

Empirical Characterization and Modeling of Electrical Loads in Smart Homes

Sean Barker, Sandeep Kalra, David Irwin, and Prashant Shenoy
University of Massachusetts Amherst

Abstract—Smart meter deployments are spurring renewed interest in analysis techniques for electricity usage data. An important prerequisite for data analysis is characterizing and modeling how electrical loads use power. While prior work has made significant progress in deriving insights from electricity data, one issue that limits accuracy is the use of general and often simplistic load models. Prior models often associate a fixed power level with an “on” state and either no power, or some minimal amount, with an “off” state. This paper’s goal is to develop a new methodology for modeling electric loads that is both simple and accurate. Our approach is empirical in nature: we monitor a wide variety of common loads to distill a small number of common usage characteristics, which we leverage to construct accurate load-specific models. We show that our models are significantly more accurate than binary on-off models, decreasing the root mean square error by as much as 8X for representative loads. Finally, we demonstrate the use of our models in data analysis to i) generate device-accurate synthetic traces of building electricity usage and ii) filter out stable min-max loads that generate rapid and random power variations in building electricity data.

I. INTRODUCTION

Computing for sustainability—where real-world physical infrastructure leverages sensing, networking, and computation to mitigate the negative environmental and economic effects of energy use—has emerged as an important new research area. As a result, in addition to improving the energy efficiency of information technology (IT) infrastructure such as mobile devices, servers, and data centers, computing researchers are expanding their focus to now include building energy efficiency. Since buildings account for nearly 40% of society’s energy use [1], compared to an estimated 1-2% for IT infrastructure [2], this research has the potential to make a significant impact. In particular, managing electricity is critical because buildings consume the vast majority (73%) of their energy in the form of electricity [1]. Existing management techniques typically employ *sense-analyze-respond control loops*: various *sensors* monitor the building’s environment (including electricity) via a smart meter, and transmit collected data in real-time to servers, which *analyze* it to reveal detailed building usage and occupancy patterns, and finally *respond* by automatically *controlling* electrical loads¹ to optimize energy consumption.

Research challenges exist at each stage of the control loop. For example, despite much prior research [3], [4], [5], [6] accurate, fine-grained, *in situ* sensing of electricity use in real time at large scales remains impractical—it is prohibitively expensive, invasive, and unreliable. Unfortunately, timely and detailed knowledge of per-load electricity use is a prerequisite for implementing sophisticated automated load control policies that increase energy efficiency. Even small residential homes require such large-scale sensing systems [7], since they often

operate hundreds of individual loads. A promising approach to address this problem is to use fewer sensors that generate less data, and compensate by employing more intelligence in the analysis phase to infer rich information from the data. For example, prior research indicates that analyzing changes in a building’s aggregate electricity usage at small time granularities, e.g., every 15 minutes or less, reveals a wealth of information: Non-Intrusive Load Monitoring (NILM) techniques use the data to infer electricity usage for individual loads [8], while recent systems use it to infer building occupancy patterns [9]. The derived insights then inform control policies. In this case, NILM might enable buildings to identify opportunities for reducing peak demand by scheduling elastic background loads [10], such as air conditioners and heaters, while occupancy patterns are critical in determining when to turn loads off without disturbing people’s lives [11].

Since electric utilities are rapidly deploying digital smart meters capable of measuring and transmitting a building’s aggregate electricity usage in real time, a substantial amount of fine-grained electricity data for buildings is now available. For example, Pacific Gas and Electric now operates over nine million smart meters in California [12]. While today’s deployed smart meters typically measure average power usage at intervals ranging from five minutes to an hour, commodity meters are now available that measure and transmit, via the Internet, energy usage at intervals as small as every second [13], [14]. Combined with the emergence of “big data” cloud storage systems, these smart meter deployments are spurring renewed interest in analysis techniques for electricity usage data. While prior research has made progress in deriving insights from smart meter data, one issue that often limits accuracy is the use of general, and often simplistic, load models. In particular, many prior techniques for analyzing and modeling building electricity data characterize loads using simple on-off models, which associate one (or a small number) of fixed power usage levels with the “on” state and either no power consumption, or some minimal amount, with the “off” state.

On-off models do have a number of advantages. For instance, they exactly capture many simple loads, including light bulbs and other low-power resistive devices with mechanical switches. In addition, on-off models allow researchers to describe buildings as state machines that associate each building state with a fixed power level that implies the set of loads that are on, and where state transitions occur whenever a load turns on or off. Characterizing buildings as state machines admits a plethora of analysis techniques. For instance, much prior work maps building state machines to Hidden Markov Models (HMMs), and applies HMM-based techniques, such as Viterbi’s algorithm [15], [16], to determine which loads are on in each state. In this case, using only a few (often two) power states per load is advantageous, since it minimizes the number of distinct power states for the entire building and reduces

¹We use the term electrical load, or simply *load*, to refer to any distinct appliance or device that consumes electricity.

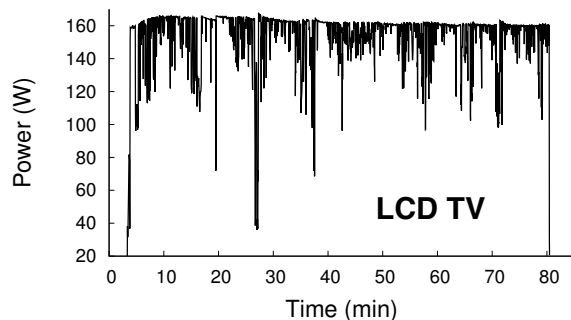


Fig. 1. An LCD TV’s power usage varies rapidly, significantly, and unpredictably while on, and does not conform to a simple on-off load model.

the complexity of analyzing the resulting state machine. Of course, even with only two power states per load, the number of building power states is still exponential in the number of loads, i.e., 2^n for n loads. Thus, even assuming simple load models, precise analysis may still be intractable, i.e., require enumerating an exponential number of states.

Unfortunately, while on-off load models are simple, they are often inaccurate, since they fail to capture the complex power usage patterns common to many loads. As a simple motivating example, Figure 1 shows a time-series of an LCD TV’s electricity usage each second. In this case, the TV’s switched mode power supply (SMPS) causes power variations as large as 120 watts (W) by rapidly switching between a full-on and full-off state to minimize wasted energy. The magnitude of these variations is effectively random—determined by the color and intensity of the TV’s pixels. An on-off model clearly does not accurately capture the TV’s power usage. As a result, modeling the TV as an on-off load may complicate higher-level analysis techniques for building data. For example, the TV may obscure the use of low-power loads, such as a 60W light bulb, since its power usage varies rapidly by $>60W$.

