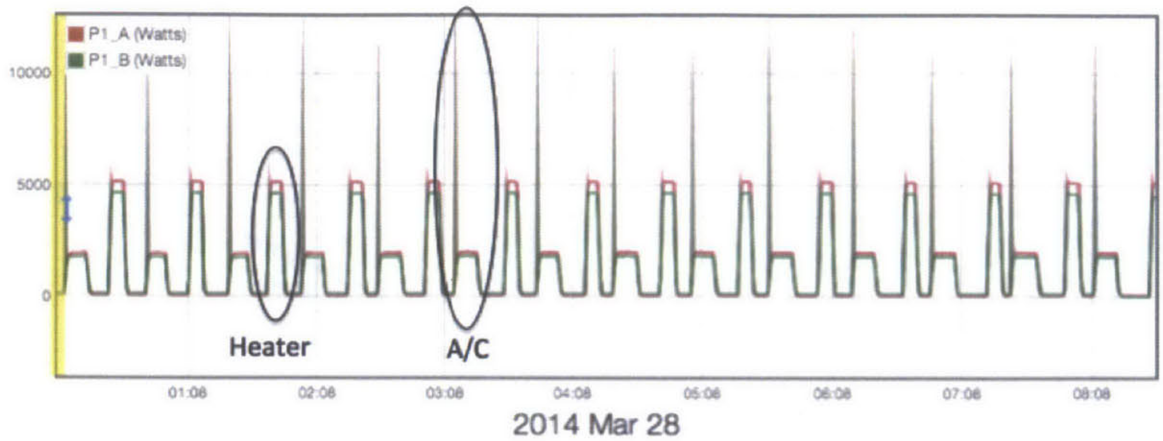

Figure 5-5: Heater and Compressor in the same room dueling over temperature

checks can easily be done remotely. NILM technology gave immediate feedback that an empty building still had ECUs running and that heating and cooling units were running simultaneously, two typical energy wasters that are easy to act upon and eliminate.

The next example puts energy savings into context. At the Ft. Polk CSH site, where generators supplied power, the non-commissioned officer (NCO) in charge of refueling did not have a flow meter on his pump nor a detailed log of each generator's daily consumption. From his experience, however, he related that he puts about 44 gallons of JP8 (aviation-quality diesel fuel) in each 100 kW generator every 12 hours. A consumption-load curve available from [37] authenticates his estimate (Fig. 5-6). The monitored CSH generators generally operated between 40-50% capacity during the monitored period, equating almost 4 gallons/hour. There is a perception among Soldiers that reducing demand while on generator power is fruitless. It is true that fuel savings are maximized under the microgrid concept, where FOBs operate fewer generators at higher capacity. Also true, however, is that reducing demand on existing generators also saves fuel. Indeed, reducing demand from 50 kW to 25 kW would mean a reduction in fuel requirements form 49 gallons to 37 gallons every 12 hours, a savings of 25%. This is not in competition but in fact complementary to savings accrued microgrid applications. Extrapolating the fuel consumption of a generator from (Fig. 5-6) for a reduction from 50 kW to 45 kW (3.68 gal/hr), this realistic 10%

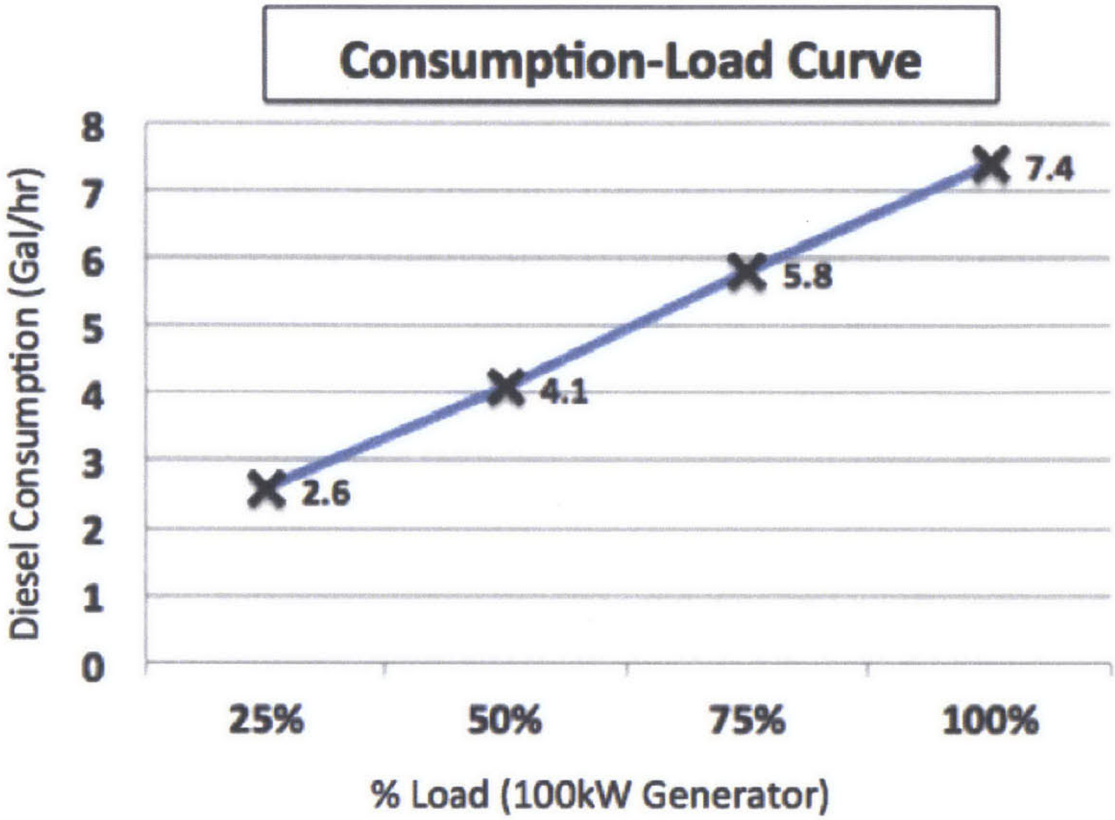

Figure 5-6: Fuel consumption by % load of a 100 kW diesel generator

reduction in electrical consumption, will result in a savings of 10 gallons per day. This corresponds to a savings of almost one 400-gallon fuel bladder per month, per generator. Note that the hospital uses 4-5 generators.

A rational operator will not knowingly waste energy. However, running equipment in ignorance of the cost allows inefficient behavior to perpetuate. NILM clarifies the cost of doing business down to the appliance level. Targeted conservation efforts will have the greatest effect when Soldiers know what they are using and how much it costs.

5.2 Early Warnings and Ex-Post Maintenance Advantages

The second category of actionable feedback relates to maintenance. Signs of machine wear and dysfunction are detectable through their electrical signals. Observable with DepNILM's high-resolution data, subtle details of machine transients can signal anomalies in the sequence of operations and trigger alarms.

Unlike traditional maintenance, which includes regular checks and parts replacements on a fixed schedule, condition-based maintenance advocates a more preemptive approach. This includes giving attention to machines that may not be broken but show signs of trouble. A key capability inherent in NILM is the ability to detect machinery faults using only their electrical signatures, and thereby know when maintenance needs to be performed.

As an example, on the morning of March 28, 2014 at Ft. Polk, an ECU in the same building (#5) that had been dueling the heater stopped behaving normally. After hours of repeatedly running 12 minutes on, 20 minutes off (Fig. 5-7), the ECU in Cool mode turned on around 8:30AM and stayed on for more than 12 hours straight. Meanwhile, a second compressor (Compressor 2) operated normally to cool building 12.

For the next several days (until maintenance personnel powered down the camp) Compressor 1 once again continued to cycle in shorter intervals, but it did so without the assistance of the supply fan. NILM detected the fault by illuminating several things. First, the power draw of Compressor 1 after the disturbance was consistently about 300W less than before, dropping from about 3.3 to just 3 kW. This difference is approximately the consumption of the supply fan. Second was the distinct absence of the "fan off" signature from all future cycles (Fig. 5-8). Third, a comparison of the differences in run times between Compressor 1 and Compressor 2 revealed significant differences. Without the supply fan circulating the air, the room will only cool through ambient heat transfers with the refrigerant, the same way a refrigerator might cool a room if the door were left ajar for long periods. For a 5-hour period,

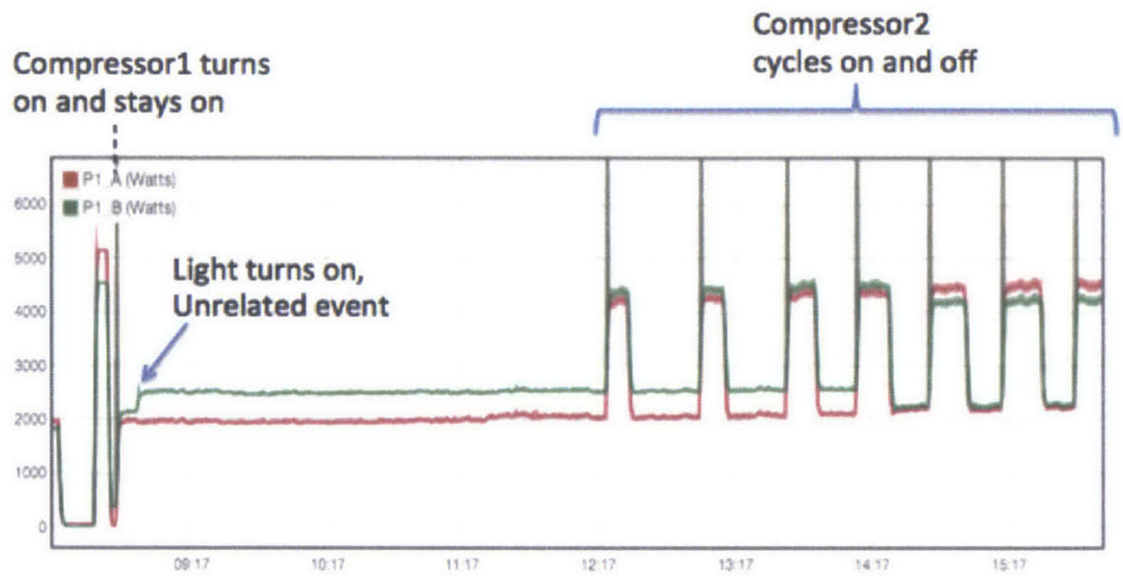

Figure 5-7: Faulty Compressor1 turns on and stays on for 12 hours

the hottest part of the day during March 29th, the broken ECU ran for 3.1 hours in intervals ranging from 7-15 minutes. The normal ECU operated for similar intervals but required only 1.4 hours, less than half the time of Compressor 1. Thus, the broken ECU used almost 2.5 times the power of a normally functioning machine. Also note that building 12 was occupied while building 5 was not, and body heat warms a room. On this particular day, the temperature peak was only 73 F). It is likely that on a hot day, the Compressor 1 would not have been able to keep the room cool.

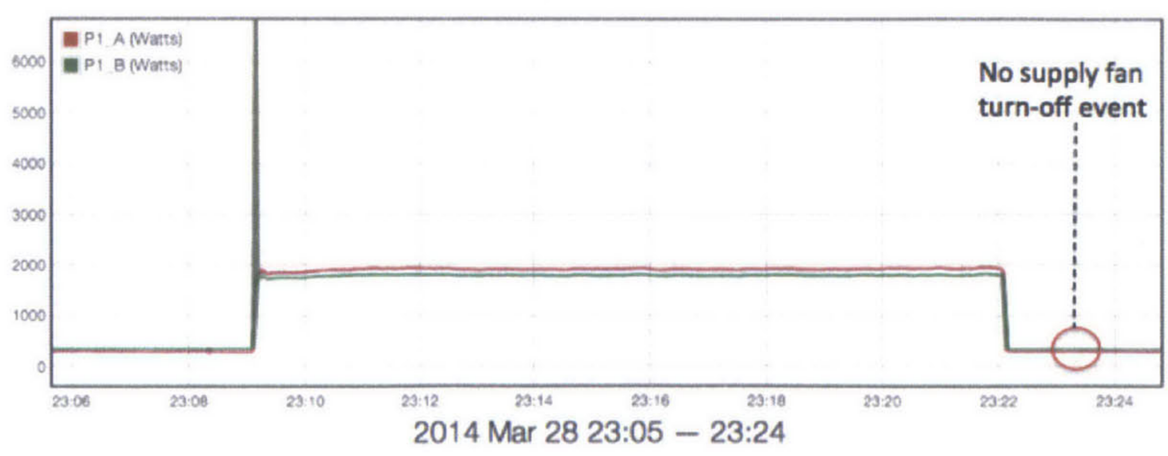

Figure 5-8: Compressor 1 turns off but the fan never turns off again

The broken ECU illustrates that monitoring with the NILM system makes waste apparent that might otherwise go unnoticed. Since the compressor continued to run and the room was maintaining temperature, the machine appeared to be working fine. However, the amount of power it consumed was far more than a functional unit would have consumed. Ultimately, the broken ECU costs double what a normal ECU costs to operate, and will continue to do so until routine maintenance finally catches it.

The field hospital at Ft. Polk shows condition-based maintenance in another light. A major feature of the system is to act like an aircraft black-box repository for historical electrical facts. During NILM monitoring, the MEP 807 3-phase 100 kW diesel-powered generator unexpectedly cut off (shed) all electrical loads several times over a 1-week period (Fig. 5-9). On five of these occasions, the generator engine also shut down seconds after. After several blackouts followed by unsuccessful troubleshooting, the unit replaced the generator with another identical machine. When the replacement generator responded similarly, also shedding the load and cutting the engine, they tried a third generator with similar results. Finally, support personnel elected to relieve the distribution box of most of its loads and redistribute them to other generators. The maintenance NCOIC believed the problem to be with one of the electrical loads in the Emergency Room (ER), though he could not confirm it.

The MEP 807 is a synchronous, brushless A/C generator with a digital voltage regulator (DVR) and a digital generator set control (GSC) panel (Fig. 5-10). The interface displays per-phase voltage, current, power factor, VAR, as well as frequency, total kW consumption, and other data. It also has alarms that light when fault conditions are present [38]. For a generator operator, their goal is to keep the voltage and frequency precise and stable and ensure the % load (in kW) in an optimally efficient range (usually 80% capacity). During these faults in the field, however, "low oil pressure" was the only alarm lit when the generator engine cut. The % load never surpassed 60% according to the operators, though this was based on infrequent spot checks. NILM data confirms their assertion about % load, however.

After the field exercise, the mechanics investigated the problem using a load bank.

