

SmartCharge: Cutting the Electricity Bill in Smart Homes with Energy Storage

Aditya Mishra, David Irwin,
Prashant Shenoy, and Jim Kurose
University of Massachusetts Amherst
{adityam,irwin,shenoy,kurose}@cs.umass.edu

Ting Zhu
Binghamton University
tzhu@binghamton.edu

ABSTRACT

Market-based electricity pricing provides consumers an opportunity to lower their electric bill by shifting consumption to low price periods. In this paper, we explore how to lower electric bills without requiring consumer involvement using an intelligent charging system, called SmartCharge, and an on-site battery array to store low-cost energy for use during high-cost periods. SmartCharge’s algorithm reduces electricity costs by determining when to switch the home’s power supply between the grid and the battery array. The algorithm leverages a prediction model we develop, which forecasts future demand using statistical machine learning techniques. We evaluate SmartCharge in simulation using data from real homes to quantify its potential to lower bills in a range of scenarios. We show that typical savings today are 10-15%, but increase linearly with rising electricity prices. We also find that SmartCharge deployed at only 22% of 435 homes reduces the aggregate demand peak by 20%. Finally, we analyze SmartCharge’s installation and maintenance costs. Our analysis shows that battery advancements, combined with an expected rise in electricity prices, have the potential to make the return on investment positive for the average home within the next few years.

Categories and Subject Descriptors

J.7 [Computer Applications]: Computers in Other Systems—*Command and control*

General Terms

Design, Measurement, Management

Keywords

Energy, Battery, Electricity, Grid

1. INTRODUCTION

The cost of generating electricity is rising. The average price of electricity for residential consumers in the United States has risen 29% over the past five years [25]. Despite energy-efficiency improvements, residential electricity

demand in the U.S. has increased 49% over the last twenty years, due to a steady rise in the number of household electrical devices. These facts combined with stagnant income growth over the past decade—down 1.9% in inflation-adjusted USD [8]—have resulted in electricity costs consuming a growing share of household budgets. The average home electricity bill now accounts for 2.8% of household income, and has risen by \$300 to \$1,419 per year over the last twenty years (in inflation-adjusted USD) [4]. Since today’s prices do not incorporate negative externalities associated with electricity generation, such as air pollution and climate change, its real cost to society is likely much higher than today’s prices reflect. Studies suggest that recent price and demand increases will continue into the foreseeable future.

Of course, the most direct way for consumers to cut their electricity bill is to simply use less electricity. Unfortunately, as the trends above indicate, rising prices have not yet motivated consumers to conserve power. Another important way to cut bills is to reduce demand peaks, which have a disproportionate affect on generation costs. Peak demands drive both capital expenses—by dictating the number of power plants, transmission lines, and substations—and operational expenses, since “peaking” generators are generally dirtier and costlier to operate than baseload generators [15]. To illustrate the impact of peak demands, Figure 1 shows the marginal cost of operating generators in the southeast U.S., and demonstrates that the marginal cost for generating electricity is non-linear and increases rapidly as utilities move up the dispatch stack to satisfy increasing demand [11]. Peak demands also result in significantly higher transmission losses, since these losses are proportional to the square of current. Thus, even small reductions in peak usage have a significant impact on generation costs. Recent estimates attribute 10-20% of generation costs in the U.S. to servicing only the top 100 hours of peak demand each year [18].

In an attempt to reduce peak demand, many utilities are transitioning from conventional fixed-rate pricing models, which charge a flat fee per kilowatt-hour (kWh), to new market-based schemes, e.g., real-time or time-of-use pricing, which more accurately reflect electricity’s cost by raising and lowering prices during peak and off-peak periods, respectively. For instance, Illinois already requires utilities to provide residential customers the option of using hourly electricity prices based directly on wholesale prices [23], while Ontario charges residential customers based on a time-of-use scheme with three different price tiers (off-, mid-, and on-peak) each day [21]. We envision utilities widening the use of market-based pricing in the future to reduce generation costs, as demands and prices increase.

Unfortunately, market-based electricity pricing places a

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

e-Energy 2012, May 9-11 2012, Madrid, Spain.

Copyright 2012 ACM 978-1-4503-1055-0/12/05 ...\$10.00.




Figure 1: The marginal cost to generate electricity increases as utilities dispatch additional generators to satisfy increasing demand. Data from [11].

significant burden on consumers to continuously monitor prices, and then alter their usage to reduce costs without disrupting normal daily activities. The task is challenging, since most consumers have no idea how much power individual devices consume, and generally do not want to think about or plan their electricity usage. Thus, consumers may not respond appropriately to price changes, and the grid may not gain the cost-saving benefits of peak reduction. Further, as we show in Section 5.2, even if consumers respond appropriately, today’s market-based pricing plans may actually increase grid peaks (and costs) if demand is highly elastic and responsive to price changes. The difficulty in regulating demand may also discourage consumers from opting into market-based pricing plans. For instance, in Illinois, less than one percent of consumers have opted to switch from fixed-rate to market-based pricing [2].

To address the problem, we propose *SmartCharge*, an intelligent charging and discharging system that determines when and how much to store low-cost energy for use during high-cost periods based on expectations of future demand. *SmartCharge*’s primary benefit is that it does not require consumers to alter their electricity usage to reduce their electric bill under market-based pricing plans. Instead, *SmartCharge* reduces costs by determining when to switch a home between using (and storing) grid power and using previously stored power from a battery array. We frame the cost-minimization problem as a linear optimization that leverages knowledge of next-day electricity prices and usage patterns. Since electricity prices are largely set in day-ahead markets [19], next-day prices are well-known. We predict next-day consumption by developing statistical machine learning (ML) to build a model based on important predictive metrics, such as weather, time-of-day, day-of-week, etc.

