

The Case for Efficient Renewable Energy Management in Smart Homes

Ting Zhu
Binghamton University
tzhu@binghamton.edu

Aditya Mishra David Irwin Navin Sharma Prashant Shenoy Don Towsley
University of Massachusetts Amherst
{adityam, irwin, nksharma, shenoy, towsley}@cs.umass.edu

Abstract

Distributed generation (DG) uses many small on-site energy sources deployed 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 many homes is difficult. In this paper, we explore a different approach that combines residential TOU pricing models with on-site renewables and modest energy storage to incentivize DG. We propose a system architecture and control algorithm to efficiently manage the renewable energy and storage to minimize grid power costs at individual buildings. We evaluate our control algorithm by simulation using a collection of real-world data sets. Initial results show that the algorithm decreases grid power costs by 2.7X while nearly eliminating grid demand peaks, demonstrating the promise of our approach.

Categories and Subject Descriptors

J.7 [Computer Applications]: Computers in Other Systems

General Terms

Design, Measurement, Management

Keywords

Building Energy, Smart Grid, Renewable Energy

1 Introduction

Buildings account for 40% of U.S. energy consumption [2], with the residential sector accounting for 54% of this total. The vast majority (70%) of this energy is from

electricity, which, due to environmental concerns, is largely generated at “dirty” power plants far from population centers. As a result, nearly half (47%) of residential energy consumption is lost during transmission and distribution (T&D) from power plants to distant homes [2]. One way to decrease both T&D losses and carbon emissions is through distributed generation (DG) from many small on-site energy sources deployed at individual buildings and homes. However, in practice, DG has drawbacks that have, thus far, prevented its widespread adoption. For instance, DG uses renewable wind and solar energy sources which buildings cannot dispatch at any time to satisfy their energy demands. Since the energy consumption density (kWh/sqft) of buildings is higher than the energy generation density of solar and wind deployments at most locations, buildings still must rely heavily on the electric grid for power.

More importantly, large centralized power plants benefit from economies-of-the-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 consumers sell the unused energy they produce back to the utility company—to offset their cost relative to grid energy. DG is a much less attractive option where net metering is not available. Net metering laws and regulations vary widely across states; it is not available in at least 4 states and the regulations are weak in many others [6]. Unfortunately, even where available, states typically place caps on both the total number of participating customers and/or the total amount of energy contributed per customer [3]. After exceeding these caps, utilities are no longer obligated to accept excess power from DG deployments. As one example, the state of Washington caps the total number of participating customers 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 it by making it difficult, or impossible, for utilities to balance supply and demand.

Today’s energy prices do not make DG financially attractive enough to consumers to reach even these low state caps. However, more widespread adoption 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 [5], while California’s Execu-

*This work was funded in part by NSF under grant #1019343 to CRA for the CI Fellows Project and grants CNS-1143655, CNS-0916577, CNS-0855128, and CNS-0834243.

tive Order S- 21-09 in California calls for 33% of generation from renewables by 2020 [4]. Given current laws, if DG becomes more widespread, residential consumers will have to look beyond net metering to reduce costs and balance on-site energy production and consumption. We envision consumers using a combination of on-site renewables, modest on-site 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 time-of-use (TOU) electricity pricing for residential consumers providing an opportunity to recoup the loss of net metering revenue, while also introducing new financial incentives for DG where net metering is not available.

The primary contribution of this paper is a new system architecture and control algorithm for managing on-site renewables, on-site energy storage, and grid energy in buildings to minimize grid energy costs for TOU electricity prices at small scale. Our system determines both the fraction of power to consume from the grid versus on-site battery-based 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) a TOU pricing model. 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 TOU rate period. We evaluate our algorithm by simulation using a collection of real data sets, including power consumption data from a real home, energy harvesting data from a solar panel and wind turbine deployment, National Weather Service (NWS) forecast data, and TOU pricing data from Ontario, Canada.

