SmartCap: Flattening Peak Electricity Demand in Smart Homes

Sean Barker, Aditya Mishra, David Irwin, Prashant Shenoy, and Jeannie Albrecht†
University of Massachusetts Amherst
†*Williams College*

***Abstract*—Flattening household electricity demand reduces generation costs, since costs are disproportionately affected by peak demands. While the vast majority of household electrical loads are interactive and have little scheduling flexibility (TVs, microwaves, etc.), a substantial fraction of home energy use derives from background loads with some, albeit limited, flexibility. Examples of such devices include A/Cs, refrigerators, and dehumidifiers. In this paper, we study the extent to which a home is able to transparently flatten its electricity demand by scheduling only background loads with such flexibility. We propose a Least Slack First (LSF) scheduling algorithm for household loads, inspired by the well-known Earliest Deadline First algorithm. We then integrate the algorithm into SmartCap, a system we have built for monitoring and controlling electric loads in homes. To evaluate LSF, we collected power data at outlets, panels, and switches from a real home for 82 days. We use this data to drive simulations, as well as experiment with a real testbed implementation that uses similar background loads as our home. Our results indicate that LSF is most useful during peak usage periods that exhibit "peaky" behavior, where power deviates frequently and significantly from the average. For example, LSF decreases the average deviation from the mean power by over 20% across all 4-hour periods where the deviation is at least 400 watts.**

I. Introduction

Recent studies indicate that residential and commercial buildings account for over 75% of electricity consumption in the United States [2]. As a result, designing new "green" buildings and retrofitting existing buildings with green technologies has become both an important research challenge and societal need. In the residential sector, many techniques exist to reduce either a home's energy footprint or its energy bill. For instance, smart buildings may use motion sensors to track occupants and opportunistically disconnect *loads*[1] in empty rooms [11]. Alternatively, consumers may participate in automated demand response programs increasingly offered by electric utilities, which automatically turn off home appliances when the demand for electricity is high [10]. These intelligent load management schemes reduce a home's energy footprint and its bill by automatically disconnecting loads from power when necessary or convenient. This paper focuses on an intelligent load management scheme for flattening household electricity usage or demand.

Flattening demand implies reducing the difference between the peaks and troughs in a home's electricity usage, thereby creating a flatter usage pattern that lessens the deviation from the average usage. Demand flattening has the potential to benefit residential consumers as the electric grid becomes smarter and more efficient, since peak demands have a disproportionate affect on grid capital and operational costs, including transmission, generation, and fuel costs. For instance, demand flattening significantly reduces transmission and distribution losses, which account for nearly half (47%) of residential energy consumption [3], since these losses are proportional to the square of current.

To incentivize demand flattening, utilities are transitioning from flat pricing models to variable time-of-use or peak-load models [4], [5], [8], [17]. Since the marginal cost to generate electricity rises as demand increases, utilities are beginning to add surcharges to bills based on a consumer's peak usage. For example, a utility may determine the bill, in part, based on a customer's largest half-hour of electricity demand within a day, regardless of the total day's energy consumption. The new electricity pricing models provide consumers strong incentives to regulate not only their total energy consumption, but also their consumption profile. In particular, these new pricing models incentivize customers to lower their peak consumption by flattening their usage.

Unfortunately, while conceptually simple—to control its demand, a home need only decide when to disconnect its loads—intelligent load management has proven difficult to implement in practice. One reason is that disconnecting loads requires active consumer involvement during peak periods, such as turning off unnecessary lights, programming a thermostat, or postponing washing clothes. Prior studies have shown that compelling consumers to change their household routines is challenging [9]. While providing occupants real-time feedback of their power consumption may initially incentivize them to reduce their usage, once the novelty wears off occupants typically revert to their previous habits. Even for consumers that wish to actively manage their load, choosing which loads to disconnect and when is a complex decision that must be continuously re-evaluated based on information that is constantly changing. To address the problem, we have designed *SmartCap*, a system for automatically monitoring and controlling household loads.

As a key step in SmartCap's design, this paper studies the extent to which homes are able to flatten their home electricity demand without affecting home occupants or requiring their active involvement. We explore the impact of modifying *background* electrical loads that are completely transparent to home occupants and have no impact on their perceived comfort. While the vast majority of electrical loads in homes are *interactive* and have little scheduling flexibility

[1] We use the term load throughout the paper to refer to any appliance or device in the home that draws electricity.

(lights, TVs, microwaves, etc.), a substantial portion of home electricity demand derives from loads with some limited flexibility. These flexible loads, such as air conditioners (A/Cs), refrigerators, freezers, dehumidifiers, and heaters, typically operate in the background: while the result of their power draw is readily apparent, e.g., a comfortable room temperature and frozen food, *when* they draw power and the *magnitude* of this power draw is not important. Note that flattening demand is distinct from, and orthogonal to, conservation efforts that reduce total energy consumption over long periods. Instead of reducing total energy usage, flattening demand redistributes consumption by shifting load to decrease demand peaks while filling in troughs. A goal of our work is to quantify when and how much demand flattening is possible from background loads.

We hypothesize that homes are capable of flattening electricity demand during peak load times by intelligently scheduling only background loads. To evaluate our hypothesis, we analyze power data gathered from a real home at outlets, switches, and panels over three months. Our data shows that while background loads account for under 10% of the loads on a typical summer day, they consume nearly 60% of the energy. SmartCap's load scheduler flattens demand by scheduling background loads according to a Least Slack First (LSF) policy, inspired by the Earliest Deadline First algorithm in computing systems, where slack is a measure of how long each background load is able to remain off without affecting its objective, e.g., maintaining an environmental setpoint or fully charging a battery. We evaluate SmartCap by simulating background load scheduling using data from our home deployment. We also implement SmartCap in a smart home testbed we have built, which uses similar background loads as our home. We leverage our testbed to experiment with SmartCap using real appliances. As an example of our results, we show that LSF decreases the average absolute deviation from the mean power (a measure of flatness) by over 20% for all 4-hour periods (over the 82 day period) where the deviation is greater than 400W.