Our hypothesis is that combining *SmartCharge* with market-based pricing is capable of reducing electricity costs for consumers over the short- and long-term. Over the short-term, consumers save by storing energy during low-cost periods for use during high-cost periods. Over the long-term, as *SmartCharge* penetration increases, average prices will fall due to significant reductions in peak demand. However, as we discuss in Section 5.2, to attain maintain peak reduction at scale using *SmartCharge*, utilities will need to modify today’s market-based electricity pricing plans, which do not properly incentivize energy storage at scale. In evaluating our hypothesis, we make the following contributions:

SmartCharge Design. We detail *SmartCharge*’s architecture, which includes a battery array and charger, DC→AC inverter, and power transfer switch, as well as a gateway server and energy/voltage sensors to monitor home electricity consumption and the battery array’s state of charge.




Figure 2: A depiction of *SmartCharge*’s architecture, including its battery array and charger, DC→AC inverter, power transfer switch, energy/voltage sensors, and gateway server.

We then outline the linear optimization problem the gateway server solves each day to reduce costs by switching the home’s power source between the grid and a battery array. **ML-based Consumption Prediction.** Since solving *SmartCharge*’s optimization problem requires knowledge of next-day consumption, we develop a ML-based prediction technique that learns the unique characteristics of a home’s usage pattern over time. We show that our approach has an average error of 5.75% in a case study of a real home over a 40 day period. Our evaluation shows that solving the optimization using ML-based predictions comes within 8-12% of an oracle with perfect knowledge of next-day consumption.

Implementation and Evaluation. We evaluate *SmartCharge* in both simulation, using power data from real homes and existing market-based residential pricing plans, and with a small-scale prototype using a home UPS system and a few household appliances. Our results show that *SmartCharge* is able to reduce a typical home’s electric bill by 10-15% using realistic battery capacities. We also show that, if widely deployed, *SmartCharge* reduces grid demand peaks by 20%. Finally, we analyze *SmartCharge*’s installation and maintenance costs, and show that recent battery advancements combined with modest (and expected) price increases may make *SmartCharge*’s return on investment positive within the next few years.

2. SMARTCHARGE ARCHITECTURE

Figure 2 depicts *SmartCharge*’s architecture, which utilizes a power transfer switch that is able to toggle the power source for the home’s electrical panel between the grid and a DC→AC inverter connected to a battery array. A gateway server continuously monitors 1) electricity prices via the Internet, 2) household consumption via an in-panel energy monitor, and 3) the battery’s state of charge via voltage sensors. Before the start of each day, the server solves an optimization problem based on the next day’s expected electricity prices, the home’s expected consumption pattern, and the battery array’s capacity and current state of charge, to determine when to switch the home’s power source between the grid and the battery array. The server also determines when to charge the battery array when the home uses grid power. In §6, we provide a detailed estimate of *SmartCharge*’s installation and maintenance costs based on




Figure 3: Example TOU and hourly market-based rate plans in Ontario and Illinois, respectively.

price quotes for widely-available commercial products.

Most utilities still use fixed-rate plans for residential customers that charge a flat fee per kilowatt-hour (kWh) at all times. In the past, market-based pricing plans were not possible, since the simple electromechanical meters installed at homes had to be read manually, e.g., once per month, and were unable to record when homes consumed power. However, utilities are in the process of replacing these old meters with smart meters that enable them to monitor electricity consumption in real time at fine granularities, e.g., every hour or less. As a result, utilities are increasingly experimenting with market-based pricing plans for their residential customers. To cut electricity bills, SmartCharge relies on residential market-based pricing that varies the price of electricity within each day to more accurately reflect its cost. We expect many utilities to offer such plans in the future.

There are multiple variants of market-based pricing. Figure 3 shows rates over a single day for both a time-of-use (TOU) pricing plan used in Ontario, and a real-time pricing plan used in Illinois. TOU plans divide the day into a small number of periods with different rates. The price within each period is known in advance and reset rarely, typically every month or season. For example, the Ontario Electric Board divides the day into four periods (7pm-7am, 7am-11am, 11am-5pm, and 5pm-7pm) and charges either a off-peak-, mid-peak, or on-peak rate (6.2¢/kWh, 9.2¢/kWh, or 10.8 ¢/kWh) each period [21]. The long multi-hour periods and well-known rates enable consumers to plan their usage across reasonable time-scales and adopt low-cost daily routines, e.g., running the dishwasher after 7pm each day. However, while TOU pricing more accurately reflects costs than fixed-rate pricing, it is not truly market-based since actual prices vary continuously based on supply and demand.

TOU pricing is a compromise between fixed-rate pricing and real-time pricing, where prices vary each hour (or less) and reflect the true market price of electricity. Unfortunately, real-time pricing complicates planning. Since prices may change significantly each hour, consumers must continuously monitor prices and adjust their daily routines, which may now have different costs on different days. Illinois was the first U.S. state to require utilities to offer residential consumers the option of using real-time pricing plans. To facilitate planning, Illinois utilities provide simple web pages, e.g., www.powersmartpricing.org/chart, to view next-day prices each evening. While some utilities use real-time prices not known in advance, most utilities use day-ahead market prices, which are set one day in advance. Since utilities purchase most of their electricity in day-ahead markets, e.g., 98% in New York [19], next-day prices are well-known.

SmartCharge works well with both TOU and real-time pricing plans. In either case, SmartCharge solves the op-




Figure 4: Example from January 3rd with and without SmartCharge using Illinois prices from Figure 3.

timization problem detailed in the next section at the end of each day to determine when to switch between grid and battery power to minimize costs, based on next-day prices and expected next-day consumption. The number of periods each day—four in Ontario or twenty-four in Illinois—simply changes a parameter in the optimization’s constraints.

3. SMARTCHARGE ALGORITHM

SmartCharge cuts electricity bills by storing energy during low-cost periods for use during high-cost periods. The total possible savings each day is a function of both the home’s rate plan and its pattern of consumption. Throughout the paper, we use power data from a real home we have monitored for the past two years as a case study to illustrate SmartCharge’s potential benefits. The home is an average 3 bedroom, 2 bath house in Massachusetts with 1700 square feet. To measure electricity, we instrument the home with an eGauge energy meter [10], which installs in the electrical panel by wrapping two 100A current transducers around each leg of the home’s split-leg incoming power. We have monitored the home’s power consumption every second for the past two years. In 2010, the home consumed 8240kWh at a cost of \$1203.53 (or 22.6 kWh/day), while in 2011 it consumed 9732kWh at a cost of \$1339.51 (or 26.7 kWh/day). The costs are near the \$1419 average U.S. home electric bill.