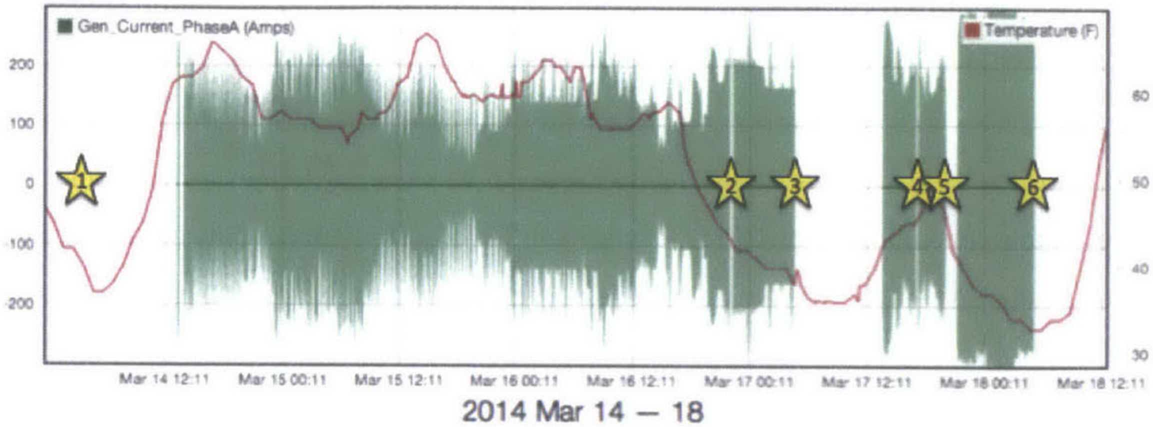

Fault #	Date	Time	Temp (F)
1	13-Mar-14	~1:00 AM	44
2	16-Mar-14	9:17 PM	48
3	17-Mar-14	3:56 AM	41
4	17-Mar-14	4:09 PM	49
5	17-Mar-14	6:51 PM	52
6	18-Mar-14	4:17 AM	38

Figure 5-9: Current (left trace) and Temperature (right trace) during Generator Faults (marked with stars) during a 1-week period

Connecting the generator directly to the full load, there were no apparent problems. Then, connecting to the generator through the distribution box, they tested each circuit breaker (CB) under load. It was discovered that one breaker, the same one that had been connected to the ER, was faulty. Three times in a row, the generator shed the load and cut the engine in the same manner as it had in the field. No fault alarm was visible except for the same "low oil pressure" light. Upon further inspection, the 3-phase CB had an open pole and was replaced, resolving the problem. In light of Fig. 5-9, it seems consistent that the generator crashed at night during cooler temperatures. Metal expands when heated and contracts when cooled. The circuit through the third pole of the malfunctioning CB could very likely have been closed when temperatures were warmer and open when temperatures cooled.

The NILM helped isolate the cause of the faults by narrowing down the possibilities. No one was watching when the faults actually occurred, so the voltage, current, frequency, and power consumption data at the final moments were unknown. 8kHz

Figure 5-10: Generator Set Control Panel of MEP 807

samples of voltage and current from the generator faults in the field provided this useful feedback. First, the generator was not overloaded, evident from the power plots during the last 10 seconds before load shed. Fig. 5-11 depicts a close-up of fault number 3 and indicates about a 50% load. Note that Fig. 5-11 shows the spectral envelope of peak current.

Second, the voltage varied widely after the load shed. Shown in Fig. 5-12, fault 2, after the current goes to zero, peak voltage in phase B dips to 80V. In this instance, the generator recovered and reconnected the load before faulting again shortly after. The DepNILM drew power from the auxiliary outlet on the generator, so when the load shed, the current sensor output went to zero. However, the generator continued to provide power to the outlet and an uninterruptible power supply (UPS) in the system kept the computer recording even after the generator altogether ceased.

Third, the magnitudes of three phases of current were relatively balanced. A

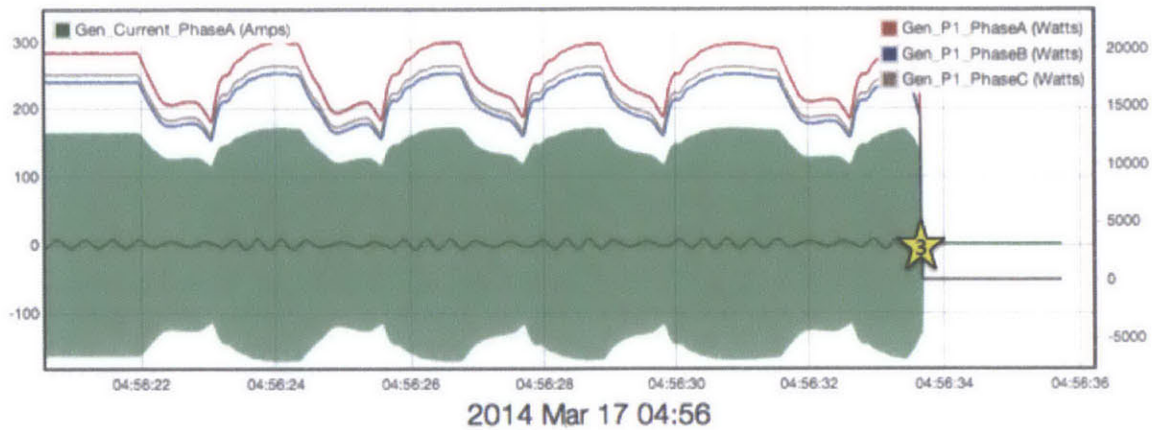

Figure 5-11: Current and 3-phase Power 10 seconds before fault 3

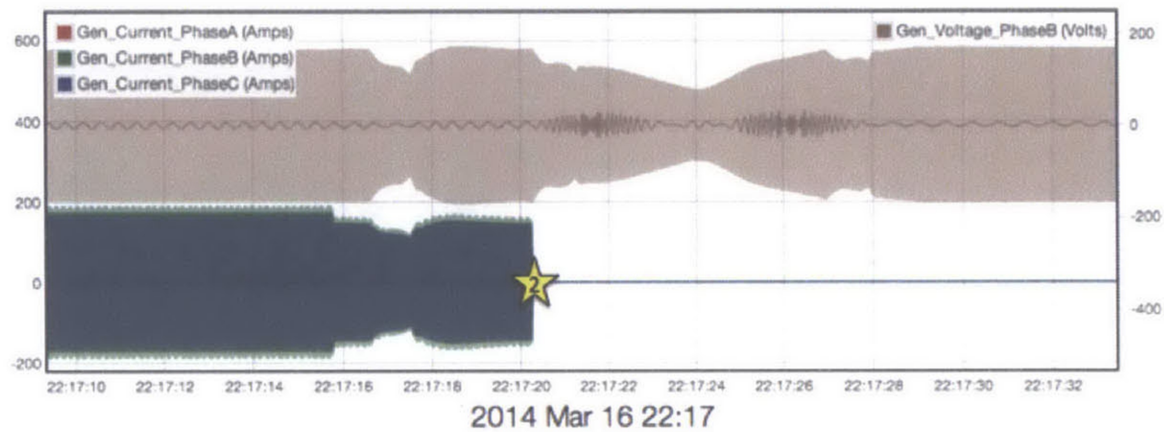

Figure 5-12: Voltage and 3-phase Current and 3-phase power 10 seconds before and after fault 2

snapshot of a few line cycles from fault 4 shows this (Fig. 5-13). Also visible, however, is some disproportionate harmonic distortion on the phases. Reasons for this are speculative. Perhaps most of the fluorescent lights are plugged into the same phase, as is the case at Ft. Devens. The electronic ballasts distort current. The same distortion could also be attributed to the 10+ laptops plugged in at any time. Another postulation is that there is a voltage imbalance due to the open phase on one of the breakers causing the current to take a similar shape. Due to safety concerns, we elected to monitor only one phase of voltage with the DepNILM through the auxiliary outlet and assume clean sine waves on the other two in order to calculate power. The absence of these other two phases, a limitation to the conclusiveness of the cause of the faults, is a problem that can be solved in future installations.

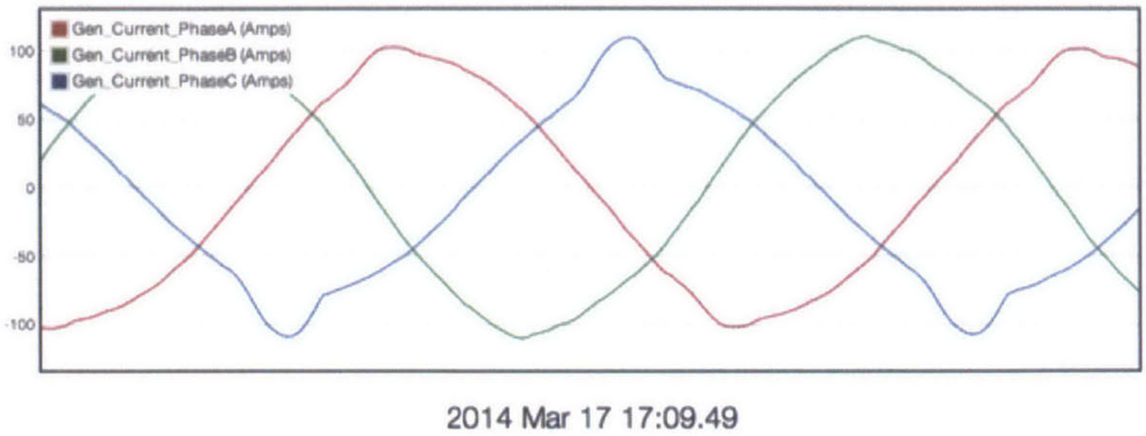

Figure 5-13: Line cycle of 3-phase Current seconds before fault 4

A final example of condition-based maintenance revealed through NILM comes from the ECUs at Ft. Devens. As previously explained, the compressor has a "soft start" function followed by periodic cycles where it changes rapidly between loaded vs. unloaded state. Apparent from the technical manual as well as theory of operation documents made available by the vendor, the compressor stays loaded or unloaded but does not turn off. Exceptions include when high or low-pressure situations pop circuit breaker #5 or #6 respectively, but the compressor in those situations will not restart automatically after pressure resolves. The breakers have to be flipped manually.

On September 13, 2013, one ECU signaled irregular behavior in the form of frequent compressor start-ups (every 5-10 minutes over several hours). These are indeed compressor cycles because zooming in reveals the familiar shape of the soft start from Fig. 5-14. The temperature at Ft. Devens at the time was 65 F). These irregularities were identified a month after the event, so fault codes available from the ECU display were unavailable. However, the NILM data and the technical manual [33] provide basis for some reasonable conjectures as to the cause. First, the outside temperature is low to be demanding cool air. Second, the Copeland Scroll Digital Fault Codes list has 8 possibilities, two of which require manual reset (disqualifying them). Code 3 is a "Compressor Protector Trip", and causes compressor lockout, presumably meaning shutting it off until the protector is reset. Code 5 is due to demand signal loss and leads to the same action. This seems unlikely as the compressor kept coming back on, apparently regaining the signal. The action of Code 6 is to reduce the capacity to 50%, which is not what happens electrically. Code 7 was un-programmed at the time of this writing. Code 8's action is to unload the compressor, also not the case here. The compressor was turning off and then on repeatedly. Code 9, a low voltage fault, causes the compressor to trip. Of the possibilities above, 3, 5, and 9 seem possible and 9 the most likely. Troubleshooting from a distance is possible through NILM Manager when these black-box facts are available, no matter where tech support may be located. It is much easier to prepare for (and carry parts for) a narrower list of possibilities, and focused recommendations can be made if site visits are unfeasible.

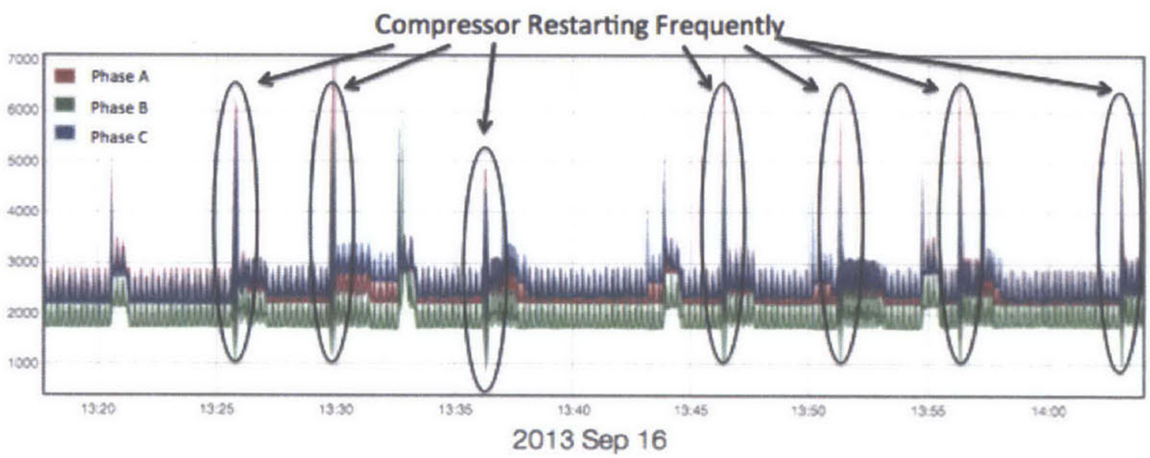

Figure 5-14: Strange behavior of ECU compressor

Compressor activity at the hospital site at Ft. Polk also showed potential for electrical diagnosis. Immediately visible from recordings was the significant variation in power between different ECUs. During the night, when temperatures were generally around 60 F), operators often left some ECUs in Cool mode and some in Heat mode. The tents all connect with narrow hallways allowing temperatures to diffuse into adjacent rooms. The ECU Compressors used by the CSH modulate using a VFD. The inputs to the VFD are six temperature sensors and two pressure sensors. The system seeks to maintain pressure within a specified range, and uses the condenser fan and VFD to do so. From the following diagrams in Fig. 5-15, it is clear to see when the ECUs are working harder to maintain pressure in that range. Compressor 1 is stable, turning on and ramping up smoothly, then finally peaking as the condenser fan turns on. The temperature in the tents is close to the thermostat temperature limiting run time to only a few minutes. Compressor 2 is less stable, with close to 500 W fluctuations accompanied by frequent cycling of the condenser fan. Note the outdoor temperature is between 65-70 F). Several possibilities for this could be considered, but one of the first should actions should be inspecting the color of the refrigerant charge through the sight glass.

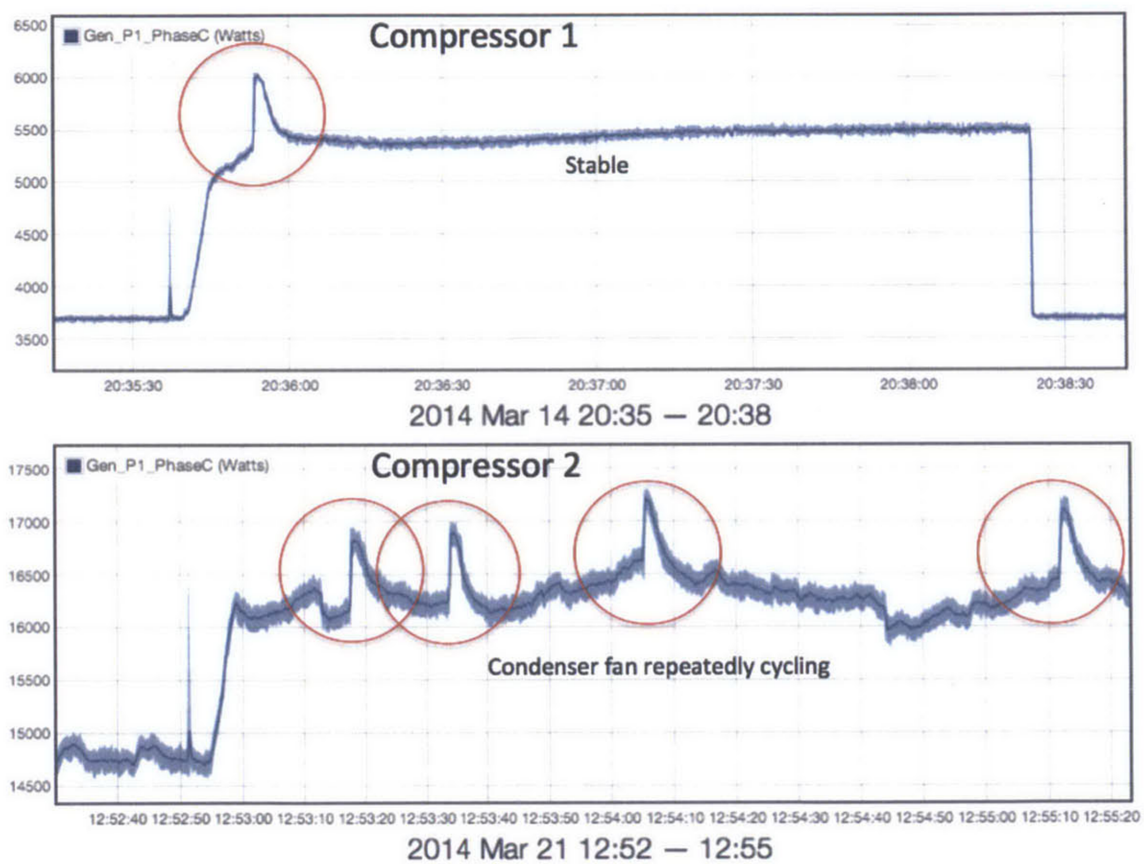

Figure 5-15: Hospital compressors 1) under stable conditions, 2) cycling VFD to maintain pressure, cycling condenser fan and VFD to maintain pressure

Another common metric signaling the need for routine maintenance is the number of hours operated. At Ft. Devens, the NILM calculated machine hours for like devices using the on and off times. Fig. 5-16 makes a comparison between hours of operation over three weekends, two of them occupied and one unoccupied. Having such data available from a single meter reduces the need to manually check each machine. Rather, maintainers could be sent a notice when some hour threshold was reached.