II. Background and Problem Formulation

The focal point of SmartCap's architecture is an intelligent smart home gateway. The home gateway serves as the interface between a smart home and the smart grid. As shown in Figure 1, the gateway receives information from multiple potential sources, including real-time electricity prices and demand-response signals from the grid, generation data from on-site renewables, and consumption data from each household load. The gateway's data sources inform its load scheduling policy. This policy determines *which* loads to power and *when* by issuing actuation commands to loads. While we focus on the problem of scheduling background loads to flatten demand without affecting occupants, our home gateway is capable of implementing scheduling policies with other objectives, such as ensuring home power

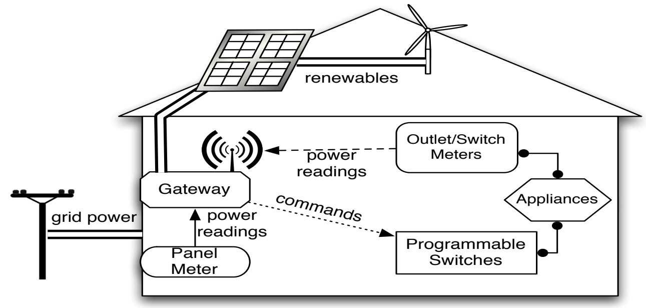

Figure 1. A graphical depiction of a SmartCap-enabled home.

demands are always less than supply when using intermittent on-site renewables [20]. SmartCap's architecture depends on loads that expose programmatic control to turn them on and off. While today's "dumb" appliances generally do not expose such remote actuation capabilities, utilities are currently testing such smart appliances for demand response initiatives [10]. As appliances begin to allow remote actuation at finer granularities, advanced techniques for controlling power will be possible. Given current standards for remote actuation, connecting loads to external programmable switches and outlets using home automation protocols, such as X10 or Insteon, is sufficient in many cases to provide programmatic load actuation in today's appliances. We currently use Insteon-enabled outlets and switches in our home deployment and testbed [12].

We divide electrical loads into two broad groups: interactive and background loads. Household occupants directly control interactive loads by toggling switches, and actively observe their behavior; examples include lights, TVs, computers, microwaves, and vacuums. The vast majority of household loads are interactive. Our LSF scheduling policy assumes that only the occupants are able to control interactive loads. In contrast, household occupants do not directly control background loads, and only passively observe their behavior; examples include refrigerators, dehumidifiers, and A/Cs. As long as these loads satisfy occupant expectations, e.g., a target temperature or humidity level, their usage pattern is neither important nor noticeable. SmartCap monitors background loads and controls when they consume power. We view transparently flattening demand from background loads as an important prerequisite in satisfying many other demand-side scheduling objectives. While disconnecting interactive loads may be necessary at certain times to strictly cap power usage, scheduling background loads without affecting occupants should always be the first priority under constraint. Note that SmartCap's architecture explicitly does not permit utilities to monitor or control household loads, since such capabilities represent an invasion of privacy [15].

III. Load Analysis and Observations

To study the extent to which scheduling background loads is able to flatten demand, we collect fine-grained power data

Load	*Peak*	*Average*	*Quantity*
Refrigerator	456W	74W	1
Freezer	437W	82W	1
HRV	1129W	24W	1
Dehumidifier	505W	371W	1
Main A/C	1046W	305W	1
Bedroom A/C 1	571W	280W	1
Bedroom A/C 2	571W	141W	1
Background	**4715W**	**1277W**	7
Interactive	**9963W**	**887W**	85

Table I
IN THE SUMMER, BACKGROUND LOADS IN OUR HOME ACCOUNT FOR 59% OF THE TOTAL ENERGY CONSUMPTION.

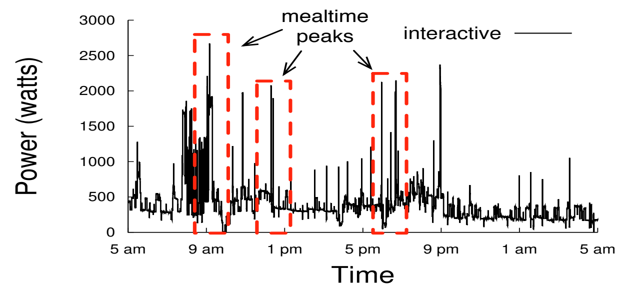

Figure 2. The power consumption of interactive loads is highly variable throughout the day. As expected, peak power consumption occurs around mealtimes in the morning, early afternoon, and early evening.

from a real home that houses three occupants. We have collected the home's aggregate power for the last 12 months and power at each outlet and switch for the past 82 days. Since our monitoring did not affect the occupants' daily routine, our data reveals realistic home power usage patterns over the monitoring period. Our home deployment continuously gathers power usage data for the entire home every second and 30 individual outlet loads every few minutes; our prototype maintains a record of the on-off state of 30 of the home's wall switches at every instant in time. SmartCap's gateway is also able to remotely (and programmatically) control the home's outlets and wall switches. More details about our home deployment are available in prior work [12].

