Energy Accountability Using Nonintrusive Load Monitoring

Mark D. Gillman, John S. Donnal, Jim Paris, Steven B. Leeb, Mohamed Ahmed Hassan El Sayed, Kenneth Wertz, and Scott Schertz

***Abstract*— Conventional power meters measure total kilowatt-hours yet reveal little about how power was used. Modern solid-state metering solutions are not necessarily taking full advantage of the inexpensive but high performance computation capability that is available. This paper contains the details of a field test of a new software architecture for nonintrusive utility monitoring that endeavors to solve the big data problem of handling interesting but high bandwidth data streams from many monitored sites, e.g., residences and commercial buildings. Results from a field test are used to illustrate the utility of this system.**

***Index Terms*— Energy efficiency, spectral analysis, monitoring.**

I. Actionable Consumption Information

ENERGY waste can and must be quantified – "You don't know what you're losing unless you measure it" [1]. Once measured, facts coupled with feedback can help reduce waste. Thorough reports on the subject estimate that energy scorekeeping will lead to savings of 4-12%, the highest percentage coming when information is real-time and detailed down to the appliance level [2], [3].

Buildings consume over 70% of generated electricity in the United States [4]. Consumer behavior and automated control systems play a significant role in energy usage. Poor equipment operation leads to waste, which can nullify the effects of greener appliances and more effective insulation. Some forms of waste include leaving devices running inadvertently, cooling/lighting unoccupied spaces, or otherwise operating equipment in ignorance of the cost. Equipment malfunctions contribute to energy loss as well. Facilities and machinery degrade in proportion to their age. Dirty filters, clogged vacuum pumps [5], or faulty sensors, if left undetected, lead to recurring energy losses long before equipment ultimately breaks down [6]. Electromechanical systems may also command substantial resources from other utilities like gas and water, complicating the detection of waste.

Are we using contemporary computation tools to our best advantage for tracking and troubleshooting patterns of energy and utility consumption? It is relatively straightforward to update the conventional watt-hour or water flow meter with a digitally tabulating version of essentially the same device. What opportunities do we miss by failing to take every advantage of the exponential growth in increasingly inexpensive computation capability? How do we reap benefits without creating expensive demands on other resources like information transmission, which involves physical expenses that may not keep pace with reductions in the cost of raw computation?

Nonintrusive electrical monitoring has benefited hugely from these same advances in computation [7]–[10]. Further, enhanced signal processing has enabled the deployment of new sensors that permit a nonintrusive load monitor (NILM) to track water consumption in parallel with electrical consumption. New opportunities for load identification have opened, and the NILM has proven a useful platform for performing diagnostics for critical loads in addition to more familiar energy scorekeeping [5], [11]. New data handling models for nonintrusive load modeling help solve the "big data" problem of monitoring many sites with a diverse and numerous collection of loads by increasing the availability of actionable information while minimizing the need for the transmission of large quantities of data. This paper describes this new data architecture and illustrates a field application at a local elementary school. The NILM is able not only to track consumption data for individual electrical loads, but also to derive associated consumption in other utilities like natural gas.

II. Modern Monitoring

Many current approaches to power monitoring involve unfortunate compromises. Sub-metering appliances provide highly discretized data, but the requisite hardware and communication requirements are substantial and invasive [12]. 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 [13]. 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.

Manuscript received December 29, 2013; revised February 10, 2014; accepted February 10, 2014. Date of publication February 12, 2014; date of current version April 16, 2014. This work was supported in part by the Grainger Foundation, in part by the Massachusetts Institute of Technology–Kuwait Partnership, and in part by the MIT SkTech Program. The associate editor coordinating the review of this paper and approving it for publication was Prof. Subhas C. Mukhopadhyay.

The authors are with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA, and also with Kuwait University, Safat 13060, Kuwait (e-mail: gillman@mit.edu; jdonnal@mit.edu; bigjim@mit.edu; sbleeb@mit.edu; elsmah@hotmail.com; wertz@sharon.k12.ma.us; scotty@sharon.k12.ma.us).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSEN.2014.2305911

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 [14]. 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.

