GreenCharge: Managing Renewable Energy in Smart Buildings

Aditya Mishra, David Irwin, Prashant Shenoy, *Fellow, IEEE,* Jim Kurose, *Fellow, IEEE,* and Ting Zhu, *Member, IEEE*

***Abstract*—Distributed generation (DG) uses many small on-site energy harvesting deployments at individual buildings to generate electricity. DG has the potential to make generation more efficient by reducing transmission and distribution losses, carbon emissions, and demand peaks. However, since renewables are intermittent and uncontrollable, buildings must still rely, in part, on the electric grid for power. While DG deployments today use net metering to offset costs and balance local supply and demand, scaling net metering for intermittent renewables to a large fraction of buildings is challenging. In this paper, we explore an alternative approach that combines market-based electricity pricing models with on-site renewables and modest energy storage (in the form of batteries) to incentivize DG. We propose a system architecture and optimization algorithm, called GreenCharge, to efficiently manage the renewable energy and storage to reduce a building's electric bill. To determine when to charge and discharge the battery each day, the algorithm leverages prediction models for forecasting both future energy demand and future energy harvesting. We evaluate GreenCharge in simulation using a collection of real-world data sets, and compare with an oracle that has perfect knowledge of future energy demand/harvesting and a system that only leverages a battery to lower costs (without any renewables). We show that GreenCharge's savings for a typical home today are near 20%, which are greater than the savings from using only net metering.**

I. Introduction

Buildings today consume more energy (41%) than either of society's other broad sectors of energy consumption—industry (30%) and transportation (29%) [1]. As a result, even small improvements in building energy efficiency, if widely adopted, hold the potential for significant impact. The vast majority (70%) of building energy usage is in the form of electricity, which, due to environmental concerns, is generated at "dirty" power plants far from population centers. As a result, nearly half (47%) of energy use in residential buildings is lost in electricity transmission and distribution (T&D) from far-away power plants to distant homes [1]. An important way to decrease both T&D losses and carbon emissions is through distributed generation (DG) from many small on-site renewable energy sources deployed at individual buildings and homes. Unfortunately, in practice, DG has significant drawbacks that have, thus far, prevented its widespread adoption. In particular, DG primarily relies on solar panels and wind turbines that generate electricity intermittently based on uncontrollable and changing environmental conditions. Since the energy consumption density, in kilowatt-hours (kWh) per square foot, is higher than the energy generation density of solar and wind deployments at most locations, buildings must still rely heavily on the electric grid for power.

Another major drawback of DG is that large centralized power plants benefit from economies-of-scale that cause their generation costs, even accounting for T&D losses, to be significantly lower than DG. As a result, today's DG deployments rely heavily on net metering—where buildings sell the unused energy they produce back to the utility company—to offset their cost relative to grid energy. DG is a much less financially attractive where net metering is not available. Net metering laws and regulations vary widely across states—it is not available in four states and the regulations are weak in many others [2]. Further, even where available, states typically place low caps on both the total number of participating consumers and the total amount of energy contributed per customer [28]. After exceeding these caps, utilities are no longer required to accept excess power from DG deployments. As one example, the state of Washington caps the total number of participating consumers at 0.25% of all customers. One reason for the strict laws limiting DG's contribution is that injecting significant quantities of power into the grid from unpredictable renewables at large scales has the potential to destabilize the grid by making it difficult, or impossible, for utilities to balance supply and demand. Large baseload power plants that produce the majority of grid energy are simply not agile enough to scale their own generation up and down to offset significant fractions of renewable generation.

Thus far, current laws have not been an issue, since today's energy prices do not make DG financially attractive enough to reach even these low state caps. However, more widespread adoption of DG is critical to meeting existing goals for increasing the fraction of environmentally-friendly renewable energy sources. For example, the Renewables Portfolio Standard targets 25% of electricity generation from intermittent renewables [8], while California's Executive Order S-21-09 in California calls for 33% of generation from renewables by 2020 [31]. Given current laws, if and when DG becomes more widespread, buildings will have to look beyond net metering to balance on-site energy generation and consumption, while also reducing DG's costs. We envision consumers using a combination of on-site renewables, on-site battery-based energy storage, and the electric grid to satisfy their energy requirements, while also balancing local supply and demand.

In parallel, we envision the adoption of market-based elec-

Manuscript received October 8, 2012; revised March 15, 2013.

Aditya Mishra, David Irwin, Prashant Shenoy, Jim Kurose are with University of Massachusetts, Amherst, MA-01003 (e-mail: {adityam,irwin,shenoy,kurose}@cs.umass.edu)

Ting Zhu is with State University of New York, Binghamton, NY, 13902-6000 (e-mail: tzhu@binghamton.edu)

tricity pricing providing a new opportunity to recoup the loss of net metering revenue, while also introducing new financial incentives for DG where net metering is not available. 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. Satisfying peak demands is significantly more expensive (∼10x) than off-peak demands, since peak demands drive both capital expenses—by dictating the number of power plants, transmission lines, and substations—and operational expenses—"peaking" generators are generally dirtier and costlier to operate than baseload generators [19]. For instance, Illinois already requires utilities to provide residential customers the option of using hourly electricity prices based directly on wholesale prices [30], while Ontario charges residential customers based on a time-of-use scheme with three different price tiers (off-, mid-, and on-peak) each day [26].