3.1 Potential Benefits

To better understand SmartCharge’s potential for savings, it is useful to consider a worst-case scenario where 100% of the home’s consumption occurs during the day’s highest rate period. Consider our home’s hourly electricity use on January 3rd, 2012, as depicted in Figure 4. On this day, the home consumed 43.7 kWh, primarily due to the occupants running multiple laundry loads after returning from a holiday trip. With Ontario’s TOU plan, if the home had consumed 100% of the day’s power during the 10.8¢/kWh on-peak period, and SmartCharge shifted it all to the 6.2¢/kWh off-peak period, then the maximum savings is 43%, or \$2.01 (from \$4.72 to \$2.71) for the day. Since the home did not consume 100% of its power during the on-peak period, the maximum realizable savings (if we shift all of the on-peak and mid-peak consumption to the off-peak period) is only 30%, a decrease of \$1.14 for the day (from \$3.85 to \$2.71). In practice, battery and inverter inefficiencies, which combined are ~80% efficient, reduce the savings further, to \$0.99 for the day. This per-day savings rate translates to a yearly savings of \$361.35, if the system achieves it every day.

Real-time pricing plans, as in Illinois, offer even more potential for savings, since the difference between the highest and lowest rate is significantly larger than a typical TOU

plan. For example, on August 1st, 2011 in Illinois, the average rate from 2pm-7pm was 10.42¢/kWh, while the average rate from 1am-6am was 2.36¢/kWh. The highest rate of 11.9¢/kWh occurs at 4pm, and is over 5X larger than the lowest rate of 2.3¢/kWh from 2am-5am. In this case, with January 3rd’s consumption pattern and battery/inverter inefficiencies, SmartCharge is still capable of reducing costs by 59%, or \$1.78 (from \$3.02 to \$1.24). However, Figure 4 demonstrates that the actual savings also depend on the on-site storage capacity. In this case, with 12kWh of usable energy storage, SmartCharge is only able to shift five hours of consumption during the highest rate daytime periods to the lowest rate nighttime periods. In particular, there is not enough capacity to store low-cost nighttime energy for use during the mid-price periods. As a result, consumption in the late morning and early evening remains unchanged. With 12kWh of storage capacity, the cost reduction falls to 32%, or \$0.96 (from \$3.02 to \$2.06) for the day.

Of course, home consumption patterns and hourly rates vary each day, which may decrease (or increase) a home’s actual yearly savings. To understand why home consumption patterns are important, consider the following scenario using the Ontario TOU pricing plan. In Ontario, while SmartCharge may fully charge its battery array during the lowest rate period (7pm-7am), it may also consume that stored energy during the day’s first high rate period (7am-11am). If the home expects to consume at least the battery array’s entire usable capacity during the day’s second high rate period (5pm-9pm), it is cost-effective, assuming ideal batteries, to fully charge the batteries during the mid-rate period (11am-5pm) when electricity costs are 17% less than in the high rate period. However, if the home only expects to use 20% of the battery’s capacity during the subsequent high rate period, it is only cost-effective to charge the battery 20% during the mid-rate period, since there will be an opportunity to charge the battery further (for 33% less cost) during the next low-rate period. In this case, charging the battery more than 20% wastes money. Introducing more price tiers, as in real-time markets, complicates the problem further. As a result, we frame the problem of minimizing the daily electricity bill as a linear optimization problem.

3.2 Problem Formulation

While batteries exhibit numerous limitations (e.g., charging rate, capacity), inefficiencies (e.g., energy conversion efficiency, self-discharge), and non-linear relationships (e.g., between capacity, lifetime, depth of discharge, discharge rate, ambient temperature, etc.), SmartCharge’s normal operation places it at the efficient end of these relationships. The system mostly charges the battery once a day during the night, which prevents stratification and extends battery lifetime by limiting the number of charge-discharge cycles. The self-discharge rate of valve-regulated absorbed glass mat (VRLA/AGM) lead-acid batteries (commonly called sealed lead-acid batteries), estimated at 1-3% per month, is insignificant, amounting to no more than \$13 per year for a 12kWh battery array with an average electricity price of 10¢/kWh. Sealed lead-acid batteries are generally 85-95% efficient, while inverters are 90-95% efficient. For SmartCharge’s battery array and inverter, we assume an energy conversion efficiency of 80%, which mirrors the efficiency rating for VRLA/AGM lead-acid batteries in a recent Department of Energy report on energy storage technolo-

gies [22]. Thus, the batteries waste 1W for every 4W they are able to store and re-use. Additionally, depth of discharge (DOD) for sealed lead-acid batteries impacts their lifetime, i.e., the number of charge-discharge cycles, due to the crystallization of lead sulfate on the battery’s metal plates. In our evaluation, we find that a DOD of 45% minimizes battery costs by balancing lifetime with usable storage capacity for a typical battery designed for home photovoltaic (PV) installations, e.g., the Sun Xtender PVX-2580L [24].

The ambient temperature and rate of discharge also have an impact on usable capacity, according to Peukert’s law. To maximize lifetime, we expect SmartCharge installations to reside in a climate-controlled room with a temperature near 25C. Rated capacity is typically based on a C/20 discharge rate, i.e., the rate of discharge necessary to deplete the battery’s capacity in 20 hours. A discharge rate higher or lower than C/20 results in less or more usable capacity, respectively. The home in our case study has averaged near 1kW per hour over the last two years, so a 20kWh battery capacity approaches this rating. As we show in §5, reasonable battery capacities for SmartCharge with a 45% DOD are near or above 20kWh. Finally, sealed lead-acid batteries are capable of fast charging up to a C/3 rate, i.e., charges to full capacity in three hours [17]. In §5, we use a maximum charge rate of C/4 for the usable storage capacity, which translates to a C/8 rate for a battery used at 45% DOD. As we show, faster charging rates are not beneficial, since market-based pricing plans generally offer long low-rate periods for charging at night.