1) Hardware: In the field test described below, the NILM measures voltage and current using LEM LF 305-S current transducers and LEM LV 25-P voltage transducers. The data is stored and processed on a computer located with the measuring devices. The 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 samples the LEM sensors at 8 kHz. This sample rate can be varied as needed for any given application. The 8 kHz sample rate is frequently used in our monitoring work because it is adequate to resolve utility harmonics with substantial resolution to at least the 13th harmonic, useful for transient recognition. Also, this sample 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.

2) NilmDB: A new high speed data management software algorithm has been developed called NilmDB [15]. NilmDB stores data collected at high speeds, including 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. NilmDB incorporates robust synchronization to the utility to compute spectral envelopes in the presence of noise and voltage waveform distortion. Preprocessed data like spectral envelopes, i.e., harmonic content in the observed current waveforms, can be stored as separately accessible database streams from raw data. 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. 1, 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.

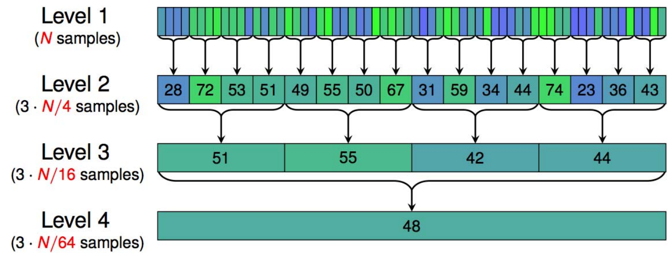

Fig. 1. Original signal of N samples decimated multiple times.

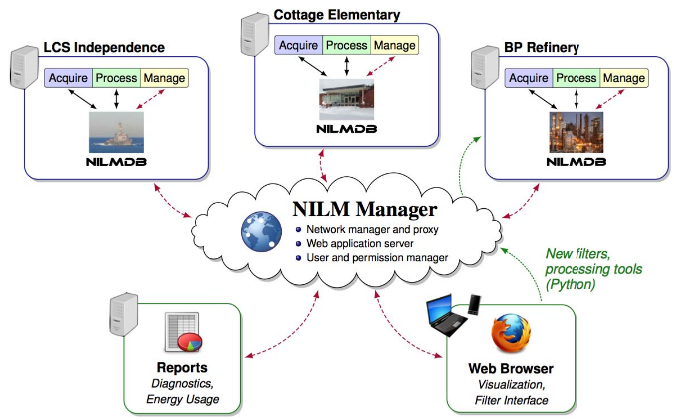

Fig. 2. Nilm Manager controlling several installations.

3) NILM Manager: These decimated files can be accessed by the second key software suite in the nonintrusive energy monitoring system: NILM Manager. NILM Manager [16] is a secure central portal serving NILM computers, communicating over the web. It supports remote access to information from NilmDB installations anywhere in the internet-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. 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. 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 [16].

4) Load Discrimination: Load discrimination filters numerically analyze preprocessed data. Edge detection filters created and applied in this paper search for events, or changes in the power signature that reveal the activity of individual devices. The filters detect ON/OFF events using two criterions: the first difference and the change of mean. In a given data set of n samples, the i-th sample is subtracted from the (i-1) sample to yield the first difference. The change of mean is calculated by taking the difference between a multiple-sample average (mean) after the transient and a similar average before the transient. At turn-on, a given device draws increased apparent (S) power in some combination of real (P) and reactive (Q) components. An event is detected when the average magnitude of P and/or Q increase above a set threshold and the derivative is sufficiently large. Threshold values should narrowly exceed noise levels and will establish which transients will be ignored. Also, considering only the first difference has shown to be an effective criteria, differing from other methods that are more computationally intensive [17]. In practice, many device transients vary too widely to make point-by-point comparisons practical. However, the instantaneous response to changes in input power is sufficiently similar for the devices we consider in this paper.

Once an edge is found, the characteristics of that edge are recorded, namely the time, utility phase, steady state delta-kW and delta kVAR, and the first difference magnitude of P and Q at the fundamental, 3rd, 5th and 7th harmonic frequencies. Settling into steady state typically takes less than a second, though there are some exceptions such as variable speed driven loads.