Our premise is that simple on-off models discard a significant amount of information that is potentially useful in analyzing data. As a result, in this paper, we focus explicitly on accurately characterizing and modeling a variety of common household loads. Our methodology is empirical: we i) gather fine-grained electricity usage data from dozens of loads across multiple homes, ii) characterize their behavior by distilling a small number of common usage attributes, and then iii) derive accurate load-specific models based on these attributes. One of our contributions is to show that a small number of model types accurately describe nearly all household loads. Many of our findings stem from basic knowledge of power systems that has not yet been fully exploited in data analysis. Thus, one of our goals is to highlight how many identifiable load attributes, which are well-known in power systems, manifest themselves in electricity data collected by smart meters. Our hypothesis is that accurate load models, which leverage domain knowledge from power systems, provide a foundation for designing new electricity data analysis techniques. In evaluating our hypothesis, we make the following contributions:

Empirical Data Collection and Characterization. We instrument a wide variety of common electrical loads in multiple homes, and collect electricity usage data for each load, every second, for over two years. We show empirically that homes operate similar types of loads, e.g., lighting, AC motors,

heating elements, electronic devices, etc., which results in significant commonality in power usage profiles across loads. We then characterize the data to identify distinguishing attributes in per-load power usage, forming the building blocks of our models. While many of these attributes are well-known in power systems, we show how they manifest in sensor data.

Modeling Methodology. We use our empirical characterization to construct a small number of load-specific model types. We show that our basic model types, or a composition of them, capture nearly all of our household loads. Our models go beyond on-off models, by capturing power usage characteristics that i) decay over time, ii) have frequent variations (as with the TV in Figure 1), iii) exhibit complex repetitive patterns of simpler internal loads, and iv) are composites of two or more simpler loads. We show that our models are significantly more accurate than on-off models, decreasing the root mean square (RMS) error by as much as 8X for representative loads. Since our methodology is general, it applicable to modeling other loads or appliances beyond those in this paper.

Case Studies. Finally, while we expect our models to have numerous uses, we illustrate two specific examples of novel applications we have designed using them. First, we use our models to generate device-accurate synthetic traces of building electricity usage for use by NILM researchers. One barrier to evaluating NILM techniques is deploying sensors to collect ground truth data from every building load for evaluation. Using our models, NILM researchers could quickly generate different types of synthetic building traces by composing together a representative collection of load models. Second, we design filters capable of identifying and removing loads, such as the LCD TV, that exhibit rapid power variations. Since these loads introduce numerous spurious building power states, removing them may improve the accuracy of HMM-based analysis techniques that describe buildings as state machines.

II. EMPIRICAL DATA COLLECTION AND CHARACTERIZATION

A typical home consists of dozens of electrical loads, including heating and cooling equipment, lights, appliances of various kinds and electronic equipment. A partial lists includes:

- *Heating, cooling, and climate control equipment* such as a central air conditioner, window air conditioner, space heater, electric water heater, dehumidifier, fan, air purifier;
- *Kitchen appliances* such as an electric range, microwave, refrigerator, coffee maker, toaster, blender, dishwasher;
- *Laundry appliances* such as a washing machine and dryer;
- *Lighting* including incandescent and fluorescent lights;
- *Miscellaneous electronic devices* such as a television, audio receiver, radio, battery charger, laptop and desktop computer, and gaming console; and
- *Other appliances* such as a vacuum and carpet cleaner.

Below, we briefly describe the data collection infrastructure we use to gather data from these common household loads. We then derive several insights from our data, which we use to design different types of load models in the next section.

A. Data Collection Infrastructure

Since our methodology is empirical, we instrument three homes with a large number of energy monitoring sensors to

gather ground truth electricity usage data from a wide variety of loads. Each instrumented home consists of a smart home gateway in the form of an embedded Linux server that queries each sensor to collect data. We have deployed several different types of energy monitoring sensors, as described below.

We use current transducer (CT) sensors to monitor electricity usage for large loads wired to dedicated circuits, such as air conditioners, washing machines, dryers, dishwashers, and refrigerators. These sensors connect to in-panel meters, such as The Energy Detective (TED) [14] or eGuage [13] and sample per-circuit electricity usage each second. We use plug-level energy sensors to track energy use for smaller loads plugged into wall outlets. Our plug-level sensors are commodity Insteon iMeters [17] and Z-Wave Smart Energy Switch meters [18], which use powerline and wireless communication, respectively, to transmit readings to our smart home gateway. These sensors also collect electricity usage data up to every second. Finally, we use Insteon-enabled wall switches to monitor switched loads wired directly into the electrical system. These switches replace normal wall switches, and transmit on-off-dim events to the gateway whenever a user manually toggles the switch.

Our system has been continually monitoring hundreds of individual loads every second for nearly two years in each of our three instrumented homes. Since our level of instrumentation is time-consuming and expensive to replicate, we have made much of our collected data available to benefit other researchers [7]. We leverage our data to characterize various loads based on a few elemental types, described below.

B. Characterizing Different Types of Loads

Despite their tremendous variety, most residential loads fall into one of a few elemental load types based on how they consume power in an alternating current (AC) system. In particular, loads are categorized as either resistive, inductive, capacitive, or non-linear based on how they draw current in relation to voltage, which in an AC system varies along a smooth sinusoidal pattern. These categories reveal properties of the loads that we leverage in our models. Since many researchers outside of power systems may be unfamiliar with these load types, for each type of load we first review its salient characteristics. We then empirically characterize data from multiple representative loads of each type to observe how their specific characteristics manifest themselves in the data.

Resistive Loads. Loads that consist of any type of heating element are resistive. Incandescent lights, toasters, ovens, space heaters, coffee makers, etc., are examples of common resistive loads in a home. Formally, if a load draws current along a sinusoidal pattern in the same phase as the voltage, i.e., the maximum, minimum, and zero points of the voltage and current sine waves align, then the load is purely *resistive*.

Figure 2 depicts a time-series of the power usage for four different resistive loads with heating elements: an incandescent light bulb, a toaster oven, a coffee maker, and a sandwich press. In general, the power usage of these loads resembles a “step” when turned on, with usage that remains relatively stable and flat. The incandescent light acts as a nearly perfect resistive load with a power usage that equals the bulb’s wattage. While the toaster oven, coffee maker, and sandwich press act similarly to the light bulb, they each experience an initial higher power

usage that slowly decays to a relatively stable usage, as highlighted in Figure 2. This initial higher power is due to the large inrush current that occurs until the device warms up and the resistance decreases, after which it stabilizes.

Observation 1: Resistive loads exhibit stable power usage when turned on, with high-power heating elements exhibiting an initial surge followed by a slow decay to stable power.

Inductive Loads. AC motors are the most common and widely-used examples of inductive loads. Motors are the primary component of many household devices, including fans, vacuum cleaners, dishwashers, washing machines, and compressors in refrigerators and air conditioners. Formally, if a load draws current along a sinusoidal pattern that peaks *after* the voltage sine wave, i.e., the current waveform lags the voltage waveform, then the load is purely *inductive*.

Figure 3 depicts a time-series of the power usage for four inductive loads: two refrigerators, a freezer, and a vacuum cleaner. All four loads operate AC motors. Unlike the resistive loads above, each inductive load experiences a significant, but brief, initial power usage. The surge is also due to inrush (or surge) current that occurs when starting an AC motor, although it is typically much higher than for heating elements. Intuitively, the underlying reason is that, while heating elements heat up slowly, the rotor inside a motor must transition from completely idle to full speed within seconds. As before, power usage exhibits a decay that eventually stabilizes. In contrast to resistive loads, however, motors exhibit small variations even during this “stable” phase. For instance, the first refrigerator shown in Figure 3(a) exhibits small fluctuations that repeat during each cycle of the compressor. The second refrigerator, the freezer and the vacuum cleaner depicted in Figures 3(b), (c) and (d) all show an initial spike and a sharper, smoother decay, with small variations as the usage stabilizes. These patterns demonstrate that, unlike resistive loads, modeling inductive loads using simple on-off step functions is problematic.