A. *Interactive vs. Background Loads*

To quantify the potential benefits of scheduling background loads, we separate the power consumption of background loads from that of interactive loads. In our prototype home, we monitor seven background loads at outlets: a refrigerator, a freezer, a dehumidifier, three window air conditioning units (A/Cs), and a heat recovery ventilation (HRV) system. By contrast, we estimate that the home used 85 distinct interactive loads over the past year. Thus, SmartCap does not attempt to schedule the vast majority of household loads, since it would affect the home's occupants.

Interactive loads that we do not schedule include lights, entertainment appliances (e.g., TV, cable box, gaming console), computing equipment (e.g., routers, laptops, desktops), kitchen appliances (e.g., microwave, toaster oven, espresso maker, garbage disposal), and miscellaneous devices (e.g., clocks, vacuums, hair dryers). In most cases, disconnecting any of these loads from power when in use is readily apparent to occupants. We also group clothes dryers, washing machines, and dishwashers with interactive loads. While we could schedule the start time of these appliances, we do not include them because adjusting the start time affects occupants. To see why, consider that a scheduler may be able to decide when an appliance executes, but occupants must ultimately initialize the appliance, e.g., fill it with clothes or dishes, before its scheduled start time. Changing the start time may force occupants to initialize the appliance at an inopportune time. Further, for clothes dryers and washing machines, their operation is often pipelined, with households washing multiple laundry loads back-to-back.

Observation #1: While background loads comprise 7.5% of the total loads over our monitoring period, they account for 59% of the average energy consumption. Table I shows the peak and average power consumption for each background load we monitor during a representative week in the summer, as well as the peak and average power consumption for all background and interactive loads. During this week, background loads consume 209 kWh, while interactive loads consume 146 kWh. The three window A/C units significantly increase the fraction of energy consumed by background loads, since each A/C draws between 400W and 1kW when the compressor is on. On hot days, the compressor may run as much as half the day, depending on the comfort level the occupants desire. Note that during the winter the A/Cs do not run, since the home uses a gas furnace for heat. As a result, background load is lower in the winter. In this case, the duct heater for the HRV system, which heats incoming air from the outside, dominates background energy consumption, accounting for 70% of the total, while the refrigerator, freezer, and dehumidifier account for the remaining 30%. Below, we highlight other observations from our home's data that influences our approach to scheduling.

B. *Interactive Variability*

Observation #2: The power consumption of interactive loads varies due to the actions of occupants throughout the day, and is not readily predictable. Figure 2 highlights this point by showing the power consumption of the interactive loads in isolation on a typical day. Additionally, Figure 3 shows consumption patterns for four interactive loads. Notice that the power draws of these loads vary considerably throughout the day, with the peak periods occurring during the morning between 6am and 10am and in the early evening between 5pm and 9pm. These periods coincide with food preparation and are partially the result of using high-power kitchen appliances, such as a coffee pot, garbage disposal, microwave, dishwasher, or toaster oven. During the night, the minimum steady state power consumption is roughly

200W, while during the morning and evening it frequently rises above 2kW for frequent short periods.

The kitchen appliances tend to induce peaks by using large amounts of power for relatively short time periods, such as the coffee pot in Figure 3. Our observation also holds for meal preparation at breakfast, lunch, and dinner. Accurately predicting the power consumption of interactive loads at fine time scales is difficult. While the home's occupants typically eat dinner between 4pm and 8pm, if and when they use a microwave, toaster oven, dishwasher, or garbage disposal is highly variable during this four hour time window each day. Additionally, the occupants have flexible work schedules, and often work from home during the day—on this day one of the occupants ate lunch at home, which accounts for the spike in power around noon. Since interactive loads are not readily predictable, our scheduler must be able to react to drastic and sudden changes in their power consumption.

C. Background Variability

Observation #3: The operating period of background loads varies due to both environmental conditions and external events, and is also not readily predictable. Figure 4 highlights the point by graphing the power consumption of four of the background loads we monitor. Each background load is clearly periodic: it alternates between distinct 'on' and 'off' states. While it is possible to design these loads with variable drive controllers, all the background loads in our home use simple on-off controllers that toggle between an on and off state [18]. In this case, the on-off periodicity is a result of each background load maintaining an environmental *setpoint*: in this example, the refrigerator and freezer maintain their internal temperature within a fixed guardband, the dehumidifier maintains a humidity level within a fixed guardband, and the HRV heats outside air to a pre-specified temperature. The *guardband* defines the acceptable maximum and minimum levels for the load's target environmental metric. Common household loads use simple control loops to stay within the guardband. For example, when the load's metric reaches a maximum allowable value, the load turns on until the metric reaches a minimum value, at which point the load turns off.

Since environmental conditions vary, neither the length nor the magnitude of a load's on-off period is entirely regular. To illustrate, the figure shows that the refrigerator (upper-right) and freezer (upper-left) exhibit longer on periods in the early evening between 5pm and 9pm, along with some transient usage spikes. In both cases, the longer on periods are the result of the occupants opening the refrigerator and freezer doors, which increases the internal temperature and causes them to turn on their compressors to lower the temperature. Tasks other than maintaining temperature also contribute to the transient spikes in power consumption. For example, both the refrigerator and freezer power multiple 60W incandescent light bulbs when the door is open and also

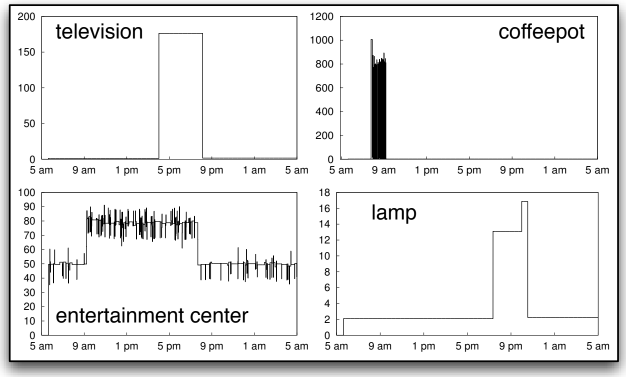

Figure 3. Power data for example interactive loads. Occupant behavior, which is not readily predictable, determines when these loads draw power.

periodically make ice; the refrigerator also cools a separate freezer compartment. The refrigerator exhibits a much more irregular consumption pattern, since it resides in the kitchen and the occupants open its door more frequently than the basement freezer. The HRV and dehumidifier exhibit irregular periods for similar reasons.