Our simulation results demonstrate the promise of our approach, by showing that our control algorithm reduces energy costs by 3.9X compared to homes without DG and 2.7X compared to homes with DG that do not control their renewable generation. Section 2 motivates our approach through a simple example using real-world data, while Section 3 provides an overview of our system and presents our control algorithm. Section 4 then evaluates our algorithm using the data sources mentioned above. Finally, Section 5 discusses related work, and Section 6 concludes.

2 Motivation

A key element of recent smart grid initiatives is the introduction of variable market-based TOU pricing to residential consumers. TOU pricing incentivizes consumers to lower their consumption when demand is high by increasing the price of electricity. While market-based pricing based on instantaneous supply and demand is common in wholesale electricity markets where utilities buy and sell energy, utilities are only beginning to introduce such pricing to residential consumers. The primary goal of TOU pricing is to reduce strain on the grid during demand peaks by incentivizing consumers to shift their consumption in time. TOU pricing may also lower generation costs, since they are disproportionately affected by the peak, rather than the average, electricity demand. As one example of TOU pricing today, the Ontario

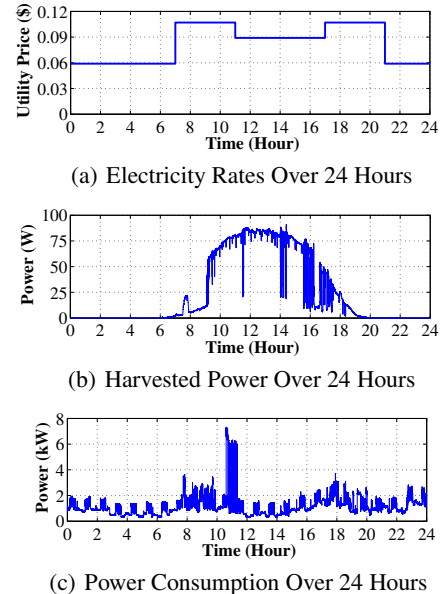


Figure 1. Electricity rates, solar harvesting, and energy consumption over 24 hours.

Electricity Board (OEB) introduced a simple pricing model for residential consumers in 2009.

Figure 1(a) shows one example of the electricity rates for a particular month, which the OEB divides into three categories: on-peak, mid-peak, and off-peak. The highest (on-peak) rate of 10.7 ¢/kWh is from 7am to 11am in the morning and from 5pm to 9pm in the evening, while the second highest (mid-peak) rate of 8.9 ¢/kWh is from 11am to 5pm in the middle of the day. The lowest (off-peak) rate of 5.9 ¢/kWh is from 9pm to 7am in the late evening and early morning. The pricing model is much simpler than in wholesale markets, where spot prices are not pre-set and vary as often as every 5 minutes. The OEB sets a different fixed ratio for on-, mid-, and off-peak rates in the summer (May 1st-October 31st) and winter (November 1st-April 30th), and on weekends and holidays. However, the exact rates change on a monthly basis according to generation costs and demand.

For instance, Ontario's electricity rate over the last two years increased by as much as 15.7% over one particular 6 month period. While renewable energy sources are able to offset rising costs, their power output also varies significantly over time. For instance, Figure 1(b) shows the power harvested from a solar panel over 24 hours during April in Amherst, MA. (within 400 miles of Ottawa, Ontario). Additionally, the power consumption of a home also varies considerably over time: Figure 1(c) shows a single day's power consumption for one 3-bedroom, 2-bath home at a nearby location. The figures show that local power generation and consumption is variable and not well-matched.

One way to reduce grid power costs is to change residential consumption patterns to align with low prices and plentiful local generation. However, the price elasticity of electricity demand is generally not high for residential consumers, i.e., price fluctuations do not readily alter consumer demand, in part, because consumers have little knowledge about how much power they are consuming or how much it costs. A key