Observation 2: Inductive loads with AC motors exhibit a large initial spike in power usage, followed by a decay to a stable power level. The decay rate is load-dependent, with the stable power level also exhibiting smaller time-scale variations.

Capacitive Loads. Capacitive loads are the dual of inductive loads. While many loads have capacitive elements, inductive and resistive characteristics dominate their overall behavior. Thus, there are no significant capacitive loads in buildings. Formally, if a load draws current along a sinusoidal pattern that peaks *before* the voltage sine wave, i.e., the current waveform leads the voltage waveform, then the load is purely *capacitive*.

Non-linear Loads. Finally, any load that does not draw current along a sinusoidal pattern is called *non-linear*. Non-linear loads may also be resistive, inductive, or capacitive according to when their current waveform peaks. The most predominant non-linear (and largely inductive) loads are electronic devices, including desktop computers and TVs. The non-linear nature of these loads is primarily due to the use of switched-mode power supplies (SMPS). Fluorescent lights are another example of a non-linear (inductive) load. Smaller electronic devices that convert AC to low-voltage DC, such as battery chargers for portable devices and digital clocks, are also non-linear.

Figure 4 shows the power usage of four different non-

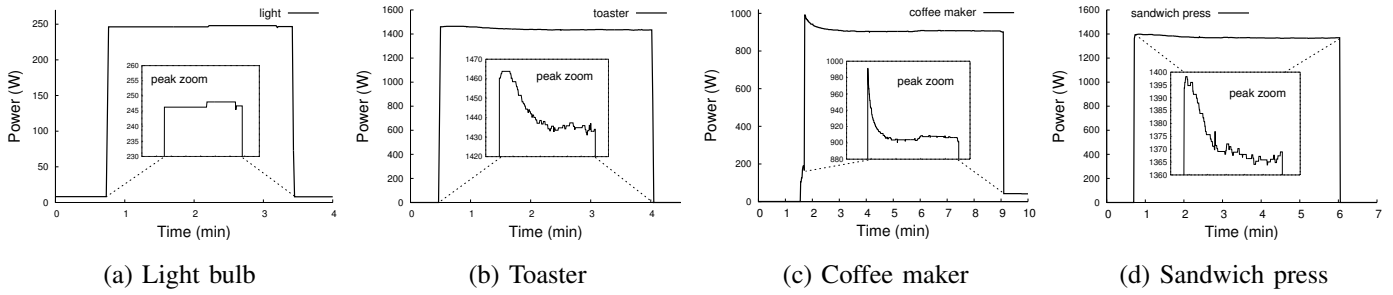


Fig. 2. Example resistive loads, demonstrating “step” behavior with a possible initial surge.

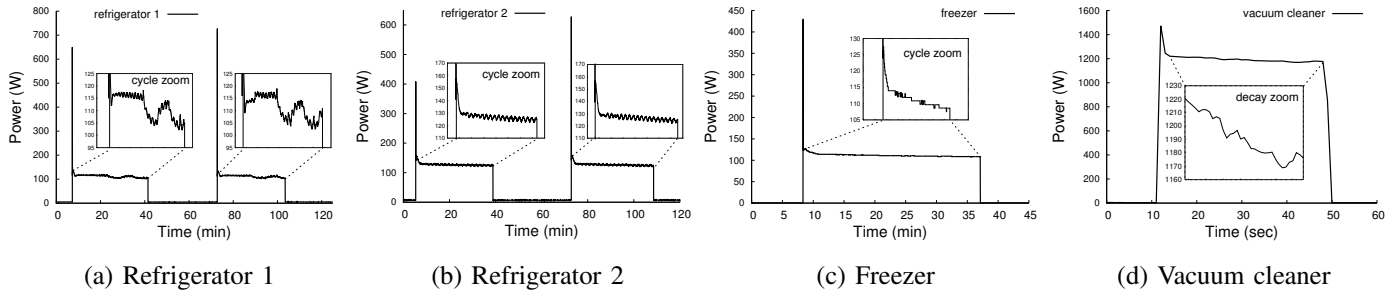


Fig. 3. Example inductive loads, demonstrating significant surge current and subsequent decay.

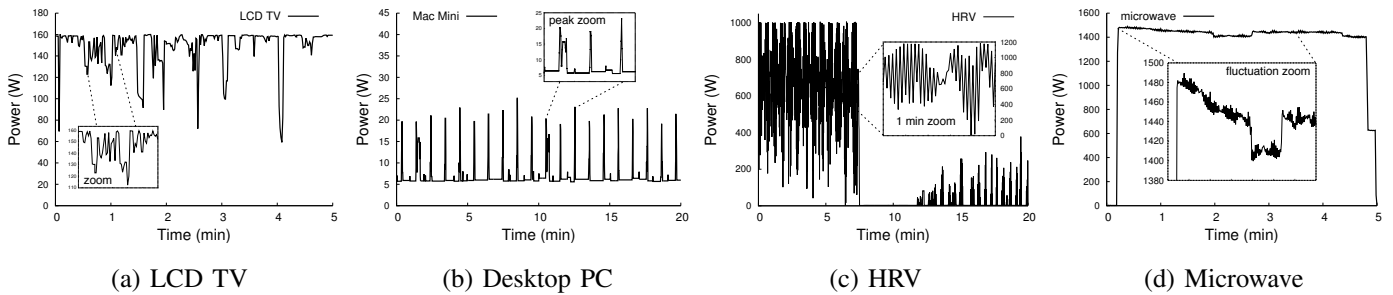


Fig. 4. Example non-linear loads, demonstrating random variations and possible ceilings and/or floors.

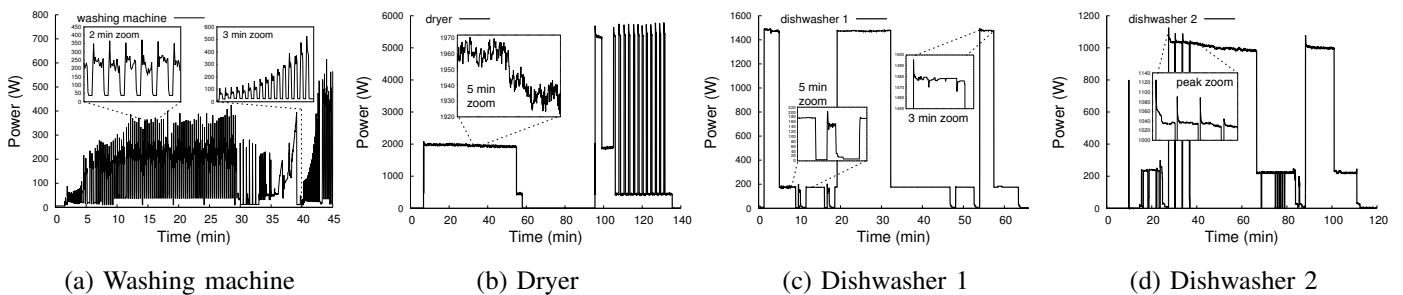


Fig. 5. Example composite loads, demonstrating combinations of simpler loads arranged in phases.