The dehumidifier's operating cycle dictates that it runs until it reaches a setpoint humidity—in our case 50%—or until it has run for two consecutive hours, at which point it remains off for 2 hours to cool down. Thus, on hot and humid summer days, the dehumidifier will run for 2 hours every 4 hours if it cannot reach its setpoint humidity, and consume a significant fraction of power (1.8 kWh). On moderately humid days, the dehumidifier will come on and off according to its setpoint humidity, causing an irregular on-off period. On this day, the environmental humidity was high, so the dehumidifier ran regularly. Similar to the refrigerator/freezer, the window unit A/Cs exhibit irregular periods based on changing outdoor temperatures and the frequency with which exterior doors open and close. While some environmental factors may be partially predictable, such as temperature or humidity, interactive events such as doors opening and closing also affect the period and power consumption of background loads. Thus, scheduling background loads must take into account these difficult to predict changes in their periodicity.

IV. Load Scheduler

SmartCap's background load scheduler leverages the well-known concept of *slack*, which quantifies the extent to which a scheduler is able to advance, defer, raise, or lower a load's power consumption without affecting its operational goal [6], [13], [14], [19]. Before detailing the LSF algorithm, we first discuss different types of load controllers to understand the available dimensions of scheduling freedom.

A. Load Controllers

Simple on-off controllers encompass the vast majority of controllers found in residential loads, since they are

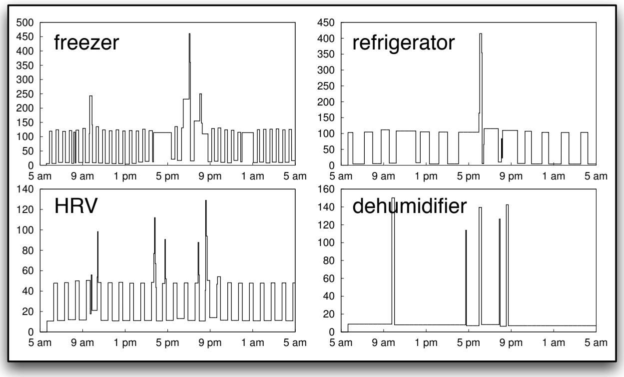

Figure 4. Power signatures for four background loads in our home. The on-off period varies with environmental conditions, and is not regular.

cheap and reliable. As discussed earlier, on-off controllers often maintain an environmental metric, e.g., temperature or humidity, within a specified guardband. For these loads, slack arises from the fact that the load is able to remain off until its metric reaches the guardband's maximum (or minimum) value, at which point the load must turn on. In effect, these loads indirectly store power in their contained environment by increasing (or decreasing) a target metric, which then slowly decreases (or increases) due to leakage with the outside environment. On-off controllers are also commonly driven by timers, which dictate fixed-length on-off periods. While a scheduler is able to advance or defer when these loads turn on or off, as long as they do not violate their guardband or fixed-length on-off period, it is not able to raise or lower power consumption when the loads are on.

Battery chargers are another example of a load with slack, since they are capable of raising or lowering their power consumption by adjusting the charging rate. While most household batteries are small, e.g., phones, laptops, and tablets, the emergence of plug-in electric vehicles (EVs) is poised to introduce a large load with substantial slack to homes. EVs that plug into standard 120V/15A outlets are able to charge at a rate of up to 1.8kW, while a dual-pole 240V/30A circuit that uses both legs of a home's split-phase input power is able to charge at a rate of up to 7.2kW. In either case, advanced chargers are capable of varying the rate of charge up to these maximums. For battery chargers, the primary scheduling constraint is fully charging the battery over some duration, or charging to an acceptable capacity,

While not present in our prototype, variable drive controllers are capable of raising and lowering their power consumption when on. These controllers offer clear benefits over on-off controllers, but they are typically not found in household appliances due to cost and reliability issues. As a result, our experiments do not study their impact.

B. Scheduler

We define a load's slack at any time t as the remaining length of time the load can be off, i.e., disconnected from power, without assuring that it will violate its objective. For a load that maintains an environmental condition with an on-off controller, it must turn on when its environmental metric reaches a guardband boundary. For a battery charger, it must turn on when only the maximum charging rate over the remaining plug-in duration is sufficient to fully charge the battery, or to charge it to an acceptable capacity. We define slack in units of time, rather than energy as in [19], only for ease of exposition—slack time is proportional to slack energy for stable load and environmental conditions. We assume each load is able to maintain an estimate of its remaining slack time based on its current power state and by monitoring the state of its internal and external environment. As shown in Figure 5, slack estimates may change over time based on both the load's power state—when the load is off slack increases—and environmental conditions, such as a refrigerator door opening or the humidity increasing. Since these changes in slack may be unpredictable, our scheduler is reactive and online, continually adjusting which loads receive power based on their available slack. Finally, we assume that our gateway is able to query the slack of each load at any time using simple models as in [19].