The primary contribution of this paper is a new system architecture and control algorithm, called GreenCharge for managing on-site renewables, on-site energy storage, and grid energy in buildings to minimize grid energy costs for market-based electricity prices. Our system determines both the fraction of power to consume from the grid versus on-site battery-based energy storage, as well as when and how much to charge battery-based storage using grid energy. The primary inputs to our control algorithm are 1) the battery's current energy level, 2) a prediction of future solar/wind energy generation, 3) a prediction of future energy consumption patterns, and 4) market-based electricity prices. The output is the amount of power to consume from the grid, as well as the power to discharge or charge the battery from renewables or the grid, over each rate period. We evaluate our system using a collection of real data sets, including power consumption data from a real home, energy harvesting data from a solar and wind deployment, National Weather Service (NWS) forecast data, and TOU pricing data from Ontario, Canada.

We compare GreenCharge with two other approaches: i) an approach from initial work, called SmartCharge [23], that only uses energy storage without renewables to reduce prices and ii) an oracle with perfect knowledge of future energy consumption and generation. GreenCharge extends our initial work on SmartCharge in multiple ways. First, SmartCharge only optimized prices by determining when and how much to charge a battery at off-peak hours. GreenCharge extends this idea to account for intermittent renewable generation, e.g., by using forecast-based models to predict future energy harvesting—a major enhancement to SmartCharge. In addition, this paper includes new material describing our use of communication protocols in implementing a GreenCharge prototype, as well as a revised linear programming formulation and algorithm that accounts for renewable generation. Finally, our work includes substantial experiments to understand the impact of adding renewables to SmartCharge. Our results show that GreenCharge saves an additional 10-15% on electric bills beyond SmartCharge, which only uses a battery, and is near the performance of an oracle with perfect future knowledge.

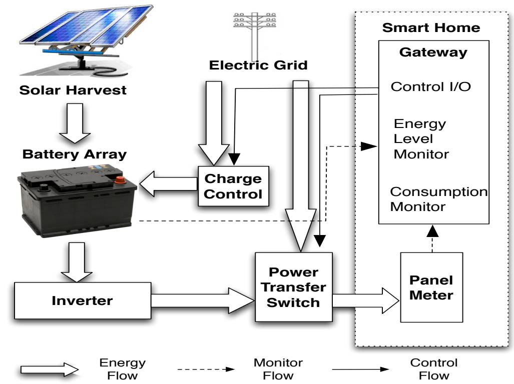

Fig. 1. A depiction of GreenCharge's architecture, including its battery array and charger, DC→AC inverter, solar and/or wind energy sources, power transfer switch, energy/voltage sensors, and gateway server.

II. GreenCharge Architecture

Figure 1 depicts GreenCharge'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. On-site solar panels or wind turbines connect to, and charge, the battery array. A smart gateway server continuously monitors 1) electricity prices via the Internet, 2) household consumption via an in-panel energy monitor, 3) renewable generation via current transducers, and 4) the battery's state of charge via voltage sensors. Our SmartCharge system, which we compare against in this work, utilizes the same architecture, but does not use renewables [23].

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 and generation 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 §VI, we provide a detailed estimate of GreenCharge's installation and maintenance costs based on price quotes for widely-available commercial products.

A. Network Communication and Sensing

One challenge with instantiating GreenCharge's architecture is transmitting sensor data about energy consumption, energy generation, and battery status to GreenCharge's smart gateway server in real time. The simplest way to measure energy consumption and generation is to wrap current transducers (CT) around wires in the building's electrical panel. In this case, two CTs are necessary to cover both legs of a building's split leg input power from the grid, as well as a CT for each connection to a renewable source. Note that CTs use the Hall Effect [14] for measuring voltage and current, and only require wrapping a sensor around a wire without cutting any wires. CTs must be installed in the panel, since this is the only place in the building that has the incoming grid lines exposed for sensors. Since electrical panels are often in remote corners of a building, transmitting readings wirelessly is difficult. While wired Ethernet is an attractive option, it requires running an Ethernet cable from GreenCharge's gateway server to the electrical panel. Instead, to overcome wireless interference and prevent running an Ethernet cable into the panel, GreenCharge uses a powerline-based communication protocol to transmit readings to the server.

Multiple types of powerline-based communication protocols exist. The most common are X10, Insteon, and HomePlug. X10 is by far the oldest protocol, having been developed in 1975; it is primarily used for controlling applications, which only requires sending brief, short control messages. Unfortunately, X10 has severe bandwidth limitations (a maximum of 20bps) and reliability problems, which make it undesirable for continuous real-time sensing. The bandwidth limitations alone prevent X10 from being used to continuously sense multiple data sources. Since powerline is a broadcast network, the 20bps bandwidth is across *all* devices. In addition to the bandwidth limitations, the protocol has no acknowledgements, so it is impossible to detect packet losses and retransmit. Further, powerline noise caused by switched mode power supplies results in substantial losses with X10 in most buildings. In our own prototype, we initially used the Energy Detective (TED) power meter for monitoring electricity consumption and generation at the electrical panel. However, we discovered that the meter uses an unreliable X10-like protocol that experiences communication problems while sending data over the powerline due to sensitivity to noise. While the display blinks orange when the problems occur, the data masks the problem by always recording the last power reading as the current power reading.