linear loads: an LCD TV, a Mac Mini desktop computer, a microwave oven, and an actively-heated heat recovery ventilator (HRV). These loads exhibit significant power fluctuations when active, but also have a stable floor or ceiling from which these fluctuations derive. The LCD TV shown in Figure 4(a) exhibits a stable maximum usage with random power reductions from this ceiling. These fluctuations result from displaying a variety of color and pixel intensities on the screen. In contrast, the desktop computer shown in Figure 4(b) has a stable minimum power draw, with random power spikes above this floor depending on its workload, e.g., causing the CPU to ramp up, etc. Both the TV and desktop computer consist of a switched mode power supply (SMPS) that regulate the

power usage of the device and switch between a full-on and full-off state to minimize wasted energy. The HRV shown in Figure 4(c) demonstrates two regular modes of operation; an active heating mode—with instantaneous intensity managed by the HRV controller—and a passive recirculation mode. In both modes, there are large, random variations in power usage. In the active state, there is also a clear stable maximum usage. Finally, the microwave shown in Figure 4(d) has what initially appears to be a straightforward step, similar to the resistive loads. However, zooming into the figure shows the microwave’s non-linear behavior, with rapid, albeit small, variations in the second-to-second usage, along with larger periodic power shifts. These examples demonstrate that on-off

models are inappropriate for non-linear loads, since two power states cannot capture their wide range of power variations.

Observation 3: Non-linear loads exhibit significant random variations in power usage. These fluctuations are often range-bound and capped by a floor or ceiling in the power level.

C. Reactive Power and Composite Loads

Reactive Power. Another important characteristic of the elemental load types above is how they consume *reactive power*. While real power is the amount of power delivered to a load, and is often referred to as simply electricity or power (without the qualifier), reactive power is the amount of power generated, but not delivered, to the load; it is also measured in units of watts, but written as voltage-amperes reactive (VAR) to distinguish it from real power. Reactive power arises when a load draws current *out of phase* with the voltage. Thus, only non-resistive loads generate reactive power. At a high level, reactive power is the result of the instantaneous power (the product of current and voltage) occasionally becoming negative within each cycle of AC power, due to out-of-phase current and voltage. This state causes power to flow towards the generator and away from the load. Reactive power is typically dissipated as heat in power lines. For our purposes, reactive power provides additional useful information for modeling, and many commodity power meters are capable of measuring it. As a result, our models include both real and reactive power.

Figure 6 depicts companion graphs for selected loads that shows their reactive, rather than real, power usage. Figure 6(a) shows that, similar to their real power usage, dimmable incandescent lights produce a stable—zero if not dimmed—amount of reactive power when on, although the magnitude of the draw peaks at 50% dim level and decreases as the light approaches either 0% or 100% dim level. Likewise, the refrigerator in Figure 6(b) exhibits a mostly flat reactive draw, while the HRV in Figure 6(c) has a rapidly varying power usage. Both patterns are similar to each load’s pattern of real power usage.

Observation 4: While the magnitude of reactive power differs from real power, a load’s pattern of reactive power consumption is qualitatively similar to its real power consumption.

Composite Loads. Finally, many household loads, particularly large appliances, are not purely resistive, inductive or non-linear. Instead, these loads consist of multiple components, each of which may be one of the simpler load types. For instance, a central air conditioner may consist of a compressor, a fan to blow air into ducts, duct dampeners to control air flow, and central humidifiers to control humidity. A refrigerator, which has compressor that is an inductive load, may also consist of door lights, an ice maker, and a water dispenser. Similarly, electric dryers, washing machines, and dishwashers also consist of a motor—to spin clothes and circulate water via a pump—and a heating element—to dry clothes or warm water. In addition, these appliances often operate in repetitive cycles that operate each of their constituent loads differently, such as washing, draining, and then drying for a dishwasher.

Figure 5 depicts the power usage of a washing machine, a dryer, and two dishwashers. As shown, these loads exhibit distinct behavior in different parts of their cycle depending on which component of the appliance is in use. For example,

based on the observations above, distinguishing when a complex load, such as the dryer, activates its heating element versus its motor is straightforward. Finally, an appliance may activate its various components in sequence, in parallel, or both. For instance, a central air conditioner may operate the compressor, the fan and the dampeners concurrently, while a dishwasher may operate its motor, pump and heater in sequence.

Observation 5: Composite loads consist of simpler resistive, inductive and non-linear loads that operate in parallel, in sequence, or both. As a result, composite loads exhibit distinct behaviors in different operating regions of their active cycle.

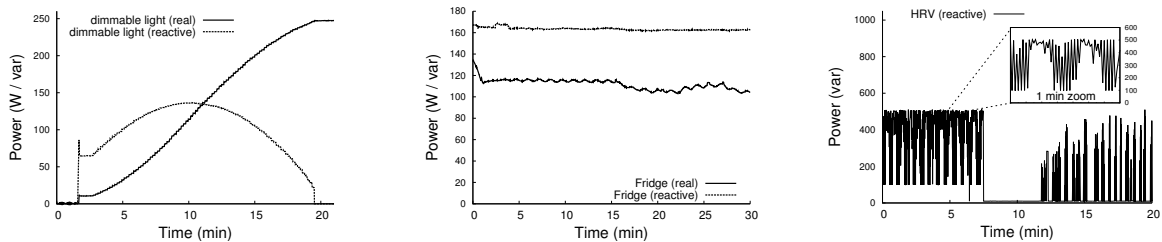
D. Summary

In classifying loads in terms of the elemental load types above, we observe that *nearly every common household electric load is a composition of one or more of the small number of resistive, inductive, and non-linear loads described above*, with heating elements and AC motors consuming the majority of electricity in homes. Further, each type of elemental load exhibits similar characteristics when active: heating elements have a stable power usage or one that decays slowly over time, AC motors have a spike in power on startup and then vary their power usage smoothly over time, while SMPSSs exhibit rapid and significant power variations. As we discuss in the next section, the presence of only a few elemental load types in homes simplifies model design, enabling us to accurately capture their behavior using a few basic types of models.

III. MODELING ELECTRIC LOADS

Based on our empirical observations from the previous section, we develop models to capture key characteristics of each load type. We first present four basic model types—*on-off*, *on-off decay*, *stable min-max*, and *random range*—to describe simple loads, and then use these models as building blocks to form compound *cyclic* and *composite* models that describe more complex loads. Ideal models describe i) *how much* real and reactive power a load uses when active, ii) *how long* a load is active, and iii) *when* a load is active. However, in many cases, users manually control loads, such that when a load is active and for how long is non-deterministic. For example, a user may run a microwave any time for either ten seconds or ten minutes. For these loads, we assume a random variable captures this non-determinism, and focus our efforts, instead, on modeling *how* each load behaves when active.

Given each model type, we employ an empirical methodology to construct accurate load-specific models: we leverage our empirical observations of the load’s power usage as a training set, and employ curve-fitting methods to map one of the model types onto the time-series data. If the best model type is not clearly evident *a priori*, we fit multiple models and then choose the one that yields the best fit. As described below, depending on the model type, we may employ simple regression or more complex curve-fitting methods, such as LMA [19], to construct a load-specific model for a given model type. As discussed in Section II, reactive power for loads exhibits similar behavior as real power, and thus constructing a model of a load’s reactive power consumption uses the same methodology as above.