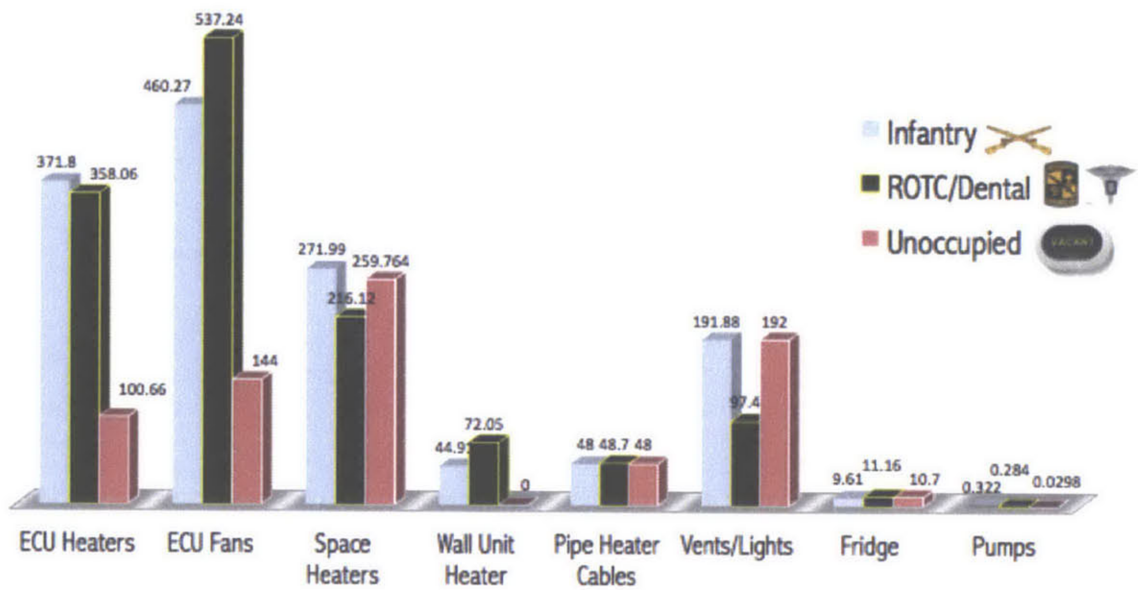

Figure 5-16: Machine Hours over three different weekends

With appliance-level resolution, equipment malfunctions can be identified and targeted efficiently. Furthermore, every NILM device becomes a black-box reference. Through any web-interface, NILM Manager permits the user to access highly detailed data regarding the electrical network, even in a bandwidth-constrained environment. To quantify the reduced data transfer requirements, we remotely accessed, analyzed, and diagnosed electrical consumption data from Ft. Polk using a mobile 3G WIFI hotspot from AT&T collocated with the three DepNILM systems. Over the following weeks of analysis, total data transfer amongst all three DepNILMs was less than 2GB. Most of the preceding accountability and maintenance examples were identified from our desks at MIT. Tech support can be anywhere and do the same.

5.3 Human Activity Awareness

NILM provides awareness of human activity within a network. Each device has an electrical fingerprint, and specific devices imply associated human actions. When electrical activity is detailed enough to tell when appliances cycle, experience and intuition can extract meaning.

Here, results require both art and science. Every electrical device has a specific purpose. Some, like the ECUs, are thermostat controlled and thus without direct human control most of the time. The On and Off patterns are closely follow fluctuations of outside air temperature. For most loads, however, their turn-on corresponds to a specific action taken by the user. The projector tends to run during briefings, for instance. The lights turn off at night when it is time to sleep, or on in the morning when it is time to wake up. The oxygen machine runs continuously when there are patients on a ventilator. For Soldiers in the field, the operating schedule is called the battle rhythm. Through some interpretation on the part of the analyzer, patterns take shape from observing what loads they use and when.

For sanitary and environmental reasons, wastewater must be managed. Fortunately, this makes it measurable electrically through pump events. On Ft. Devens, this is a multi-stage operation. Waste from the showers, sinks, and latrines is gravity fed into underground holding tanks with a capacity of 10-12 gallons each. Each double-latrine and double-shower has its own tank, and each tank has its own pump. There are a total of 4 sub-systems on the 150-Soldier camp. When full, submersible pumps eject the contents of their respective tanks into a larger reservoir managed by one larger pump.

Pump event times, therefore, correspond to when Soldiers are using water, namely the latrines and showers. During the same weekend described in Fig. 5-1, the corresponding pump events are overlaid with milestones of the timeline in Fig. 5-17 to evoke meaning. Between 1200-1660 hrs on Friday, with only 2 personnel from the unit on site, there were no pump events. Three pumps events within 1 minute at 1559 hours signaled the arrival of the main body (88 personnel). From 1600-2200 hrs, pump events were spread out about every 30 minutes. Between the hours of 2200 hrs on Friday and 0459 hrs on Saturday, not one pump cycled. This corresponds with the time the unit was probably sleeping. There were 13 events within 41 minutes early Saturday morning while the Soldiers did hygiene. Training for the day consisted of a weapons range 2km away, and 88 personnel left the base by 0700 hrs. The last two pump event of the morning, predictably, were both at 0700 exactly. One pump event during the 9-hour training can be explained by noting that two were personnel left behind to watch equipment. Scheduled to return from training at 1700 hrs that day, three pump events between 1626 and 1700 hrs indicate a somewhat earlier return. Saturday night, between 2100 and 2200 hrs, there were 10 events. After a day in the field, several took showers. That last event that night was at 2222 hrs, and the first on Sunday was at 0437 hrs. Aside from the early outlier, there were 11 events from 0556-0632 hrs where morning hygiene took place. Compared to Saturday, it appears that wake-up time was an hour later on Sunday. With only 4 scattered events after 0632 hrs, the last of which was at 0835 hrs, it is evident that most of the unit vacated the site before the scheduled 1200 hrs. Looking at the other electrical data, 7 fans and heaters were turned off, between 1100 and 1105 hrs. That is around the time the final individuals departed the site.

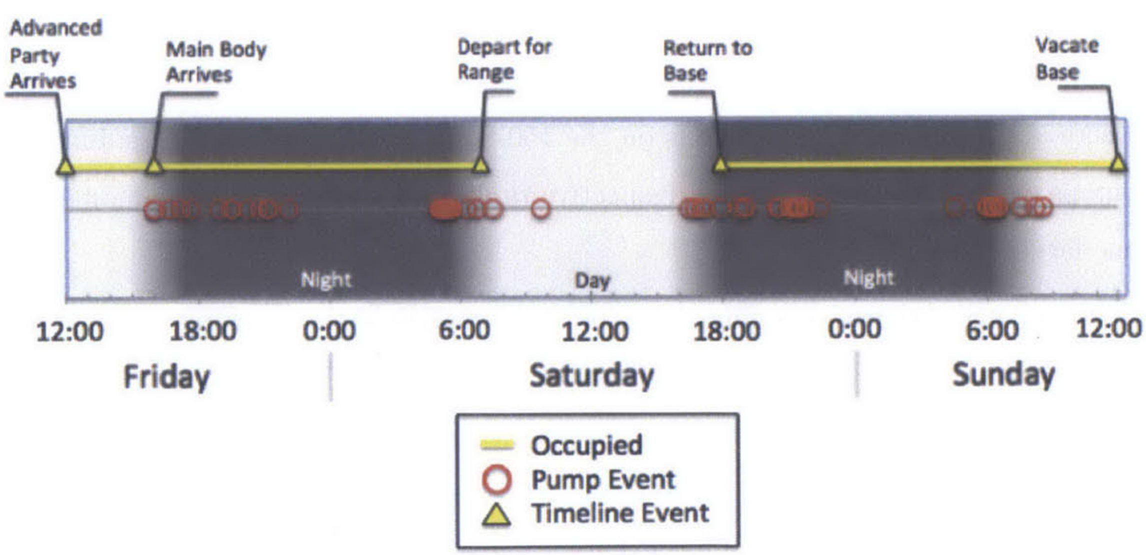

Figure 5-17: Pump events during a 48-hour training weekend indicate human activity

At Ft. Polk, during the monitored training rotation, the CSH unit received and treated notional casualties during the 10 days of the exercise. An active Army Brigade Combat Team (BCT) was conducting missions nearby against an opposing force. During their advances, the observers grading the BCT assessed casualties, and the unit had to respond appropriately, treating the notional wounds and evacuating priority injuries by helicopter to the CSH. During the monitored period, the hospital received, treated, and discharged over 100 casualties.

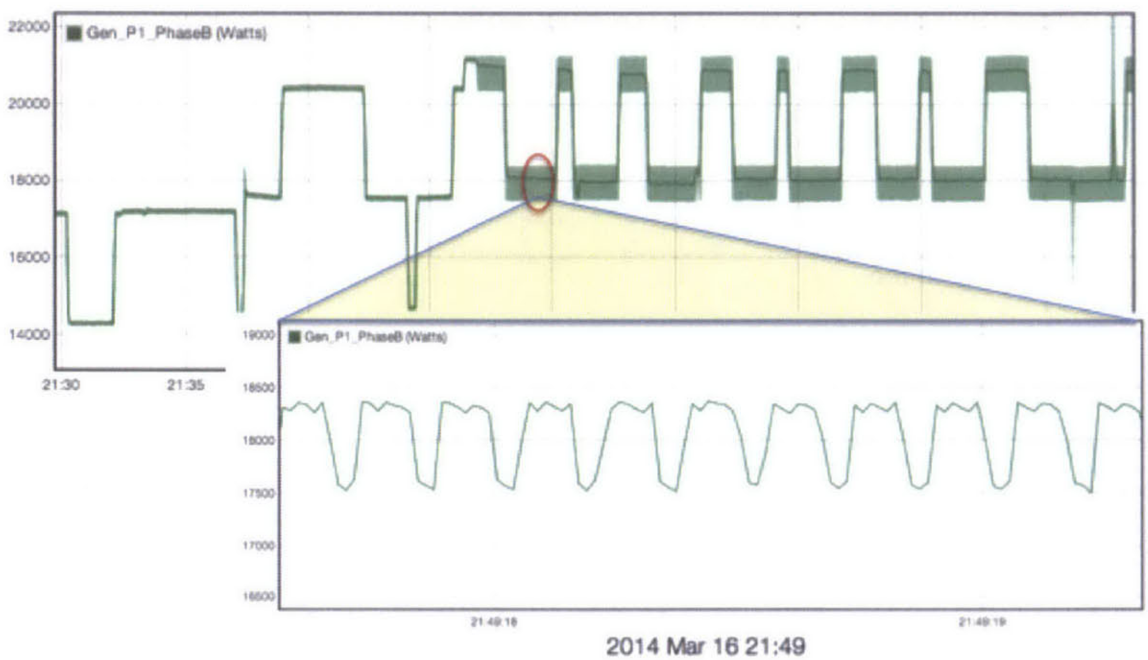

Figure 5-18: The electrical signature of an electric heating blanket may indicate the arrival of a new hypothermic patient

One tool that was used during casualty situations shows up clearly. The Bair Hugger, used for treating or preventing hypothermia, is a forced air temperature-heating unit. Hypothermia is very common amongst trauma patients, where over 50% receive some sort of treatment for it in conjunction with their other injuries. When attached to a blanket through a hose, warm air generated by the Bair device flows around the patient raising their body temperature. Several times a second, a heating filament is repeatedly cycled on and off to maintain temperature as a fan continuously blows across the coil towards the patient (Fig. 5-18). The appearance of this device in the data is a good indication that new patients have arrived.

Another indicator of human patterns is the ubiquitous coffee maker, symbolic of the beginning of a long period of work. Its resistive signature is clear, though it depends on the size of the coffee pot also. For the large pot in the hospital, there are repeated 45-second bursts of 1kW power draw followed by 15 seconds off until the coffee pot is turned off. Fig. 5-19 shows that the morning of March 15 started at 0838. On the 16th, it started at 0900. One could opine that this time corresponds to the highest-ranking person's arrival to work, though that would be circumstantial. In the sleeping quarters, where the hospital staff rested, a small coffee maker could be seen most mornings beginning shortly after 0600 hrs, corresponding to the 12-hour shift changes at 0700 hrs every morning.

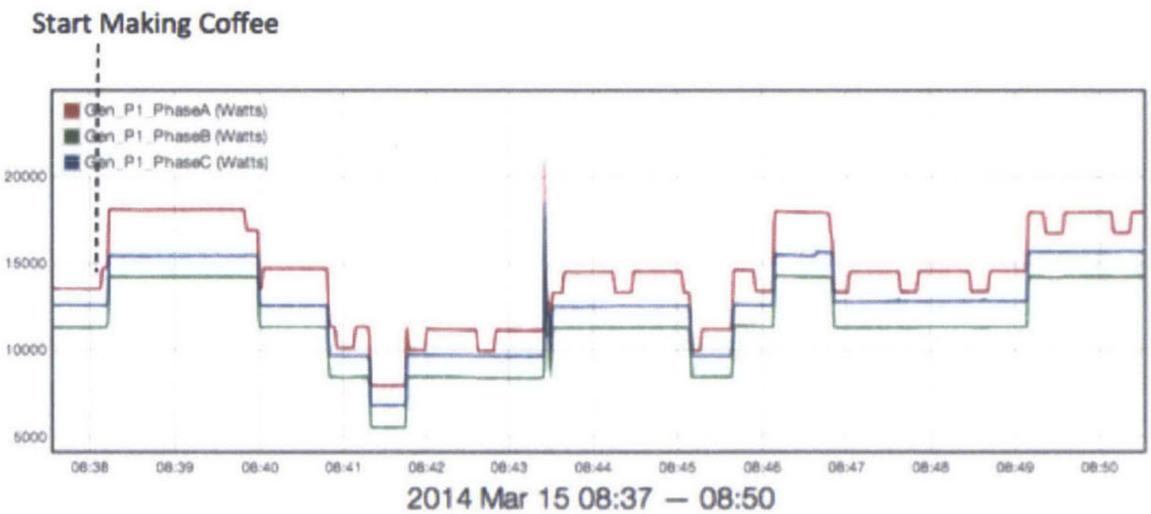

Figure 5-19: A coffee maker on Phase A turns on most mornings around 0900 hrs

Detailed electrical measurements permit another means of confirmation of human activity. Accurate arrival or departure times generally correspond with specific devices turning on or off. The presence or absence of certain electrical fingerprints during a certain period may confirm whether something is happening or not happening. Gathering this type of human intelligence requires two things: a library of electrical signatures of the loads on the network and some understanding of the size and disposition of the occupants. A skilled analyst accustomed to looking at electrical data may be able to recognize specific equipment simply from experience if actual loads are unknown. The benefit of loads on a base camp, especially the FP-150 at Ft. Devens, is that military loads tend to be highly standardized. The library of loads need only be taken once, and they apply wherever that base camp is set up.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 6

Beyond Accountability, Maintenance, and Human Activity

6.1 Other Features

Current reveals a great about an electrical system. Beyond power, it also proves useful in several other ways. We learned from Cottage School that the NILM shows promise as a sub meter for gas. If the flow rate is relatively constant during run times, then amount of gas consumed is that flow rate amount multiplied by the hours operated, which can be determined by the boiler On and Off times detected by the NILM.