Insteon is an improvement to X10 that includes acknowledgements, retransmissions, and optimizations to overcome powerline noise. However, Insteon still has bandwidth limitations that, in practice, reduce its maximum rate to near 180bps [17]. While useful for controlling devices via the powerline, it is still insufficient for continuous real-time sensing of multiple data sources. Thus, in our own prototype we chose a power meter that uses the HomePlug Ethernet-over-powerline protocol. Unlike Insteon and X10, Homeplug was initially designed to stream high definition audio and video data from the Internet to televisions. As a result, it was designed from the outset to support high-bandwidth applications. HomePlug modems exist that are capable of transmitting up to 200Mbps. Since HomePlug simply implements Ethernet over the powerline, it can support a standard TCP stack to ensure reliable communication. Our prototype uses an eGauge power meter [11], and uses HomePlug to continuously transmit power consumption and generation readings over a building's powerline to GreenCharge's gateway server. Below, we discuss how the server gets current market prices for electricity.

B. Market-based Electricity Pricing

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,

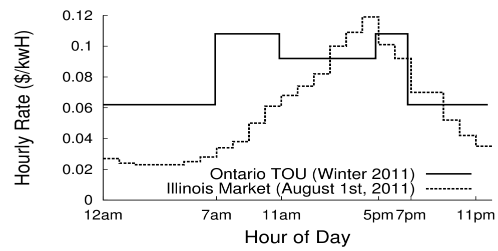

Fig. 2. Example TOU and hourly market-based rate plans in Ontario and Illinois, respectively.

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, GreenCharge 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 2 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 [26]. 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. While some utilities use real-time prices not known in advance, most utilities use day-ahead market prices, which are are set one day in advance. Since utilities purchase most of their electricity in day-ahead markets, e.g., 98% in New York [25], next-day prices are well-known.

There are many possible ways for GreenCharge's gateway server to monitor prices in real time. In the simplest case, utilities can provide simple web pages with current prices. For example, Illinois utilities are already required to do this, e.g., www.powersmartpricing.org/chart posts next-day prices each evening. Utilities may also use explicit protocols to "push" prices to GreenCharge's gateway server whenever they change.

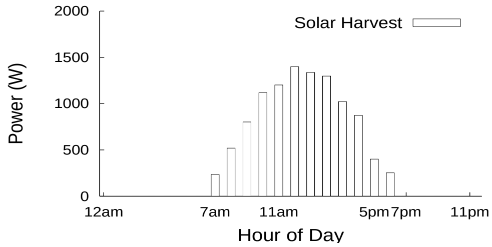

Fig. 3. Example solar harvest data from a day in August.

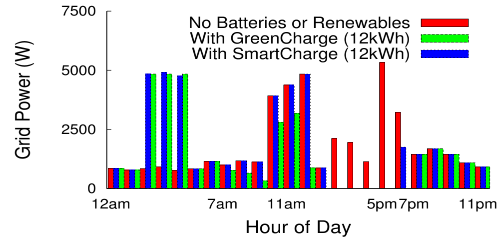

Fig. 4. Example from January 3rd with and without GreenCharge using Illinois prices from Figure 2.

For example, utilities could run publish/subscribe protocols that interact with smart meters to broadcast price changes. In this case GreenCharge's gateway server could interact with a building's local smart meter to discover prices. Authors in [18], [13] propose to combine IP multicast and publish-subscribe technologies to scale real-time price broadcast to millions of users for Ecogrid [9]. When using smart meters, utilities could disseminate prices using the smart meter's communication protocol, e.g., often cellular wireless or wired powerline, rather than the public Internet.

Transactive control system, presented in [16], proposes another way of price dissemination in smart grids. In transactive control, responsive demand assets are controlled by a single, shared, price-like value signal. It defines a hierarchical node structure and the signal path through these nodes, and includes the predicted day-ahead price values. Alternatively, IEC 61850 ([3]), which has been used between DER (Distributed Energy Resources) plants for energy and price information exchange, can be extended for price exchange in smart grids. [22] presents a survey of a set of existing communication protocols. The report also analyzes suitability of the surveyed protocols for their application in real-time price exchange.

GreenCharge is compatible with any method above for retrieving real-time prices, and works well with both TOU and real-time pricing plans. In either case, GreenCharge solves the optimization 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.

III. GreenCharge Algorithm

GreenCharge cuts electricity bills by combining on-site renewable generation with energy storage that stores energy during low-cost periods for use during high-cost periods. As discussed in §I, GreenCharge extends our SmartCharge system that only uses energy storage to cut electricity bills without renewables. The total possible savings each day is a function of both the home's rate plan and its pattern of generation and 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 GreenCharge'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 [11], 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. Separately, we have deployed solar panels to study variation in solar power generation. Figure 3 depicts power generation from a sunny day.

