

Scaling Distributed Energy Storage for Grid Peak Reduction

Aditya Mishra, David Irwin, Prashant Shenoy, and Ting Zhu[‡]
University of Massachusetts Amherst [‡]Binghamton University
{adityam,irwin,shenoy}@cs.umass.edu †tzhu@binghamton.edu

ABSTRACT

Reducing peak demand is an important part of ongoing smart grid research efforts. To reduce peak demand, utilities are introducing variable rate electricity prices. Recent efforts have shown how variable rate pricing can incentivize consumers to use energy storage to cut their electricity bill, by storing energy during inexpensive off-peak periods and using it during expensive peak periods. Unfortunately, variable rate pricing provides only a weak incentive for distributed energy storage and does not promote its adoption at scale. In this paper, we present the storage adoption cycle to describe the issues with incentivizing energy storage using variable rates. We then propose a simple way to address the issues: augment variable rate pricing with a surcharge based on a consumer’s peak demand.

The surcharge encourages consumers to flatten their demand, rather shift as much demand as possible to the low-price period. We present PeakCharge, which includes a new peak-aware charging algorithm to optimize the use of energy storage in the presence of a peak demand surcharge, and use a closed-loop simulator to quantify its ability to flatten grid demand as the use of energy storage scales. We show that our system i) reduces upfront capital costs since it requires significantly less storage capacity per consumer than prior approaches, ii) increases energy storage’s ROI, since the surcharge mitigates free riding and maintains the incentive to use energy storage at scale, and iii) uses aggregate storage capacity within 18% of an optimal centralized system.

Categories and Subject Descriptors

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

Keywords

Energy, Battery, Electricity, Grid, Peak shaving

1. INTRODUCTION

As is now well-known, a significant fraction of the electric grid’s capital and operational expenses (CapEx and OpEx)

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

e-Energy’13, May 21–24, 2013, Berkeley, California, USA.
Copyright 2013 ACM 978-1-4503-2052-8/13/05 ...\$15.00.

result from satisfying its peak power demands. For example, recent work estimates that 10%-18% of North American CapEx, in terms of energy generation capacity, is idle and wasted over 99% of the year [8]. Similarly, peak demand also influences OpEx, by i) requiring utilities to operate high cost and inefficient “peaking” generators to meet demand [1], ii) contributing to higher transmission charges, which are set based on peak demand, and iii) forcing utilities to offset supply shortages by purchasing electricity in the wholesale market at inopportune times, i.e., when it is most expensive. Thus, reducing peak demand and its impact on CapEx and OpEx is an important part of ongoing smart grid research efforts. One way to reduce peak demand that has received significant attention in the research community is leveraging energy storage to shift some demand from peak to off-peak periods. To shift demand, prior work proposes to store energy during off-peak periods, which increases off-peak demand, and use it during peak periods, which then decreases peak demand [4, 7, 11, 13, 22, 23, 24].

To implement the approach, utilities may either i) install large-scale *centralized* energy storage systems at strategic points in the grid, such as at power plants and substations [13], and directly control when they store and release energy, or ii) incentivize consumers to install and control their own small-scale energy storage systems *distributed* at buildings throughout the grid. Prior research has focused largely on the latter case, since the increasing adoption of variable rate pricing plans by utilities [6, 16, 21] provides an incentive [4, 7, 11, 22, 23, 24], and endowing buildings with energy storage has additional value-added benefits, e.g., providing power during outages and conditioning power to increase its quality. Since variable rate pricing plans charge higher rates during periods of peak demand, consumers that store energy during off-peak periods—when prices are low—and use it during peak periods—when prices are high—are able to lower their electricity bill. While many energy storage technologies exist, including pumped water storage, flywheels, and compressed air, batteries are currently the most viable option for storing energy at building-scale.

Prior research analyzes the potential savings for residential [4, 7, 23, 24] and industrial [11, 22] consumers to install batteries. The focus is largely on cost-benefit analyses using existing pricing plans, which vary electricity’s price per kilowatt-hour (kWh). Unfortunately, for the reasons below, these plans *provide only a weak incentive for distributed energy storage and do not promote its adoption at large scales. Large Upfront Capital Costs.* Since today’s pricing plans typically exhibit low prices during off-peak nighttime

periods and high prices during peak daytime periods, they incentivize consumers to shift *all of their demand* to the off-peak period. Of course, the cost of batteries limits the amount of storage capacity available to shift demand. In our prior work on SmartCharge, we show that for a residential home with near the average U.S. electricity usage, $\sim 24\text{kWh}$ of capacity¹ maximizes the return-on-investment (ROI) when taking into account battery costs [15]. Given typical battery lifetimes, we estimate the *annual* amortized cost to maintain 24kWh of energy storage to be $\$1416$ [15]. Since the annual electricity bill for an average U.S. home is $\$1419$ [5], battery costs effectively prevent (at current price levels) a positive ROI using this much energy storage.

Rebound Peaks and Grid Instability. Current pricing plans incentivize *all consumers* to charge their batteries during off-peak, low-price periods. Thus, at large scales, simultaneous battery charging during off-peak periods will trigger rebound peaks if prices do not change to reflect the resulting increases in off-peak demand. Our prior work shows that if prices do not change and 100% of consumers install 24kWh of energy storage, then the peak demand period will migrate to the (previously) off-peak period and actually *increase*, rather than decrease, peak demand by nearly 120% [15]. Note that most variable rate pricing plans in use are Time-of-Use (TOU) plans with rates that do not react quickly to changes in demand, but instead are manually reset by utilities on an infrequent basis, e.g., monthly or seasonally [16].