The DepNILM can also serve as a back-up flow meter for water. Pump events correspond to the quantity of water consumed simply by counting the times the pump cycled and multiplying by the tank size. It is not a perfect measurement. The average pump run time is 17 seconds, though the pump will run until the tank is almost empty. If water is continuously running, the pump may run longer and move more water than 10-12 gallons. While flow meters provide the same data more accurately, their displays generally must be read visually. If there is more than one water source, there will be at least that many flow meters to read. Indeed, Ft. Devens has 6 flow meters installed, one for each double-latrine and two for each double-shower (hot and cold). A test conducted on 25-27 April shows the relative accuracy of measuring water consumption electrically. A training unit occupied the camp for approximately 48 hours over the weekend. The sum of the 6 flow meters from Ft. Devens totaled 886 gallons during that period. In parallel, NILM extracted 80 pump events during the total training period. Equating each pump event to 10-12 gallons, the NILM estimate was 800-960 gallons. For a logistician, this is sufficiently accurate for planning when the unit will need another delivery of water.

A third feature of NILM is that it serves as solar sensor. The FOB buildings at Ft. Polk are equipped with outdoor lights above each door. Controlled by photo-resistors, when the sun goes down the light powers on. Fig. 6-1 shows the consistent patten of the light activation and deactivation at the beginning and end of each day around 1918 hrs. The National Oceanic and Atmospheric Administration (NOAA) records that sunrise was at 0714 hrs and sunset at 1918 hrs on March 18, 2014. The light turned off at 0710 hrs and on at 1934 hrs.

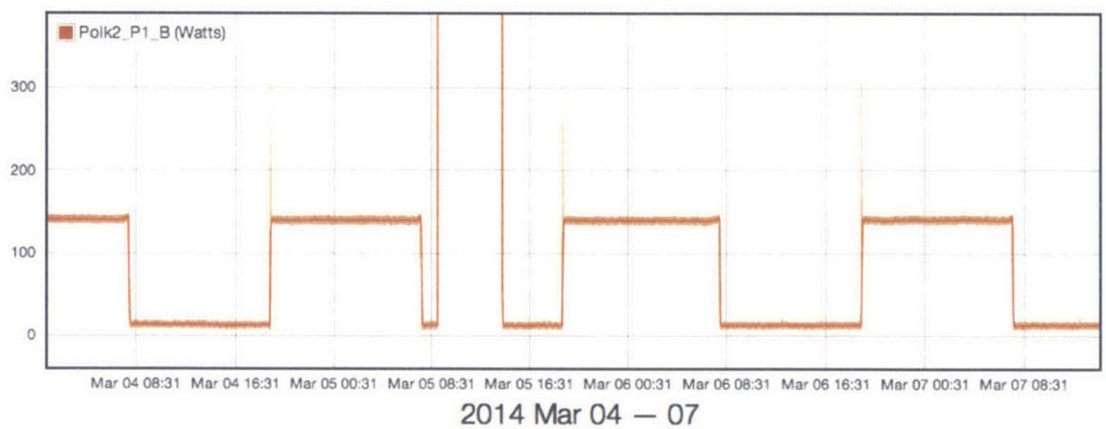

Figure 6-1: A simple display for consumer feedback provided by the NILM system

6.2 The Training Phase

The training phase can be insightful. Aside from just providing the appliance-level values of kW consumed, it can also uncover commissioning mistakes. During the training phase at Ft. Polk, the 6 identical Bard Wall-Mounted ECUs were each initially cycled through their fan, heat, and cool modes (Fig. 6-2) . The expectation is that all six signatures should be identical. However, one heater and one compressor differ in magnitude significantly from the others by nearly a factor of two. One physical difference between the wall units is the thermostat. Five of them use a White-Rodgers brand thermostat, the other uses a LuxPro. Both are digital and single stage for heat and cool. The LuxPro is programmable, though only for time of use. On page 13 of the Bard Manual #2100-398A, the unit shows it can accommodate 2-stage heat and 2-stage cool. Looking at the wiring diagram Bard provides, two thermostats are most likely wired incorrectly to the low stage. Incidentally, the six Bard ECUs connected to the other panel all had identical signatures.

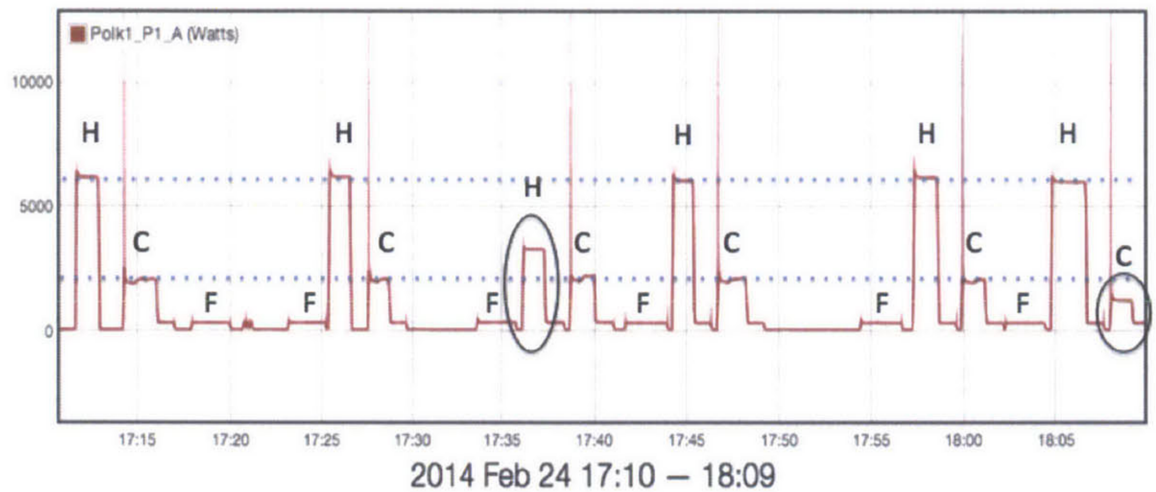

Figure 6-2: Heater (H), Compressor (C), and Fan (F) signatures show that one heater and one compressor draw significantly less power than the others

The measured effect of this different wiring scheme is interesting. When the training unit's advanced party arrived, they occupied two buildings: a TOC on panel 1 and sleeping quarters on panel 2. They used only one heater in each room. During an 11-hour period, the heater in low mode (panel 1) ran 3 hours longer than the heater in high mode (panel 2) (see Fig. 6-3). The high-mode heater cycled 24 times while the low-mode unit only turned off twice. In terms of total power consumption, the heater in low mode actually used 15% less by power, though during the cold night the low-power panel failed to reach the target thermostat setting. The trade off for lower power here was a loss of comfort.

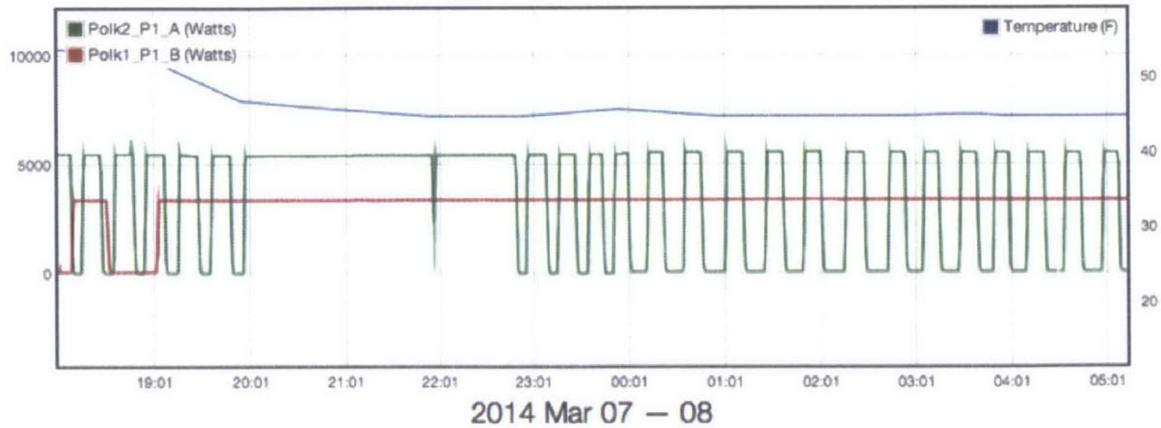

Figure 6-3: The heater on panel A drew more power but cycled more often, while the heater on panel B drew less but never reached thermostat temperature during the night

For future installations, the training phase is crucial to accurate data and complete understanding. Taking the power system down to zero is not always feasible due, so cycling loads in isolation may not happen. This is still acceptable so long as you cycle each load several times. This way, if some unknown device interferes during the training phase, at least one of the events should be "clean". Also, few people (even maintainers) actually know the full extent of the loads on a system or subsystem. The best way to ensure you are not missing anything during the training phase is to work through each sub panel. Every load is traceable to a circuit breaker by code, so if you can cycle all of the breakers, you will get all of the loads. Last, many loads have multistage operations, and it is imperative that you cycle the machine through every state.

6.3 Hardware Lessons

The design for DepNILM works well for electrical panels. However, the system needs improvement for installation on generators. There are several differences. First, generators vibrate significantly. LEM CTs and Molex connectors are not robust enough to withstand this shaking. Second, the voltage terminals are not as easily accessible as they are in the panel. Panels are highly standardized, and 15-20A CBs fit easily to each phase via the bus bars. Access and protection to the different phases of voltage in a generator are unique to each generator. Often, the 250 kcmil terminal lug connectors are the only safe access point for most users. In the case of the CSH scenario at Ft. Polk, three-phase voltage access requirements were circumvented through the use of the convenience receptacle on the outside of the generator. While safe (the primary concern), this only provided a single phase of data. The other two voltage phases had to be manufactured using the zero crossings from the accessible phase to rotate the other two appropriately. See the Sinefit publication [13] for more on this.

Another learning point had to do with the UPS. The NILM system draws power from a single-phase connection on the panel or generator. This source feed splits inside the tuff box. One branch goes directly into the NILM sensor enclosure. The other branch feeds a GFCI duplex receptacle. The UPS draws power through this receptacle and powers the computer and the wireless MODEM. When the system loses power due to source failure, the UPS provides about 30 minutes of back up. Generally this is enough to survive the outage until power is restored. DepNILM is meant to be austere. Even during outages, the data on the computer is remotely accessible. Sensor data during this period is recorded as all zeros. As it happened, the generator died and remained off for longer than half and hour. The cheaper model UPS (CyberPower CP425), does not have the feature to restart automatically upon power restoration. This also happened when the entire FOB lost power. Thus, there are days of missed data due to power failure that would have been hours.

The next version of DepNILM will include more ruggedized sensors and connections. Current research on non-contact sensors [2] may resolve this issue altogether. More intelligent UPS such as the CP825 will be better able to withstand reliability issues that tend to accompany austere power systems.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 7

The Way Forward

The NILM system provides consumers the basic feedback information they need to reduce energy usage. In a very short amount of time, it tells them what is running, how much each appliance uses in an entire system, and where energy reductions can be made most effectively. It also tells them if any appliances are faulty and need maintenance. Finally, it allows them to reasonably determine, using only electrical data, who is using what device at any point in time.

Fig. 7-1 is what the user needs to know. It is the interface we developed to display feedback. We call it a Facility Management Dashboard. The average consumer doesn't care about power factor, the number of VARs on each phase, or other technical data. They want to know what they can do to reduce their costs. This entails knowing what loads are currently on and how much power each one is using.

The bottom right section of Fig. 7-1 provides a simple list of the top 5 or six loads and how many of each are currently running. Second, users need to know what they could cut and what that would save. The top right of the dashboard shows the pie chart with each slice showing the percentage of instantaneous load itemized by load type. The radius of the pie scales to the number of total kW's being used. In this example, since the largest slice is made up of heaters, that is the highest-value target. How the costumer trims their load is an individual choice, but with this tool they can measure the effect of their decisions. The notion of swing loads, where nonessential devices can be temporarily limited for a sudden power need, becomes easy to implement when the information is complete. Third, there should be historical reference data to show the time-of-use data by device over the past 24 hours. The top left shows a stack chart (color coded by major load) with this information. That way, users know when they use the most power and what loads that consists of. Finally, the bottom left section some overall metrics such as the gallons of water consumed, the peak energy consumption, and perhaps even the estimated gallons of fuel used by the generators (if applicable).

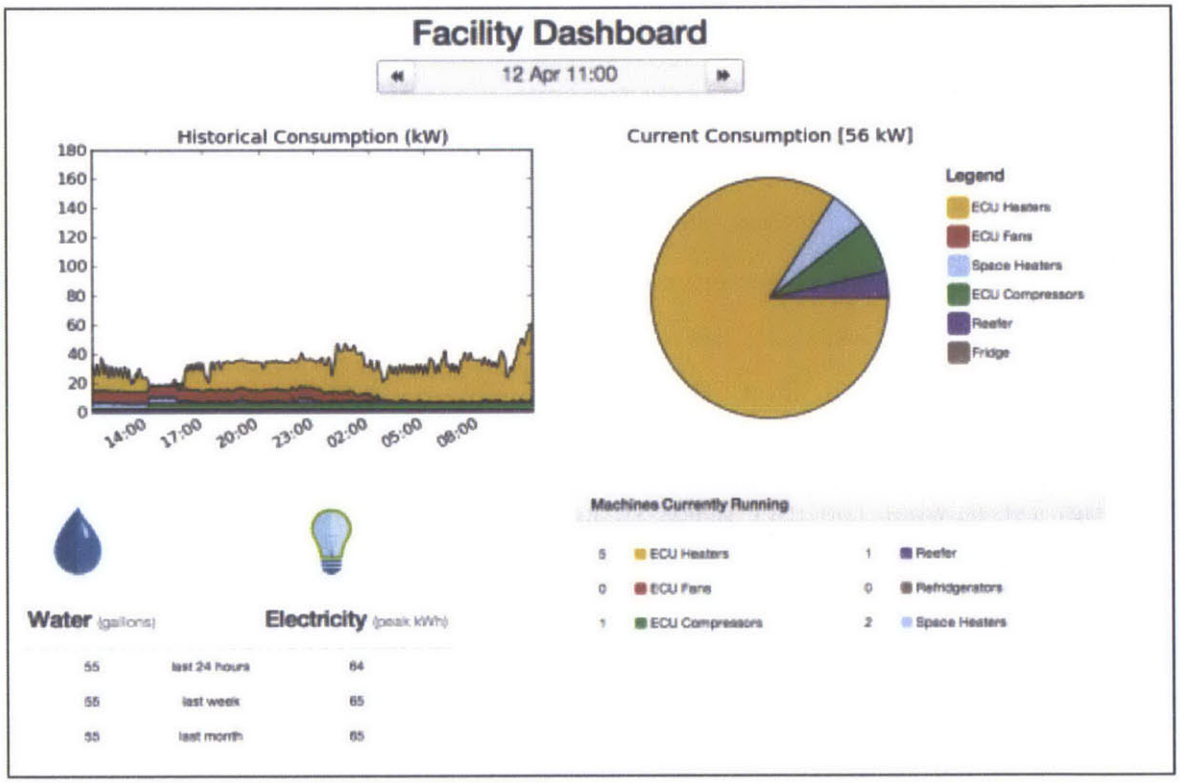

Figure 7-1: A simple display for consumer feedback provided by the NILM system

The inputs required for appliance-level feedback are voltage and current (measured at the panel), temperature, and a library of electrical load signatures in the system (obtained from the training phase). In a microgrid environment such as the military, where equipment is highly standardized, the library of loads is reasonably easy to gather. For an experienced user, the training phase may be omitted. Once accustomed to looking at transients, many loads can be reasonable inferred just from the shape and context. Some idea of the occupant's size and disposition is also helpful but not required. The output is a report on the power system that provides enough information to act upon. With no extra sensors, a good estimate of daylight, gas, and water are also possible.