A. *Potential Benefits*

To better understand GreenCharge'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. Figure 4 then compares GreenCharge using renewable production from Figure 3 with a home has only energy storage but not renewables (labeled SmartCharge), and home with no energy storage or renewables. Now consider our home's hourly electricity use on January 3rd, 2012, as depicted in Figure 4 in red. 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 all consumption was shifted 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. Finally, if we then add in the 10.5kW generated by renewables the savings increases by $0.93 to $1.92. This per-day savings rate translates to a yearly savings of $702, 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. Of course, energy consumption and generation patterns, as well as hourly rates vary each day, which may decrease (or increase) a building's actual yearly savings. To understand why energy consumption and generation patterns are important, consider the following scenario using the Ontario TOU pricing plan. In Ontario, while GreenCharge 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, even when accounting for renewable generation, 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, e.g., because renewables will generate some power during this time, 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.

B. 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.), GreenCharge'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 GreenCharge'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 technologies [27]. 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 [32].

The ambient temperature and rate of discharge also have an impact on usable capacity, according to Peukert's law. To maximize lifetime, we expect GreenCharge 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 §V, reasonable battery capacities for GreenCharge 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 [21]. In §V, 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 GreenCharge'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 I from 1 to T. We then denote the power charged to the battery from the grid during interval i as s_i, the renewable power charged to the battery as g_i, average renewable power available to the home as r_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 ith 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.

$$s_i \geq 0, \forall i \in [1, T] \tag{1}$$

$$d_i \geq 0, \forall i \in [1, T] \tag{2}$$

$$g_i \geq 0, \forall i \in [1, T] \tag{3}$$

$$g_i \leq r_i, \forall i \in [1, T] \tag{4}$$

$$s_i \leq C/4, \forall i \in [1, T] \tag{5}$$

$$g_i \leq C/4, \forall i \in [1, T] \tag{6}$$

$$\sum_{t=1}^{i} d_t \leq e * \sum_{t=1}^{i} s_t + e * \sum_{t=1}^{i} g_t, \forall i \in [1, T] \tag{7}$$

$$(\sum_{t=1}^{i} s_t + \sum_{t=1}^{i} g_t - \sum_{t=1}^{i} d_t/e) * I \leq C, \forall i \in [1, T] \tag{8}$$

$$m_i = (p_i + s_i - d_i) * I * c_i, \forall i \in [1, T] \tag{9}$$

The first second and third constraint ensure the energy charged to, or discharged from, the battery is non-negative. The fourth constraint ensures that total renewable energy charged to the battery is less than or equal to the available renewable energy. The fifth and sixth constraint limits the battery's maximum charging rate. The seventh constraint specifies that the power discharged from the battery is never greater than the total power charged to the battery multiplied by the inefficiency parameter. The eighth constraint states that the energy stored in the battery array, which is the difference

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 I
AVERAGE PREDICTION ERROR (%) OVER 40 DAY SAMPLE PERIOD FOR SVM WITH DIFFERENT KERNEL FUNCTIONS.

between the energy charged to or discharged from the battery over the previous time intervals, cannot be greater than its capacity. Finally, the ninth constraint defines the price the home pays for energy during the ith interval. The objective and constraints define a linearly constrained optimization problem that is solvable using standard linear programming techniques. GreenCharge solves the problem at the beginning of each day to determine when to switch between grid and battery power, and when to charge the battery from grid vs renewables. SmartCharge uses a similar linear programming formulation without the constraints specific to renewable energy. Since the approach uses knowledge of next-day consumption and generation patterns, we next detail techniques for predicting next-day consumption and generation, and quantify their accuracy for our case study home.

IV. PREDICTING CONSUMPTION AND GENERATION

As discussed in §III, solving GreenCharge's linear optimization problem requires *a priori* knowledge of next day consumption and generation patterns. We develop a machine learning based approach to predicting demand, and use an approach developed in prior work [29] to predict next day energy harvesting based on weather forecasts. We discuss each mode in turn.

A. ML-based Demand Prediction

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

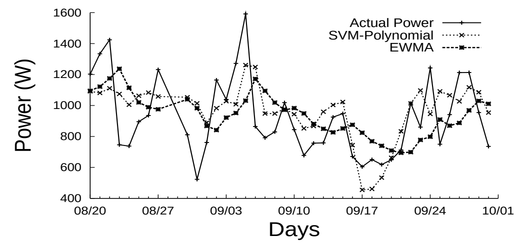

Fig. 5. Predicting energy consumption using the past does not capture day-to-day variations due to changing weather, weekly routines, holidays, etc.

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 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 [6] 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 [6]. For simplicity, we only predict consumption for the Ontario TOU rate periods in Figure 2.

Before training our model, we employed Correlation-based Feature Subset Selection (CFSS) to refine the number of input features [15]. 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

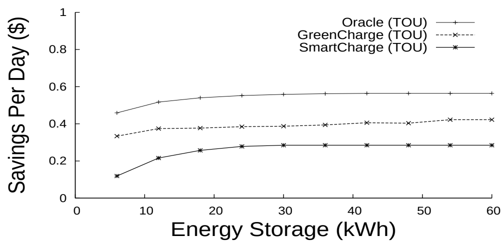

Fig. 6. Average dollar savings per day for both SmartCharge and GreenCharge in our case study home.