(a) Dimmable light (dim from 0% to 100%)

(b) Refrigerator

(c) HRV (real power shown in Figure 4c)

Fig. 6. Reactive power demonstrates the same types of patterns as real power and can help in identifying devices.

A. Basic Model Types

On-off Model. As discussed in Section I, prior work often uses simple on-off models for loads. An on-off model includes two states—an *on* state that draws some fixed power p_{active} and an *off* state that draws zero, or some minimal amount of, power p_{off} . Conventional, non-dimmable incandescent lights are the canonical example of a load that conforms to an on-off model. Dimmable lights also conform to on-off models, although p_{on} depends on the dim level. As shown empirically in Figure 6(a), a dim-level of $N\%$ yields a proportionate reduction in real power usage. In addition, while real power usage is a simple linear function of the dim level, reactive power is a quadratic function that peaks at 50% dim level.

Constructing an on-off model is simple—we use regression to determine appropriate values p_{on} and p_{off} . In particular, we partition the time series of load power usage into two mutually exclusive time-series with data for the on and off periods, to determine the best values of p_{on} and p_{off} .

On-off Decay Model. An on-off decay model is a variant of the on-off model that accounts for an initial power surge when a load starts, followed by a decrease in power usage over time. As discussed in Section II, AC motors are the most common example of a load that has this behavior, e.g., refrigerator, freezer, vacuum cleaner. Resistive loads with high-power heating elements, such as the toaster or coffee maker, also conform to an on-off decay model, although the surge and the decay in these devices is far less prominent than in AC motors. We characterize on-off decay models using four parameters: p_{active} , p_{off} , p_{peak} , and λ . The first two parameters are the same as in on-off models, while p_{peak} represents the level of inrush current when a device starts up and λ represents the rate of decrease to the stable p_{active} power level. We model decay using an exponential function as follows, where t_{active} is the length of the active interval:

$$p(t) = \begin{cases} p_{active} + (p_{peak} - p_{active})e^{-\lambda t} & , 0 \leq t < t_{active} \\ p_{off} & , t \geq t_{active} \end{cases}$$

As with the on-off models above, the length of the active interval for on-off decay models is not known *a priori* and often depends on user behavior. However, we have observed that in many cases users repeatedly operate devices in the same way, e.g., a toaster that toasts a bagel every morning. In many cases, the device determines the length of t_{active} automatically, e.g., the compressor for a refrigerator or freezer may turn on for an average of 20 minutes in each cycle. In these cases, we incorporate the mean value of t_{active} into the model.

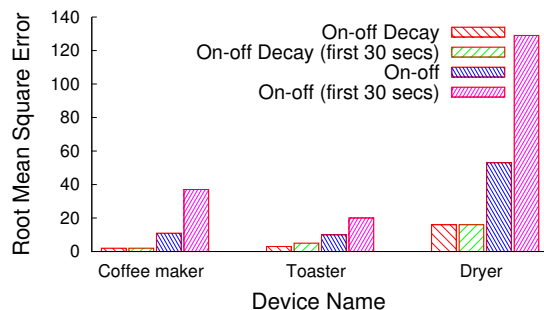


Fig. 8. On-off decay models are often more accurate than on-off models.

Constructing an on-off decay model requires fitting an exponentially decaying function onto the time-series data, in addition to determining p_{peak} , p_{active} , and p_{off} . We employ the well-known LMA algorithm [19] to numerically find the exponential decay function that best fits, i.e., based on a least-squares nonlinear fit, the time-series data.

Figure 7(a) shows the specific the on-off decay model for a coffee maker in parallel with its real power data. The figure demonstrates that the exponential decay is a highly accurate approximation of the coffee maker’s power usage. In this case, $p_{active} = 905$, $p_{peak} = 990$, $p_{off} = 0$, and $\lambda = 0.045$. Likewise, Figures 7(b) and (c) show on-off decay models and real power data for a toaster and a portion of a dryer cycle. For each of the three loads in Figure 7, we also fitted an on-off model. After fitting the best on-off model, we calculated root mean square error (RMSE) for both the on-off and on-off decay models for each load’s duration of activity and for its first 30 seconds of activity (since capture ‘on’ events is particularly important). Figure 8 shows the results: the decay model decreases the error in the on-off model by as much as 8X, particularly in the first 30 seconds where the on-off model is unable to capture the rapid decay behavior.

Stable Min-Max Model. While on-off and on-off decay models accurately capture the behavior of resistive and inductive loads, they are inadequate for modeling non-linear loads. As seen in Section II, many non-linear loads maintain a stable maximum or minimum power draw when active, but often vary randomly and frequently from this stable state. These variations are due to the device regulating their electricity usage at a fine grain to instantaneously “match” the needs of the tasks the device is performing. Our stable min-max model captures this behavior. We model stable min-max devices as having a stable maximum or minimum power denoted by p_{active} when active. The power usage then deviates, or “spikes,” up or down from this stable value at some frequency. The magnitude of the spike is a random value uniformly distributed between p_{active} and

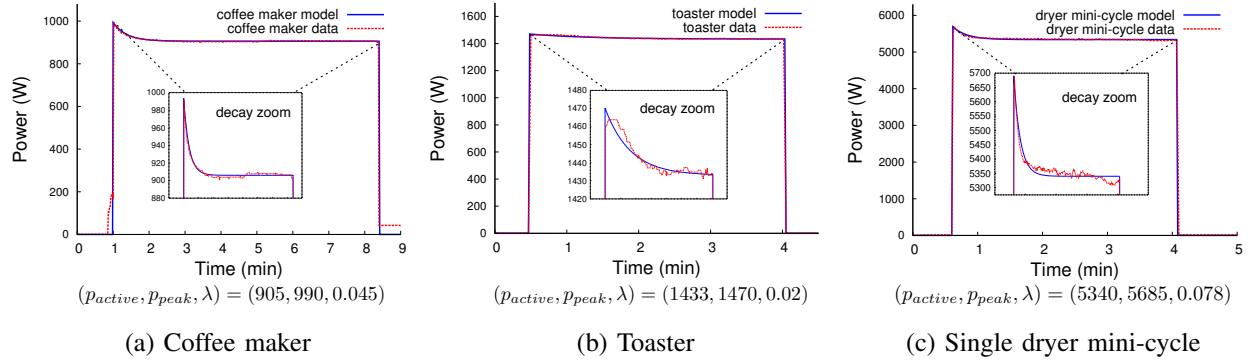


Fig. 7. An on-off decay model closely matches the average power usage measured each second for a variety of resistive and inductive loads.

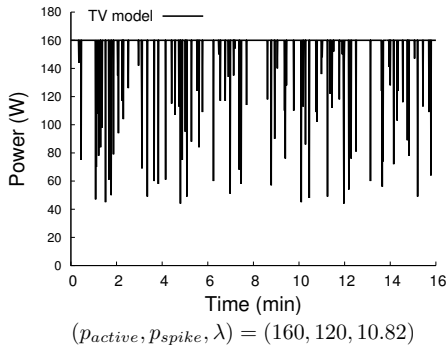


Fig. 9. A stable-max model of the LCD TV from Figure 1.

p_{spike} , where p_{spike} denotes the maximum deviation in power per spike. The interarrival time of the spikes are exponentially distributed with mean λ . Thus, the stable min-max model has three parameters: a stable max or min power usage p_{active} , the maximum deviation p_{spike} for each spike, and λ , which governs the interarrival time of spikes. Note that p_{active} may denote either the stable maximum or minimum power, with power deviations decreasing or increasing, respectively.