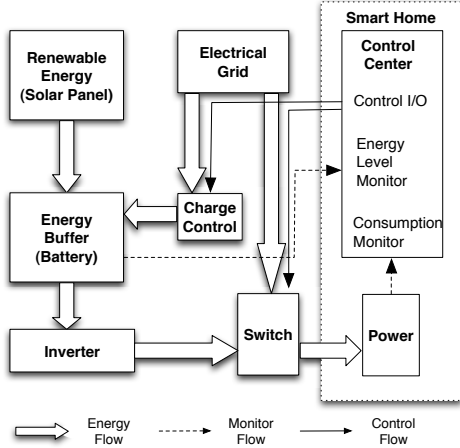


Figure 2. System Architecture

goal of our approach is to minimize grid energy costs without requiring a building’s occupants to change their energy consumption patterns or even think about them. Prior studies have shown that compelling consumers to change their routines is challenging [8]. Even for consumers that wish to alter their consumption to decrease costs, choosing which loads to disconnect and when is a complex decision that must be continuously re-evaluated based on information, e.g., energy generation, prices, etc., that is constantly changing. For Ontario, where price varies by a factor of 1.8 between the off-peak and on-peak rates, there is a strong incentive to adjust electricity consumption to reduce costs.

Rather than change consumption patterns, we propose using a modest amount of local battery storage combined with a control algorithm to minimize grid costs in the presence of an intermittent renewable energy source. To see how, again consider Figure 1(a) and Figure 1(b), which show that the highest electricity rates (from 7am to 11am) are within the period where renewable energy is not sufficient to power our example home. Thus, the building must use energy from the electric grid when rates are highest. However, given a modest size battery, the building is able to store and buffer grid energy during low rate periods to supplement renewable generation during high rate periods. The challenge is determining when and how much energy to store: if the building stores more energy than required, it wastes renewable energy due to limited battery capacity, while if it stores less energy than required, it must purchase grid energy at high rates.

3 Design

Figure 2 depicts the general architecture of our smart home energy management system. The heart of the system is the control center, which periodically records the home’s aggregate energy consumption and, based on historical use patterns and other information, predicts the expected energy consumption for the next 24 hours. Additional information that could aid predictions include data from other in-home sensors, e.g., occupancy or motion sensors, mobile phones, e.g., occupant GPS coordinates, or the Internet, e.g., online calendar programs. The control center also predicts the renewable energy the home expects to harvest over the next 24 hours based on the time of year and the next day’s weather

forecast. The control center then uses the predictions to determine how much energy it should store in the on-site battery based on the current battery level, the expected renewable energy, and the expected energy consumption.

3.1 Renewable Energy Prediction

For renewable energy, we use a prediction model similar to Sharma et al. [13] that translates a weather forecast from the National Weather Service (NWS) into a solar or wind energy harvesting prediction. We focus our experiments on solar energy, since it is the predominant renewable energy source in residential DG deployments, although the prediction model applies equally well to wind energy. We briefly summarize the model below, which uses the forecasted sky condition—as a percentage of cloud cover between 0% and 100%—to predict solar energy harvesting. The NWS releases a sky condition forecast, in addition to other weather metrics, every hour for the next 24 hours. At any time instance t , based on the sky condition percentage $C(t)$, we compute the solar array’s energy harvesting power $P_S(t)$ as:

$$P_S(t) = P_{max} \cdot (1 - C(t)) \quad (1)$$

Where P_{max} is the solar array’s maximum possible harvesting power. Sharma et al. [13] quantify the accuracy of Equation 1 and show that it is more accurate than existing techniques that use the past to predict the future. Thus, based on Equation 1, at any time instance t , we predict the solar energy harvesting within the next 24 hours as follows:

$$\hat{E}_S(t+T) = \int_t^{t+T} P_S(\tau) d\tau \quad (2)$$

Where T equals 24 hours. Without loss of generality, we assume that $t = kT$, since we run our algorithm every evening at the start of the 9pm low rate period, rather than at midnight $t = 0$. For simplicity, we use $\hat{E}_S(k+1)$ to represent $\hat{E}_S((k+1)T)$. Thus, we rewrite Equation 2 as follows:

$$\hat{E}_S(k+1) = \int_{kT}^{(k+1)T} P_S(\tau) d\tau \quad (3)$$

3.2 Energy Consumption Prediction

To predict the home’s energy consumption, we use a simple model based on an Exponentially Weighted Moving Average (EWMA). The EWMA exploits the diurnal nature of home consumption, while also adapting to seasonal variations. On a typical day, we expect the total energy consumption to be similar to the total energy consumption of previous days with slight deviations due to weather, e.g., a mild day that does not require A/C, or daily activities, e.g., use of the clothes dryer on laundry day. More sophisticated models are possible that take into account changing weekend activity patterns, weather conditions, or other information. One goal of this work is to quantify how much cost reduction we are able to achieve with a simple and straightforward prediction model. Let $E_C(k)$ denote the amount of energy consumed in the k -th day and $\hat{E}_C(k+1)$ denote the predicted energy consumed on the $(k+1)$ -th day, which is given by:

$$\hat{E}_C(k+1) = \alpha \cdot \hat{E}_C(k) + (1 - \alpha) \cdot E_C(k) \quad (4)$$

Where α is a weighting factor based on the observed prediction error over previous days. Since a TOU pricing model has different electricity rates at different time intervals within each individual day, we further predict the energy consumption at the lower rate and higher rate in the $(k+1)$ -th day by using Equations 5 and 6, respectively.

$$\hat{E}_{CL}(k+1) = \alpha \cdot \hat{E}_{CL}(k) + (1 - \alpha) \cdot E_{CL}(k) \quad (5)$$

$$\hat{E}_{CH}(k+1) = \alpha \cdot \hat{E}_{CH}(k) + (1 - \alpha) \cdot E_{CH}(k) \quad (6)$$

Where $E_{CL}(k)$ and $E_{CH}(k)$ are the actual energy consumption at the lower rate and higher rate on the k -th day, respectively.

Finally, we examine the lower rate and higher rate compared with the energy conversion efficiency of our system, which is the product of the battery's charging efficiency and the grid-tie inverter's efficiency [1]. Our goal is to charge the battery when the electricity rate is low, and discharge the battery to power the home when the rate is high. However, if the energy conversion efficiency is less than the ratio of the low rate and high rate values, then storing energy in the battery during low rate periods actually wastes more energy than directly using it from the grid during high rate periods.

As an example, using Ontario's TOU model, the ratio of the lowest rate (5.9 ¢/kWh) and the second highest rate (8.9 ¢/kWh) is $\frac{5.9}{8.9} = 66.29\%$. If the energy conversion efficiency is less than 66.29%, directly using grid energy during the 8.9 ¢/kWh period is more efficient than charging the battery at 5.9 ¢/kWh and discharging the battery during the 8.9 ¢/kWh time period. In this case, 8.9 ¢/kWh is not high enough to incentivize battery-based storage during the high rate period.

Since most lead-acid batteries have charging efficiencies greater than 85% [14] and most grid-tie inverters have efficiencies greater than 94% [1], the energy conversion efficiency is greater than $85\% \cdot 94\% = 79.9\%$, which is greater than the ratio of the lowest rate (5.9 ¢/kWh) and the second highest rate (8.9 ¢/kWh). Therefore, both the highest rate and the second highest rate for Ontario's model incentivizes battery-based storage. We use $\hat{E}_{CH}(k+1)$ as the expected total energy consumption in the $(k+1)$ -th day when the electricity rates are 8.9 ¢/kWh and 10.7 ¢/kWh.

3.3 An Efficient Control Algorithm

Given the simple harvesting and consumption prediction models above, we propose a simple control algorithm for minimizing grid power costs in DG deployments that decides how much energy to store in the battery based on the battery's remaining energy, predicted available environmental energy, and predicted energy consumption in the next 24 hours. Pseudo-code for the control algorithm, which we explain below, is shown in Algorithm 1. Here, $\hat{E}_r(k+1)$ is the expected energy remaining inside the battery that can be consumed in the $(k+1)$ -th day. We compute $\hat{E}_r(k+1)$ as follows:

$$\hat{E}_r(k+1) = \eta \cdot E_r(k) \quad (7)$$

Where η is the efficiency of the grid-tie inverter and $E_r(k)$ is the remaining energy inside the battery at the beginning of the lowest rate period of the k -th day.

To summarize, our control algorithm accounts for the following three cases to make charging decisions for each rate period within each day.