Given the constraints above, we frame SmartCharge’s linear optimization problem as follows. The objective is to minimize a home’s electricity bill using a battery array with a usable capacity (after accounting for its DOD) of C kWh. We divide each day into T discrete intervals of length l from 1 to T . We then denote the power charged to the battery during interval i as s_i , the power discharged from the battery as d_i , and the power consumed from the grid as p_i . We combine both the battery array and inverter inefficiency into a single inefficiency parameter e . Finally, we specify the cost per kWh over the i th interval as c_i , and the amount billed as m_i . Formally, our objective is to minimize $\sum_{i=1}^T m_i$ each day, given the following constraints.

$$s_i \geq 0, \forall i \in [1, T] \quad (1)$$

$$d_i \geq 0, \forall i \in [1, T] \quad (2)$$

$$s_i \leq C/4, \forall i \in [1, T] \quad (3)$$

$$\sum_{t=0}^i d_t \leq e * \sum_{t=0}^i s_t, \forall i \in [1, T] \quad (4)$$

$$\left(\sum_{t=0}^i s_t - \sum_{t=0}^i d_t / e \right) * I \leq C, \forall i \in [1, T] \quad (5)$$

$$m_i = (p_i + s_i - d_i) * I * c_i, \forall i \in [1, T] \quad (6)$$

The first and second constraint ensure the energy charged to, or discharged from, the battery is non-negative. The third constraint limits the battery’s maximum charging rate. The fourth constraint specifies that the power discharged from the battery is never greater than the power charged to the battery multiplied by the inefficiency parameter. The fifth constraint states that the energy stored in the battery

Model	12am-7am	7am-11am	11am-5pm	5pm-7pm	7pm-12am	Average (%)
SVM-Linear	14.77	27.32	46.72	18.49	47.03	29.5
SVM-RBF	22.44	63.77	71.93	17.84	35.01	42.51
SVM-Polynomial	4.74	4.62	6.48	7.99	5.14	5.75

Table 1: Average prediction error (%) over 40 day sample period for SVM with different kernel functions.

array, which is the difference between the energy charged to or discharged from the battery over the previous time intervals, cannot be greater than its capacity. Finally, the sixth constraint defines the price the home pays for energy during the i th interval. The objective and constraints define a linearly constrained optimization problem that is solvable using standard linear programming techniques. SmartCharge solves the problem at the beginning of each day to determine when to switch between grid and battery power. Since the approach uses knowledge of next-day consumption patterns, we next detail statistical machine learning techniques for predicting next-day consumption and quantify their accuracy for our case study home.

4. ML-BASED DEMAND PREDICTION

As discussed in §3, solving SmartCharge’s linear optimization problem requires *a priori* knowledge of next day consumption patterns. A simple approach to predicting consumption is to use past-predicts-future models that assume an interval’s consumption will closely match either that interval’s consumption from the previous day or the prior interval’s consumption. As we show, the approach does not work well for the multi-hour intervals in Ontario’s TOU pricing plan. Instead, we develop statistical machine learning (ML) techniques to accurately predict consumption each interval. While our techniques have numerous applications, e.g., dispatch scheduling in microgrids, we focus solely on their application to SmartCharge in this paper.

We experimented with a variety of prediction techniques, including Exponentially Weighted Moving Averages (EWMA), Linear Regression (LR), and Support Vector Machines (SVMs) with various kernel functions, including Linear, Polynomial, and Radial Basis Function (RBF) kernels. EWMA is a classic past-predicts-future model that predicts consumption in the next interval as a weighted sum of the previous interval’s consumption and an average of all previous intervals’ consumption. More formally, EWMA predicts the energy consumption for each interval on day k as $\hat{E}_C(k+1) = \alpha E_C(k) + (1-\alpha)\hat{E}_C(k)$, where α is a configurable parameter that alters the weight applied to the most recent interval versus the past. Note that since each interval’s power consumption is different, we apply EWMA to each interval independently on a daily basis. As might be expected, since home consumption patterns vary largely around mealtimes, we found that predicting consumption based on the preceding interval to be highly inaccurate.

Both LR and SVM are regression techniques that combine and correlate numerous indicators (or features) of future power consumption to predict next-day usage. We experimented with a total of nine features: outdoor temperature and humidity, month, day of week, previous day power, previous interval power, as well as whether or not it is a weekend day or a holiday. We also included the EWMA prediction as an additional feature. To predict next-day temperature and humidity, we used weather forecasts from the National




Figure 5: Predicting energy consumption using the past does not capture day-to-day variations due to changing weather, weekly routines, holidays, etc.

Weather Service available from the National Digital Forecast Database (<http://www.nws.noaa.gov/ndfd/>). To evaluate our techniques we used power data collected every second from our case study home over a period of four months from June to September 2011. For the LR and SVM models, we used the first 70 days of the data set for model training, and the last 40 days for evaluating the model’s accuracy. We use the LibSVM library [5] to implement our LR and SVM models. Our SVM models use the *nu*-SVR regression algorithm, which we found always performed better than the ϵ -SVR algorithm [5]. For simplicity, we only predict consumption for the Ontario TOU rate periods in Figure 3.

Before training our model, we employed Correlation-based Feature Subset Selection (CFSS) to refine the number of input features [14]. CFSS evaluates the predictive ability of each individual feature along with the degree of redundancy between features. We apply CFSS separately for each of the five intervals, since the pattern of power consumption varies each interval. CFSS reduces the number of features in prediction model from nine to: four for 12am-7am, seven for 7am-11am, seven for 11am-5pm, six for 5pm-9pm, and five for 9pm-12am. In general, we find that more features are useful during periods with high, variable consumption.