Uncertain Return-on-Investment. One way to prevent rebound peaks is to alter electricity rates in real-time as peak and off-peak demand changes. Although not widespread, some utilities are experimenting with real-time pricing (RTP) plans for residential consumers, where rates vary dynamically each hour based on demand [6, 21]. Unfortunately, consumers only benefit from energy storage by exploiting the difference between peak and off-peak prices. With RTP plans, as peak demand declines and off-peak demand rises due to the increasing use of energy storage, the difference between the peak and off-peak price narrows, reducing energy storage’s benefits [24]. In the extreme, if grid demand is near flat then the price of electricity will be similar at all times [4, 15, 24]. Once the peak/off-peak price differential is not large enough to compensate for the conversion losses from storing energy in batteries, there is no benefit to using energy storage. Our prior work estimates that grid demand would be near flat once just 22% of consumers install 24kWh of energy storage [15], which is consistent with related work [4, 24]. Consumers are unlikely to invest in energy storage with such uncertain future long-term benefits.

Socialized Benefits and Free Riders. For residential consumers, the annual cost to install and maintain battery-based energy storage is much higher—around 10X for average consumers in the U.S.—than the annual savings on an electric bill using current battery costs, electricity rates, and pricing plans [15]. However, prior work does not consider the grid-wide reductions in generation costs from lowering the grid’s aggregate peak demand. Unfortunately, with existing pricing plans, these cost savings are distributed (or socialized) across all consumers, since they manifest themselves as cheaper electricity rates. Thus, variable rate pricing plans provide a weak, non-optimal incentive for energy storage. Strengthening the incentive requires eliminating free riders

to ensure that the consumers that invest in energy storage reap its full benefits, especially given the large capital costs.

The problems above arise from the interaction between current pricing plans and battery charging algorithms that minimize cost. We argue that solving these problems requires re-designing both pricing plans and charging algorithms to explicitly encourage energy storage adoption. In particular, any charging algorithm should prevent grid instability regardless of the pricing plan, similar to how TCP prevents Internet congestion even though it does not maximize end-user bandwidth. Likewise, pricing plans should sustain, not eliminate, the incentive to use energy storage as capacity scales. Finally, the charging algorithm and pricing plan should work together to ensure a stable grid, while also maximizing energy storage’s ROI at scale.

1.1 Contributions

Ideally, energy storage distributed at buildings throughout the grid would behave like centrally-controlled energy storage of equal capacity. That is, the “right” fraction of buildings would i) charge their batteries whenever grid demand is below average and ii) discharge their batteries whenever grid demand is above average, such that aggregate grid demand remains flat and constant at the average. Of course, ensuring the behavior of any self-organizing distributed system emulates that of an equivalently-sized centralized system is challenging. In this case, determining when and how many batteries should charge requires explicit feedback from the grid and coordination among all buildings, which does not scale. This paper targets an alternative approach: designing a charging algorithm and pricing plan where individual consumers (partially) flatten their own demand. As we discuss, our distributed approach does not require global coordination between consumers and the utility, and addresses each of the issues with scaling distributed energy storage.

The main drawback to incentivizing consumers to flatten their own demand is that it may require more aggregate energy storage capacity to flatten grid demand than the minimum required using a centralized approach. Since batteries are expensive, minimizing overall storage capacity and distributing it as widely as possible among consumers is critical to reducing per-consumer capital costs and increasing ROI. Our hypothesis is that, when consumers’ peak demand is well-aligned, a charging algorithm and pricing plan that flattens each consumer’s demand uses aggregate storage capacity near the optimal centralized approach. In evaluating our hypothesis we make the following contributions.

Scalable Design. We describe the *storage adoption cycle* that arises as energy storage scales. We show that existing charging algorithms and pricing plans cannot simultaneously minimize an electric bill and ensure grid stability at scale. In particular, preventing rebound peaks requires some (explicit or implicit) feedback from the grid to signal algorithms to rate-limit charging as demand rises. To resolve the cycle, we propose augmenting variable rate plans with a peak demand surcharge, and then modifying charging algorithms to account for it. Our system, called PeakCharge, is a complete redesign of our SmartCharge system [15] that optimizes a consumer’s electricity costs in the presence of a peak demand surcharge.

Closed-loop Experimentation. We implement a closed-loop simulator, which replays traces of real household demand, using a representative generator dispatch stack,

¹Operated at a maximum of 45% depth-of-discharge.




Figure 1: Prior switch-based architectures do not significantly lower an individual building’s peak demand. Figure from [15].

which specifies the cost to generate electricity as demand rises, to dynamically compute electricity rates based on demand. Our simulator is closed-loop since our charging algorithm reacts to the rates, which in-turn alters demand and then changes the rates. In contrast, prior work has evaluated energy storage using only open-loop simulations, where consumer behavior does not affect prices. Using our simulator, we experimentally verify the undesirable behavior of existing charging algorithms and pricing plans at scale.

Grid- and Consumer-scale Evaluation. We evaluate both the grid- and consumer-scale effects of PeakCharge, comparing it with prior “greedy” approaches that store as much energy as possible during low-price periods. Our analysis shows that, when compared with these systems, PeakCharge i) reduces upfront capital costs since it requires significantly less storage capacity per consumer and ii) increases ROI, since a peak surcharge mitigates free riding and maintains energy storage’s incentive at large scales, while requiring aggregate storage capacity within 18% of optimal.

2. OVERVIEW AND APPROACH

Our work leverages the use of battery-based energy storage systems to reduce electricity costs. We assume an intelligent battery-based energy storage system that is capable of determining when, and how much, to charge and discharge batteries based on variable electricity rates over time to minimize electricity costs. To be cost-effective, these systems must i) limit energy storage capacity due to battery costs, which, amortized over their lifetime, are currently \$100-\$200 per year per kWh of usable capacity for the VRLA/AGM lead acid variety widely used in stationary energy storage systems, and ii) account for the ~20% conversion loss from storing energy in batteries. Note that, since a lead-acid battery’s lifetime is a function of its depth-of-discharge (DOD), a 24kWh battery operated at 50% DOD has only 12kWh of usable capacity. As in past work, we consider both the savings from batteries and their cost (20% energy loss and capital cost) when considering a system’s ROI.