Algorithm 1: Efficient Control

```

1 if  $\hat{E}_r(k+1) + \hat{E}_S(k+1) \geq \hat{E}_{CH}(k+1) + \hat{E}_{CL}(k+1)$ 
  then
2   | Directly use the battery to power the house;
3 else if  $\hat{E}_r(k+1) + \hat{E}_S(k+1) \geq \hat{E}_{CH}(k+1)$  then
4   | while  $\hat{E}_r(k+1) + \hat{E}_S(k+1) - \hat{E}_{CH}(k+1) > 0$  do
5     |   | Directly use the battery to power the house;
6 else if  $\hat{E}_r(k+1) + \hat{E}_S(k+1) < \hat{E}_{CH}(k+1)$  then
7   |   | while  $\hat{E}_r(k+1) < \hat{E}_{CH}(k+1) - \hat{E}_S(k+1)$  do
8     |     | Charge the battery;
```

Case 1: If the sum of the battery's expected remaining energy $\hat{E}_r(k+1)$ and the predicted solar energy $\hat{E}_S(k+1)$ is greater than or equal to the total expected energy consumption during both the high rate and low rate periods, the home does not need to use any grid power. Instead, the control center uses the energy stored inside the battery, which includes the harvested solar energy, to directly power the home (lines 1 and 2). This case is unlikely, since our solar array is nowhere near large enough to power the home, and the majority of the harvesting occurs during a high rate period.

Case 2: If the sum of the battery's expected remaining energy $\hat{E}_r(k+1)$ and the predicted solar energy $\hat{E}_S(k+1)$ is greater than or equal to the expected energy consumption during the high rate period $\hat{E}_{CH}(k+1)$, then the battery has extra energy after the period ends, i.e., $\hat{E}_r(k+1) + \hat{E}_S(k+1) - \hat{E}_{CH}(k+1)$. The control center uses the extra energy to power the house during the low rate period (lines 3 to 5). We note that the home only uses battery power during the low rate period, if it is able to power the home during the high rate periods.

Case 3: If the sum of the battery's expected remaining energy $\hat{E}_r(k+1)$ and the predicted solar energy $\hat{E}_S(k+1)$ is less than the expected energy consumption during the high rate period $\hat{E}_{CH}(k+1)$, then the control center charges the battery during the low rate periods until the charged power and the predicted solar energy equals the expected consumption, i.e., $\hat{E}_r(k+1) + \hat{E}_S(k+1) = \hat{E}_{CH}(k+1)$ (lines 6 to 8).

While there are numerous ways to improve the simple control algorithm above, our goal in this paper is to quantify a lower-bound for potential cost reductions using simple approaches for prediction and control.

4 Implementation and Evaluation

We evaluate the performance of our control algorithm by simulation using empirical traces of (i) a solar array's harvesting power, (ii) a real home's power consumption, (iii) local weather forecasts, and (iv) Ontario's TOU pricing model. We set up and recorded the solar harvesting power using solar panels shown in Figure 3(a), which has a maximum power output of 185 watts in full sunlight. We measure the power consumption of the home by installing a TED 5000 in the home's electrical panel. The TED 5000 wraps 200Amp current transducers (CTs) around each leg of the home's split-phase input power from the grid (shown in Figure 3(b)), and records consumption every second. The man-




Figure 3. Experiment Setup




Figure 4. Energy Harvesting Power over 5 days




Figure 5. Power Consumption over 5 days

ual states that measurements are accurate to within 2% of the aggregate power down to a single watt. The measurement units in the panel transmit power data over the home’s power network to a server that stores and publishes it to a web interface. Our setup has been collecting consumption data for over a year.

In our evaluation, we compare our efficient control (EC) algorithm with the following two approaches:

No Renewable Energy (NRE): In this approach, the home purchases all of its energy from the grid, and does not use any renewable energy sources. This approach represents today’s homes without DG.

No Efficient Control (NEC): In this approach, the home uses renewable energy only as a supplement to the grid, by consuming it whenever it is available. The home stores any extra renewable energy in its battery, but never charge the battery from the grid. This approach represents today’s homes with DG, but without net metering available.