Live data is being collected and reports generated at Ft. Devens as of this publication using the software filters I wrote onto the NILM Manager platform. Experiments into real-time report generation of in the format of Fig. 7-1 are being conducted to refine the product and verify the accuracy. This chart is our vision for the how to inform the consumer who desires actionable feedback about their electrical network.

NILM is scalable through the NILM Manager tool. Every base camp and facility deploys a DepNILM on each panel, and each monitoring device should be networked via LAN or WAN, depending on the circumstances. When the computers located with each NILM are online, they automatically phone home to NILM Manager. If not online, they will report to a central computer within the facility. Everything on their hard drive is immediately visible, including historical data as far back as the installation day. After a day of data collection, the major loads will likely stand out. The training phase could be done remotely, though it is important to get precise times that each device is turning on or off. Once the load signatures are collected, the software to process reports can updated remotely. Once complete, the base camp manager will have the data from Fig. 7-1 available any time. Moreover, the logistician, higher-level leader, or authorized technician has access to the same data.

Appliance-level feedback helps provide awareness of human activity, saves energy, and reduces maintenance costs. The NILM system offers a unique, accurate, and inexpensive method to provide this feedback because it is built for the consumer.

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix A

Ft. Devens Load Disaggregation Code

```
"""    DEVENS 1  12APR14
  Testing Live Functionality
"""
import collections


def filter(data, interval, args, insert_func, state):
  """--[Auto-Generated: Do Not Remove or Modify]
    data is a python array where each element
    is a numpy array of timestamped values
    (timestamps are 64 bit microseconds)

    Access data in input array with these names:
            A_P1:  _i_A_P1
            A_Q1:  _i_A_Q1
            A_P3:  _i_A_P3
            A_Q3:  _i_A_Q3
            A_P5:  _i_A_P5
```

A_Q5: _i_A_Q5

A_P7: _i_A_P7

A_Q7: _i_A_Q7

B_P1: _i_B_P1

B_Q1: _i_B_Q1

B_P3: _i_B_P3

B_Q3: _i_B_Q3

B_P5: _i_B_P5

B_Q5: _i_B_Q5

B_P7: _i_B_P7

B_Q7: _i_B_Q7

C_P1: _i_C_P1

C_Q1: _i_C_Q1

C_P3: _i_C_P3

C_Q3: _i_C_Q3

C_P5: _i_C_P5

C_Q5: _i_C_Q5

C_P7: _i_C_P7

C_Q7: _i_C_Q7

Access data in output array with these names:

fan_on: _o_fan_on

fan_off: _o_fan_off

heater_on: _o_heater_on

heater_off: _o_heater_off

space_on: _o_space_on

space_off: _o_space_off

pump_on: _o_pump_on

pump_off: _o_pump_off

fridge_on: _o_fridge_on

fridge_off: _o_fridge_off

```
            reefer_on: _o_reefer_on
            reefer_off: _o_reefer_off
            hot_water_on: _o_hot_water_on
            hot_water_off: _o_hot_water_off
            compressor_on: _o_compressor_on
            compressor_off: _o_compressor_off
    __"""
    #TODO <your code here>

# from Manager, the columns are:
# [0] Unix Time stamp

#PA
# [1] P1
# [2] Q1
# [3] P3
# [4] Q3
# [5] P5
# [6] Q7
# [7] P7
# [8] Q7

# PB
# [9] P1
# [10] Q1
# [11] P3
# [12] Q3
# [13] P5
# [14] Q5
# [15] P7
```

```
# [16] Q7

#PC
# [17] P1
# [18] Q1
# [19] P3
# [20] Q3
# [21] P5
# [22] Q5
# [23] P7
# [24] Q7


##########################NOTES
#none

    data_reshape=np.concatenate((data[_i_A_P1],data[_i_A_Q1][:,[1]],
                                 data[_i_A_P3][:,[1]],data[_i_A_Q3][:,[1]],
                                 data[_i_A_P5][:,[1]],data[_i_A_Q5][:,[1]],
                                 data[_i_A_P7][:,[1]],data[_i_A_Q7][:,[1]],
                                 data[_i_B_P1][:,[1]],data[_i_B_Q1][:,[1]],
                                 data[_i_B_P3][:,[1]],data[_i_B_Q3][:,[1]],
                                 data[_i_B_P5][:,[1]],data[_i_B_Q5][:,[1]],
                                 data[_i_B_P7][:,[1]],data[_i_B_Q7][:,[1]],
                                 data[_i_C_P1][:,[1]],data[_i_C_Q1][:,[1]],
                                 data[_i_C_P3][:,[1]],data[_i_C_Q3][:,[1]],
                                 data[_i_C_P5][:,[1]],data[_i_C_Q5][:,[1]],
                                 data[_i_C_P7][:,[1]],data[_i_C_Q7][:,[1]]),
                                 axis=1)
    data=data_reshape
```

```
print np.shape(data)
new_array = state.retrieveSlot('prev_values')
i=1
suppress_lines = 0
num_events_TOTAL = 0      # counter for total events
num_events_ON = 0         # counter for ON events

## Initial Conditions
pumpA_toggle = False
fan_count = 0
heater_count = 0
space_count = 0
pump_count = 0
fridge_count = 0
reefer_count = 0
hot_water_count = 0
compressor_count = 0
event_time = 0

# other variables for tracking things (obsolete?)
line_counter = 0
deriv_counter   0


# definition of AVERAGE; we care about the first value,
how many to average, and which column from prep data
def average(start, how_many, column):
  x = 1;
  sub_total = 0;
```

```
    while x <= how_many:
        point = new_array[start][column]
        sub_total = sub_total + point
        start += 1
        x += 1

    total = sub_total/how_many
    return total

# definition of DERIVATIVE; care about how many samples
  to take derivative of and which column from prep data
def derivative(how_many, column):
    x = 1;
    sub_total = 0;
    start = 34
    result = 0

    while x<= how_many:
        deriv = new_array[start][column] - new_array[start - 1][column]
        x += 1
        start += 1
        sub_total = sub_total + deriv

    result = sub_total/how_many
    return result

# definition of event maximum (over 10 samples)
def maxi(column):
    maxi = max(new_array[30][column], new_array[31][column],
               new_array[32][column],new_array[33][column],
```

```
                  new_array[34][column], new_array[35][column],
                  new_array[36][column],new_array[37][column],
                  new_array[38][column],new_array[39][column],
                  new_array[40][column])
      return maxi

   # definition of event minimum (over 10 samples)
   def mini(column):
      mini = min(new_array[30][column], new_array[31][column],
                  new_array[32][column],new_array[33][column],
                  new_array[34][column], new_array[35][column],
                  new_array[36][column],new_array[37][column],
                  new_array[38][column],new_array[39][column],
                  new_array[40][column])
      return mini

##   ## open the file
##   phaseX = open(X,'r')

   #####################################################Start the Loop
   ## THE BEGINNING

   for line in data:
      new_array.append(line)
      i += 1

      ## if array is not long enough, go back to the beginning
      if len(new_array) < 120:
         continue
```

```
fan_on = 0
fan_off = 0
heater_on = 0
heater_off = 0
space_on = 0
space_off = 0
pump_on = 0
pump_off = 0
fridge_on = 0
fridge_off = 0
reefer_on = 0
reefer_off = 0
hot_water_on = 0
hot_water_off = 0
compressor_on = 0
compressor_off = 0

A_delta_kW = 0
B_delta_kW = 0
C_delta_kW = 0

A_peak_P1 = 0
B_peak_P1 = 0
C_peak_P1 = 0

## if you are suppressing after an event, go back to the beginning
if suppress_lines > 0:
  suppress_lines = suppress_lines - 1
  continue                    # means nothing below will execute
```

```
time = new_array[33][0] # unix time
######################################
# Code not working ??????
# Make sure resample is turned off!!!!!!!!
######################################

## calculate the numbers you need to determine whether
    event happened;

## derivative of each column (first difference);
    see definition above
A_deriv = derivative(1,1)
A_Qderiv = derivative(1,2)
A_deriv3 = derivative(1,3)
A_Qderiv3 = derivative(1,4)
A_deriv5 = derivative(1,5)
A_Qderiv5 = derivative(1,6)
A_deriv7 = derivative(1,7)
A_Qderiv7 = derivative(1,8)


## averages of power...
# before potential event (P1, P3, P5, P7, Q1, Q3, Q5, Q7)
A_prev_power = average(29,2,1)
A_prev_power3 = average(29,2,3)
A_prev_power5 = average(29,2,5)
A_prev_power7 = average(29,2,7)
A_prev_reactive_power = average(29,2,2)
A_prev_reactive_power3 = average(29,2,4)
A_prev_reactive_power5 = average(29,2,6)
```

```
A_prev_reactive_power7 = average(29,2,8)

# after potential event
# P1
A_on_new_power = average(108,12,1)
A_on_short_power = average(38,5,1)
A_on_medium_power = average(55,5,1)
A_delta_kW_on = A_on_new_power - A_prev_power
A_delta_kW_short_on = A_on_short_power - A_prev_power
A_delta_kW_medium_on = A_on_medium_power - A_prev_power

A_off_new_power = average(35,2,1)  # start sooner after off event
    for average
A_off_short_power = average(38,5,1)
A_off_medium_power = average(55,5,1)
A_delta_kW_off = A_off_new_power - A_prev_power
A_delta_kW_short_off = A_off_short_power - A_prev_power
A_delta_kW_medium_off = A_off_medium_power - A_prev_power

# P3
A_on_new_power3 = average(108,12,3)
A_off_new_power3 = average(35,2,3)  # start sooner after off event
    for average
A_delta_kW3_on = A_on_new_power3 - A_prev_power3
A_delta_kW3_off = A_off_new_power3 - A_prev_power3

# P5
A_on_new_power5 = average(108,12,5)
A_off_new_power5 = average(35,2,5)  # start sooner after off event
    for average
```

```
A_delta_kW5_on = A_on_new_power5 - A_prev_power5
A_delta_kW5_off = A_off_new_power5 - A_prev_power5

# P7
A_on_new_power7 = average(108,12,7)
A_off_new_power7 = average(35,2,7)  # start sooner after off event
    for average
A_delta_kW7_on = A_on_new_power7 - A_prev_power7
A_delta_kW7_off = A_off_new_power7 - A_prev_power7

# Q1
A_on_new_reactive_power = average(108,12,2)
A_off_new_reactive_power = average(35,2,2)
A_delta_kVAR_on = A_on_new_reactive_power - A_prev_reactive_power
A_delta_kVAR_off = A_off_new_reactive_power - A_prev_reactive_power

# Q3
A_on_new_reactive_power3 = average(108,12,4)
A_off_new_reactive_power3 = average(35,2,4)
A_delta_kVAR3_on = A_on_new_reactive_power3
    - A_prev_reactive_power3
A_delta_kVAR3_off = A_off_new_reactive_power3
    - A_prev_reactive_power3

# Q5
A_on_new_reactive_power5 = average(108,12,6)
A_off_new_reactive_power5 = average(35,2,6)
A_delta_kVAR5_on = A_on_new_reactive_power5
    - A_prev_reactive_power5
A_delta_kVAR5_off = A_off_new_reactive_power5
```

```
    - A_prev_reactive_power5

# Q7
A_on_new_reactive_power7 = average(108,12,8)
A_off_new_reactive_power7 = average(35,2,8)
A_delta_kVAR7_on = A_on_new_reactive_power7
    - A_prev_reactive_power7
A_delta_kVAR7_off = A_off_new_reactive_power7
    - A_prev_reactive_power7

## derivative of each column (first difference);
   see definition above
B_deriv = derivative(1,9)
B_Qderiv = derivative(1,10)
B_deriv3 = derivative(1,11)
B_Qderiv3 = derivative(1,12)
B_deriv5 = derivative(1,13)
B_Qderiv5 = derivative(1,14)
B_deriv7 = derivative(1,15)
B_Qderiv7 = derivative(1,16)

## averages of power...
# before potential event (P1, P3, P5, P7, Q1, Q3, Q5, Q7)
B_prev_power = average(29,2,9)
B_prev_power3 = average(29,2,11)
B_prev_power5 = average(29,2,13)
B_prev_power7 = average(29,2,15)
B_prev_reactive_power = average(29,2,10)
B_prev_reactive_power3 = average(29,2,12)
B_prev_reactive_power5 = average(29,2,14)
```

```
B_prev_reactive_power7 = average(29,2,16)

# after potential event
# P1
B_on_new_power = average(108,12,9)
B_on_short_power = average(38,5,9)
B_on_medium_power = average(55,5,9)
B_delta_kW_on = B_on_new_power - B_prev_power
B_delta_kW_short_on = B_on_short_power - B_prev_power
B_delta_kW_medium_on = B_on_medium_power - B_prev_power

B_off_new_power = average(35,2,9)  # start sooner after off event
    for average
B_off_short_power = average(38,5,9)
B_off_medium_power = average(55,5,9)
B_delta_kW_off = B_off_new_power - B_prev_power
B_delta_kW_short_off = B_off_short_power - B_prev_power
B_delta_kW_medium_off = B_off_medium_power - B_prev_power

# P3
B_on_new_power3 = average(108,12,11)
B_off_new_power3 = average(35,2,11)  # start sooner after off event
    for average
B_delta_kW3_on = B_on_new_power3 - B_prev_power3
B_delta_kW3_off = B_off_new_power3 - B_prev_power3

# P5
B_on_new_power5 = average(108,12,13)
B_off_new_power5 = average(35,2,13)  # start sooner after off event
    for average
```

```
B_delta_kW5_on = B_on_new_power5 - B_prev_power5
B_delta_kW5_off = B_off_new_power5 - B_prev_power5

# P7
B_on_new_power7 = average(108,12,15)
B_off_new_power7 = average(35,2,15)  # start sooner after off event
    for average
B_delta_kW7_on = B_on_new_power7 - B_prev_power7
B_delta_kW7_off = B_off_new_power7 - B_prev_power7

# Q1
B_on_new_reactive_power = average(108,12,10)
B_off_new_reactive_power = average(35,2,10)
B_delta_kVAR_on = B_on_new_reactive_power - B_prev_reactive_power
B_delta_kVAR_off = B_off_new_reactive_power - B_prev_reactive_power

# Q3
B_on_new_reactive_power3 = average(108,12,12)
B_off_new_reactive_power3 = average(35,2,12)
B_delta_kVAR3_on = B_on_new_reactive_power3
    - B_prev_reactive_power3
B_delta_kVAR3_off = B_off_new_reactive_power3
    - B_prev_reactive_power3

# Q5
B_on_new_reactive_power5 = average(108,12,14)
B_off_new_reactive_power5 = average(35,2,14)
B_delta_kVAR5_on = B_on_new_reactive_power5
    - B_prev_reactive_power5
B_delta_kVAR5_off = B_off_new_reactive_power5
```

```
    - B_prev_reactive_power5

# Q7
B_on_new_reactive_power7 = average(108,12,16)
B_off_new_reactive_power7 = average(35,2,16)
B_delta_kVAR7_on = B_on_new_reactive_power7
    - B_prev_reactive_power7
B_delta_kVAR7_off = B_off_new_reactive_power7
    - B_prev_reactive_power7

## derivative of each column (first difference); see definition above
C_deriv = derivative(1,17)
C_Qderiv = derivative(1,18)
C_deriv3 = derivative(1,19)
C_Qderiv3 = derivative(1,20)
C_deriv5 = derivative(1,21)
C_Qderiv5 = derivative(1,22)
C_deriv7 = derivative(1,23)
C_Qderiv7 = derivative(1,24)


## averages of power...
# before potential event (P1, P3, P5, P7, Q1, Q3, Q5, Q7)
C_prev_power = average(29,2,17)
C_prev_power3 = average(29,2,19)
C_prev_power5 = average(29,2,21)
C_prev_power7 = average(29,2,23)
C_prev_reactive_power = average(29,2,18)
C_prev_reactive_power3 = average(29,2,20)
C_prev_reactive_power5 = average(29,2,22)
```

```
C_prev_reactive_power7 = average(29,2,24)

# after potential event
# P1
C_on_new_power = average(108,12,17)
C_on_short_power = average(38,5,17)
C_on_medium_power = average(55,5,17)
C_delta_kW_on = C_on_new_power - C_prev_power
C_delta_kW_short_on = C_on_short_power - C_prev_power
C_delta_kW_medium_on = C_on_medium_power - C_prev_power

C_off_new_power = average(35,2,17)  # start sooner after off event
     for average
C_off_short_power = average(38,5,17)
C_off_medium_power = average(55,5,17)
C_delta_kW_off = C_off_new_power - C_prev_power
C_delta_kW_short_off = C_off_short_power - C_prev_power
C_delta_kW_medium_off = C_off_medium_power - C_prev_power

# P3
C_on_new_power3 = average(108,12,19)
C_off_new_power3 = average(35,2,19)  # start sooner after off event
     for average
C_delta_kW3_on = C_on_new_power3 - C_prev_power3
C_delta_kW3_off = C_off_new_power3 - C_prev_power3

# P5
C_on_new_power5 = average(108,12,21)
C_off_new_power5 = average(35,2,21)  # start sooner after off event
     for average
```

```
C_delta_kW5_on = C_on_new_power5 - C_prev_power5
C_delta_kW5_off = C_off_new_power5 - C_prev_power5

# P7
C_on_new_power7 = average(108,12,23)
C_off_new_power7 = average(35,2,23)  # start sooner after off event
    for average
C_delta_kW7_on = C_on_new_power7 - C_prev_power7
C_delta_kW7_off = C_off_new_power7 - C_prev_power7

# Q1
C_on_new_reactive_power = average(108,12,18)
C_off_new_reactive_power = average(35,2,18)
C_delta_kVAR_on = C_on_new_reactive_power - C_prev_reactive_power
C_delta_kVAR_off = C_off_new_reactive_power - C_prev_reactive_power

# Q3
C_on_new_reactive_power3 = average(108,12,20)
C_off_new_reactive_power3 = average(35,2,20)
C_delta_kVAR3_on = C_on_new_reactive_power3
    - C_prev_reactive_power3
C_delta_kVAR3_off = C_off_new_reactive_power3
    - C_prev_reactive_power3

# Q5
C_on_new_reactive_power5 = average(108,12,22)
C_off_new_reactive_power5 = average(35,2,22)
C_delta_kVAR5_on = C_on_new_reactive_power5
    - C_prev_reactive_power5
C_delta_kVAR5_off = C_off_new_reactive_power5
```

```
        - C_prev_reactive_power5

  # Q7
  C_on_new_reactive_power7 = average(108,12,24)
  C_off_new_reactive_power7 = average(35,2,24)
  C_delta_kVAR7_on = C_on_new_reactive_power7
        - C_prev_reactive_power7
  C_delta_kVAR7_off = C_off_new_reactive_power7
        - C_prev_reactive_power7


# EDGE DETECTION ALGORITHM
#PA
  ##############################################################
  # CHECK 4A: Is there an ON event?
  line_counter += 1

  if A_deriv > 150 and A_delta_kW_short_on > 150:

    time_ON = timestamp_to_human(time)
    event_time = time_ON

    # FOR SMALL EVENTS
    if A_delta_kW_short_on < 500:
      A_delta_kW = A_delta_kW_short_on
      suppress_lines = 10

    # FOR QUICK EVENTS THAT COME BACK DOWN
    elif A_delta_kW_on < 150:          # no change in long
      A_delta_kW = A_delta_kW_short_on
```

```
      suppress_lines = 5

    # FOR MOST COMMON EVENTS
    else:
      A_delta_kW = A_delta_kW_on
      suppress_lines = 10

  ##################################################
  # CHECK 4B: Is there an OFF event?

  elif A_deriv < -100 and A_delta_kW_off < -100:
    A_delta_kW = A_delta_kW_off
    time_OFF = timestamp_to_human(time)
    event_time = time_OFF
    suppress_lines = 3


#PB
  ##################################################
  # CHECK 4A: Is there an ON event?
  elif B_deriv > 150 and B_delta_kW_short_on > 150:

    time_ON = timestamp_to_human(time)
    event_time = time_ON

    # FOR SMALL EVENTS
    if B_delta_kW_short_on < 500:
      B_delta_kW = B_delta_kW_short_on
      suppress_lines = 10
```

```
  # FOR QUICK EVENTS THAT COME BACK DOWN
  elif B_delta_kW_on < 150:         # no change in long
    B_delta_kW = B_delta_kW_short_on
    suppress_lines = 5

  # FOR MOST COMMON EVENTS
  else:
    B_delta_kW = B_delta_kW_on
    suppress_lines = 10

##################################################
# CHECK 4B: Is there an OFF event?

elif B_deriv < -100 and B_delta_kW_off < -100:
  B_delta_kW = B_delta_kW_off
  time_OFF = timestamp_to_human(time)
  event_time = time_OFF
  suppress_lines = 3

  #PC
##################################################
# CHECK 4A: Is there an ON event?
elif C_deriv > 150 and C_delta_kW_short_on > 150:

  time_ON = timestamp_to_human(time)
  event_time = time_ON

  # FOR SMALL EVENTS
  if C_delta_kW_short_on < 500:
    C_delta_kW = C_delta_kW_short_on
```

```
      suppress_lines = 10

    # FOR QUICK EVENTS THAT COME BACK DOWN
    elif C_delta_kW_on < 150:          # no change in long
      C_delta_kW = C_delta_kW_short_on
      suppress_lines = 5

    # FOR MOST COMMON EVENTS
    else:
      C_delta_kW = C_delta_kW_on
      suppress_lines = 10

  ##########################################################
  # CHECK 4B: Is there an OFF event?

  elif C_deriv < -100 and C_delta_kW_off < -100:
    C_delta_kW = C_delta_kW_off
    time_OFF = timestamp_to_human(time)
    event_time = time_OFF
    suppress_lines = 3

  ################ NO EVENT, go back to beginning
  else:
    continue

##############################################################################################

# IF there WAS an event....

  A_maxP1 = maxi(1)
```

```
A_maxP3 = maxi(3)
A_maxP5 = maxi(5)
A_maxP7 = maxi(7)
A_maxQ1 = maxi(2)
A_maxQ3 = maxi(4)
A_maxQ5 = maxi(6)
A_maxQ7 = maxi(8)
A_minP1 = mini(1)
A_minP3 = mini(3)
A_minP5 = mini(5)
A_minP7 = mini(7)
A_minQ1 = mini(2)
A_minQ3 = mini(4)
A_minQ5 = mini(6)
A_minQ7 = mini(8)

A_peak_P1 = A_maxP1 - A_minP1
A_peak_P3 = A_maxP3 - A_minP3
A_peak_P5 = A_maxP5 - A_minP5
A_peak_P7 = A_maxP7 - A_minP7
A_peak_Q1 = A_maxQ1 - A_minQ1
A_peak_Q3 = A_maxQ3 - A_minQ3
A_peak_Q5 = A_maxQ5 - A_minQ5
A_peak_Q7 = A_maxQ7 - A_minQ7

B_maxP1 = maxi(9)
B_maxP3 = maxi(11)
B_maxP5 = maxi(13)
B_maxP7 = maxi(15)
B_maxQ1 = maxi(10)
```

```
B_maxQ3 = maxi(12)
B_maxQ5 = maxi(14)
B_maxQ7 = maxi(16)
B_minP1 = mini(9)
B_minP3 = mini(11)
B_minP5 = mini(13)
B_minP7 = mini(15)
B_minQ1 = mini(10)
B_minQ3 = mini(12)
B_minQ5 = mini(14)
B_minQ7 = mini(16)

B_peak_P1 = B_maxP1 - B_minP1
B_peak_P3 = B_maxP3 - B_minP3
B_peak_P5 = B_maxP5 - B_minP5
B_peak_P7 = B_maxP7 - B_minP7
B_peak_Q1 = B_maxQ1 - B_minQ1
B_peak_Q3 = B_maxQ3 - B_minQ3
B_peak_Q5 = B_maxQ5 - B_minQ5
B_peak_Q7 = B_maxQ7 - B_minQ7

C_maxP1 = maxi(17)
C_maxP3 = maxi(19)
C_maxP5 = maxi(21)
C_maxP7 = maxi(23)
C_maxQ1 = maxi(18)
C_maxQ3 = maxi(20)
C_maxQ5 = maxi(22)
C_maxQ7 = maxi(24)
C_minP1 = mini(17)
```

```
C_minP3 = mini(19)
C_minP5 = mini(21)
C_minP7 = mini(23)
C_minQ1 = mini(18)
C_minQ3 = mini(20)
C_minQ5 = mini(22)
C_minQ7 = mini(24)

C_peak_P1 = C_maxP1 - C_minP1
C_peak_P3 = C_maxP3 - C_minP3
C_peak_P5 = C_maxP5 - C_minP5
C_peak_P7 = C_maxP7 - C_minP7
C_peak_Q1 = C_maxQ1 - C_minQ1
C_peak_Q3 = C_maxQ3 - C_minQ3
C_peak_Q5 = C_maxQ5 - C_minQ5
C_peak_Q7 = C_maxQ7 - C_minQ7

##############################################################
# LOAD CLASSIFICATION ALGORITHM

# CHECK 1: Fan ON Event?

if 500 < A_delta_kW < 700 and  -600 < A_delta_kVAR_on < -250 and
        -50 < A_delta_kW3_on < 35 and 3700 < A_peak_P1 < 5100:
    fan_on = 1
    suppress_lines = 60



# CHECK 2: Fan OFF Event?
```

```
elif -450 > A_delta_kW > -700 and 200 < A_delta_kVAR_off < 600 and
     -35 < A_delta_kW3_off < 35 and A_delta_kVAR7_off < -3 and
     pumpA_toggle  == False and A_peak_P1 < 800:
  fan_off = 1


##################################################
# CHECK 3: Heater ON Event?

## double on
elif 5800 < A_delta_kW <7000:
  heater_on = 1
  what_happened = "Double Heater_ON"
###adjusted by John
  if(time<interval.start):
    time=interval.start+5
    suppress_lines = 120
  output = [time-1, fan_on, fan_off, heater_on, heater_off, space_on,
            space_off, pump_on, pump_off, fridge_on, fridge_off,
            reefer_on, reefer_off, hot_water_on, hot_water_off,
            compressor_on, compressor_off, event_time,
            what_happened, A_delta_kW,B_delta_kW,
            C_delta_kW]
  insert_func([output[:17]])
####
  print '\t'.join(["%s"%x for x in output])


# single on
elif 2000 < A_delta_kW < 5800 and A_peak_P1 > 2900:
```

```
        heater_on = 1
    elif 2900 < B_delta_kW < 3700 and B_peak_P1 > 2900:
        heater_on = 1

    # CHECK 4: Heater OFF Event?

    elif -2000 > A_delta_kW > -3900 and A_peak_P1 > 2800:
        heater_off = 1
    elif -2000 > B_delta_kW > -3900 and B_peak_P1 > 2800:
        heater_off = 1


    ###########################################################
    # CHECK 5: Space ON Event?

    # PA
    elif 1200 < A_delta_kW < 1600 and A_peak_P1 < 1600 and
            C_delta_kW_on < 200:
          space_on = 1

    # PB
    elif 900 < B_delta_kW < 1600 and B_peak_P1 < 1800 and
            C_delta_kW_on < 200:
        space_on = 1

    # CHECK 6: Space OFF Event?

    # PA
    elif -1600 < A_delta_kW < -900 and pumpA_toggle == False and
            C_peak_P1 < 500:
```

```
    space_off    1



  # PB
  elif  -1600 < B_delta_kW < -900 and  pumpA_toggle == False and
        C_peak_P1 < 500:
    space_off = 1

  ############################################################
  # CHECK 6a: Hot Water ON Event?

  # PA
  elif 300 < A_delta_kW < 450 and 2600 < A_peak_P1 < 3400:
    hot_water_on = 1
    suppress_lines = 60

  # CHECK 6b: Hot Water OFF Event?

  # PA
  elif -300 > A_delta_kW > -450 and -300 < A_delta_kVAR_off < 800
        and 0 > A_delta_kVAR3_off > -70 and C_peak_P1 < 100:
    hot_water_off = 1

  ############################################################
  # CHECK 7: Fridge ON Event?

  elif 450 < C_delta_kW < 700 and -800 < C_delta_kVAR_on < -100
        and 2800 < C_peak_P1 < 3400:
    fridge_on - 1
    suppress_lines    60
```

```
# CHECK 8: Fridge OFF Event?

elif -450 > C_delta_kW > -600 and -100 < C_delta_kVAR_off < 800
    and C_peak_P1 < 400:
  fridge_off = 1

##############################################################
# CHECK 7: Reefer ON Event?

    # single on
elif 1000 < A_delta_kW < 1600 and A_delta_kVAR3_on > 100:
  reefer_on = 1
elif 1000 < B_delta_kW < 1600 and B_delta_kVAR3_on > 100:
  reefer_on = 1
elif 1000 < C_delta_kW < 1600 and C_delta_kVAR3_on > 100:
  reefer_on = 1

# CHECK 8: Reefer OFF Event?

elif -1100 > A_delta_kW > -1200 and A_delta_kVAR_off > 400:
  heater_off = 1
elif -1100 > B_delta_kW > -1200 and A_delta_kVAR_off > 400:
  heater_off = 1
elif -1100 > C_delta_kW > -1200 and C_delta_kVAR_off > 400:
  heater_off = 1

##############################################################
# CHECK 9: Pump ON Event?
```

```
elif 300 < A_delta_kW < 1700 and -200 > A_delta_kVAR_on > -1000
        and A_peak_P1 > 4500:
    pump_on = 1
    pumpA_toggle = True
elif 300 < B_delta_kW < 1700 and -200 > B_delta_kVAR_on > -1000
        and B_peak_P1 > 4500:
    pump_on = 1
    pumpA_toggle = True
elif 300 < C_delta_kW < 1700 and -200 > C_delta_kVAR_on > -1000
        and C_peak_P1 > 4500:
    pump_on = 1
    pumpA_toggle = True

    # CHECK 10: Pump OFF Event?

elif -400 > A_delta_kW > -1600 and -100 < A_delta_kVAR_off < 1000
        and pumpA_toggle == True:
    pump_off = 1
    pumpA_toggle = False
elif -400 > B_delta_kW > -1600 and -100 < B_delta_kVAR_off < 1000
        and pumpA_toggle == True:
    pump_off = 1
    pumpA_toggle = False
elif -400 > C_delta_kW > -1600 and -100 < C_delta_kVAR_off < 1000
        and pumpA_toggle == True:
    pump_off = 1
    pumpA_toggle = False

    ################################################################
# CHECK 11: Compressor ON Event?
```

```
elif 1000 < A_peak_P1 < 1100 and 700 < B_peak_P1 < 800 and
        800 < C_peak_P1 < 900:
    compressor_on = 1
    suppress_lines = 900

# CHECK 12: Compressor OFF Event?

elif -500 > A_delta_kW > -800 and 800 < A_delta_kVAR_off < 1100
        and pumpA_toggle == False:
    compressor_off = 1

##########################################################
## If event matches nothing above

else:
    if 1200 < A_peak_P1 < 1600 or 900 < B_peak_P1 < 1600:
        if A_delta_kW < 0 or B_delta_kW < 0:
            space_off = 1
        else:
            space_on = 1
    else:
        print "Unknown Event", timestamp_to_human(time), A_delta_kW,
                B_delta_kW, C_delta_kW, A_delta_kVAR3_on
    continue

if fan_on == 1:
    what_happened = "Fan_ON"
    fan_count += 1
elif fan_off == 1:
```

```
    what_happened = "Fan_OFF"
    fan_count -= 1
elif heater_on == 1:
    what_happened = "Heater_ON"
    heater_count += 1
elif heater_off == 1:
    what_happened = "Heater_OFF"
    heater_count -= 1
elif space_on == 1:
    what_happened = "Space_ON"
    space_count += 1
elif space_off == 1:
    what_happened = "Space_OFF"
    space_count -= 1
elif pump_on == 1:
    what_happened = "Pump_ON"
    pump_count += 1
elif pump_off == 1:
    what_happened = "Pump_OFF"
    pump_count -= 1
elif fridge_on == 1:
    what_happened = "Fridge_ON"
    fridge_count += 1
elif fridge_off == 1:
    what_happened = "Fridge_OFF"
    fridge_count -= 1
elif reefer_on == 1:
    what_happened = "Reefer_ON"
    reefer_count += 1
elif reefer_off == 1:
```

```
    what_happened = "Reefer_OFF"
    reefer_count -= 1
  elif hot_water_on == 1:
    what_happened = "Hot_Water_ON"
    hot_water_count += 1
  elif hot_water_off == 1:
    what_happened = "Hot_Water_OFF"
    hot_water_count -= 1
  elif compressor_on == 1:
    what_happened = "Compressor_ON"
    compressor_count += 1
  elif compressor_off == 1:
    what_happened = "Compressor_OFF"
    compressor_count -= 1
  else:
    print "column error, no ones"

  if(time<interval.start):
      time=interval.start+5
      suppress_lines = 120

  output = [time, fan_on, fan_off, heater_on, heater_off, space_on,
      space_off, pump_on, pump_off, fridge_on, fridge_off, reefer_on,
      reefer_off, hot_water_on, hot_water_off, compressor_on,
      compressor_off, event_time, what_happened, A_delta_kW,
      A_delta_kW3_ON, A_delta_kW5_ON, A_delta_kW7_ON,
      A_delta_kVAR, A_delta_kVAR3_on, A_delta_kVAR5_on,
      A_delta_kVAR7_on, B_delta_kW, C_delta_kW, A_peak_P1,
      A_peak_P3, A_peak_P5, A_peak_P7, A_peak_Q1, A_peak_Q3,
      A_peak_Q5, A_peak_Q7]
```

```
insert_func([output[:17]])

print '\t'.join(["%s"%x for x in output])
continue
```