2.1 PeakCharge Architecture

Previously proposed architectures for leveraging energy storage [15] use a programmatic power transfer switch, which allows them to toggle a building’s power supply between the grid and a battery. Thus, in addition to a charging algorithm that decides when and how much to charge batteries, the system also decides when to toggle the building’s power supply between the grid and the battery, based on

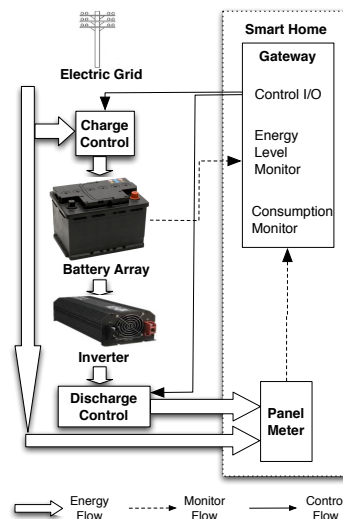


Figure 2: PeakCharge architecture, which includes a battery capable of programmatically controlling the rate of discharge wired in parallel with the grid.

expectations of future prices and demand. Of course, when batteries supply power, the building’s load dictates the rate of discharge due to Kirchhoff’s laws. Although not programmatic, such switches are common in commercial standby UPS systems, which automatically switch to battery power when grid voltage falls below a preset threshold. The coarse switching architecture works well in previous systems, since they connect to the grid and charge batteries during lengthy low-price periods at night before switching to battery power during lengthy high-price periods during the day.

In contrast, we assume a system architecture that is capable of controlling a battery’s rate of discharge independent of the building’s load. For example, if a building is consuming 1kW of power, the system is able to control the fraction of the 1kW the battery supplies, with the grid supplying the remainder. Thus, the system may choose to satisfy 1kW of demand using 500W via the battery and 500W via the grid, or using 200W via the battery and 800W via the grid. Controlling the rate of discharge is necessary for PeakCharge’s approach, which encourages buildings to flatten their demand rather than simply shift large amounts of demand from daytime to nighttime. As Figure 1 demonstrates, for individual buildings, the simple switching architecture does not significantly reduce (or flatten) an individual building’s peak demand. The figure (from [15]) illustrates how, due to off-peak battery charging, our prior switch-based SmartCharge system simply shifts the original peak demand to the off-peak period to minimize electricity costs.

There are two primary ways to control a battery’s rate of discharge. A simple approach is to install multiple switches capable of switching separate fractions of a building’s load between grid and battery power. For example, the system may be able to individually switch each circuit. In this case, the system controls the rate of discharge by monitoring the load on each circuit and switching some subset of circuits to the battery to achieve a specific rate of discharge. An alternative, cleaner approach depicted in Figure 2 is to connect the battery in parallel to the grid and use a discharge controller to programmatically limit the rate of discharge. These controllers use pulse-width modulation (PWM) to

control the charge or discharge rate by connecting and disconnecting the battery at rapid frequencies. Unfortunately, controllers capable of programmatically setting the rate of discharge are not widely available, since their primary purpose today is in testing equipment [26]. However, programmatic control may become more widespread in the future, since recent work beyond our own also requires this capability [14, 25]. We assume this latter method is available to control the discharge rate in PeakCharge.

Finally, both our work and prior work derives from the fact that the marginal cost for a utility to generate each additional watt of power increases non-linearly as utilities activate additional generators to satisfy increasing demand. Utilities maintain a dispatch stack of generators: as grid demand rises utilities activate, or “dispatch,” additional generators in ascending order of their marginal cost. Figure 3 shows the demand-cost function we use to compute generation costs based on demand in our closed-loop simulator, and demonstrates the non-linear relationship between cost and demand. To derive our function, we scaled real demand-cost data from the Southeastern U.S. from a 2008 report [9] by the Federal Energy Regulatory Commission (FERC) to match the peak demand in our traces, discussed in Section 5, while also ensuring a median electricity cost of 10¢/kWh, which is near the average cost of electricity in the U.S. We then fitted an exponential function to this scaled data for use in our simulations.

2.2 The Storage Adoption Cycle

Figure 4 depicts the *storage adoption cycle* that arises from the use of distributed energy storage at large scales to minimize electricity costs in the presence of variable rate electricity prices. At the top of the figure, variable demand for power first causes the price of electricity over time to change based on the demand-cost function from Figure 3. Variable pricing, in turn, incentivizes consumers to adopt energy storage to reduce their costs by shifting demand to low price periods. However, as more consumers shift demand using energy storage, the difference between the grid’s peak and off-peak demand narrows resulting in a flatter grid demand profile. As a result, prices also flatten to reflect the new demand distribution. Unfortunately, flat prices eliminate the incentive to use energy storage, which causes demand to vary again and the cycle to repeat. Of course, our depiction is idealistic. It assumes, first, a high enough price differential to warrant energy storage, which is not the case today, as outlined in Section 1. Second, grid/utilities may not explicitly want to incentivize energy storage. In practice, completing each step would take a long time, potentially requiring significant regulatory changes and large capital investments. Further, the phenomenon depicted in Figure 4, and described below is only for illustrative purposes to demonstrate potential trends as the use of energy storage scales.