Empirically constructing a load-specific stable mix-max model requires both determining the stable power level p_{active} and characterizing the magnitude and frequency of the power spikes. We employ simple regression to determine the stable power level p_{active} from the time series data, e.g., after filtering out the data for spikes and finding the fit for p_{active} . The mean observed duration between spikes then yields the parameter λ . Figure 9 shows our stable-max model for the LCD TV (from Figure 1) using a maximum p_{active} of 160W and a λ of 10.82, which we derive from the TV’s real power usage data. Importantly, as we discuss in Section IV, both the model and the raw data have similar statistical properties, which simple filters can recognize by detecting when power variations are significant, frequent, and symmetric, e.g., a decrease and then immediate increase in power of similar magnitude.

Random Range. Finally, we found that some devices draw a seemingly random amount of power within a fixed range when active. This is likely due to the fact that taking average power readings each second is too coarse a frequency to capture the device’s repetitive behavior. We model such loads by determining upper and lower bounds on their power usage, denoted by p_{max} and p_{min} . When active, our model randomly

varies power within these bounds using a random walk. Note that the random range model is similar to the stable min-max model in that both employ upper and lower bounds on power usage. However, while the deviations in the stable min-max model are spikes from a stable value, those in the random range model are power variations within a range. The microwave is an example of a load that exhibits this behavior. As shown in Figure 4(d), when turned on, the power usage of the microwave fluctuates continuously between 1400W and 1480W.

Random range models require determining the minimum and maximum of the load’s range of power usage. We determine these values by simply choosing the minimum and maximum power values observed in training data, or by deriving a distribution of power values from the data and choosing a high and low percentile of the distribution to be the minimum and maximum, p_{min} and p_{max} . We then model the variations with a random walk within the range.

B. Compound Model Types

While the models above accurately capture the behavior of simple loads, many loads, including large appliances, exhibit complex behavior from operating a variety of smaller constituent loads. We devise two types of compound models for complex loads that use the basic building blocks above.

Cyclic Model. Cyclic loads repeat one of the basic model types in a regular pattern, often driven by timers or sensors. For example, the HRV heater employs a timer that activates for 20 minutes each hour. Similarly, a refrigerator duty-cycle is based on sensing its internal temperature, which rises and falls at regular intervals and fits our model well, as shown in Figures 3(a) and (b). A cyclic model augments a basic model by specifying the length of the active and inactive period, t_{active} and $t_{inactive}$, each cycle. Constructing cyclic models is straightforward, since it only requires extracting the duration of the active and inactive periods from the empirical time-series power data. We currently use the mean of the active and inactive periods from the time-series observations to model t_{active} and $t_{inactive}$.

Composite Model. Composite loads exhibit characteristics of multiple basic model types either in sequence or parallel. Example composite loads include dryers, washing machines, and dishwashers, as shown in Figure 5. *Sequential composite loads* operate a set of basic load types in sequence; we model them as simple piecewise functions that encode the

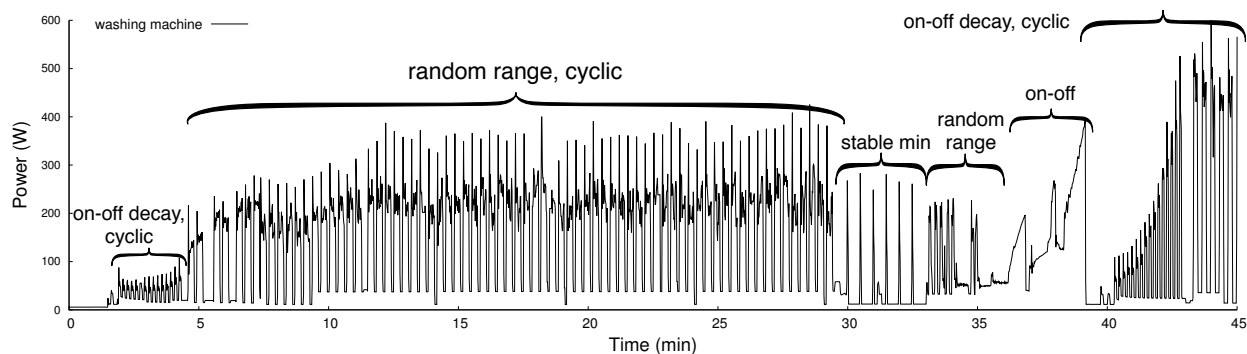


Fig. 10. A single complete cycle of a washing machine.

sequence of basic load models, including how long each load operates. For instance, a model for a dishwasher is a sequence of stages: modeled as the operation of the motor (wash stage), pump (drain stage), motor (rinse stage), pump (drain stage) and heater (dry stage), where each individual stage uses an inductive or resistive load. Some loads also exhibit characteristics of two or more basic models in *parallel* if two basic loads operate simultaneously. For example, a refrigerator may simultaneously activate both a compressor and an interior light. We model *parallel composite loads* by adding together the power usage for two or more of the basic model types. Finally, composite loads may also be cyclic, referred to as *cyclic composite loads*, which repeat a pattern of individual model types at regular intervals. Our methodology permits arbitrary compositions of sequential, parallel or cyclic loads.

Constructing load-specific composite models is more complex and requires additional manual inputs. For example, constructing a *sequential composite* model requires manually partitioning and isolating load time-series data into individual sequences that reflect the activation of the various load components. We must then construct basic models from above for each component in the sequence. The composite model is then simply a concatenation of these piecewise models in sequence.

As an example, Figure 10 shows an extended operating cycle of the washing machine with the annotations for different basic load model types in the sequence. We represent these models as large piecewise functions of the basic models that describe each constituent load. We omit these functions here for brevity. In addition, many of the large appliances that have composite models also have numerous operating states. For example, the washing machine and dryer in one of our homes each has over 25 different types of cycles. In this case, an accurate model requires a different piecewise functions for each type of cycle. However, in the homes we monitor we have found that most residents operate devices using only a small number states—in most cases one.

Constructing parallel composite models poses additional challenges. Since the time-series data for a load captures the power usage for all components that are concurrently active, there is no straightforward general-purpose technique to extract individual models from the composite time-series data. In practice, however, extracting basic models is often possible through exogenous means. For instance, many loads permit operating individual components to isolate them for profiling, e.g., such as a running a dryer on tumble mode without any heat or using an air conditioner’s fan without any cooling. After

separately profiling a constituent load, such as the tumbler or fan, it is possible to operate the compressor and the fan, and then infer the compressor power usage by “filtering out” the tumbler or fan usage from the aggregate. In some devices, such as a refrigerator, it also might be possible to deploy additional sensors that monitor important events, such as a door opening that triggers lights, to filter them out.