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix B

DepNILM Assembly Instructions

DepNILM is the deployable version of the NILM sensor system. Fig. B-1 shows the design and itemized component list, and the instructions to put one together is below. Fig. B-3 is a list of parts, vendors, and approximate cost.

B.1 Hardware

B.1.1 Drill

1. In open space, lay out major components on plywood. Disconnect all NILM cables, and don't wire outlet boxes yet
    - NILM
    - UPS
    - Computer
    - 4" square outlet box
    - Two (2) single-gang outlet boxes
2. Determine hole placement for plywood; mark with pencil
    - NILM - 4 holes
    - UPS - 2 wood screws

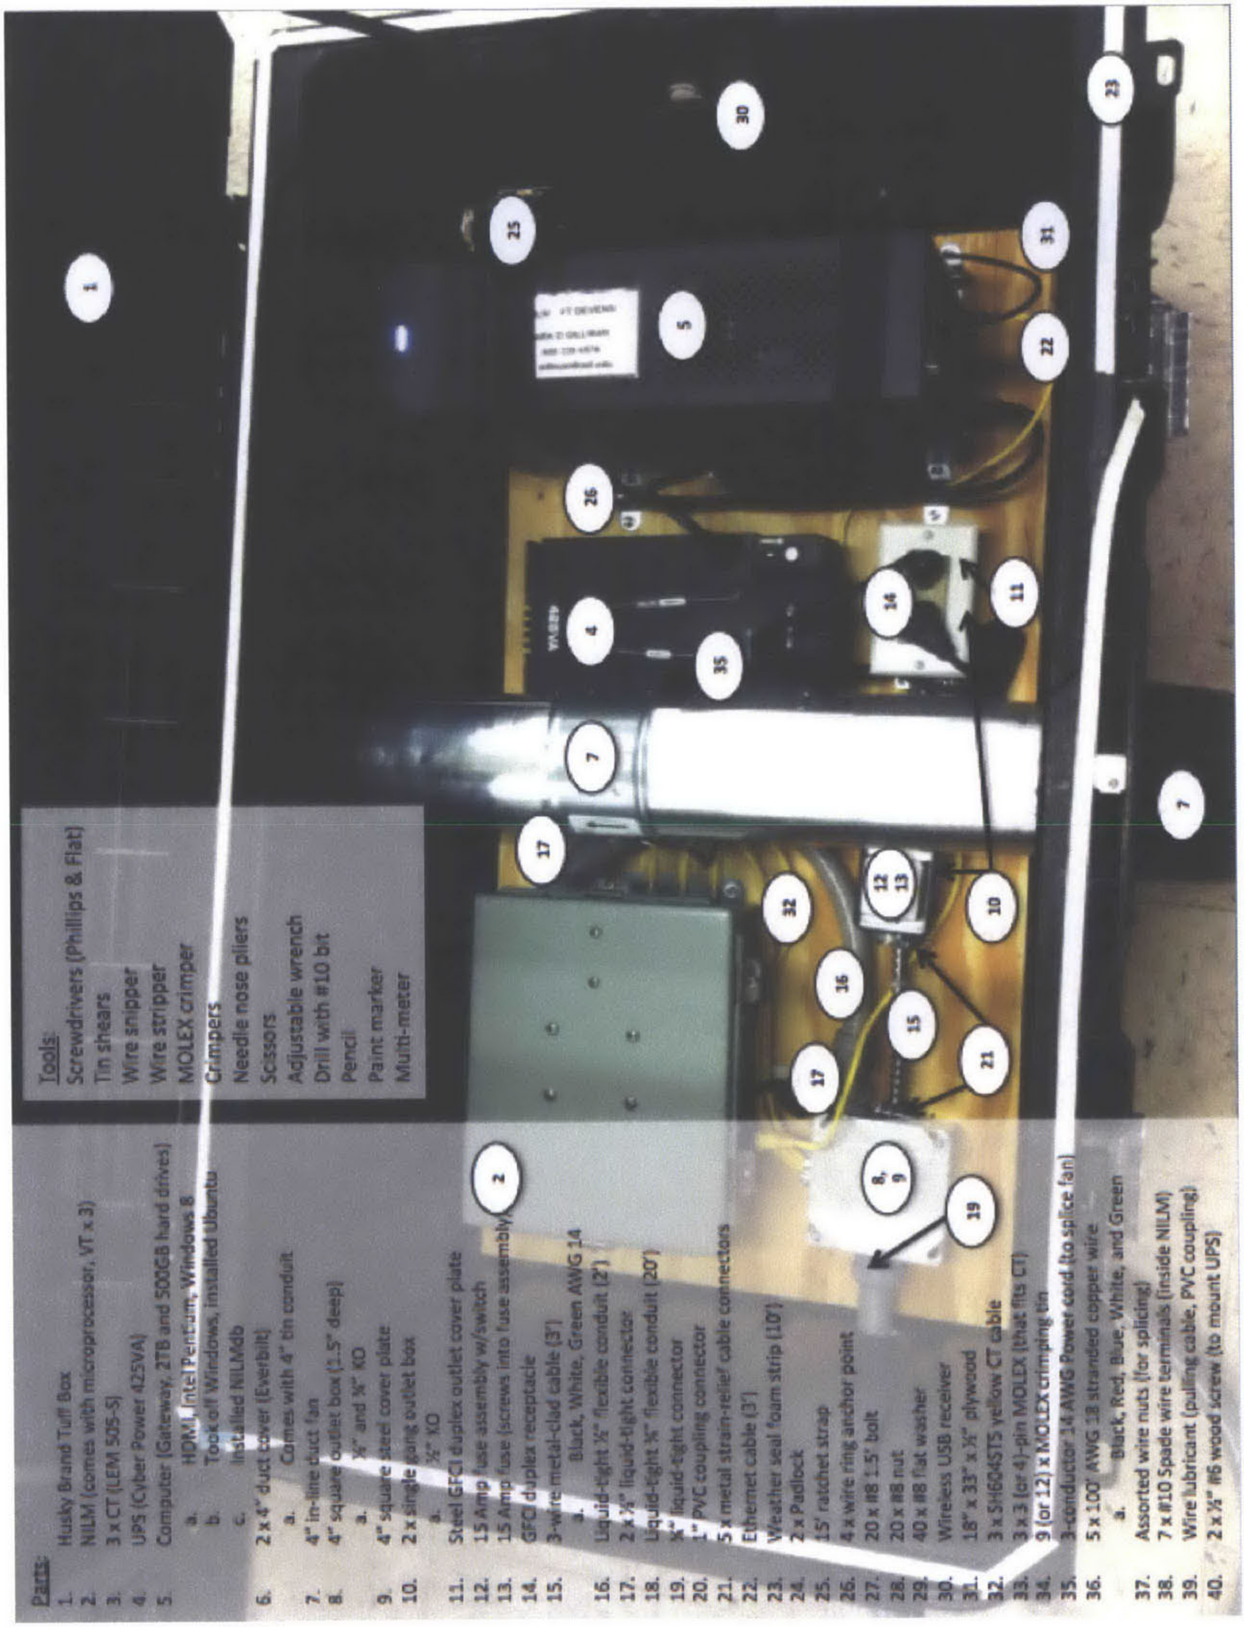

Figure B-1: DepNILM Assembly

- Computer - 8 holes for wire ring anchor points
- 4" box - 4 holes
- single-gang boxes - 2 holes each

3. Determine hole placement for tuff box

   - 2 x 4 $\frac{1}{4}$" holes (1 inlet, 1 outlet) for duct covers centered on front and back; make sure it is low enough so the duct cover can sit flush with the side of box
   - 1 x 1" hole for $\frac{3}{4}$" conduit to access 4" square box

4. Drill all holes

5. Bolt all outlet boxes to plywood using #8 screws, washers, and nuts

B.1.2 Wire (working backwards from GFCI)

1. GFCI

   - Cut section of metal clad cable to size (from GFCI to fuse/switch assembly); strip with pliers (pull to unravel) to expose 5" of conductors on either end; wire GFCI outlet according to manufacturer's instructions; continuity test on outlet will not work until you energize GFCI and press reset; affix metal cable to outlet box through bottom KO using strain reliever clamp; screw GFCI to outlet box

2. Fuse/Switch Assembly

   - Cut another section of metal clad cable to size (from fuse/switch to 4" box); expose 5" on both ends; wire fuse/switch assembly according to manufacturer's instructions; the ground screw should have both the incoming and outgoing green (ground) wires; orientation does not matter, but don't mix up the incoming and outgoing black (hot) wires; White (neutral) wires will be spliced with wire nut; affix incoming and outgoing metal cable to outlet box through top and bottom KO's using strain reliever clamps; screw fuse/switch assembly to outlet box; terminate loose end of metal clad cable in center hole of 4" square outlet box (facing fuse/switch) using strain relief clamp

3. Panel to Tuff box cable

   - Yellow cables: Inside each yellow cable are 4 colored conductors shielded in metal; crimp the appropriate ends and snip excess conductors; insert the crimped ends into the MOLEX connector appropriately. Use $\frac{3}{4}$" heat shrink to finish end. Do continuity checks. Lay cables out on the ground

   - Colored wires: Cut 5 sections of 18 AWG wire (one each of Black, Red, Blue, Green, and White); length should be 6" longer than yellow cables; use zip ties to combine entire bundle (8 strands total), and lay bundle flat

   - Cut $\frac{3}{4}$" liquid tight conduit to size; it will contain entire bundle; leave 18" of yellow cable exposed at both ends; feed all cables through conduit (use lubricant); panel side will have MOLEX connectors and extra 6" of colored wires, tuff box side gets the round connectors that plug into NILM box

   - Order is important! Thread conduit through outside PVC coupler; thread cable through 1" hole in tuff box; thread conduit through inside PVC coupler; thread conduit through $\frac{3}{4}$" liquid-tight connector; affix $\frac{3}{4}$" liquid-tight connector into appropriate KO in 4"box; screw PVC coupling together now

4. 4" outlet box

   - Bolt NILM to plywood ensuring all wire ports are accessible; cut $\frac{1}{2}$" liquid-tight flexible conduit to size (from NILM to 4" square box); cut 5 sections of 18 AWG wire (one each of Black, Red, Blue, Green, and White) and route through conduit into NILM and into 4" box; Black, Red, and Blue terminate at closest respective fuse ends using #10 Spade wire terminals; White and Green terminate in black screw terminals; leave 5" of each wire exposed in 4" box; next affix $\frac{1}{2}$" conduit to NILM and 4" box using $\frac{1}{2}$" liquid-tight connectors

   - There are now 3 sets of cables terminating inside the 4" box; route the yellow cables through appropriate KO towards NILM and plug them in, affixing them to 4" box with strain relief clamp; splice all like-colored ends together using wire nuts. There should be 3 blacks spliced together, 3 whites, 3 greens, 2 reds, and 2 blues

- Conduct continuity checks from panel end to NILM and from panel end to GFCI screw terminals (outlet ports not connected until GFCI is energized); remember to put fuse in and turn switch on; put cover on 4" box

5. Fan

   - Wire according to instructions

B.1.3 Final Assembly

1. Finish wiring the NILM correctly and close cover
2. Affix UPS to plywood
3. Bolt computer to plywood using ratchet strap routed through anchors
4. Put cover on GCFI
5. Cut holes in tin conduit to bring in fresh air and blow out box air; attach tin to duct covers, inset through 4" holes in box, and place fan in the middle (oriented appropriately)
6. Place weather seal around opening edge of tuff box
7. Seal duct covers to box
8. Plug everything in; fan and UPS in GFCI, Computer in UPS, wireless receiver into computer, ethernet into NILM and computer

B.1.4 Testing

1. Use 3-phase HP 6834B (Agilent) power source; wire an appropriate 3-phase plug using the panel-side colored wires to plug into source; set voltage to 120V RMS, set current limit to 5 A, set frequency to 60 Hz; turn on

2. Ensure computer temperature does not exceed 80 ; ensure UPS is working; plug in CTs on loads of interest and ensure they are reading; calibrate results

B.2 Software

B.2.1 Computer

1. Follow the instructions in Jim Paris' Thesis [1] to setup Ubuntu and NilmDB

2. The disaggregation filters for Ft. Devens, Sharon, and Polk are available through the website `http://www.nilmdb.com`

B.2.2 Labjack

1. from the Labjack website and using a Windows computer, ensure the Labjack UE9 has version of the software (see Fig. B-2)