The storage adoption cycle may also cause grid instability if prices do not react fast enough to changes in demand. To demonstrate the potential for instability, we ran a simple experiment using our trace-driven closed-loop simulator to show how grid power demand could experience significant oscillations even if utilities alter prices each day based on the previous day’s demand. Day-ahead planning is common, since consumers require some pricing feedback to adjust their behavior and utilities require some advance notice




Figure 3: The model we use in our simulator of the marginal cost to generate electricity as demand increases. The fitted function we use is based on scaled data from a recent FERC report [9].

to activate generators. The simulator, which we discuss in Section 5, takes traces of demand as input; in this case, we use demand traces we collected of 194. For our experiment, 49 of the 194 homes have usable energy storage capacity that is 50% of their average daily demand, where each home uses a “greedy” charging algorithm similar to prior work [15], which minimizes electricity costs by charging as much as possible during the lowest price periods. Our simulator is closed-loop, since it computes the next day’s prices each hour using Figure 3’s demand-cost function and the previous day’s demand.

Figure 5 shows how the peak demand periods change dramatically each day. On the first day, everyone charges during the low-price period at night (12am-6am), which increases demand during that period and, hence, also increases the price of electricity during that time on the second day. As a result, on the second day the lowest-price period shifts to the morning (6am-12pm), which is the low-demand period from the previous day, and causes peak demand to shift dramatically from the nighttime 12am-6am period to the morning 6am-12pm period. Since generators require lead time to activate, utilities carefully plan generator dispatch schedules each day based on the previous day’s demand. If demand were to change dramatically each day, as in this scenario, the grid would be incapable of balancing supply and demand. This simple example highlights how naive battery charging-discharging algorithms can potentially cause grid instability for some of the existing pricing plans, like [6], [16], under certain conditions. Note that, in this paper, we do not evaluate the cost-benefit tradeoff between the excessive distributed storage (caused by our approach) and the potential affects of the storage adoption cycle.

2.3 An Optimal Approach

Before describing PeakCharge’s charging algorithm, we first define and consider an optimal centralized battery charging scheme. Ideally, to minimize generation costs based on the demand-cost function from Figure 3, the optimal approach would shift aggregate grid demand such that it was the same—equal to average demand—all the time. If we assume a centrally controlled battery array, then an optimal algorithm simply charges and discharges batteries whenever grid demand is below or above average, respectively, such that demand is always equal to the average. With this algorithm, the minimum energy capacity necessary to flatten demand is equal to the maximum capacity ever required




Figure 4: Illustration of the storage adoption cycle.

to charge or discharge the batteries to sustain the average. Of course, as mentioned in Section 1, a centralized system has drawbacks compared to the distributed approach, which provides value-added benefits to consumers.

As an example, Figure 6 depicts a grid demand profile from the first day of our trace of 194 homes, as well as the average demand for the day. In this case, the maximum capacity required to charge or discharge the battery occurs between hour 16 and 23 and equals 392kWh (equivalent to the area between the instantaneous demand and the average demand from hour 16 to 23). With the optimal approach this 392kWh of storage would reduce generation costs by 23% based on our demand-cost function in Figure 3. If this storage capacity were distributed evenly among all 194 homes, then each home would need only 2.02kWh of usable energy storage (or 4.5kWh of rated capacity used at 45% depth-of-discharge to maximize lifetime). This capacity is over 5X less than the 24kWh of rated capacity each home requires to maximize energy storage’s ROI based on our previous SmartCharge work [15], which uses existing variable rate pricing plans and a “greedy” charging algorithm. Qualitatively, this result holds for any greedy battery charging-discharging approach for price saving, and not just SmartCharge. Since battery costs scale linearly with capacity, maintaining 5X less capacity decreases costs by 5X (from \$1416 amortized per year to maintain 24kWh to \$266 per year to maintain 4.78kWh). The example demonstrates how minimizing capacity, and distributing it as widely as possible among consumers reduces the ROI per consumer of energy storage.

3. SCALABLE DESIGN

The storage adoption cycle discourages distributed energy storage from scaling to a large fraction ($> 20\%$) of consumers. Unfortunately, variable rate electricity prices incentivize consumers with energy storage to use greedy battery charging algorithms, which charge batteries as much as possible during the lowest price periods to minimize electricity costs. At large scales, the use of greedy charging algorithms results in either i) large rebound peaks (if prices do not react to changing demand), ii) grid instability (if prices react slowly to changing demand as in Figure 5), or iii) no benefit to the consumer (if prices react quickly to changing demand by flattening). None of these outcomes is desirable. Variable rate pricing is effective at reducing peak demand today only because electricity’s price elasticity of demand is typically low, i.e., consumers do not react strongly to




Figure 5: Load oscillations in our simulated microgrid, in presence of day-ahead real time pricing.

changes in electricity’s price. As a result, only a small fraction of consumer demand shifts to low price periods. In contrast, large-scale distributed energy storage makes electricity’s price completely elastic with demand, causing a large fraction of demand to shift to the lowest price period.

Properly incentivizing distributed energy storage at scale requires rethinking electricity pricing plans. Our premise is that augmenting existing variable rate pricing plans with a peak demand surcharge (or penalty) is a simple and effective way of addressing the storage adoption cycle. A peak demand surcharge bills consumers $\$X/\text{kWh}$ based on their peak demand over an N minute interval within some billing period M . Typical values are $N = 1$ hour and $M = 1$ day; for example, in this case, the consumer in Figure 1 would incur an additional charge for using $\sim 5\text{kWh}$ during their peak hour of that day. Utilities already use such a peak demand surcharge for large, primarily industrial, consumers. Put simply, a large peak demand surcharge incentivizes consumers to flatten their own demand to minimize their peak, rather than simply shift as much demand to the lowest price period. Of course, if consumers flatten their own demand, then grid demand will also flatten. We discuss below, how penalizing peak usage addresses the problems from Section 1. While the incentive to flatten due to a peak surcharge disappears once a home reaches peak demand, we design our proposed peak-aware charging algorithm to avoid peaks when possible. Our results indicate that the algorithm is successful most of the time.