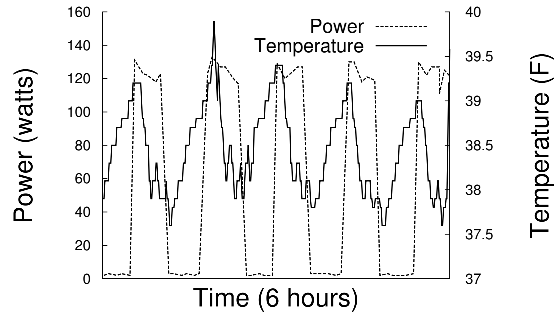

Figure 5. A depiction of slack in our refrigerator's simple on-off control loop. The compressor turns on once the internal temperature reaches an upper threshold, and turns off once it reaches a lower threshold.

Before describing our scheduler, we first illustrate a simple example using ideal background loads with well-defined on and off periods in isolation, and without uncontrollable interactive loads. The illustration demonstrates how shifting power usage is able to flatten demand. Figure 6(a) depicts an extreme example, where the slack for three window A/C units that draw 1kW when on dictates that they must turn on for 15 minutes anytime within each hour to maintain their respective setpoint temperatures. In the worst case, without any scheduling, these units may be nearly synchronized and cause power usage to reach 3kW for close to 15 minutes over the hour, while drawing 0W for the remaining 45 minutes. In the best case, with appropriate scheduling, it is possible to shift the on periods such that only a single A/C is on at any given time, resulting in a peak usage of only 1kW (Figure 6(b)); since the on periods of the A/Cs interleave with room to spare, we are able to perfectly flatten demand.

To quantify flattening over an interval, we use the *average*

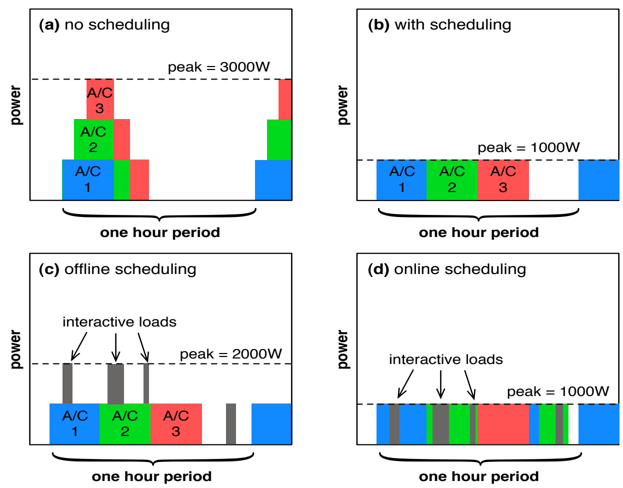

Figure 6. A background load scheduler is capable of flattening demand, but must account unpredictable interactive and background loads.

absolute deviation from the mean power, which is an average of the absolute difference between power at every time t and the average power. We use this metric instead of the standard deviation simply because it is more intuitive; standard deviation exhibits the same trends but is greater than or equal to our metric. The magnitude of the deviation quantifies how much demand varies; a lower deviation indicates flatter demand and a better schedule. In our example, the worst-case no scheduling scenario has a deviation of 1125W from the mean power, while the best-case scenario has a deviation of 375W due to 15 minutes of no power consumption at the end of the period. In this scenario, interleaving the A/Cs results in a 3x reduction in the deviation and, thus, a significantly flatter demand profile.

As noted in prior work [6], [13], the scheduling problem for ideal background loads with regular known on-off periods distills to a simple offline optimization problem in the absence of interactive loads. Figure 6(c) demonstrates how interactive loads alter scheduling by inserting into our previous example four 5 to 15 minute peaks of 1000W during the hour-long period, as could be expected from heating up food in a microwave. Even though A/Cs have enough slack within the hour to defer their power consumption whenever the microwave turns on (Figure 6(d)), an algorithm that determines the schedule in advance will not know about these microwave events. While this is a simple idealized example, it illustrates that load scheduling in the presence of unpredictable interactive loads is an online, and hence heuristic, process. Sudden and unpredictable changes to a load's slack, such as from opening doors or changes in weather, introduce similar issues that warrant an online approach. As we discuss in Section VI, and in contrast to Figure 6(a) and (b), we find that scheduling background loads is most advantageous during "peaky" periods with many short, but high power, interactive loads.

SmartCap's scheduler executes every interval T to determine which background loads receive power (and how much for the battery charger). In our simulator and testbed, we choose T's length to be significantly less (one minute) than the typical on-off periods of our background loads; the setting also ensures that background loads are not quickly turned on and off, which may degrade their reliability. We assume that once a load's slack reaches zero, the scheduler must provide it the necessary power regardless of the increase in peak usage. We call our basic load scheduling policy *Least Slack First* (LSF), since it supplies power to loads in ascending order of their current slack value. Thus, loads with a lower slack have a higher priority. LSF is a direct adaption of the Earliest Deadline First (EDF) scheduling policy common in real-time operating systems. We combine LSF with a target capacity threshold to determine how many loads to power, and how much power to supply to battery chargers. Once the sum of the background loads' power usage reaches the capacity threshold, the scheduler stops powering additional background loads. Figure 7 depicts how LSF scheduling flattens demand for a real power signal, assuming three A/Cs turn on near each other as in Figure 6. As in our example, LSF flattens the demand profile by interleaving the on periods.