Figures 4 and 5 show the environmental energy harvested by the solar array, as well as the home’s power consumption, respectively, over a 5 day period. In our simulation, we scale up the solar generation from our small-scale deployment by 17 to closely align its aggregate energy consumption with that of the home. We choose a battery capacity of 12kWh,

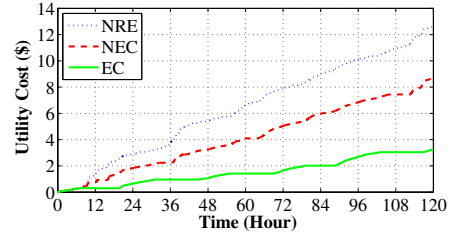


Figure 6. Utility Cost over 5 days

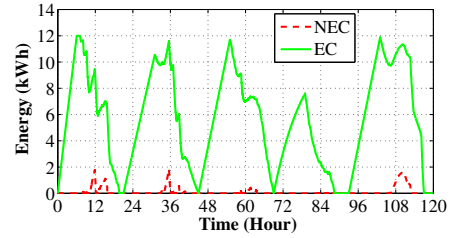


Figure 7. Remaining Energy inside Battery over 5 days




Figure 8. Power Requested from Grid (NEC)




Figure 9. Power Requested from Grid (EC)

which equals half the battery capacity of the recently introduced Nissan Leaf plug-in electric car. Our simulator assumes the initial energy inside the battery is zero.

Figure 6 compares the cumulative utility cost over time using the Ontario TOU pricing model for the three approaches—EC, NRE, and NEC—for the same 5 day period. Note that since our algorithm only considers a low and high rate period, we classify Ontario’s mid-rate period the same as high rate period. The figure shows our main result, which is that the control algorithm reduces grid power costs over the period by a factor of 3.9X, compared to the NRE approach that does not use renewable energy, and by a factor of 2.7X, compared to the NEC approach that does not control the use of renewable energy. Since a 12kWh lead-acid battery (with a 5 year warranty) costs roughly \$2,000, at current electricity rates the savings from our system recoups the battery’s cost in less than three years. Further, the long-term rising trend in electricity rates ensures that the savings will increase in the future.

To better understand how our EC algorithm attains these

cost reductions, consider Figure 1(a), which shows the change in the price of grid power over 1 day window, and Figure 7, which compares the battery's remaining energy when using the NEC and EC approaches. The figures show that by pre-charging the battery at the lowest rate, the energy stored in the battery combined with the harvested energy is capable of powering the house for significant periods of time during the high-rate periods. As a result, the home does not pay for grid power when it is most expensive. Figures 8 and 9 show the effect on the grid by comparing the grid power consumed using NEC and EC approaches.

By using the battery as an energy buffer, the EC approach requests zero energy from electrical grid during most of the high rate period. Additionally, the EC approach dramatically reduces the number of power consumption peaks. These peaks have a disproportionate effect on generation costs. Thus, reducing them has the potential to reduce prices across the entire the grid by reducing the cost to generate electricity. The peaks that the EC algorithm does allow are a function of the limited battery capacity in our simulation. The result shows that not only does EC have the potential to lower the cost of power for individual homes, but it also has the potential to lower generation costs across the entire grid.

5 Related Work

Using energy storage to reduce generation costs is normally conducted at large scales. Utilities use pumped-storage hydroelectricity, where they pump water uphill during periods of low demand and release it to generate electricity during periods of high demand. Pumped-water storage enables utilities to decrease the operating time and cost of the most expensive and inefficient generators, which only run during times of peak demand. While similar in principal, using energy storage at individual buildings to reduce costs presents different challenges than utility-scale systems.

Demand Forecasts. Across many homes electricity demand is highly predictable: utilities are able to develop accurate demand forecasts based on historical demand data and weather conditions. Individual homes exhibit more stochastic behavior based on the actions of a small group of occupants. There is room for significant improvement to the simple day-ahead demand forecasts we use in this paper. Developing accurate forecasts requires monitoring both building-wide and per-load power usage. Much prior work exists on load monitoring for a variety of building types [11, 12]. Demand forecasts may also incorporate auxiliary data from occupancy and other types of sensors [9, 10] to aid prediction.