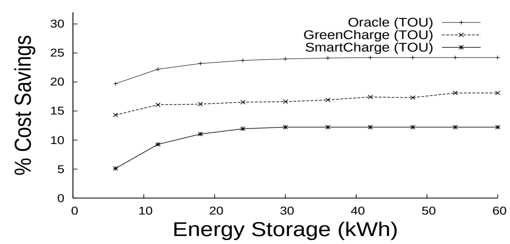

Fig. 7. Average percentage savings for both SmartCharge and GreenCharge in our case study home.

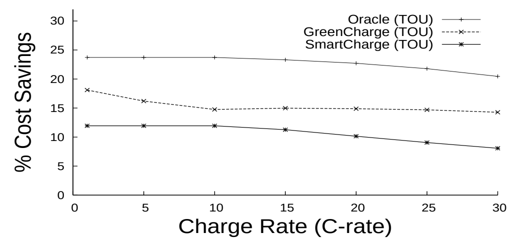

Fig. 8. SmartCharge's and GreenCharge's savings as a function of the charging rate for a 24kWh storage capacity.

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

B. Predicting Energy Harvesting from Weather Forecasts

For predicting the harvested solar energy we use the prediction model presented in [29]. For a given solar panel deployment this model translates the forecasted sky cover, by National Weather Service (NWS), into solar energy harvesting prediction. The NWS publishes weather forecast including sky condition forecast, every hour. The forecast contains predicted sky condition for next 24 hours. The model computes predicted solar harvesting power for every hour as:

$$Power = MaxPower * (1 - SkyCondition) \qquad (10)$$

$Power$ in (10) is the predicted solar harvesting power, $MaxPower$ is the maximum possible solar power that can be harvested from the given solar panel in a given hour of day assuming perfectly sunny day, and $SkyCondition$ is the fraction of sky that is covered with clouds.

V. Experimental Evaluation

To illustrate GreenCharge's potential for savings, we use the home described in §III to evaluate the savings using Ontario's TOU rate plans in simulation from Figure 2. 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 and generation 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 in §IV. Similarly, to predict next-day generation, we use the forecast-based model from §IV. Finally, to quantify the optimal savings, we compare with an oracle that has perfect knowledge of next-day consumption and generation.

Unless otherwise noted, our experiments use home power data from the same 40 day period in late summer as the previous section, and generation data from our own solar panel installation scaled up to generate equal to the home's average power consumption. We use CPLEX, a popular integer and linear programming solver, to encode and solve GreenCharge's (and 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 §III, we use an energy conversion efficiency of 80% for the battery and a C/4 charging rate for the usable storage capacity.

A. Household Savings

Figure 6 shows the average savings per day in USD for the TOU rate plan 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 and 20% for GreenCharge. If we extrapolate the savings over an entire year, we estimate that GreenCharge with 24kWh of storage is capable of saving $200, while SmartCharge is capable of saving $100. Finally, the graphs show that GreenCharge's performance is close to that of an oracle with perfect knowledge of future consumption and generation: mispredictions only cost a few dollars each year

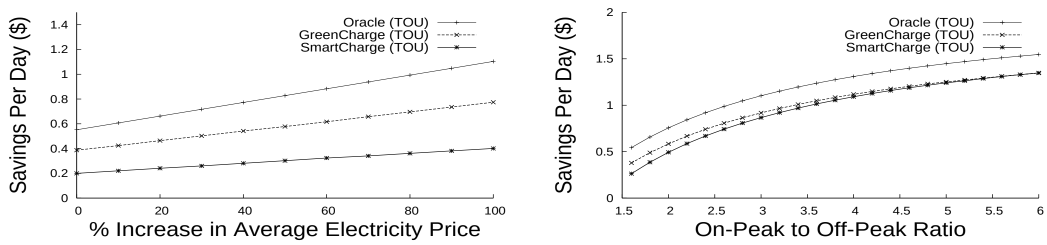

(a) (b)

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

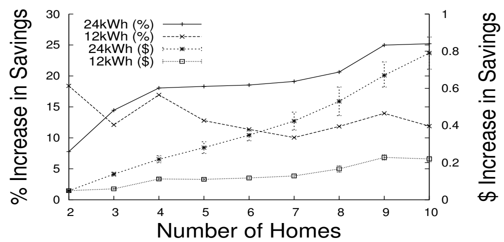

Fig. 10. Additional savings (in % and $) from sharing 12kWh and 24 kWh between homes.

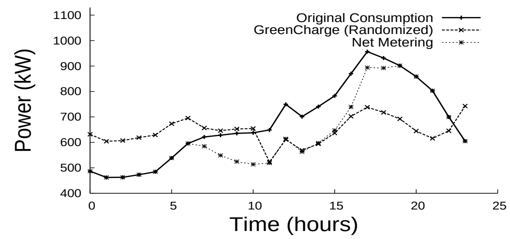

Fig. 12. Demand flattening with Net Metering.

with 24kWh storage capacity, or under 10% of the total savings.