We then experimented with multiple variations of LR models, including least squares and different regularized models (LASSO, ElasticNet, and Ridge Regression), since we found that temperature, humidity, and past data were approximately linear with respect to power consumption. However, our best performing LR model (ElasticNet) had an average error of 37%. EWMA performed much better, although Figure 5 demonstrates its limitations in predicting future consumption. The figure shows actual power consumption each day during the first interval (12am-7am), as well as EWMA ($\alpha = 0.35$) and the SVM-Polynomial model. EWMA is unable to predict large spikes or dips in consumption before they occur. Instead, EWMA’s predictions never vary too far from the mean usage. In contrast to EWMA, the SVM approach is able to partially predict many of the spikes and dips in consumption. Over our 40 day testing period, we found that SVM-Polynomial had an average error of only 5.75%. The SVM model with the Linear and RBF kernel performed worse than EWMA, as Table 1 shows, with




Figure 6: Average dollar savings per day for both real-time and TOU prices in our case study home.




Figure 7: Average percentage savings for both real-time and TOU prices in our case study home.

a 29.5% and 42.5% average error, respectively. As a result, in §5 we use SVM-Polynomial to evaluate SmartCharge.

5. EXPERIMENTAL EVALUATION

To illustrate SmartCharge’s potential for savings, we use the home described in §3 to evaluate the savings using real hourly real-time and TOU rate plans in simulation. We also implement a small-scale SmartCharge prototype using a home UPS system and a few household appliances. For real-time prices, we use rates from June to September 2011 in the hourly day-ahead market run by the New England Independent System Operator (ISO), which operates the electricity market in our home’s region. We use historical market data publicly available that ISOs are required to publish [20]. Since we use day-ahead market prices, we have perfect knowledge of next-day prices. For TOU pricing, we use the Ontario rate plan from Figure 3. While our home is not located in Ontario, it lies at the same latitude and experiences a similar climate. Thus, the prices are not entirely mismatched to our home’s consumption profile. In our experiments, we vary the pricing plans and battery characteristics to see how future price trends and battery technology impact savings. To predict next-day usage, we use the SVM-Polynomial model described §4. Finally, to quantify the optimal savings, we compare with an oracle that has perfect knowledge of next-day consumption.

Unless otherwise noted, our experiments use home power data from the same 40 day period in late summer as the previous section. We use CPLEX, a popular integer and linear programming solver, to encode and solve SmartCharge’s optimization problem, given next-day prices and expected consumption levels. Note that we consider only usable storage capacity in kWh in this section, which is distinct from (and typically much less than) battery capacity. In the next section, we discuss the battery capacity necessary to attain a given storage capacity. As mentioned in §3, we use an energy conversion efficiency of 80% for the battery and a C/4 charging rate for the usable storage capacity.





Figure 8: SmartCharge’s savings as a function of the charging rate for a 12kWh storage capacity.

5.1 Household Savings


Figure 6 shows the average savings per day in USD for both the real-time and TOU rate plans over the 40 day period, as a function of storage capacity, while Figure 7 shows the savings as a percentage of the total electricity bill. The graphs show that a storage capacity beyond 30kWh does not significantly increase savings. Further, smaller storage capacities, such as 12kWh, are also capable of reducing costs, near 10% for SmartCharge. If we extrapolate the savings over an entire year, we estimate that SmartCharge with 24kWh of storage is capable of saving \$101.59. Finally, the graphs show that SmartCharge’s performance is close to that of an oracle with perfect knowledge of future consumption: mispredictions only cost an estimated \$12.09 each year with 24kWh storage capacity, or near 12% of the total savings. Due to different price levels, the TOU plan saves slightly more dollars per day, while the real-time plan saves a larger percentage of the bill. As we show next, both the pricing plan and battery characteristics impact the savings. Since the savings for both the real-time and TOU rate plan are similar, for clarity we focus our remaining results on the TOU rate plan, which is more widely used today.

The experiments above assume that we use today’s battery characteristics and price levels. Of course, a more efficient battery and inverter would increase the usable storage capacity in a battery array. As the experiments above indicate, increasing storage capacity increases the savings up to a 30kWh capacity. Figure 8 demonstrates that the maximum charging rate has a minimal effect on savings, since the TOU rate plan (as well as the real-time plan) offer a long period of relatively low rates during the night for charging. The charging rate need only be high enough, e.g., a C/10 rate, to charge the battery over these periods. Figures 9(a) and (b) show how the savings change if we vary either the average price (while keeping price ratios constant) or the peak-to-off-peak price ratio (while keeping the average price constant) for a 12kWh capacity. The graphs demonstrate that, as expected, rising prices or ratios significantly impact the savings. In the former case, the relationship is linear, with a doubling of today’s average price resulting in a doubling of the savings for SmartCharge. Thus, if average electricity prices continue to rise 5% per year, as in the past, SmartCharge’s expected savings should also increase at 5% per year. In the latter case, while the savings rate decreases slowly as the ratio increases, the savings nearly doubles (up 88%) if the current ratio increases slightly from 1.6 to 2.

Finally, Figure 10 shows the additional savings homes are able to realize by sharing battery capacity with neighbors. Sharing is beneficial when homes exhibit peaks at different times by allowing them to share the available storage ca-



(a)



(b)

Figure 9: Varying the average electricity price (a) and the peak-to-off-peak price ratio (b) impacts savings.




Figure 10: Additional savings (in % and \$) from sharing 12kWh and 24 kWh between homes.

capacity. For the experiment, we use power data for a single day from a pool of 353 additional homes we monitor (described below), such that each point is an average of twenty runs with a set of k randomly chosen homes. We report both the additional dollar and percentage savings per home. We include 90% confidence intervals for the dollar savings. The experiment shows that sharing a battery array between homes results in additional savings as we increase the number of homes. As expected, more homes require more storage capacity to reap additional benefits. With 10 homes sharing 24kWh per home, the additional savings is 25%. However, with 12kWh per home the percentage savings does not increase beyond 15% when sharing with more than four homes.

5.2 Grid Peak Reduction

The purpose of real-time and TOU rate plans is to lower peak electricity usage across the entire grid. We evaluate the potential grid-scale effect of SmartCharge using power data from a large sampling of homes. We gather power data at scale from thousands of in-panel energy meters that anonymously publish their data to the web. Since we do not know if the meters are installed in commercial, industrial, or residential buildings, we filter out sources that do not have typical household power levels and profiles, i.e., peak power less than 10kW and average power less than 3kW. We also filter out sources with large gaps in their data. After filtering, we select 435 homes from the available sources.