Energy Harvesting Forecasts. Weather differs significantly at individual homes based on localized weather conditions. Thus, automatically developing per-site models that translate weather forecasts into energy harvesting predictions is important. We use a model by Sharma et al. [13] that predicts solar harvesting using the next day's sky condition forecast. Since today's residential TOU pricing models only change price 2 or 3 times per day (8 to 12 hours), buildings must be able to accurately predict energy harvesting at hour-to-day time-scales to significantly reduce costs.

Control Algorithm. Finally, large pumped-water storage has larger capacities and is more efficient than battery-based

storage. As a result, storing too much energy in batteries at the wrong time may actually increase costs. In general, a home control algorithm is more challenging than a utility-scale algorithm, since it must operate at finer time-scales, using more stochastic data sources, and with less margin for error. Additionally, the same approach must be applicable to homes of varying sizes and demand profiles, as well as different harvesting capabilities and TOU pricing models.

Our results show that our techniques are capable of reducing power drawn from the grid on average and during times of peak demand. Related work achieves similar goals by automating home energy management [7] using programmatic switches that disconnect devices during times of peak use. However, this requires new intelligent devices that communicate with the utility to determine their duty cycle. Our approach only uses local battery storage and renewables, and does not require replacing or augmenting existing loads.

6 Conclusions

In this paper, we propose a novel system architecture and control algorithm to efficiently manage renewable energy, battery storage, and grid power for DG deployments. Given a TOU pricing model, our control algorithm makes the decision based on the predicted future renewable energy generation and energy consumption. We evaluate our algorithm in simulation using real-world data sets, and show that the approach reduces grid power costs 3.9X relative to homes without DG deployments and 2.7X relative to today's homes with DG deployments (but without net metering).

7 References

- [1] Grid-tie Inverter. http://en.wikipedia.org/wiki/Grid-tie_inverter.
- [2] U.S. Department of Energy. Building Energy Data Book. <http://buildingsdatabook.eren.doe.gov>.
- [3] Freeing the Grid: Best and Worst Practices in State Net Metering Policies and Interconnection Procedures. <http://www.newenergychoices.org/uploads/FreeingTheGrid2009.pdf>, 2009.
- [4] State of California Executive Order S-21-09. <http://gov.ca.gov/executive-order/13269>, 2009.
- [5] Database of State Incentives for Renewables and Efficiency. <http://www.dsireusa.org>, 2010.
- [6] SERC: State Environmental Resource Center. <http://www.serconline.org/netmetering/stateactivity.html>, 2011.
- [7] N. Banerjee, S. Rollins, and K. Moran. Automating Energy Management in Green Homes. In *HomeNets*, 2011.
- [8] S. Davidoff, M. Lee, C. Yiu, J. Zimmerman, and A. Dey. Principles of Smart Home Control. In *UbiComp*, 2006.
- [9] V. Erickson and A. Cerpa. Occupancy Based Demand Response HVAC Control Strategy. In *BuildSys*, 2010.
- [10] V. Erickson, Y. Lin, A. Kamthe, R. Brahme, A. Surana, A. E. Cerpa, M. D. Sohn, and S. Narayanan. Energy Efficient Building Environment Control Strategies Using Real-Time Occupancy Measurements. In *BuildSys*, 2009.
- [11] X. Jiang, M. V. Ly, J. Taneja, P. Dutta, and D. Culler. Experiences with a High-Fidelity Wireless Building Energy Auditing Network. In *SenSys*, 2009.
- [12] Y. Kim, T. Schmid, Z. Charbiwala, and M. Srivastava. Viridroscope: Design and Implementation of a Fine Grained Power Monitoring System for Homes. In *UbiComp*, 2009.
- [13] N. Sharma, J. Gummesson, D. Irwin, and P. Shenoy. Cloudy Computing: Leveraging Weather Forecasts in Energy Harvesting Sensor Systems. In *SECON*, 2010.
- [14] J. Stevens and G. Corey. A Study of Lead-Acid Battery Efficiency Near Top-of-Charge and the Impact on PV System Design. In *PVSC*, 1996.