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. We evaluate the effect of maximum battery charging rate on home savings using TOU pricing plan over 40 day traces in presence of 24kWh battery capacity. Figure 8 demonstrates that the maximum charging rate has a minimal effect on savings, since the TOU rate plan offers 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 24kWh capacity, assuming C/4 charging rate for the usable storage capacity, for both GreenCharge and SmartCharge. 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 both GreenCharge and SmartCharge. Thus, if average electricity prices continue to rise 5% per year, as in the past, the expected savings for both systems 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%) for both GreenCharge and SmartCharge 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 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.

B. Grid Peak Reduction

The purpose of market-based rate plans is to lower peak electricity usage across the entire grid. We evaluate the potential grid-scale effect of GreenCharge 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. Power consumption trace for each home is at the granularity of one hour. 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 GreenCharge and SmartCharge with energy storage. For these experiments we assume that each home has an available energy storage equal to half the home's average daily consumption. Charging rate of C/4 for the usable storage capacity is assumed. The figure shows that GreenCharge and SmartCharge are capable of reducing peak power by roughly 20% when little more

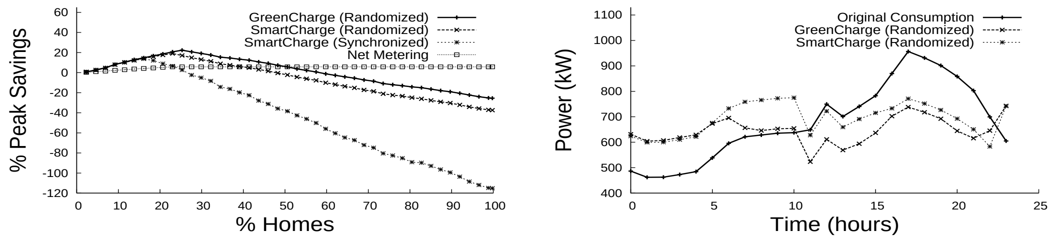

(a) (b)

Fig. 11. With 25% of homes using GreenCharge, the peak demand decreases by 22.5% (a) and demand flattens significantly (b).

than 20% 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 GreenCharge or 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. The same point occurs when only 25% of homes use the system without randomized charging. ’Net Metering’ represents those homes which have on-site renewable deployments, however, they don’t have on-site battery installations for storing this energy. Hence, the renewable energy is consumed as soon as it is generated. In contrast to GreenCharge and SmartCharge the peak savings from ’Net Metering’ increase from 0% to 5.75% and then flattens out. The reason being, net metering does not use any on-site battery storage, it simply uses the renewable energy whenever it is available else the power is drawn from the grid. Also, as can be seen from figure 12 net metering effectively flattens out the mid day peaks between 11am and 2pm, however, it does poorly to shave the evening peak which occurs after 5pm. This is because solar energy harvest reduces significantly towards sunset. Clearly, battery storage is required to shave the evening peaks. Another important observation from figure 12 is that net metering increases the difference between the minimum and maximum power drawn from the grid during day time, i.e., between 7am to 7pm, hence making load on the grid less predictable and sporadic.

All our experiments assume that prices do not change in response to homes installing battery-based energy storage, 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, the gradual load shifting would alter prices in each rate period. At some point, as Vytelingum et al.[34] 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 GreenCharge’s and SmartCharge’s economics at scale further in §VI. Figure 11(b) shows grid power usage over time, with 0% and 22% of the homes using GreenCharge and SmartCharge with randomized charging, and demonstrates how both approaches cause demand to “flatten” significantly. Such a peak reduction would have a profound effect on generation costs, likely lowering them by more than 20% [24]. Finally, with 20% of homes using GreenCharge or 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.

VI. Cost-Benefit Analysis

The previous section shows that GreenCharge cuts an electric bill by 20% with today’s market-based pricing plans, compared to around a 10% decrease with SmartCharge. In this section, we first discuss GreenCharge’s return on investment (ROI), including its installation and maintenance costs. We ground our discussion using price quotes, primarily from the altE store (http://www.altestore.com), for widely-available commercial products.

A. Return-on-Investment

In many instances, homes already have the necessary infrastructure to implement GreenCharge. For example, many homes in developing countries already utilize UPSs because of instability in the power grid. In addition, homes with photovoltaic (PV) systems 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 GreenCharge. 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 II shows cost estimates for purchasing and installing GreenCharge’s components. For the inverter, we assume Apollo Solar’s True Sinewave Inverter, which combines an inverter, battery charger, and transfer switch into a single 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 GreenCharge’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, GreenCharge’s largest

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 II
ESTIMATED COST BREAKDOWN FOR INSTALLING SMARTCHARGE'S SUPPORTING INFRASTRUCTURE.

expenses are its battery array and solar panel installation. We discuss each below.

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 [32] 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 GreenCharge'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.

While solar panel prices are dropping dramatically, current prices are $7-$9 per watt for installing solar generation. Since both the average consumption in our example home (and the average across the U.S.) is 1kW, it would cost $4000 for a system capable of producing half the home's electricity. Of course, since a solar installation does not produce its maximum power all the time, our home would likely need a installation with at least a 4x larger capacity than our desired output. As a result, to generate half the home's electricity from solar panels would cost $16,000-$20,000.

In the U.S., GreenCharge likely qualifies for a Residential 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 [8], which we estimate lower costs by 20%. Despite the advantages, today's lead-acid batteries and solar panels 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 [10], [12], [27]. 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) [27]. In addition, solar panel prices per installed watt are predicted to drop to $1 per watt over the next decade.