Figure 11(a) plots the peak power over all the homes as a function of the fraction of homes using SmartCharge with 12kWh of energy storage. The figure shows that SmartCharge is capable of reducing peak power by 20% when 22% of homes use the system, as long as the homes randomize when they begin overnight charging. If everyone begins charging at the same time, e.g., at 12am at night, the peak reduction decreases to a maximum of only 8%. Even using randomized charging, if more than 22% of consumers install SmartCharge, then the peak reduction benefits begin

to decrease, due to a nighttime “rebound peak”. Once 45% of consumers use the system the evening rebound peak actually becomes larger than the original peak without SmartCharge. The same point occurs when only 24% of homes use the system without randomized charging. Of course, the experiments assume that prices do not change in response to homes installing SmartCharge, i.e., a large fraction of homes install the system simultaneously. A more plausible and realistic scenario is that the rate of adoption slowly rises with the differential between the peak and off-peak prices. In this scenario, SmartCharge’s load shifting would alter prices in each rate period. At some point, as Vytelingum et al.[28] formally show, the price changes would make the system increasingly less attractive for new users, as the difference between peak and off-peak prices would approach zero.

We discuss SmartCharge’s economics at scale further in §6. Figure 11(b) shows grid power usage over time, with 0% and 22% of the homes using SmartCharge with randomized charging, and demonstrates how SmartCharge causes demand to “flatten” significantly. Such a peak reduction would have a profound effect on generation costs, likely lowering them by more than 20% [18]. Finally, with 22% of homes using SmartCharge, the increase in total energy usage is only 2%. The result demonstrates that the benefits of flattening likely outweigh the increased energy consumption due to battery/inverter inefficiencies.

5.3 Lab Prototype Results

We constructed a small-scale proof-of-concept prototype using a home UPS connected to a few common household appliances. While not typically designed for entire homes, today’s UPSs include the inverters, transfer switches, charge controllers, battery enclosure, cabling, and battery sensors necessary for a SmartCharge system in a single appliance. We chose the APC Smart-UPS 2200VA XL as our UPS, which includes software to monitor its capacity and charge/discharge state. The UPS has a usable capacity of 450Wh, but is expandable to 16kWh, at a discharge rate of 100W/s. The UPS switches to battery in roughly 25ms, which is less than the holdup time, i.e., the duration a device is able to sustain operation without power, in modern power supplies. We experimented with both charging and discharging the UPS. The unit charges from 45% to 100% capacity in 80 minutes at a linear rate, and discharges in 35 minutes with an average load of 384W. We connect a refrigerator, freezer, dehumidifier, and two laptops to the UPS system. We then emulate a TOU rate plan over a two hour period, where the first hour corresponds to a peak period and the second corresponds to an off-peak period. Figure 12 shows that in this simple case SmartCharge uses battery power




Figure 11: With 22% of homes using SmartCharge, the peak demand decreases by 20% (a) and demand flattens significantly (b).




Figure 12: Our UPS-based prototype reduces peak usage by 69% when using a few common appliances.

during the peak period and then switches to grid power during the off-peak period. Without SmartCharge, during the peak period the grid load was on average 298W and during the off-peak period it was 128W. With SmartCharge, the peak period has an average grid load of only 91W while the off-peak period has an average load of 324W, resulting in a 69% reduction in peak electricity consumption.

6. COST-BENEFIT ANALYSIS

The previous section shows that SmartCharge cuts an electric bill by 10-15% with today’s market-based pricing plans. In this section, we first discuss SmartCharge’s return on investment (ROI), including its installation and maintenance costs, and then discuss its advantages over centralized energy storage. We ground our discussion using price quotes, primarily from the altE store (<http://www.altestore.com>), for widely-available commercial products.

6.1 Return-on-Investment

In many instances, homes already have the necessary infrastructure to implement SmartCharge. For example, many homes in developing countries already utilize UPSs because of instability in the power grid. As we discuss below, in the future, homes with photovoltaic (PV) systems may require on-site energy storage to balance an intermittent supply with demand without the aid of net metering. Batteries in electric vehicles (EVs) could also serve as energy storage. In each case, the homes already include the required infrastructure and battery capacity to implement SmartCharge. Since the homes would not need new infrastructure, the ROI is positive in these cases. Below, we discuss the ROI for homes that do not already have the necessary infrastructure.

Table 2 shows cost estimates for purchasing and installing SmartCharge’s components. For the inverter, we assume Apollo Solar’s True Sinewave Inverter, which combines an inverter, battery charger, and transfer switch into a sin-

Component	Total
Inverter	\$2099.00
Battery Charger	-
Transfer Switch	-
Inverter Gateway	\$287.00
Energy Monitor	\$200.00
Server	\$159.00
Battery Enclosure	\$1426.00
Cabling	\$200.00
Labor	\$500.00
Total	\$4871.00

Table 2: Estimated cost breakdown for installing SmartCharge’s supporting infrastructure.

gle appliance. To read battery state and control the appliance, we attach an additional communications gateway available for the inverter. Numerous home energy meters are available: The Energy Detective (TED) is a popular choice and costs \$200. Nearly any server is adequate to support SmartCharge’s software. We use an embedded DreamPlug server at a cost of \$159 as the gateway in the homes we now monitor. To hold the battery array, we assume two MNEBE-C 12-battery modular enclosures. Finally, we estimate \$200 for cabling and a day’s labor at \$500 for installation. The total estimated cost, excluding batteries, is \$4871.

Of course, SmartCharge’s largest expense is its battery array. Sealed VRLA/AGM lead-acid batteries are the dominant battery technology for stationary home UPSs and PV installations, due to their combination of low price, high efficiency, and low self-discharge rate. By contrast, lithium ion batteries, while lighter and more appropriate for EVs, are much more expensive. We use, as an example, the Sun Xtender PVX-2580L with a 3kWh rated capacity (at a C/20 discharge rate), which costs \$570 [24] and is designed for deep-cycle use in home PV systems. The battery’s manual specifies its lifetime as a function of its number of charge-discharge cycles and the DOD each cycle. We use the data to estimate the yearly cost of batteries—in \$/kWh of *usable* storage capacity—as a function of the depth of discharge (Figure 13) amortized over their lifetime, assuming SmartCharge’s typical single charge-discharge cycle per day. The usable storage capacity takes DOD into account: a battery rated for 10kWh operated at 50% depth of discharge has a usable capacity of only 5kWh. Figure 13 demonstrates that cost begins to increase rapidly after a 45% DOD, with an estimated cost of \$118/kWh of usable capacity.