To illustrate power data and transient events, Fig. 3 depicts an hour of preprocessed data from an elementary school machine room. Several transients are visible above the 3kW baseline load. Inductive-load turn-on events are relatively easy to detect due to their distinctive inrush currents. 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. Sensor dynamic range is adaptable to the anticipated current range at a site. 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 10 W to 100 kW range, where most load transients show significant features, targeting the larger machines that would have the greatest impact on consumption. Noise in the system on the order of 150W (peak to peak) exists due to the number and type of loads. Averaging

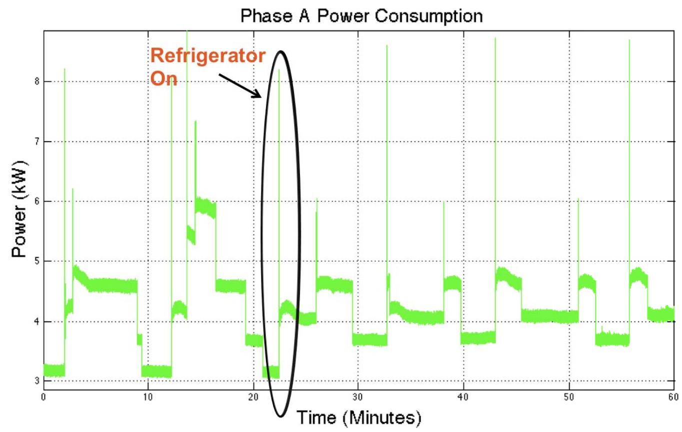

Fig. 3. One hour of real power data at Cottage. A refrigerator turning on is one of several events during this period.

techniques help discriminate smaller load transients, but this study focused primarily on the larger, more costly machinery whose signatures are readily extractable.

The classifier filter matches events to a corresponding device. In more detail, a comparison is made between each edge detected to a list of known values. A library of load characteristics for each device in the system must first be created. This 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, a topic of recent interest [18], [19]. To ensure reliable identifications, rich streams of relevant data may be used to create exemplars, including real and reactive transient power shapes, higher harmonic content, and state models of load behavior.

Many load identification methods have been researched, a brief summary of which can be found in [20]. Most methods are implemented on systems with 15 or less loads, are residential in nature [7]–[10], and benefit from load classification techniques generally defined by two features, 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 have similar kW and kVAR characteristics at the fundamental frequency to use only fundamental components. To these features, we added harmonics, first differences, and utility phase.

III. Case Study

The Cottage Elementary School in the Sharon School District in Massachusetts has served as a fascinating and representative test bed to demonstrate the NilmDB/NILM Manager 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. 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.

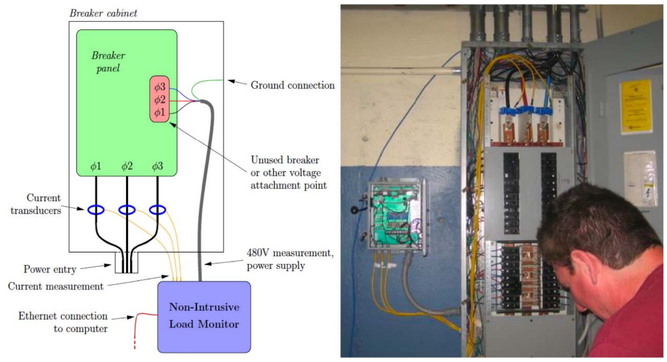

Fig. 4. NILM schematic and installation.

An electrician installed the NILM system on a 3-phase subpanel known as the emergency panel (EBPP). Fig. 4 shows the connection scheme. In the interests of student education, a website was also established to graphically display the preprocessed power consumption data from that panel (http://www.nilmdb.com, under the Cottage Elementary drop down menu). 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 subpanels at Cottage, but the EBPP accounts for about 1/4th of the school's total electrical power consumption during winter months.