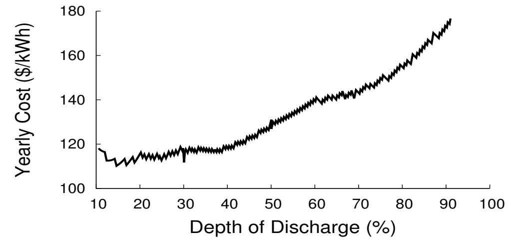

Fig. 13. Amortized cost per kWh as a function of depth of discharge.

Lead-carbon batteries combined with modest and expected price increases (25%) and peak-to-off-peak ratios (25%), as well as a decrease in solar panel prices, would produce a positive ROI for GreenCharge in a few years. As Figure 11 shows, enabling only 20% of homes with GreenCharge 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 20% of homes would lower costs by nearly 5X.

Alternatively, utilities might consider modifying their pricing plans to incentivize GreenCharge (and 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 [4], 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 the system 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 GreenCharge-enabled homes approaches 100%.

A full discussion of GreenCharge'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. GreenCharge fundamentally changes this fact by making demand nearly fully elastic with price.

B. Distributed vs. Centralized

Utilities have already begun to deploy large, centralized battery arrays to reduce peak usage and integrate more wind

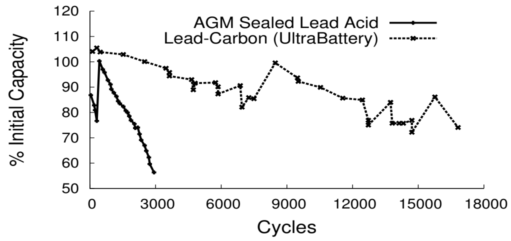

Fig. 14. Comparison of sealed lead-acid and lead-carbon battery lifetime. Data from [27].

and solar farms, which require substantial energy storage to match an intermittent supply with variable demand. However, distributing battery storage and energy harvesting throughout the grid has a number of inherent advantages over a centralized approach. In particular, local energy storage and generation serves 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 and generation may be a catalyst for implementing microgrids, where matching supply and demand is difficult without an energy buffer. Storing and generating 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 and energy harvesting deployments 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. GreenCharge provides an alternative to net metering to offset costs in home PV systems that use batteries instead of net metering.

VII. Related Work

Daryanian et al. [7] 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. In addition, we also 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. Finally, we include renewables into the system, as well as use a model for predicting renewable generation, which has not been considered in prior work to the best of our knowledge.

More recent work explores a similar problem as ours, but from different perspectives and without renewable generation. For example, van de ven et al. [33] 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, solar panel generation, 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. [34] and Carpenter et al. [5] 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. [20] 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 increases super-linearly as utilities move up the dispatch stack to satisfy increasing demand. 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 §VI-B.

VIII. Conclusion

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

Acknowledgements

This work has been supported by NSF grants CNS-1143655, CNS-0916577, CNS-0855128, CNS-0834243, CNS-0845349, CNS-1217791.

References

[1] U.S. Department of Energy. Building Energy Data Book. http://buildingsdatabook.eren.doe.gov.

[2] SERC: State Environmental Resource Center. http://www.serconline.org/netmetering/stateactivity.html, 2011.

[3] IEC 61850. http://en.wikipedia.org/wiki/IEC_61850, 2012.

[4] 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.

[5] 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.

[6] C. Chang and C. Lin. LIBSVM: A Library for Support Vector Machines. *ACM Transactions on Intelligent Systems and Technology*, 2(27), April 2011.

[7] B. Daryanian, R. Bohn, and R. Tabors. Optimal Demand-side Response to Electricity Spot Prices for Storage-type Customers. *TPS*, 4(3), August 1989.

[8] Database of State Incentives for Renewables and Efficiency. http://www.dsireusa.org, 2010.

[9] EcoGrid EU FP7 project. http://www.eu-ecogrid.net/, 2012.

[10] The Economist, Lead-acid Batteries Recharged: A 150-year-old Technology Looks to the Future, Jan. 14 2009.

[11] eGauge Monitoring. http://www.egauge.net/, 2012.

[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] D. Gantenbein, C. Binding, B. Jansen, A. Mishra, and O. Sundstrom. EcoGrid EU: An Efficient ICT Approach for a Sustainable Power System. In *SustainIT*, October 2012.

[14] E. Hall. On a new action of the magnet on electric currents. *American Journal of Mathematics*, 2(3):287–292, November 1879.

[15] M. Hall. Correlation-based Feature Selection for Discrete and Numeric Class Machine Learning. In *ICML*, June 2000.

[16] D. Hammerstrom, T. Oliver, R. Melton, and R. Ambrosio. Standardization of a Hierarchical Transactive Control System. In *Grid Interop*, September 2009.

[17] 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.

[18] B. Jansen, C. Binding, and A. Mishra. Input on the Real-Time Price distribution protocol for Ecogrid EU WP 3 Task 3.6. http://www.zurich.ibm.com/pdf/ecogrid/price_distribution_protocol_1.1.pdf, July 2012.