In the U.S., SmartCharge likely qualifies for a Residen-




Figure 13: Amortized cost per kWh as a function of depth of discharge.

tial Renewable Energy Tax Credit, reducing its cost by 30%. Additionally, U.S. state and local governments offer an assortment of tax incentives for energy-efficiency improvements [7], which we estimate lower costs by 20%. Despite the advantages, today’s lead-acid batteries are still too expensive to produce a positive ROI at current electricity prices. For instance, while 24kWh of usable storage capacity saves \$91.25 per year using the Ontario TOU rate plan, batteries alone would cost \$1416 per year assuming the take breaks above. However, recent advancements in battery technology promise to dramatically reduce battery costs in the near future. Lead-carbon batteries have an expected lifetime 10x longer than today’s sealed lead-acid batteries at roughly the same cost [9, 12, 22]. Figure 14 shows the extended lifetime using data from recent tests conducted at Sandia National Labs comparing today’s sealed lead-acid battery and a new lead-carbon battery (the UltraBattery) [22].

Lead-carbon batteries combined with modest and expected price increases (25%) and peak-to-off-peak ratios (25%) would produce a positive ROI for SmartCharge in a few years. Assuming this scenario, Figure 15 plots SmartCharge’s yearly expense, including battery and infrastructure costs (amortized over 20 years), along with its estimated yearly savings for our case study home, as a function of usable storage capacity. Note that our ROI estimates do not include the savings from lowering generation costs for *all* homes by reducing peak demands. As Figure 11 shows, enabling only 22% of homes with SmartCharge would dramatically reduce peak demands, and, hence, generation costs for *all* homes, even those that have not invested in the system. Since all homes benefit from lower prices, utilities may consider subsidies that spread costs across all consumers, which for 22% of homes would lower costs by nearly 5X.

Alternatively, utilities might consider modifying their pricing plans to incentivize SmartCharge in all homes by increasing the fraction of the bill based on peak usage. While many utilities charge large consumers based on their peak usage over a day or month [1], residential bills typically do not include such a charge. Incorporating a substantial peak usage charge in electric bills would prevent the large rebound peaks in Figure 11 by directly incentivizing homes to flatten demand, rather than shift as much demand as possible to low-cost periods (causing the rebound peak). With market-based plans that only charge per-kWh, as more consumers install SmartCharge and shift their demand to low-cost periods, the price difference between the low-cost and high-cost periods would lessen to reflect the new demand distribution, thus lowering the ROI and discouraging additional homes from installing the system. A substantial peak-usage charge would maintain the financial incentives and continue to flatten demand (and prevent rebound peaks) as the fraction of




Figure 14: Comparison of sealed lead-acid and lead-carbon battery lifetime. Data from [22].




Figure 15: SmartCharge’s projected yearly expense and savings assuming recent battery advancements.

SmartCharge-enabled homes approaches 100%.

A full discussion of SmartCharge’s impact on the economics of electricity generation is outside the scope of this paper. However, it is clear that today’s market-based pricing plans assume that the price elasticity of electricity demand is low, i.e., changes in price do not have a significant impact on demand. SmartCharge fundamentally changes this fact by making demand nearly fully elastic with price.

6.2 Distributed vs. Centralized

Utilities have already begun to deploy large, centralized battery arrays to reduce peak usage and integrate more wind and solar farms, which require substantial energy storage to match an intermittent supply with variable demand. However, distributing battery storage throughout the grid has a number of inherent advantages over a centralized approach. In particular, home energy storage may serve as backup power during extended blackouts, lessening the economic impact of power outages and promoting a more stable grid. A centralized system also introduces a single point of failure. Further, substantial home energy storage may be a catalyst for implementing microgrids, where matching supply and demand is difficult without an energy buffer. Storing energy at its point-of-use also reduces transmission losses by eliminating losses incurred from generator to battery array.

Finally, perhaps the most important argument for installing many distributed battery arrays in homes, rather than large centralized arrays, is to encourage distributed generation without relying on net metering. While today’s PV installations typically use net metering to offset costs by selling energy back to the grid, it is not a scalable long-term solution. Injecting significant quantities of power into the grid from unpredictable and intermittent renewables has the potential to destabilize the grid by making it difficult to balance supply and demand. SmartCharge provides an alternative to net metering to offset costs in home PV systems that use batteries instead of net metering. We are currently studying how to include renewables in SmartCharge’s al-

gorithm. Our initial results suggest that homes with PV installations also benefit from SmartCharge.

7. RELATED WORK

Daryanian et al. [6] first identified the opportunity to exploit energy storage in real-time electricity markets using a linear programming formulation similar to ours. However, their problem formulation ignores many of the battery inefficiencies that influence the realizable savings. Further, the work does not address stochastic demand in residential settings, whereas we develop machine learning techniques to accurately predict next-day consumption. Finally, we conduct experiments to analyze the peak reduction effects of energy storage in the grid using real data, as well as analyze the ROI for installing and maintaining the system.