IV. USING THE MODELS

While we expect our models to have numerous uses in a variety of analysis and optimization tasks in building energy management, below, we illustrate two initial examples of novel applications we have designed: i) generating device-accurate synthetic data of a building’s aggregate electricity usage and ii) designing filters to identify and remove random and frequent power variations from stable min-max loads. As we discuss, both applications should prove useful to researchers.

A. Device-accurate Synthetic Building Data

Evaluating new techniques for analyzing electricity data requires actual building data for testing. Unfortunately, while recording a building’s aggregate electricity usage is simple, requiring only a single smart meter, recording detailed aspects of the building’s environment is not. For instance, evaluating the accuracy of a NILM algorithm, which disaggregates building electricity data into power data for individual loads, requires power data from *both* the entire building and each of its constituent loads. However, NILM’s entire purpose is to prevent the need for recording such ground truth data at each load, since setting up even a test infrastructure for this is prohibitively expensive, invasive, and time-consuming. While there are now a small number of data sets for a few buildings available for NILM researchers to use in evaluation [7], [20], [21], they typically do not instrument *every* load nor do they cover a wide range of building types or load characteristics.

To address this problem, we use our models to automatically generate device-accurate synthetic electricity data for buildings. Being device-accurate means that we include *both the synthetic aggregate time-series power data for a building, as well as time-series power data for each of the constituent loads* in the building generated using our models. While prior work targets generating synthetic traces of building power usage [22], we are not aware of any previous work that focuses on being device-accurate. We expect device-accurate modeling to enable researchers to evaluate new techniques for data analytics on a variety of building types without requiring them

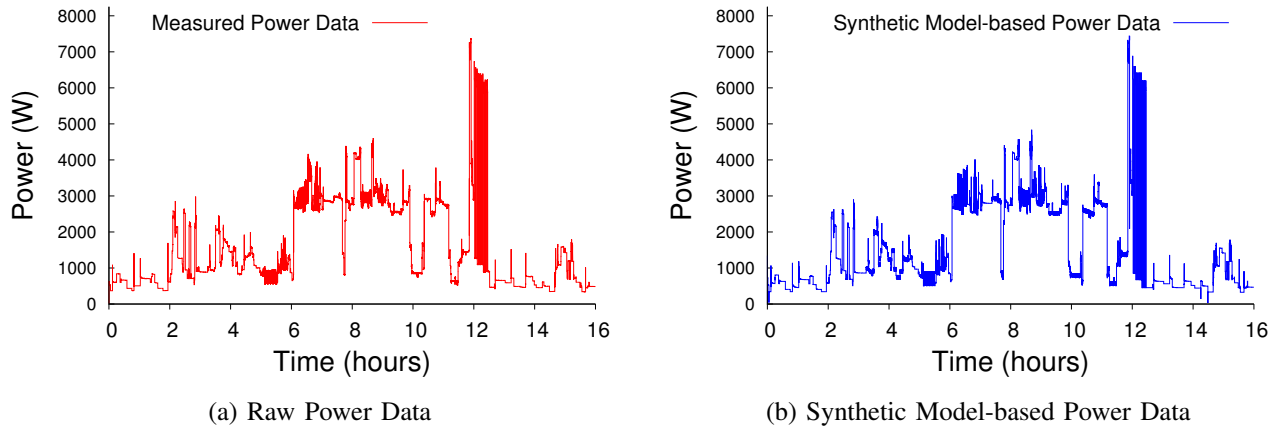


Fig. 11. Aggregate power usage in a home from measured power data (a) and from our corresponding devices models (b).

to deploy a large number of power meters. Importantly, our device-accurate synthetic building data has similar statistical properties as real data. Figure 11 shows both a real trace of the aggregate power usage of one of the homes we monitor (a), and a synthetically derived trace using our models (b).

We generate the synthetic trace from the “on” events for the home’s loads, where we insert our model for a load whenever it comes on. Not only do the traces in Figure 11 look qualitatively similar, but they also have similar properties: the real data has an average power of 1200W, a standard deviation of 1072W, and 5591 changes in power >15 W, while the synthetic data has an average power of 1165W, a standard deviation of 1073W, and 5833 changes in power >15 W. Since the synthetic data is composed of data from models of individual loads, it is useful for analysis techniques that look for patterns in the aggregate usage data. By comparison, if we generate on-off models that include at most 4 power states per load (as in recent work [20]), there are only 1985 changes in power >15 W, which eliminates many identifiable load-specific characteristics. We are currently using our models to design a general workload generator that automatically produces synthetic trace data for different types of buildings, while also providing users control over the trace’s statistical properties, e.g., number of loads, average power, variance, etc., since some analysis techniques may be appropriate for some types of buildings and not for others. We plan to release our device-accurate synthetic data and workload generator as part of our public Smart* data set [7].

B. Stable Min-max Filters

Since stable min-max loads account for the large majority of power variations in a home, filtering out the variations can significantly improve subsequent data analysis techniques, including HMM-based techniques that model buildings as state machines composed of on-off loads. Figure 12 shows a bar graph of the number of power variations in a single day for each of 19 active circuits in one of the homes we monitor. The graph demonstrates that a small number of circuits account for the vast majority of the power variations in the home—notably resulting from two non-linear loads (the active HRV and the living room TV) and two composite loads with non-linear components (the washing machine and dryer), each of which is highlighted in Figure 4. Our stable min-max filter

works by scanning through the power data time-series for a home and maintaining a *stable power* parameter, which only updates if power deviates from the current parameter setting for more than time T by some threshold power $P_{threshold}$. Time T and threshold $P_{threshold}$ are based on the load’s model. For T , the goal is to select a value large enough to ensure that a power variation is a stable load (and not a sequence of random variations in the stable min-max model), but short enough to prevent filtering out legitimate short-lived loads, e.g., a lamp that is only on for 30 seconds. For $P_{threshold}$, the goal is to select a value large enough to capture significant variations.

Figure 13 applies the filter to aggregate data from a home that includes the TV and HRV (from Figure 4) combined with a 60W light bulb. Since the TV and the HRV exhibit rapid power variations every few seconds (see Figure 4), due to controllers that rapidly switch between a fully-on and fully-off state, the value of T need only be slightly greater than the typical frequency of the variations. In both cases we set $T = 5$ seconds. In addition, we select a value of 10W for $P_{threshold}$ since both loads have a narrow maximum power usage that varies by less than 10W. As the figure shows, the filter makes the HRV and TV appear to be simpler loads with discrete power states, making it possible to easily identify a 60W light bulb turning on or off by observing changes in power in the filtered data. However, without filtering, any data analysis technique, e.g., HMMs, that uses changes in aggregate power usage to identify when the light bulb turns on or off would have difficulty, since operating either the TV or HRV would obscure changes in the light bulb’s power by introducing significant and frequent power variations. In fact, the HRV’s power variations alone are large enough—greater than 800W in some cases—to obscure all but the largest loads in the home.

V. RELATED WORK

In this paper, we focus explicitly on modeling the power usage of common electrical loads. While recent work targets modeling for specific appliances, e.g., a particular brand of refrigerator [23], it does not generalize to a broad range of devices. Much of the prior research on modeling power usage for individual loads has been done in the context of Non-Intrusive Load Monitoring (NILM). While we expect our models to be broadly useful for data analysis, including, but not limited to, NILM, we survey related work in NILM below.