A. 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. 5). Operation of the heat system depends on several user-established inputs. If the outside air temperature is below 55 °F, the boilers will operate according to water temperature settings in the loop. If the return-loop temperature is below 170 °F, the boilers will operate until it reaches 185 °F. 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. 6. The minimum and maximum kW

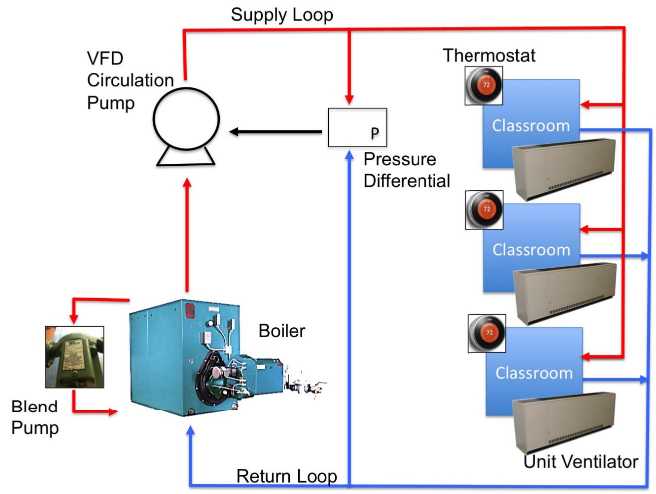

Fig. 5. Cottage heating system.

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 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 4.5kW 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, to include smaller pumps and many of the electronic equipment in the IT room, were purposely set aside to draw attention to the larger, more energy-consuming equipment.

B. 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 depicts one hour of the collective power signal on phase A on Monday, March 26, 2013, from 12:00-1:00 PM. Two Python programming scripts developed for Cottage filtered the preprocessed data. The following figures will demonstrate how the filters decompose this signal into its individual loads. With the DC 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. 7, 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

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

Fig. 6. Loads monitored by NILM at Cottage Elementary School.

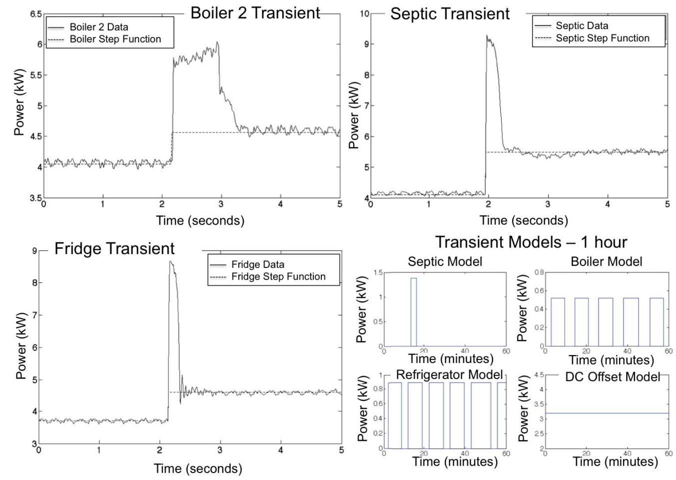

Fig. 7. Load transients modeled with step functions.

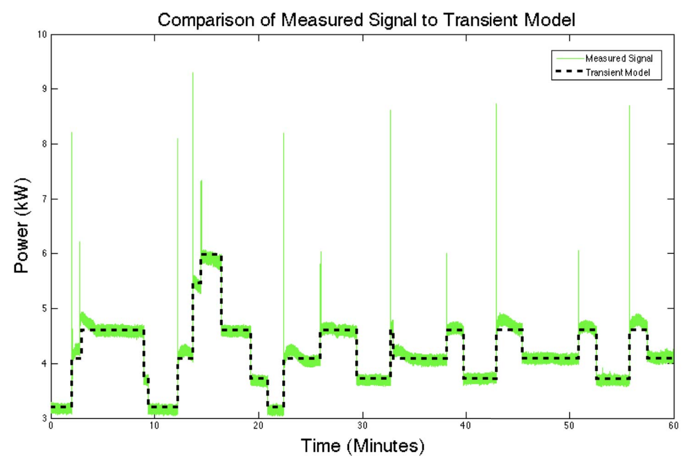

Fig. 8. Comparison of original signal with modeled signal.

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 [21]. In Cottage, most loads were distinguishable using relatively simple characterizations of the "load transient," i.e., just the change in steady power consumption.

For example, see Fig. 7. The Boiler Pump must physically move water that is initially static and thus requires more force at first to overcome inertia. As more laminar flow is reached, however, the power requirements on the pump quickly approach steady state operation. Inrush current peaks at about 2 kW for fractions of a second. Power fluctuates for another few milliseconds before leveling off at a level that is somewhere in the range of 0.46−0.58 kW at steady state. The other transients follow similar patterns for moving air 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. 7.