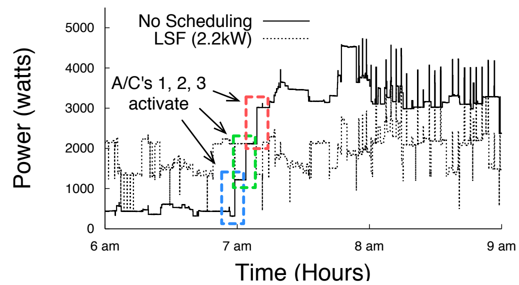

Figure 7. Example of how LSF flattens demand.

Our experiments use an adaptive threshold based on an exponentially weighted moving average of the home's power consumption over the previous hour. Setting the capacity threshold presents a trade-off. A threshold too low causes the scheduler to defer too many loads, resulting in their slack values approaching zero in tandem. This induces large peaks by ultimately forcing the scheduler into simultaneously powering many loads with zero slack. A threshold too high causes the scheduler to power too many background loads at a time, resulting in a peak that is higher than necessary.

V. Prototype: Design and Implementation

Our SmartCap home deployment is in an average 3-bedroom, 2-bath house with 1700 sq. ft. and a total of 8 rooms across three floors, including a basement. Since the prototype is a real home with three occupants that went about their daily routines during the monitoring period, our data reflects real-life home usage patterns. The home does not have central air, and its furnace and water heater use

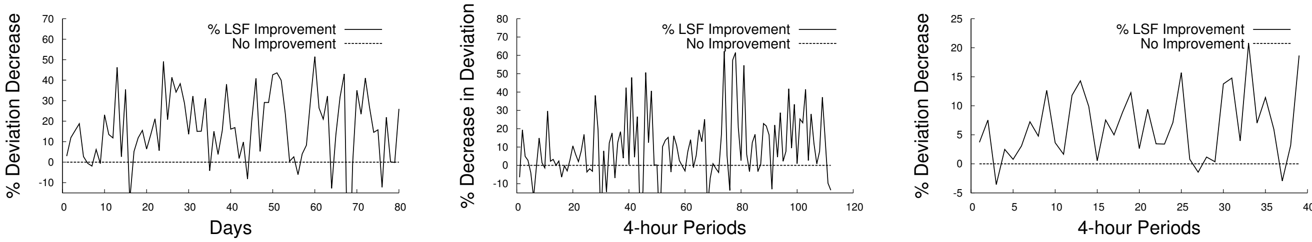

(a) Each Day (b) Mid-range 4-hour Periods (c) High-range 4-hour Periods

Figure 8. LSF decreases the absolute average deviation from the mean power (with no scheduling) on the vast majority of days (91%), as well as over peak 4-hour periods with mid-range and high-range deviations.

natural gas, which removes three potentially large consumers of electricity. During the summer, the occupants use three window A/C units to cool the home—one large unit in the living room and a smaller unit in each upstairs bedroom.

We provide a brief summary of our SmartCap deployment. A TED 5000 measures power consumption for the entire home every second using meter-like measurements of the wires supplying grid power to the home's main circuit breaker panel. The TED specification claims accuracy within 2%; we found the TED to be within 1% of the utility power readings during the monitoring period. We use Insteon-enabled switches to monitor and control loads; Insteon is a common, commercially-available home automation protocol that uses power line communication. In particular, we use the Insteon iMeter Solo to monitor power at background load outlets, and the Insteon ApplianceLinc to control power to our background loads from our gateway. Our gateway connects to an Insteon Power Line Modem (PLM), which is able to inject Insteon commands and listen for responses over the home's power lines. The gateway both polls the iMeters for their power usage and issues on-off commands to the appliances through the PLM. For background load scheduling, SmartCap only requires power data for the whole home and at the seven background loads. However, our prototype is capable of remotely monitoring and controlling each outlet and wall switch in the home [12]. To monitor environmental metrics and compute slack, we deploy eight temperature and humidity sensors inside or near each background load, as well as outside, using an Oregon Scientific WMR200A weather station.

In addition to our in-home SmartCap deployment, we also setup a smart home testbed to mimic our home's background loads. The testbed enables us to perform repeatable experiments, such that we do not disturb home occupants. It uses the same SmartCap system as our real deployment: Insteon-enabled power meters and switches to monitor and control background loads. The background loads include a humidifier, dehumidifier, multiple electric heaters, a freezer, and a refrigerator – we use heaters rather than A/Cs, since our testbed resides within a window-less room and A/Cs require outside drainage. Since we use external load control switches that are not integrated with the appliance to connect and disconnect power, we use appliances that remember and restart in the same setting after a power outage. For experiments, we are able to replay traces using our home data both with and without LSF scheduling.

VI. Evaluation

We evaluate LSF in simulation and in our smart home testbed to explore its performance in realistic settings. Our simulator, written in Java, uses input traces of household load events to simulate background load scheduling using LSF. Each load event corresponds to a change in the power level for a single load. The simulator also associates both a maximum and minimum slack value with each background load every period, which includes a single off interval and its subsequent on interval. At each period boundary, the simulator assumes the load is at its maximum slack value if it has just transitioned to the off state, and assumes the load has zero slack if it has just transitioned to the on state.

The simulator uses a simple linear model for computing per-period slack: when a load is on its slack increases linearly, and when it is off it decreases linearly. To always ensure that the load reaches its maximum slack by the end of each period, the simulator determines the slope of the linear increase or decrease in slack using the ratio of the on and off durations for the current period. Note that, due to environmental changes, each background load may exhibit different period durations, as well as per-period on and off durations, throughout the day. In practice, SmartCap may use an environmental model to compute slack in real time; linear models tend to perform well, as Figure 5 demonstrates for the refrigerator and its inside temperature. Since we automatically generate input traces from the home's power data, our per-period slack computation is an indirect way of accounting for environmental changes in simulation.