More recent work explores a similar problem as ours, but from different perspectives. For example, van de ven et al. [27] model the problem as a Markov Decision Process and claim that there is a threshold-based stationary cost-minimizing policy. The policy is optimal assuming that consumption is independent and identically distributed (i.i.d.). A preliminary evaluation with simulated demands following an i.i.d. distribution shows cost savings up to 40%. In contrast, we take a more experimental approach using traces of real home power usage and market-based rate plans. For the home in our case study, which has an aggregate power usage close to the average U.S. home, we show that the optimal savings is never more than 20% with realistic energy storage capacities (< 60kWh). Rather than solving the problem with respect to a particular demand distribution, we distill the problem to a linear program that uses our prediction model of future consumption levels

Vytelingum et al. [28] and Carpenter et al. [3] both focus on the economics of storage at scale, which we also discuss. Vytelingum et al. show that for sufficiently low adoption rates, the difference between the peak and off-peak prices approaches zero, reducing the financial incentives for installing energy storage. Similarly, in parallel with our work, Carpenter et al. also show that today's pricing schemes may increase the grid's peak demand at scale if prices do not adjust to demand. The work studies the profitability of a variety of different pricing schemes, and their effectiveness in decreasing grid demand peaks at scale. Koutsopoulos et al. [16] explore the problem from the perspective of a utility operator. In this case, the utility controls when to charge and discharge battery-based storage to minimize generation costs, assuming the marginal cost to dispatch generators is similar to Figure 1 from §1. In contrast to our problem, the approach is more applicable to large centralized energy storage facilities. We discuss the trade-offs between distributed and centralized energy storage in §6.2.

Prior work has also explored combining energy storage with market-based electricity prices to cut electricity costs in data centers [13, 26]. The approach is different, since the authors assume that neither future prices nor consumption is known or predictable. As mentioned above, we assume the use of day-ahead electricity prices, where next-day prices are well-known. Household consumption patterns are also more predictable and regular than data center workloads, which vary widely based on application and server characteristics. Our ROI analysis is also not directly comparable, since data centers already employ double-conversion UPSes to ensure operation during grid outages. Thus, data centers already

incur the 10-20% loss from battery inefficiency, which increases their ROI by increasing savings relative to their baseline and eliminating additional installation and maintenance costs.

8. CONCLUSION

In this paper, we explore how to lower electric bills using SmartCharge by storing low-cost energy for use during high-cost periods. We show that typical savings today are 10-15% per home with the potential for significant grid peak reduction (20% with our data). Finally, we analyze SmartCharge's costs, and show that recent battery advancements combined with an expected rise in electricity prices may make SmartCharge's return on investment positive for the average home within the next few years.

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

References

- [1] 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*, June 2008.
- [2] Bringing Residential Real-Time Pricing to Scale in Illinois: Policy Recommendations. Technical report, CNT Energy, 2010.
- [3] T. Carpenter, S. Singla, P. Azimzadeh, and S. Keshav. The Impact of Electricity Pricing Schemes on Storage Adoption in Ontario. In *e-Energy*, May 2012.
- [4] D. Cauchon. USA Today, Household Electricity Bills Skyrocket, December 13th 2011.
- [5] C. Chang and C. Lin. LIBSVM: A Library for Support Vector Machines. *ACM Transactions on Intelligent Systems and Technology*, 2(27), April 2011.
- [6] B. Daryanian, R. Bohn, and R. Tabors. Optimal Demand-side Response to Electricity Spot Prices for Storage-type Customers. *TPS*, 4(3), August 1989.
- [7] Database of State Incentives for Renewables and Efficiency. <http://www.dsireusa.org>, 2010.
- [8] C. DeNavas-Walt, B. Proctor, and J. Smith. Income, Poverty, and Health Insurance Coverage in the United States: 2010. Technical report, U.S. Census Bureau, September 2011.
- [9] The Economist, Lead-acid Batteries Recharged: A 150-year-old Technology Looks to the Future, Jan. 14 2009.
- [10] eGauge Monitoring. <http://www.egauge.net/>, 2012.
- [11] State of the Markets Report 2008. Technical report, Federal Energy Regulatory Commission, August 2009.
- [12] J. Furukawa, T. Takada, T. Kanou, D. Monma, L. Lam, N. Haigh, O. Lim, and R. Louey. Development of UltraBattery. Technical report, CSIRO, December 2006.
- [13] S. Govindan, A. Sivasubramaniam, and B. Urgaonkar. Benefits and Limitations of Tapping into Stored Energy for Datacenters. In *ISCA*, June 2011.
- [14] M. Hall. Correlation-based Feature Selection for Discrete and Numeric Class Machine Learning. In *ICML*, June 2000.
- [15] S. Kaplan. Power Plants: Characteristics and Costs. Technical report, Congressional Research Service Report to Congress, November 2008.
- [16] I. Koutsopoulos, V. Hatzi, and L. Tassiulas. Optimal Energy Storage Control Policies for the Smart Power Grid. In *SmartGridComm*, September 2011.
- [17] D. Linden. *Linden's Handbook of Batteries*, pages 17.1–17.39. McGraw Hill, Fourth edition, 2011.
- [18] National Public Radio. How Smart is the Smart Grid?: Interview with Dan Delurey, President Demand Response Smart Grid Coalition, July 7th, 2010.

- [19] New York ISO. Understanding the Markets. http://www.nyiso.com/public/about_nyiso/understanding_the_markets/energy_market, 2012.
- [20] New England ISO. <http://www.iso-ne.com/>, 2012.
- [21] Ontario Energy Board: Electricity Prices. <http://www.ontarioenergyboard.ca/OEB/Consumers>, 2012.
- [22] S. Schoenung. Energy Storage Systems Cost Update: A Study for the DOE Energy Storage Systems Program. Technical report, Sandia National Laboratories, 2011.
- [23] Smart Grid Information Clearinghouse: Legislation and Regulation. <http://www.sgiclearinghouse.org/Legislation?q=node/1705>, 2011.
- [24] Technical Manual for Sun Xtender Batteries, July 2011.
- [25] Electric Power Annual 2010. Technical report, U.S. Energy Information Administration, November 2011.
- [26] R. Uргаonkar, B. Uргаonkar, M. Neely, and A. Sivasubramaniam. Optimal Power Cost Management Using Stored Energy in Data Centers. In *SIGMETRICS*, March 2011.
- [27] P. van de ven, N. Hegde, L. Massoulie, and T. Salonidis. Optimal Control of Residential Energy Storage Under Price Fluctuations. In *ENERGY*, May 2011.
- [28] P. Vytelingum, T. Voice, S. Ramchurn, A. Rogers, and N. Jennings. Agent-based Micro-storage Management for the Smart Grid. In *AAMAS*, May 2010.