3.1 Benefits of a Peak Demand Surcharge

First, flattening a consumer’s demand takes significantly less energy storage capacity than shifting all of it to the lowest price period. For example, Figure 7 shows that, while 12kWh of usable energy storage is only capable of shifting a fraction of demand to the low price period, it is more than enough to completely flatten the original demand from Figure 1. As a result, the approach encourages distributing aggregate storage capacity widely across consumers, requiring less storage capacity per consumer, and resulting in lower upfront capital costs and higher per-consumer ROI. In effect, to flatten grid demand, the approach incentivizes a large number of consumers to install a small amount of energy storage (and make a small investment), rather than incentivizing a small number of consumers to install a large amount of energy storage (and make a large investment). Second, the approach prevents rebound peaks and grid instability, since consumers are (partially) flattening, rather than shifting, their demand. Third, the approach maintains the incentive to use energy storage as capacity scales, since con-




Figure 6: Instantaneous and average grid demand for 194 homes in our trace.

sumers always benefit from not paying an additional peak demand surcharge, regardless of other consumers’ behavior. Finally, utilities can mitigate free riding by altering the peak demand surcharge, in addition to the electricity rate, as generation costs change, since only the set of consumers with energy storage are able to automatically optimize for peak demand. Thus, a higher peak demand surcharge and lower rates will penalize consumers with energy storage less than consumers without it.

Of course, utilities must use a peak demand surcharge in conjunction with existing variable rate schemes, since only charging based on peak demand would encourage *more* energy use. For example, with only a peak demand surcharge, if a residential consumer runs a dryer that causes their peak load to be 7kW over an hour, the consumer has *no incentive to ever reduce their demand below 7kW*. Since the home in Figures 1 and 7 has an average electricity usage of ~ 1 kW, billing solely based on peak demand would allow the consumer to use 7X more electricity at no extra cost. Thus, utilities must balance the size of the peak demand surcharge with the electricity rates to encourage flattening without incentivizing consumers to use significantly more electricity. We examine how the size of the peak demand surcharge affects our battery charging algorithm in Section 5.

While this paper focuses primarily on how a peak demand surcharge addresses the storage adoption cycle, it also has other benefits. For instance, homes without energy storage could reduce their electricity costs using automated load scheduling techniques that flatten demand, e.g., via SmartCap [3] or nPlug [10]. Consumers have little monetary incentive to use these techniques today, since most deferrable loads, e.g., refrigerators, air conditioners, heaters, dehumidifiers, are unable to defer their usage (by up to 12 hours) to low-price nighttime periods without causing significant harm, e.g., spoiled food or an uncomfortable environment. In addition, as recent work shows, flattening demand using a battery preserves privacy [14, 25], since it removes power variations that Non-Intrusive Load Monitoring (NILM) algorithms use to identify appliance usage and behavioral patterns. Unfortunately, with existing variable rate plans consumers with a battery must choose to either use it to reduce their electricity bill or preserve privacy, but not both. A peak demand surcharge could enable consumers to minimize their electricity bill *and* preserve their privacy.

3.2 Drawbacks of a Peak Demand Surcharge

The primary drawback to encouraging consumers to flatten their own demand is that, in aggregate, it may require consumers to install more energy storage capacity than nec-




Figure 7: While 12kWh of energy storage is capable of shifting only a fraction of demand to the low price period, it is more than enough to completely flatten the demand from Figure 1.

essary to flatten grid demand. To understand why, consider a simple grid with only two homes, where each day the first home uses 1kW from 12am-12pm and 2kW from 12pm-12am, while the second home uses 2kW from 12am-12pm and 1kW from 12pm-12am. In this case, to flatten their own demand, each home requires 6kWh of energy storage for a total of 12kWh of capacity, which the homes would charge at a rate of 500W/hour when usage is 1kW and discharge at a rate of 500W/hour when usage is 2kW. However, in aggregate, the two homes’ demand is already flat—using exactly 3kW all the time—without any energy storage. Thus, in this case, energy storage is not necessary. The waste occurs because the peak periods of the two homes are not aligned with each other; if their peak periods were exactly aligned then they would each require 6kWh to flatten demand (and 12kWh would be the minimum capacity necessary to flatten aggregate demand). In general, the aggregate energy storage capacity necessary to flatten grid demand by flattening each home’s demand will diverge more from the optimal amount the more the peak and off-peak periods of the homes become less aligned. Of course, in practice, homes in the grid exhibit peak demand at similar times, which naturally reduces the divergence from optimal. We quantify this divergence using our closed-loop simulator and demand traces in Section 5.

As we discuss, augmenting existing variable rate plans with a peak demand surcharge requires rethinking the greedy charging algorithms used in prior work. Below, we present PeakCharge’s peak-aware charging algorithm, which minimizes a consumer’s electricity bill in the presence of a peak-demand surcharge.