The superposition of these three individual transient models closely approximates the original power signal (Fig. 8), 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.

IV. Case Study Result

The NILM system detected some 5100 events over a period of 6 winter days in 2013 (11-13 March and 24-26 March). 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,

	Boiler 1 Pump	Boiler 1 Fan	Boiler 2 Pump	Boiler 2 Fan	Circulation Pumps
Times used per day	15.3	15.3	63.5	63.5	1.0
Duration ON per cycle (min)	9.7	9.7	7.2	7.2	1440.0
Duration ON per day (min)	289.7	289.7	544.2	544.2	1440.0
Daily Cost	**$0.15**	**$0.37**	**$0.44**	**$0.63**	**$8.71**

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

Fig. 9. Average usage and cost per machine over 6-day period.

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**

Fig. 10. Cost comparison of Boiler 1 to Boiler 2.

more that 98% of the events were classified without error. Most often, the issue was 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. 9. From the monthly power bill, Cottage paid just over 9 cents per kWh to the utility company.

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.

In this experiment, NILM showed promise as a plausible sensor for natural gas sub-metering. Since the burner specifications 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

$$2136\text{MBTU/hr} \times 10.1\text{hr} = 647.208\text{MBTU/month} \quad (1)$$

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 whose combined capacity is about 25% of a boiler burner, but it is interesting to note that the math is in the ballpark and merits further study.

Another result made possible by the NILM is a comparison between the boilers. 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 between boilers. 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.

NILM also measured the effect of a major change to the system. On 25 March, the weather turned warmer. This led

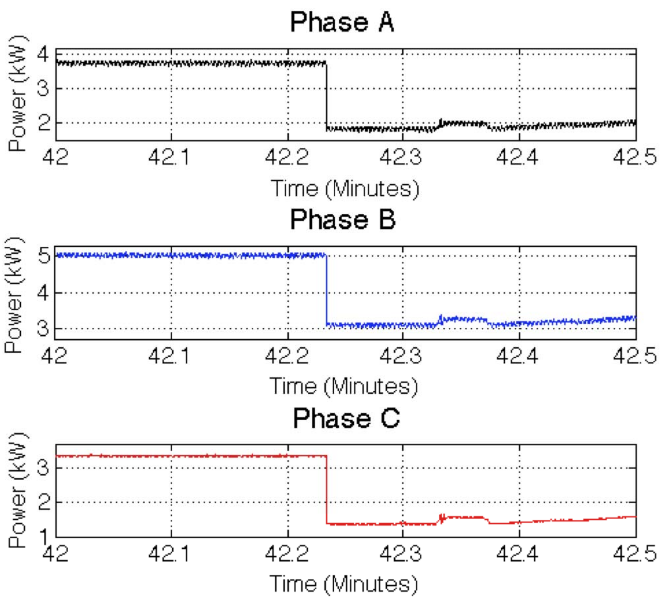

Fig. 11. Significant three-phase power drop when the VFD demand drops suddenly.

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. 11). 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.

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 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.

The reconstructed model (Fig. 8) 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 [22]. 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) [21]. 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.

V. Mission Critical Monitoring

The study at Cottage demonstrated a new software architecture for nonintrusive power system monitoring that takes advantage of low-cost *in-situ* or on-site computing. Detailed appliance-level consumption feedback is possible through NILM. The hardware footprint is minimal, essentially a "Raspberry-PI-style" computing resource and associated sensors. Network bandwidth requirements are very small. A commercial or industrial setting with or without automation could benefit greatly from the energy scorekeeping provided by the NILM software suite. This field test also demonstrated the ability of the NILM to "derive" details of other utility consumptions like natural gas. This directly points out the value of increased local "intelligence" or signal processing in unraveling the "big data" problem associated with consumption feedback and diagnostic monitoring.

The NILM system is uniquely suited for austere electrical networks where loads are standardized. It may be a spectacular tool for assisting with micro grid control and economization. For small or islanded networks, the library of loads can theoretically be pre-set, reducing the extent or possibly eliminating the need for a training phase. In terms of network requirements, the bandwidth required for communication is very tractable. We have already begun to examine this approach in oil refineries, and other possible examples include oilrigs, solar plants, wind farms, industrial parks, and other micro grid installations such as military forward operating bases. In the military's case, where there is already a mandate to reduce consumption [23], [24], accountability is made available quickly and inexpensively.