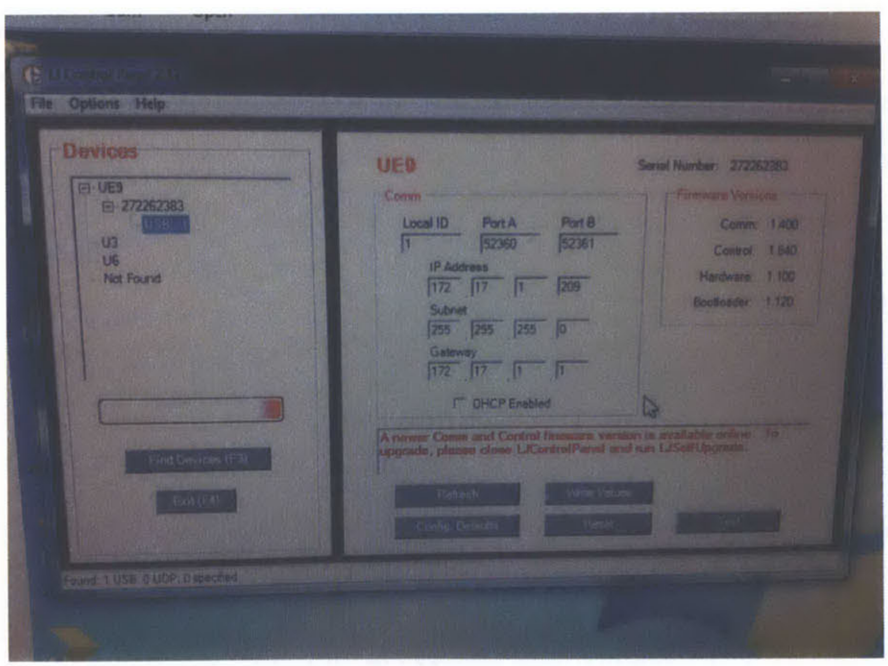

Figure B-2: Labjack Software Setup

Parts	Vendor	Price	Comments
1. Husky Brand Tuff Box	Home Depot		
2. NILM (comes with microprocessor, LEM LV-25P VT x 3)		[redacted]	Contact MIT
3. 3 x CT (LEM 505-S)	Digikey	$450	
4. UPS (Cyber Power 425VA)	Amazon	$50	
5. Computer (Dell Inspiron i660s-2313BK Desktop)	Amazon	$400	No refurbs, please
b. 2TB additional hard drive	Western Digital, Black Model	$200	
a. HDMI, Intel Pentium, Windows 8			
b. Take off Windows, install Ubuntu		[redacted]	Contact MIT
c. Installed NILMdb			Contact MIT
d. May need hard drive bay insert	Amazon	$20	
6. 2 x 4" duct cover (Everbilt)	Home Depot		
a. Comes with 4" tin conduit			
7. 4" in-line duct fan	Home Depot		
8. 4" square outlet box (1.5" deep)	Home Depot		
a. ½" and ¾" KO			
9. 4" square steel cover plate	Home Depot		
10. 2 x single gang outlet box	Home Depot		
a. ½" KO			
11. Steel GFCI duplex outlet cover plate	Home Depot		
12. 15 Amp fuse assembly w/switch	Home Depot		
13. 15 Amp fuse (screws into fuse assembly)	Home Depot		
14. GFCI duplex receptacle	Home Depot		
15. 3-wire metal-clad cable (3')	Home Depot		
a. Black, White, Green AWG 14			
16. Liquid-tight ½" flexible conduit (2')	Home Depot		
17. 2 x ½" liquid-tight connector	Home Depot		
18. 2 x Liquid-tight ¾" flexible conduit (25')	Home Depot		
19. ¾" liquid-tight connector	Home Depot		
20. 2 x 1" PVC coupling connector	Home Depot		
21. 5 x metal strain-relief cable connectors	Home Depot		
22. 2 x Ethernet cable (3')	Home Depot		
23. Weather seal foam strip (10')	Home Depot		
24. 2 x Padlock	Home Depot		
25. 2 x15' ratchet strap	Home Depot		
26. 8 x wire ring anchor point	Home Depot		
27. 20 x #8 1.5' bolt	Home Depot		
28. 20 x #8 nut	Home Depot		
29. 40 x #8 flat washer	Home Depot		
30. Wireless bridge		[redacted]	Unique to each install
31. 18" x 33" x ½" plywood	Home Depot		
32. 6 x SH6045TS yellow CT cable	Digikey	$50	
33. 6 x 3 (or 4)-pin MOLEX (that fits CT)	Digikey	$5	
34. 24 x MOLEX crimping tin	Digikey	$5	
35. 3-conductor 14 AWG Power cord (to splice fan)	Home Depot		
36. 5 x 100' AWG 18 stranded copper wire	Amazon/Home Depot		
a. Black, Red, Blue, White, and Green			
37. Assorted wire nuts (for splicing)	Home Depot		
38. 14 x #10 Spade wire terminals (inside NILM)	Digikey	$5	
39. Wire lubricant (pulling cable, PVC coupling)	Home Depot		
40. 2 x ½" #6 wood screw (to mount UPS)	Home Depot		
41. 20A Circuit Breakers (Square D, FY14020A, B, C), 1 each	Granite City	$700	
42. Ethernet to USB adapter	Amazon	$20	
43. Ethernet Cable (20')	Home Depot		
TOTAL Home Depot Cost Estimate		$1,000	
TOTAL Estimated Cost		$ 2,905.00	

Figure B-3: DepNILM Parts List

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix C

Points of Contact

C.1 Ft. Devens BCIL

Gary McMichael

Configuration Control Board Chair

gary.p.mcmichael.civ@mail.mil

Bill Singleton

Engineering Chief

william.m.singleton1.civ@mail.mil

C.2 Ft. Polk JRTC

LTC Joseph Brown

RDECOM

joseph.d.brown.mil@mail.mil

Robert Wilson

Operations Group Engineer

robert.a.wilson1.civ@mail.mil

Major Kirk Reed

Operations Officer, CSH

kirk.a.reed.mil@mail.mil

C.3 Cottage Elementary School

Scotty Schertz

Chief Maintenance Officer

scotty@sharon.k12.ma.us

Bibliography

[1] J. Paris. "A Comprehensive System for Non-Intrusive Load Monitoring and Diagnostics," Ph.D. Dissertation, Dept. Elect. Eng. and Comp. Sci., Massachusetts Institute of Technology, Cambridge, MA, 2013.

[2] J. Donnal. "Home NILM: A Comprehensive Non-Intrusive Load Monitoring Toolkit," M.S. Thesis, Dept. Elect. Eng. and Comp. Sci., Massachusetts Institute of Technology, Cambridge, MA, 2013.

[3] M. Gillman, J. Donnal, J. Paris, S. Leeb, M. El Sayed, K. Wertz, S. Schertz, "Energy Accountability Using Nonintrusive Load Monitoring." IEEE Sensors Journal, vol 14, no. 6, pp. 1923-1931, June 2014.

[4] H.Y. Iris Cheung, S. Lanzisera, J. Lai, R. Brown, S. Dawson-Haggerty, J. Taneja, D. Culler, "Detailed Energy Data Collection for Miscellaneous and Electronic Loads in a Commercial Office Building," presented at the 2012 ACEEE Summer Study on Energy Efficiency in Buildings, Asilomar, CA, Aug. 2012.

[5] (2013) The Power Standards website. [Online]. Available: `http://www.powerstandards.com/PQube.php`.

[6] `http://www.theenergydetective.com`

[7] `http://www.trendpoint.com/branch-circuit-monitoring/enersure/`

[8] `http://www.brandelectronics.com/meters.html`

[9] http://www.p3international.com/products/p4400.html

[10] C. Schantz. "Method for Nonintrusive Sensing and System Monitoring," Ph.D. Dissertation, Dept. Elect. Eng. and Comp. Sci., Massachusetts Institute of Technology, Cambridge, MA, 2014.

[11] S.M. Neuman, "Non-Intrusive Water Utility Monitoring and Free-Space Load Monitoring," Ph.D. dissertation, Dept. Elect. Eng. and Comp. Sci., Massachusetts Institute of Technology, Cambridge, MA, 2009.

[12] C. Schantz, J. Donnal, B. Sennett, M. Gillman, S. Muller, S. Leeb, "Water Non-intrusive Load Monitoring," submitted for publication in IEEE Sensors Journal as of June 2014.

[13] J. Paris, J. Donnal, Z. Remscrim, S. Leeb, and S. Shaw, "The Sinefit Spectral Envelope Preprocessor," submitted for publication in IEEE Sensors Journal, 2014.

[14] S. R. Shaw, "System identification techniques and modeling for nonintrusive load diagnostics," PhD, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, February 2000.

[15] S. Shaw, C. Laughman, "A Kalman-filter Spectral Envelope Preprocessor," IEEE Transactions on Instrumentation and Measurement, vol. 56, no. 5, pp. 2010-2017, October 2007.

[16] G.W. Hart, "Nonintrusive appliance load monitoring," Proceedings of the IEEE 80, no. 12, pp. 1870-1891, 1992.

[17] Farinaccio, R. Zmeureanu, "Using a pattern recognition approach to disaggregate the total electricity consumption in a house into the major end-uses," Energy and Buildings, vol. 30(3), pp. 245-259, 1999.

[18] K. Anderson, A. Ocneanu, D. Benitez, D. Carlson, A. Rowe, M. Bergés, "BLUED: a fully labeled public dataset for Event-Based Non-Intrusive load monitoring research," in Proceedings of the 2nd KDD Workshop on Data Mining Applications in Sustainability, Beijing, China, 2012, pp. 12-16.

[19] M.L. Marceau, R. Zmeureanu, "Nonintrusive load disaggregation computer program to estimate the energy consumption of major end uses in residential buildings," Energy Conversion and Management, vol 41, no. 13, pp.1389-1403, 2000.

[20] J.Z. Kolter, M.J. Johnson, "REDD: A Public Data Set for Energy Disaggregation Research," presented at the SustKDD, San Diego, CA, 2011.

[21] Z. Wang, G. Zheng, "Residential appliances identification and monitoring by a nonintrusive method." Smart Grid, IEEE Transactions on, vol 3, no. 1, pp. 80-92, 2012.

[22] Ming Dong; Meira, P.C.M.; Wilsun Xu; Chung, C.Y., "Non-Intrusive Signature Extraction for Major Residential Loads," Smart Grid, IEEE Transactions on , vol.4, no.3, pp.1421,1430, Sept. 2013.

[23] M. Figueiredo, A. De Almeida, B. Ribeiro, "Home electrical signal disaggregation for non-intrusive load monitoring (NILM) systems." Neurocomputing, vol. 96, pp. 66-73, 2012.

[24] C. Laughman, K. Lee, R. Cox, S. Shaw, S. Leeb, L. Norford, P. Armstrong, "Power signature analysis," IEEE Power and Energy Magazine, pp. 56-63, Mar. 2003.

[25] H. Friedman, A. Potter, T. Haasl, D. Claridge. (October 2003). "Persistence of Benefits from New Building Commissioning," California Energy Commission Technical Report. [Online]. Available: `http://www.energy.ca.gov/2003publications/CEC-500-2003-097/CEC-500-2003-097F-A18.PDF`.

[26] Gordon-Piatt Model R & S Burner Instruction Manual, version 1-RA-40.1, Manufactured by John Zink Company, LLC, May 2004.

[27] Report of the Defense Science Board Task Force on DoD Energy Strategy, "More Fight, Less Fuel", Office of the Under Secretary of Defense For Acquisition, Technology, and Logistics, Washington, D.C. 20301-3140, February 2008.

[28] Schwartz, M., K. Blakeley, R. O'Rourke, "Department of Defense Energy Initiatives: Background and Issues for Congress," Congressional Research Service, December 10, 2012.

[29] D. Robyn, Deputy Under Secretary of Defense for Installations and Environment, quoted in http://www.defense.gov/news/newsarticle.aspx?id=65480, American Forces Press Service, September 29, 2011.

[30] Report by the United States Department of Defense, "Operational Energy Strategy: Implementation Plan," March 2012, Office of the Assistant Secretary of Defense for Operational Energy Plans and Programs, Washington, DC 20301-3700.

[31] S. Burke, C. Parthemore, "A Strategy for American Power: Energy, Climate, and National Security," Center for a New American Security, *Solarium Strategy Series*, Washington, DC, June 2008.

[32] B. Reinart, "Base Camp Integration Lab opens at Fort Devens," USAG-Natick Public Affairs Article, Available: http://www.army.mil/article/60473/Base_Camp_Integration_Lab_opens_at_Fort_Devens/, June 24, 2011.

[33] Operations and Maintenance Manual, Model F100-60K Environmental Control Unit, TM-2003841 Rev G, HDT Engineered Technologies, Salon, OH.

[34] Installation Instructions, Wall Mounted Package Air Conditioners, Model WA484, Manual 2100-398A, Vol III, Tab 16, May 28, 2002, Bard Climate Control Solutions, Bryan, OH.

[35] Technical Manual, Operator, Field, and Sustainment Maintenance Manual for 60K BTU/HR Improved Environmental Control Unit (IECU), TM 9-4120-431-14, July 1, 2010, DRS Technologies, Arlington, VA.

[36] Power generation calculator provided by Global Power Supply, Available: http://www.globalpwr.com/power-calculator/

[37] Generator Efficiency Chart, Available: `http://www.dieselserviceandsupply.com/Diesel_Fuel_Consumption.aspx`

[38] Technical Manual, Operator's Manual for Generator Set, Skid Mounted, Tactical Quiet 100kW, 50/60 Hz, MEP-807A, Army TM 9-6115-729-10, 24, 24P, January 15, 2006.

[39] S. Darby, "The effectiveness of feedback on energy consumption," A Review for DEFRA of the Literature on Metering, Billing and Direct Displays, Oxford Environmental Change Institute, April 2006.

[40] K. Ehrhardt-Martinez, K. Donnelly, and S. Laitner, "Advanced metering initiatives and residential feedback programs: a meta-review for household electricity-saving opportunities," American Council for an Energy-Efficient Economy, Washington, DC, 2010.

[41] U. S. D. of Energy. (2009, August). "The smart grid: An introduction." [Online]. Available: http://energy.gov/oe/downloads/smart-grid-introduction.

[42] G.R. Mitchell, R.Cox, J. Paris, and S.B. Leeb, "Shipboard Fluid System Diagnostic Indicators Using Non-Intrusive Load Monitoring." Naval Engineers Journal, vol.119, no. 2, pp. 109-119, 2007.

[43] M. A. Piber, "Improving shipboard maintenance practices using non-intrusive load monitoring." Ph.D. dissertation, Dept. Elect. Eng. and Comp. Sci., Massachusetts Institute of Technology, Cambridge, MA, 2007.

[44] Orji, Uzoma A., Z. Remscrim, C. Laughman, S.B. Leeb, W. Wichakool, C. Schantz, R. Cox, J. Paris, J.L. Kirtley, and L.K. Norford. "Fault detection and diagnostics for non-intrusive monitoring using motor harmonics," presented at the Applied Power Electronics Conference and Exposition (APEC), Palm Springs, CA, 2010, pp. 1547-1554.

[45] T. DeNucci, R. Cox, S. B. Leeb, J. Paris, T. J. McCoy, C. Laughman, W. Greene, "Diagnostic indicators for shipboard systems using non-intrusive load monitoring,"

presented in the IEEE Electric Ship Technologies Symposium, Philadelphia, PA, July 2005.

[46] S.B. Leeb, S.R. Shaw, J.L Kirtley, "Transient event detection in spectral envelope estimates for nonintrusive load monitoring," Power Delivery, IEEE Transactions on, vol. 10, no. 3, pp. 1200-1210, 1995.

[47] S.R. Sanders, J.M. Noworolski, X.Z. Liu, G.C. Verghese, "Generalized averaging method for power conversion circuits," Power Electronics, IEEE Transactions on, vol 6, no. 2, pp. 251-259, 1991.

[48] J. Paris, Z. Remscrim, K. Douglas, "Scalability of Non-Intrusive Load Monitoring for Shipboard Applications," presented at the ASNE Day Conference, National Harbor, Maryland, Apr. 2009.

[49] Dawei He; Weixuan Lin; Nan Liu; Harley, R.G.; Habetler, T.G., "Incorporating Non-Intrusive Load Monitoring Into Building Level Demand Response," Smart Grid, IEEE Transactions on , vol.4, no.4, pp.1870,1877, Dec. 2013.

[50] J. Donnal, S. Leeb, "Wattsworth: Monitor Electric Power Anywhere," in *Applied Power Electronics Conf. (APEC) and Expo.*, Ft. Worth, TX, 16-20 March 2014, pp. 2223-2230.