4. PEAK-AWARE CHARGING

Our initial approach to designing PeakCharge’s battery charging algorithm was to simply modify the algorithm from our prior work on SmartCharge [15]. SmartCharge uses a linear program (LP) that executes at the beginning of each day and takes as input next-day electricity prices, which are typically well-known, and expected demand each hour to determine how to charge and discharge a battery throughout the day. The LP is optimal, i.e., minimizes costs, if future demand is known. For completeness, we include a description of this LP and its constraints in the Appendix. Since future demand is not known, we use machine learning, specifically a support vector machine with a polynomial kernel, to predict next-day demand over the five multi-hour

periods each day in Ontario’s TOU pricing plan [16]. On average, our predictions were within 6% of demand for a representative home, and SmartCharge’s LP achieved cost savings within 10-15% of an oracle with perfect future knowledge. As a result, for PeakCharge, we initially added constraints to SmartCharge’s LP to account for the new peak demand surcharge. As with SmartCharge, given perfect future knowledge, the LP is optimal at minimizing costs. Unfortunately, our experimental results were far from optimal: the PeakCharge variant *did not* result in flatter consumer demand and *did not* minimize electricity costs. In this case, rates were based on the Ontario TOU scheme, and the peak demand surcharge was applied to the peak hour of each day.

We found the reason for the poor results to be that, with a peak-demand surcharge, the LP is highly sensitive to the prediction of the peak demand hour each day. Unfortunately, predicting next-day demand at the granularity of an hour is much less accurate than over the multi-hour periods used in SmartCharge. Further, ensuring high prediction accuracy for the peak hour is more difficult than ensuring a high average accuracy. In contrast, SmartCharge, which only optimized for variable electricity prices, is much less sensitive to prediction accuracy. While there are corner cases the LP is able to optimize for with accurate predictions, in general, it will always charge a battery as much as possible during the lowest-rate nighttime periods. Thus, simply charging the battery at its maximum rate overnight, regardless of the predictions of next-day demand, accounts for the vast majority of the savings in SmartCharge and other systems. One implication of a sensitivity to demand prediction accuracy is that optimizing for a peak demand surcharge becomes more difficult the longer the time interval M the peak is evaluated over, since predictions are generally less accurate over longer time horizons. For instance, many utilities charge industrial consumers a surcharge for their peak demand hour within an entire month, requiring them to accurately predict demand (including the peak) each hour of the month to determine when to charge and discharge the battery using our LP above.

4.1 Optimizing for the Peak

Based on our experiences with the LP, rather than retrofit a greedy algorithm originally designed to minimize costs for variable electricity rates, to account for a peak demand surcharge, we instead began by designing an algorithm to minimize costs for a peak demand surcharge in the absence of variable electricity rates. Our starting point is the same algorithm we outlined in Section 2.3 for flattening grid demand, but applied to individual buildings. The system selects a target average power and then simply charges and discharges batteries whenever demand is below or above the target, respectively, such that demand is always equal to the average. Note that rather than run our algorithm once per day (at the beginning of the day) using predictions of next-day demand, as with SmartCharge, this algorithm naturally operates in an online manner, adjusting the charging and discharging of the battery in real time based on changing demand. This peak-centric algorithm works well as long as i) the target average is near the actual average power, and ii) the storage capacity is large enough to flatten demand.

If the target average is too small, then the approach will not store enough energy to reduce the peak by its maximum amount; if the target is too large, then it will store more

energy than necessary throughout the day. However, importantly, while the algorithm is sensitive to a prediction of average power, it does not require shorter time-scale predictions of future demand, e.g., hourly day-ahead predictions, as in the LP approach. Average power predictions over long time periods tend to be much more accurate than demand predictions over short time-scales far into the future. In fact, when predicting average power, the longer the time-scale, generally the more accurate the prediction [19], e.g., the average power of a home each year tends to vary less than each day. In addition, accurate predictions of average power over long periods, e.g., a day or month, do not require sophisticated methods [19, 20]. In this paper, to predict average demand over an interval, we simply use the average demand over the previous interval.

If the available storage capacity is too small, then the approach may discharge batteries when demand is only slightly above average, causing there to be little energy left for the highest peaks. In this case, short time-scale predictions of future demand are necessary to optimize use of the available storage capacity, i.e., save stored energy for the highest peaks each day or month. Thus, with a peak demand surcharge, the less storage capacity a consumer has, the more fine-grained and accurate the predictions required to minimize cost. However, importantly, as discussed earlier, flattening consumer demand requires much less energy storage capacity than shifting it to take advantage of variable rates.

4.2 Optimizing for Peaks and Variable Rates

Our peak-centric algorithm focuses only on flattening demand. As a result, it minimizes a consumer’s electricity costs when their bill is based solely on a peak demand surcharge in absence of variable rates charged per kWh of energy use. Given our basic algorithm, we must modify it to optimize for cost in the presence of both a peak demand surcharge and variable rates. With a high peak demand surcharge the algorithm should behave like the peak-centric algorithm, and with a low peak demand surcharge the algorithm should behave greedily, i.e., by charging at its maximum rate during low-price periods. To understand the decision of whether to behave greedily or peak-centric, consider the inequality below, which compares the benefit of greedily taking advantage of variable rates versus the cost of raising peak demand. In this case, we consider only two rate periods: high and low, where C_{high} is the cost per kWh during the high rate period and C_{low} is the cost per kWh during the low rate period. In addition, T is the length of the low-price period, P is the cost per kWh of usage during the peak hour each day, e_{loss} is the energy conversion loss as a percentage stored energy (typically 80% in practice), and X_{max} is the maximum charging rate of the battery.

$$X_{max}e_{loss}C_{high}T - X_{max}C_{low}T > X_{max}P \quad (1)$$

The left side of the inequality is the maximum monetary benefit of greedily charging the battery at its maximum rate during the low-price period and then discharging it during the high-price period, while the right side is the cost of the peak demand surcharge from charging the battery at its maximum rate. If the inequality holds then the benefit of charging greedily during the lowest-price period is greater than the cost, signaling that a consumer should act greedily. If not, then a consumer may benefit from acting peak-centric by charging (or discharging) at less than the

maximum rate during low-price periods. Unfortunately, determining exactly how much less to charge (or discharge) than the maximum is challenging, requiring the same accurate short time-scale, e.g., hourly, predictions of future demand that SmartCharge’s LP requires. As a result, we adopt a heuristic approach using four simple cases, as outlined below, based on whether the electricity rate is high or low and the demand is above or below average.