In machine rooms like at Cottage, we don't expect loads to be moved around. However, in future work, we will examine what happens when new loads are added or portable loads are connected to different phases or outlets. Ongoing research efforts will engage such scenarios with more sophisticated machine learning algorithms.

REFERENCES

[1] Rockwell Automation, Inc. New York, NY, USA. (2011, Aug.). *Owens Corning's Sustainability Plan Includes Energy-Efficient Variable Frequency Drive Solutions for Annual Savings of 538 MW-HR* [Online]. Available: http://www.rockwellautomation.com.au/Applications/CSS_Artilce.nsf/0/c92928abd85d5dc251862575f20072414a/$FILE/FIBER-AP001C-EN-P.pdf

[2] S. Darby, *The Effectiveness of Feedback on Energy Consumption: A Review for DEFRA of the Literature on Metering, Billing and Direct Displays*. Oxford, U.K.: Environmental Change Institute, Apr. 2006.

[3] K. Ehrhardt-Martinez, K. Donnelly, and S. Laitner, *Advanced Metering Initiatives and Residential Feedback Programs: A Meta-Review for Household Electricity-Saving Opportunities*. Washington, DC, USA: American Council for an Energy-Efficient Economy, 2010.

[4] (2009, Aug.). *The Smart Grid: An Introduction* [Online]. Available: http://energy.gov/oe/downloads/smart-grid-introduction

[5] M. A. Piber, "Improving shipboard maintenance practices using non-intrusive load monitoring," Ph.D. dissertation, Dept. Electr. Eng. Comput. Sci., Massachusetts Inst. Technol., Cambridge, MA, USA, 2007.

[6] U. A. Orji, Z. Remscrim, C. Laughman, S. B. Leeb, W. Wichakool, C. Schantz, *et al.*, "Fault detection and diagnostics for non-intrusive monitoring using motor harmonics," in *Proc. 25th Annu. IEEE APEC Exposit.*, Palm Springs, CA, USA, Feb. 2010, pp. 1547–1554.

[7] K. Anderson, A. Ocneanu, D. Benitez, D. Carlson, A. Rowe, and M. Bergés, "BLUED: A fully labeled public dataset for event-based non-intrusive load monitoring research," in *Proc. 2nd KDD Workshop Data Mining Appl. Sustainabil.*, Beijing, China, 2012, pp. 12–16.

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

[9] Z. Wang and G. Zheng, "Residential appliances identification and monitoring by a nonintrusive method," *IEEE Trans. Smart Grid*, vol. 3, no. 1, pp. 80–92, Mar. 2012.

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

[11] T. DeNucci, R. Cox, S. B. Leeb, J. Paris, T. J. McCoy, C. Laughman, *et al.*, "Diagnostic indicators for shipboard systems using non-intrusive load monitoring," in *Proc. IEEE Electr. Ship Technol. Symp.*, Philadelphia, PA, USA, Jul. 2005, pp. 413–420.

[12] H. Y. I. Cheung, S. Lanzisera, J. Lai, R. Brown, S. Dawson-Haggerty, J. Taneja, *et al.*, "Detailed energy data collection for miscellaneous and electronic loads in a commercial office building," in *Proc. ACEEE Summer Study Energy Efficiency Buildings*, Asilomar, CA, USA, Aug. 2012, pp. 1–12.

[13] (2013). *The Power Standards* [Online]. Available: http://www.powerstandards.com/PQube.php

[14] C. Schantz, "Method for nonintrusive sensing and system monitoring," Ph.D. dissertation, Dept. Electr. Eng. Comput. Sci., Massachusetts Inst. Technol., Cambridge, MA, USA, 2014.

[15] J. Paris, "A comprehensive system for non-intrusive load monitoring and diagnostics," Ph.D. dissertation, Dept. Electr. Eng. Comput. Sci., Massachusetts Inst. Technol., Cambridge, MA, USA, 2013.