Since our scheduler only controls background loads, our input traces represent all interactive loads as a single load with many frequent load events. To get power readings for the interactive loads, we subtract each background load from the home's aggregate power consumption. Note that since we collect aggregate power consumption every second, our trace includes a new event nearly every second to represent the changing consumption of the interactive loads.

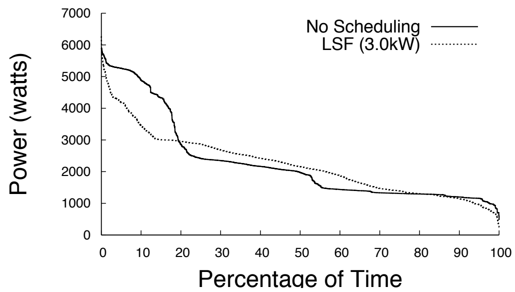

Figure 9. Load duration curves for a typical summer day with and without scheduling when using an electric vehicle.

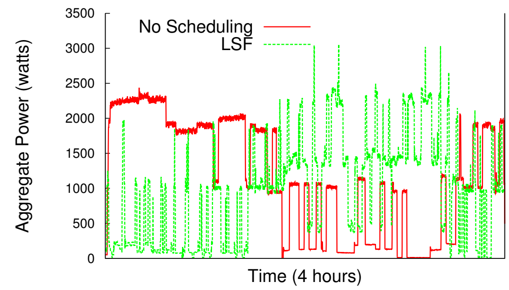

Figure 10. Power usage with and without LSF scheduling using using our smart home testbed with real background loads over a 4-hour period.

A. Simulation Results

We first evaluate LSF for flattening peak power usage in our home deployment. We focus on the last 82 days during the summer. Flattening is most important during summer months, since peak demands typically occur in these months [1]. Figure 8 shows the percentage decrease in average absolute deviation from the mean power using LSF scheduling for different periods. Recall from Section 4 that we use the average absolute deviation to quantify the flatness of the demand profile. Figure 8(a) plots the percentage decrease in deviation over each day, and demonstrates that LSF flattens the profile on over 91% of days, resulting in a 16% flatter profile on average. LSF does not flatten the profile on 9% of days, since those days already have a low deviation without scheduling. On 33% of days, LSF decreases the deviation by more than 20%. These results are significant, since each day includes long periods of relatively little activity, e.g., all night, where the average deviation is not high due to minimal occupant activity. Despite the long periods of inactivity that occur each day, LSF is still able to flatten the day-long demand profile.

We also examine how LSF performs for shorter 4-hour intervals that correspond to peak usage times, since these are the periods where demand flattening is most important. We divide 4-hour periods throughout the summer by the magnitude of their average absolute deviation (or "peakiness"). We find that LSF does not provide much improvement ($<3\%$) for periods that do not exhibit "peaky" behavior, since the demand profile is already flat. We find that over 69% of the 4-hour periods throughout the 82 days have average deviations less than 400W; these periods generally correspond to nighttime or when the home is unoccupied. The remaining 31% of the periods exhibit deviations from 400W to 1000W (22%) and over 1000W (9%). Figure 8(b) shows that LSF works well for the mid-range (400W-1000W) and high-range ($>$1000W) 4-hour periods, decreasing the respective average deviations by 23% and 21%, on average.

The data indicates that many flat 4-hour periods exist throughout our trace, which suggests that most of LSF's improvement stems from scheduling background loads around interactive loads that cause brief, but significant, power peaks. If the background loads themselves interleaved to cause significant peaks in power, we would expect more improvement during periods with few interactive loads, e.g., nighttime. Since our home has many background loads that operate based on different environmental conditions, they rarely all turn on simultaneously. Thus, without LSF, the background loads already exhibit a great deal of statistical multiplexing, and there is little LSF can do to flatten their peak usage. Our results also indicate that LSF works well during "peaky" periods, which typically occur during peak demand periods, where the average deviation is high.

B. Impact of Electric Vehicles

We also studied the impact of EVs on LSF's ability to flatten peaks. Today's grid was not provisioned for the increased power consumption from widespread EV adoption. As a result, the grid must either add capacity or use better load scheduling, e.g., through new pricing models, to force EVs to multiplex their charging over time. For instance, in our home, charging an EV on a typical summer day increases the home's total power consumption by 52%. SmartCap and LSF represent a possible avenue for flattening demand with EVs. As in the simulations above, we use data from our prototype home, but add an EV charger based on the Chevy Volt, with a battery capacity of 16kWh plugged in at night between 7pm and 6am that takes 5 hours to charge.

Figure 9 shows the results for an average summer day by plotting a load duration curve both without scheduling and using our LSF scheduler. Load duration curves are a common method for visualizing the flatness of power distributions. The curve shows the percentage of time on the x-axis during the day that electricity demand was at the corresponding power value on the y-axis. An ideal load duration graph is a completely horizontal line at the average power usage. On this day, LSF reduces the average absolute deviation by 22%. In particular, LSF reduces the peak time periods where demand is highest (on the left side of the graph) significantly, and shifts their power consumption across many of the lower power periods throughout the day.