[19] S. Kaplan. Power Plants: Characteristics and Costs. Technical report, Congressional Research Service Report to Congress, November 2008.

[20] I. Koutsopoulos, V. Hatzi, and L. Tassiulas. Optimal Energy Storage Control Policies for the Smart Power Grid. In *SmartGridComm*, September 2011.

[21] D. Linden. *Linden's Handbook of Batteries*, pages 17.1–17.39. McGraw Hill, Fourth edition, 2011.

[22] A. Mishra, C. Binding, D. Gantenbein, B. Jansen, and O. Sundström. Initial Inputs to Ecogrid Task 3.5 on DER and Aggregator Communication and Standards Contributions. http://www.zurich.ibm.com/pdf/ecogrid/SOTA_price-protocols.pdf, July 2012.

[23] A. Mishra, D. Irwin, P. Shenoy, J. Kurose, and T. Zhu. SmartCharge: Cutting the Electricity Bill in Smart Homes with Energy Storage. In *eEnergy*, May 2012.

[24] National Public Radio. How Smart is the Smart Grid?: Interview with Dan Delurey, President Demand Response Smart Grid Coalition, July 7th, 2010.

[25] New York ISO. Understanding the Markets. http://www.nyiso.com/public/about_nyiso/understanding_the_markets/energy_market, 2012.

[26] Ontario Energy Board: Electricity Prices. http://www.ontarioenergyboard.ca/OEB/Consumers, 2012.

[27] S. Schoenung. Energy Storage Systems Cost Update: A Study for the DOE Energy Storage Systems Program. Technical report, Sandia National Laboratories, 2011.

[28] SERC: State Environmental Resource Center. http://www.serconline.org/netmetering/stateactivity.html, 2012.

[29] N. Sharma, J. Gummeson, D. Irwin, and P. Shenoy. Cloudy Computing: Leveraging Weather Forecasts in Energy Harvesting Sensor Systems. In *SECON*, June 2010.

[30] Smart Grid Information Clearinghouse: Legislation and Regulation. http://www.sgiclearinghouse.org/Legislation?q=node/1705, 2011.

[31] State of California Executive Order S-21-09. http://gov.ca.gov/executive-order/13269, 2009.

[32] Technical Manual for Sun Xtender Batteries, July 2011.

[33] P. van de ven, N. Hegde, L. Massoulie, and T. Salonidis. Optimal Control of Residential Energy Storage Under Price Fluctuations. In *ENERGY*, May 2011.

[34] P. Vytelingum, T. Voice, S. Ramchurn, A. Rogers, and N. Jennings. Agent-based Micro-storage Management for the Smart Grid. In *AAMAS*, May 2010.

Aditya Mishra is a computer science Ph.D. student at the University of Massachusetts, Amherst. He received his M.Tech in Information Technology from Indian Institute of Technology, Bombay in 2007. His research interests are in the area of Smart Grids, Smart Homes, Green Computing, Computer Networks, Traffic Engineering, Quality of Service in Networks.

David Irwin is an Assistant Professor in the Department of Electrical and Computer Engineering at the University of Massachusetts Amherst. He received his Ph.D. and M.S. in Computer Science from Duke University in 2007 and 2005, respectively, and his B.S. in Computer Science and Mathematics from Vanderbilt University in 2001. His research interests are broadly in experimental computing systems with a particular emphasis on sustainability.

Prashant Shenoy received the B.Tech degree in computer science and engineering from the Indian Institute of Technology, Bombay, in 1993, and the M.S and Ph.D degrees in computer science from the University of Texas, Austin, in 1994 and 1998, respectively. He is currently a Professor of computer science at the University of Massachusetts. His current research focuses on cloud computing and green computing. He is a distinguished member of the ACM and a Fellow of the IEEE.

Jim Kurose received a B.A. degree in physics from Wesleyan University a Ph.D. degree in computer science from Columbia University. He is currently Distinguished University Professor in the Department of Computer Science. He has been a Visiting Scientist at IBM Research, INRIA, Institut EURECOM , the University of Paris, and Technicolor Research. He has served as Editor-in-Chief of the IEEE Transactions on Communications and was the founding Editor-in-Chief of the IEEE/ACM Transactions on Networking. He has been Program Co-chair of IEEE Infocom, ACM SIGCOMM, ACM SIGMETRICS and ACM e-Energy. He currently serves on the Board of Directors of the CRA, and the advisory council of the NSF CISE Directorate. He has received a number of research, education and service awards, including the IEEE/CS Taylor Booth Education Medal. He is a Fellow of the IEEE and the ACM.

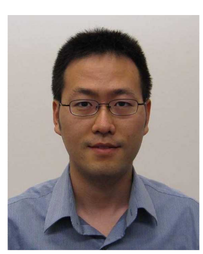

Ting Zhu is an assistant professor in the Department of Computer Science at the State University of New York Binghamton. He received his Ph.D. from the Department of Computer Science and Engineering at the University of Minnesota in 2010. He received his B.S. and M.E. degree from the Department of Information Science and Electronic Engineering at the Zhejiang University in 2001 and 2004, respectively. His research interests include wireless networks, renewable energy, embedded systems, distributed systems, and security.