[16] J. Donnal, "Home NILM: A comprehensive non-intrusive load monitoring toolkit," M.S. thesis, Dept. Electr. Eng. Comput. Sci., Massachusetts Inst. Technol., Cambridge, MA, USA, 2013.

[17] J. Paris, Z. Remscrim, and K. Douglas, "Scalability of non-intrusive load monitoring for shipboard applications," presented at the ASNE Day Conf., National Harbor, MD, USA, Apr. 2009.

[18] M. Dong, P. C. M. Meira, W. Xu, and C. Y. Chung, "Non-intrusive signature extraction for major residential loads," *IEEE Trans. Smart Grid*, vol. 4, no. 3, pp. 1421–1430, Sep. 2013.

[19] D. He, W. Lin, N. Liu, R. G. Harley, and T. G. Habetler, "Incorporating non-intrusive load monitoring into building level demand response," *IEEE Trans. Smart Grid*, vol. 4, no. 4, pp. 1870–1877, Dec. 2013.

[20] L. Farinaccio and R. Zmeureanu, "Using a pattern recognition approach to disaggregate the total electricity consumption in a house into the major end-uses," *Energy Buildings*, vol. 30, no. 3, pp. 245–259, Aug. 1999.

[21] C. Laughman, K. Lee, R. Cox, S. Shaw, S. Leeb, L. Norford, *et al.*, "Power signature analysis," *IEEE Power Energy Mag.*, vol. 1, no. 2, pp. 56–63, Apr. 2003.

[22] H. Friedman, A. Potter, T. Haasl, D. Claridge, and S. Cho, "Persistence of benefits from new building commissioning," in *Proc. 11th Nat. Conf. Building Commissioning*, Oct. 2003, pp. 20–22.

[23] *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, DC, USA, Feb. 2008.

[24] M. Schwartz, K. Blakeley, and R. O'Rourke, "Department of defense energy initiatives: Background and issues for congress," vol. 42558, Congressional Res. Service, Tech. Rep., Dec. 2012.

Mark D. Gillman received the B.S. degree in electrical engineering from the United States Military Academy in 2004. He is an active duty Army Major, Professional Engineer in the state of Missouri, and currently pursuing the M.S. degree in electrical engineering from the Massachusetts Institute of Technology. His current research interests include the measurement and analysis of microgrid power systems in military applications.

John S. Donnal received the B.S. degree in electrical engineering from Princeton University in 2007 and the M.S. degree in electrical engineering from the Massachusetts Institute of Technology in 2013, where he is currently pursuing the Ph.D. degree. His current research interests include nonintrusive load monitoring synthesis, energy harvesting, and communications systems.

Jim Paris received the B.S., M.S., and Ph.D. degrees from the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology in 2003, 2006, and 2013, respectively. He is currently doing independent consulting work in Boston, MA, USA.

Steven B. Leeb (S'89–M'91–SM'01–F'07) received the B.S., M.S., and Ph.D. degrees from the Massachusetts Institute of Technology in 1987, 1989, and 1993, respectively. He is currently a Professor with the Laboratory for Electromagnetic and Electronic Systems, Massachusetts Institute of Technology. His current research interests include the design, analysis, construction, control, and monitoring of servomechanical actuators and mechatronic systems as well as the application of exotic materials, including gel polymers, to actuator construction.

Mohamed Ahmed Hassan El Sayed received the B.S. and M.S. degrees in electrical engineering from Cairo University, Egypt, in 1968 and 1972, respectively, and the Ph.D. degree from the RWTH University, Aachen, Germany, in 1978. He is currently a Professor with the Electrical Engineering Department, Kuwait University, Kuwait City, Kuwait. His current research interests include power system analysis, reliability, protection, and artificial intelligence. Dr. Mohamed is a recipient of the Egyptian National Prize of Engineering Science in 1989.

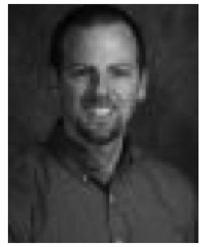

Kenneth Wertz received the Associates Degree in plumbing technology from the New England Institute of Technology in 1992. He is currently the Director of Maintenance and Operations with the Sharon School District in Massachusetts.

Scott Schertz is the Chief Maintenance Officer with the Sharon School District in Massachusetts.