C. Testbed Results

Finally, to demonstrate LSF's performance in a realistic setting we use our smart home testbed. Figure 10 shows the power usage, as measured by our Insteon power meters, for a representative 4-hour period. We use data from our home on June 15th from 2pm to 6pm to replicate the same sequence of background load on and off periods in our testbed. As discussed earlier, our gateway sends commands to Insteon ApplianceLincs to connect and disconnect background loads from power. The experiment demonstrates how LSF shifts the power usage of the background loads forward to compensate for the interactive loads early in the period. As a result, on this day, LSF decreases the absolute average deviation from the mean power by 23%.

VII. Related Work

Increasing the penetration of demand-side load management in residential settings is a key goal of smart grids. Thus, SmartCap's general architecture, which includes the home gateway, an array of real-time power meters, and programmable switches, is similar to other proposed architectures for programmatically regulating home electricity demand [6], [13], [16], [7]. While space constraints preclude a full survey of prior work, past approaches focus on using these architectures for a range of scheduling objectives, such as reducing total consumption, reducing costs based on variable prices, or varying usage to match renewable generation or make use of a battery. Our work differs in its focus on flattening demand without affecting occupants by scheduling background loads. We do not explore scheduling to satisfy other objectives, since it requires disturbing occupants by periodically disconnecting interactive loads.

Prior work also recognizes that loads with on-off controllers present a unique scheduling opportunity [6], [13], [19]. For instance, Taneja et al. [19] present an algorithm for scheduling a single refrigerator with slack that operates off wind power. Both Keshav and Rosenberg [13] and Bakker et al. [6] present offline optimization approaches to scheduling multiple on-off loads in isolation, assuming that loads with on-off controllers have well-known and regular periods. In contrast, our work quantifies the benefits of scheduling background loads in a real home. Data from our home reveals that background loads do not exhibit regular periods, due to environmental changes, while interactive loads are difficult to predict in advance. As a result, we eschew offline optimization scheduling algorithms in favor of an online approach that uses each load's current slack as a heuristic to determine its priority at any time.

VIII. Conclusion

Demand-side management is challenging, since it often requires active, and often burdensome, consumer involvement. Forcing people to think about how they use power is simply not effective in encouraging broader adoption of demand-side management. Thus, we focus on quantifying the benefits of scheduling transparent background loads. We show that LSF is able to flatten household demand over each day, despite long periods of inactivity at night. Importantly, we also show that LSF is useful over shorter (4-hour) peak usage periods, where demand is "peaky" and deviates frequently and significantly from the average.

Acknowledgements. This work is supported by NSF grants CNS-1143655, CNS-0916577, CNS-0855128, CNS-0834243, and CNS-0845349.

References

[1] U.S. Energy Information Adminstration. Annual Energy Review. www.eia.gov/totalenergy/data/annual/pdf/aer.pdf, 2009.
[2] U.S. Department of Energy. Building Energy Data Book. buildingsdatabook.eere.energy.gov/, 2010.
[3] U.S. Department of Energy. Building Energy Data Book. buildingsdatabook.eere.energy.gov/ChapterIntro2.aspx, 2010.
[4] Dynamic Pricing and Smart Grid. www.cntenergy.org/pricing/, 2011.
[5] Ontario Energy Board. www.ontarioenergyboard.ca/OEB/Consumers, 2011.
[6] V. Bakker, M. Bosman, A. Molderink, J. Hurink, and G. Smit. Demand Side Load Management Using a Three Step Optimization Methodology. In *SmartGridComm*, October 2010.
[7] A. Bar-Noy, Y. Feng, M. Johnson, and O. Liu. When to Reap and When to Sow - Lowering Peak Usage With Realistic Batteries. In *WEA*, May 2008.
[8] A. Bar-Noy, M. Johnson, and O. Liu. Peak Shaving Through Resource Buffering. In *WAOA*, September 2008.
[9] S. Davidoff, M. Lee, C. Yiu, J. Zimmerman, and A. Dey. Principles of Smart Home Control. In *UbiComp*, September 2006.
[10] D. H. et al. Grid Friendly Appliance Project. *Pacific Northwest GridWise Testbed Demonstration*, October 2007.
[11] R. Harle and A. Hopper. The Potential of Location-Aware Power Management. In *UbiComp*, September 2008.
[12] D. Irwin, A. Wu, S. Barker, A. Mishra, P. Shenoy, and J. Albrecht. Exploiting Home Automation Protocols for Load Monitoring in Smart Buildings. In *BuildSys*, November 2011.
[13] S. Keshav and C. Rosenberg. Direct Adaptive Control of Electricity Demand. Technical Report CS-2010-17, University of Waterloo, September 2010.
[14] S. Keshav and C. Rosenberg. How Internet Concepts and Technologies Can Help Green and Smarten the Electrical Grid. In *GreenNet*, August 2010.
[15] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin. Private Memoirs of a Smart Meter. In *BuildSys*, November 2010.
[16] A. Schülke, J. Bauknecht, and J. Häussler. Power Demand Shifting with Smart Consumers: A Platform for Power Grid friendly Consumption Control Strategies. In *SmartGridComm*, October 2010.
[17] T. Sickinger. PGE to Test Peak-pricing for Electricity. *The Oregonian*, September 22, 2009.
[18] L. Sonneborn and F. V. Vleck. The Bang-Bang Principle for Linear Control Systems. *SIAM Control*, Series A(2), 1965.
[19] J. Taneja, D. Culler, and P. Dutta. Towards Cooperative Grids: Sensor/Actuator Networks for Renewables Integration. In *SmartGridComm*, October 2010.
[20] T. Zhu, A. Mishra, D. Irwin, N. Sharma, P. Shenoy, and D. Towsley. Efficient Renewable Energy Management for Smart Homes. In *BuildSys*, November 2011.