- If the electricity rate is **low** and demand is **below average**, then greedily charge at the maximum rate if (1) holds, else charge at a rate to sustain the target average demand.
- If the electricity rate is **low** and demand is **above average**, then greedily charge at the maximum rate if (1) holds, else discharge at a rate to sustain the target average demand.
- If the electricity rate is **high** and demand is **below average**, then greedily discharge at the full rate (bounded by the building’s demand) if (1) holds, else do nothing.
- If the electricity rate is **high** and demand is **above average**, then greedily discharge at the full rate (bounded by the building’s demand) if (1) holds, else discharge at a rate to sustain the target average demand.

Rather than add more cases, to extend the approach to multiple rate periods, we simply divide each period into two bins, based on whether its price is higher and lower than average, and compute C_{high} and C_{low} by taking the average of the cost per period (weighted by the length of the period) in each respective bin. Based on the cases, if the inequality holds then the algorithm simply acts greedily by charging at the maximum rate when the electricity price is low, while discharging at the maximum rate (bounded by the building’s demand) when the electricity price is high. In contrast, if the inequality does not hold, then the algorithm simply toggles to using the peak-centric algorithm, with one exception. If the electricity rate is high and demand is below average, it balances the objective of the greedy algorithm, i.e., to discharge, and the peak-centric algorithm, i.e., to charge, by doing nothing. Note that, in the extreme, since variable rates are based on the grid’s demand, as grid demand flattens the rates will equal each other and the algorithm will become entirely peak-centric.

Using our peak-aware algorithm above, when the peak demand surcharge is high relative to the electricity rates, the algorithm above charges and discharges the battery to hit the expected average demand; in contrast, when it is low, the algorithm devolves to a greedily charges the battery at the maximum rate during the lowest price periods.

4.3 Summary

Our *peak-aware* algorithm above optimizes for a peak demand surcharge by using inequality (1) to determine when to act greedily and when to optimize for the peak. In the next section, we compare the peak-aware algorithm with an online *greedy* algorithm that is conceptually similar to our previous LP-based approach, but operates in an online manner by charging at the maximum rate during the lowest-price periods at night and discharging during the highest price periods during the day. Since battery capacity is typically much lower than each day’s energy usage, this simple variant performs similarly to our previous LP-based approach.

As an additional point of comparison, we also experiment

with a variant of the greedy algorithm with an additional congestion parameter, which limits the maximum charging rate of the battery by a factor P_{limit} , which is between 0 and 1. Enforcing a limit on the battery charging rate is a simple way to ensure grid stability and prevent rebound peaks, even using greedy charging. Of course, in practice, this parameter requires feedback and enforcement from a utility, which could either directly disseminate P_{limit} to consumers or allow them to indirectly infer it, e.g., by using subtle changes in line voltage as a signal of grid demand as in nPlug [10]. We show that while our *congestion-aware greedy variant* prevents rebound peaks and grid instability, without a peak demand surcharge, it reduces the savings (and ROI) of energy storage for consumers.

5. EVALUATION

To evaluate the charging algorithms from the previous section, we built a closed-loop simulator that takes as input traces of building energy usage. The simulator is closed-loop, since it determines the price of electricity each hour of each day based on the demand from i) the same hour on the previous day and ii) the demand-cost function in Figure 3; we call this Day-Ahead Real-Time (or DART) pricing, since each day’s prices are known at the beginning of the day. In addition to DART, our simulator also supports open-loop TOU pricing, where prices do not change based on demand. As in our prior work [15], we use TOU prices based on Ontario’s rates [16]; specifically, 6.3¢ per kWh from 11pm to 6am (off-peak period), 11.8¢ per kWh between 6am to 10am and 4pm to 11pm (peak periods), and 9.9¢ per kWh from 10am to 4pm (mid-peak period).

In addition to the rate plans above, the simulator also supports a peak demand surcharge in \$/kW of peak usage. The surcharge applies to the highest average demand over a 30 minute sliding window across each day. Our default surcharge in the experiments below, unless otherwise noted, is \$3/kW. This surcharge is high relative to the rates, i.e., inequality (1) does not hold, although with DART pricing rates may rise (since they vary every day based on demand). As in our prior work [15], we use a maximum charge rate of $C/4$ for the usable storage capacity, i.e., the battery charges to full capacity in 4 hours, which translates to a $C/8$ rate for a battery used at 45% DOD. We use power demand traces from 194 homes, which have an in-panel energy meter to record usage each minute, for ten consecutive days. While our traces are not at utility scale, i.e., with tens of thousands of residential homes as well as commercial and industrial buildings, they are sufficient to verify the trends in using energy storage at scale and to explore the behavior of our algorithm. However, the benefits of storage at scale will certainly vary based on the characteristics of each grid’s (and building’s) demand profile.

As with our previous work [2], we plan to make our traces available for download from our Smart* data repository located at <http://smart.cs.umass.edu>. Finally, we experiment with the algorithms from the previous section—greedy, peak-aware, and congestion-aware greedy—using DART (closed-loop) and TOU (open-loop) rate plans, examining both the grid-scale and consumer-scale effects.

5.1 Grid-scale Effects

We first examine the effect of rebound peaks when consumers use energy storage at large scales. Figure 8 shows the




Figure 8: Generation cost savings compared to using no energy storage for both closed-loop DART (a) and open-loop TOU (b) pricing plans. Zoom-in of generation cost savings for peak-aware algorithm (c).