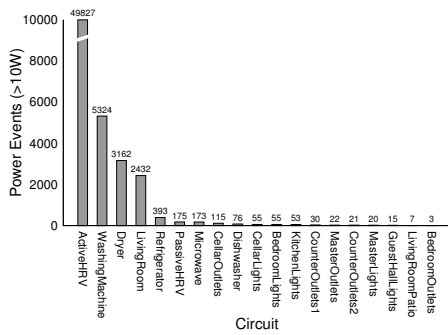


Fig. 12. A few highly variable (non-linear) loads are responsible for the majority of changes in a home’s aggregate power.

NILM techniques differ significantly based on the granularity of the current, voltage, and power readings. For instance, fine-grain readings that sample current or voltage frequently within each cycle, e.g., $\gg 60\text{Hz}$, differ significantly from the models we present, since they attempt to capture the behavior of the AC current and voltage waveform. In addition, gathering data at such high frequencies presents challenges: it requires expensive and highly calibrated equipment, while storing and transmitting the data in real time is beyond the capabilities of today’s embedded power meters. Instead, our models target a data granularity of one reading per second, since this is the finest granularity that commodity low-cost power meters support. We could also build load models using the coarse-grain data, e.g., 5 minutes to an hour, supported by today’s utility-installed smart meters [24]. Unfortunately, coarse-grain data measured every minute or more eliminates important details of each load’s operation that are useful in analysis.

Recently, a number of researchers have focused on NILM approaches for the per-second power data we use to build our models. Most of these approaches employ generic on-off load models that, as we show, are not accurate. The techniques generally use these simple models to either i) detect changes in load power states by observing changes in building power, [8] or ii) use Viterbi-style algorithms [15] to determine the most likely set of “hidden” states, e.g., combinations of power states for multiple loads, from a sequence of changes in building power [25]. These prior techniques generally do not scale to the large numbers of, often low-power, loads found in typical buildings. For instance, we are not aware of any prior approach that focuses on large-scale scenarios— > 100 loads—with many low-power loads $< 50\text{W}$, which is a common characteristic of many homes. The lack of research may be due, in part, to the inaccuracy of the underlying load models.

VI. CONCLUSION

This paper presents a new methodology for modeling common electrical loads. We derive our methodology empirically by collecting data from a large variety of loads and showing the significant commonalities between them. Finally, we illustrate examples of how to use our models for data analysis.

REFERENCES

[1] J. Kelso, Ed., *2011 Buildings Energy Data Book*. Department of Energy, March 2012.
 [2] J. Koomey, “Data Center Electricity Use 2005 to 2010,” Analytics Press, Tech. Rep., August 2011.

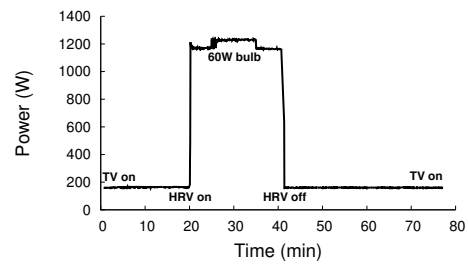


Fig. 13. The stable maximum power enables a filter that removes variations in aggregate electricity usage data, making it easier to detect on-off transitions.

[3] S. Dawson-Haggerty, S. Lanzisera, J. Taneja, R. Brown, and D. Culler, “@scale: Insights from a Large, Long-Lived Appliance Energy WSN,” in *IPSN*, April 2012.
 [4] T. Hnat, V. Srinivasan, J. Lu, T. Sookoor, R. Dawson, J. Stankovic, and K. Whitehouse, “The Hitchhiker’s Guide to Successful Residential Sensing Deployments,” in *SenSys*, November 2011.
 [5] X. Jiang, M. V. Ly, J. Taneja, P. Dutta, and D. Culler, “Experiences with a High-fidelity Building Energy Auditing Network,” in *SenSys*, November 2009.
 [6] Y. Kim, T. Schmid, Z. Charbiwala, and M. Srivastava, “ViridiScope: Design and Implementation of a Fine Grained Power Monitoring System for Homes,” in *UbiComp*, September 2009.
 [7] S. Barker, A. Mishra, D. Irwin, E. Cecchet, P. Shenoy, and J. Albrecht, “Smart*: An Open Data Set and Tools for Enabling Research in Sustainable Homes,” in *SustKDD*, August 2012.
 [8] M. Zeifman and K. Roth, “Nonintrusive Appliance Load Monitoring: Review and Outlook,” *IEEE Transactions on Consumer Electronics*, vol. 57, no. 1, February 2011.
 [9] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin, “Private Memoirs of a Smart Meter,” in *BuildSys*, 2010.
 [10] S. Barker, A. Mishra, D. Irwin, P. Shenoy, and J. Albrecht, “SmartCap: Flattening Peak Electricity Demand in Smart Homes,” in *PerCom*, March 2012.
 [11] J. Lu, T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field, and K. Whitehouse, “The Smart Thermostat: Using Occupancy Sensors to Save Energy in Homes,” in *SenSys*, November 2010.
 [12] “Pacific Gas and Electric, leveraging Smart Meter Installation Progress,” <http://www.pge.com/myhome/customerservice/smartmeter/deployment/>, February 2013.
 [13] “eGauge Energy Monitoring Solutions,” <http://www.egauge.net/>, 2013.
 [14] “Energy, Inc.” <http://www.theenergydetective.com/>, July 2012.
 [15] G. Forney, “The Viterbi Algorithm,” *Proceedings of the IEEE*, vol. 61, no. 3, pp. 268–278, 1973.
 [16] A. Viterbi, “Error Bounds for Convolutional Codes and an Asymptotically Optimum Decoding Algorithm,” *IEEE Transactions on Information Theory*, vol. 13, no. 2, pp. 260–269, 1967.
 [17] iMeter Solo, <http://www.insteon.net/2423A1-iMeter-Solo.html>.
 [18] <http://aeotec.com/z-wave-plug-in-switch>.
 [19] K. Levenberg, “A Method for the Solution of Certain Non-Linear Problems in Least Squares,” *Quarterly of Applied Mathematics*, 1944.
 [20] J. Kolter and M. Johnson, “REDD: A Public Data Set for Energy Disaggregation Research,” in *SustKDD*, August 2011.
 [21] K. Anderson, A. Ocneanu, D. Benitez, D. Carlson, A. Rowe, and M. Berges, “BLUED: A Fully Labeled Public Dataset for Event-Based Non-Intrusive Load Monitoring Research,” in *SustKDD*, August 2012.
 [22] O. Ardakanian, S. Keshav, and C. Rosenberg, “Markovian Models for Home Electricity Consumption,” in *GreenNets*, August 2011.
 [23] J. Taneja, D. Culler, and P. Dutta, “Towards Cooperative Grids: Sensor/Actuator Networks for Renewables Integration,” in *SmartGridComm*, 2010.
 [24] J. Kolter and A. Ng, “Energy disaggregation via discriminative sparse coding,” in *NIPS*, December 2010.
 [25] H. Kim, M. Marwah, M. Arlitt, G. Lyon, and J. Han, “Unsupervised Disaggregation of Low Frequency Power Measurements,” in *SDM*, April 2011.