Figure 9: Peak reduction as a percentage compared to using no energy storage for both closed-loop DART (a) and open-loop TOU (b) pricing plans. Zoom-in of peak reduction for peak-aware algorithm (c)

savings in generation costs across the entire grid compared to no energy storage, as the percentage of homes using energy storage scales up with both DART (a) and TOU (b) pricing. In this case, each home has usable energy storage capacity that is 50% of their average daily demand. The graph shows that as the number of homes using energy storage scales up, the greedy algorithm, akin to our SmartCharge algorithm, increases generation costs, i.e., the savings are negative, due to simultaneous battery charging and large rebound peaks after 20% of homes use energy storage (with DART).

TOU pricing scales slightly better than DART because the closed-loop pricing results in very low rates when consumers are not charging, which reflect in the next day’s prices. Since TOU rates do not change, it does not suffer from oscillations in peak demand or prices. The congestion-aware greedy variant (with a limit on the charging rate of 60% the maximum rate) in both cases is more scalable, but still results in rebound peaks once enough homes adopt energy storage (35% for DART and 45% for TOU). In contrast, the peak-aware algorithm steadily decreases the grid’s generation costs as more homes use energy storage, signaling that homes and the grid are successfully flattening demand. The congestion-aware greedy variant demonstrates an important point: consumers could prevent rebound peaks by rate-limiting their charging, but they would reduce their cost savings.

Figure 8(c) zooms in on the Peak-aware(DART), Peak-aware(TOU) lines from Figure 8(a) and (b), respectively. It demonstrates the steadily increasing, rather than decreasing, cost savings as more homes use energy storage. Notice maximum savings from the three variants is similar: this reflects that in each case the aggregate energy storage is sufficient to flatten demand. The difference with the peak-aware algorithm is that it distributes this capacity across 100% of consumers, while the other algorithms distribute it across a much smaller set of consumers. Figure 9 shows the corresponding reduction in peak demand for the same experiment, which, as expected, shows similar trends as the generation cost savings. Namely, the grid’s peak demand steadily decreases, rather than increases, as more homes use

energy storage with the peak-aware algorithm. As an example of this decrease in peak demand, Figure 10(a) and (b) show the time-series of power usage of an example day in our trace both with and without energy storage (for DART and TOU pricing).

We also run a similar experiment, but rather than vary the percentage of homes using energy storage we vary the amount of energy storage each home has as a fraction of its demand. In this case, 100% of homes have energy storage. Figures 11(a) and (b) show the generation cost savings and peak reduction, respectively. In both cases, the results show that each home only needs energy storage capacity that is a small fraction of its average demand. In Figure 11(a), generation cost savings stop increasing once homes have energy storage capacity that is 20% of their average demand. Similarly, Figure 11(b) shows the grid’s peak not decreasing further at the same 20% threshold. At the 20% threshold, the aggregate storage capacity across homes is near the optimal storage capacity required to flatten demand (from Section 2.3). In contrast, using our previous work on SmartCharge without a peak demand surcharge, a representative home required 50% of their average demand in storage capacity to maximize their ROI [15].

The previous experiments used the same, relatively high, peak demand surcharge of \$3/kW, such that inequality (1) does not hold and our peak-aware algorithm focuses on flattening demand. Figure 12 demonstrates the percentage peak reduction across the 194 homes in our traces as we vary the peak demand surcharge for both TOU and DART pricing plans. In this case, all homes use energy storage with usable capacity that is 50% of their average demand. As expected, for low surcharge values the homes are greedy resulting in large rebound peaks—larger than the original peaks—from simultaneous battery charging during low-price periods. However, once the peak demand surcharge passes the threshold defined by inequality (1) the homes switch to flattening their demand. In this case, for high peak demand surcharges the algorithm reduces the peak 10-15%. The threshold near \$0.60/kW represents the tipping point




Figure 10: Time series of aggregate grid demand for TOU (a) and DART (b) pricing for both without energy storage and using our peak-aware algorithm with each home having energy storage.

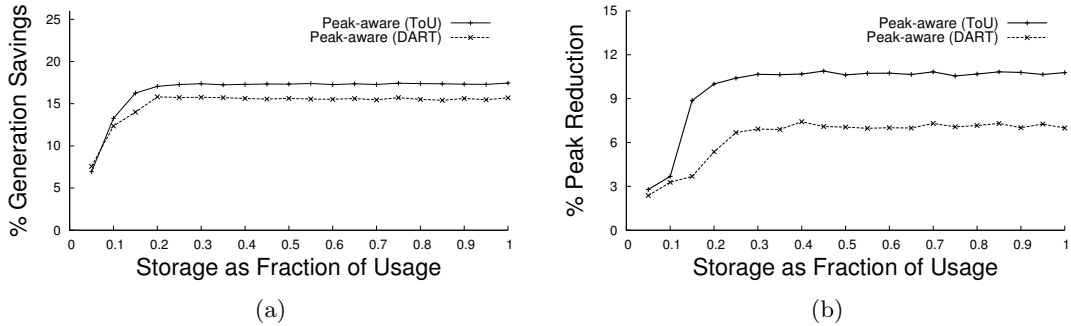


Figure 11: Cost savings (a) and grid peak reduction (b) as a function of each home's energy storage capacity.

where the benefit of charging the battery at its maximum rate during the lowest-price hour is not worth the cost of the peak demand surcharge.

Lastly, we revisit the analysis of the optimal minimum amount of energy storage capacity necessary to flatten grid demand in a centrally controlled system (from Section 2.3). For that trace, the minimum capacity required was 393kWh. For the same trace, our peak-aware algorithm requires 481kWh of aggregate capacity across the grid to maximize its peak reduction and generation savings from flattening demand. Figure 10 shows the high-level trend of aggregate demand achieved by the peak-aware algorithm relative to raw consumption. As the figure demonstrates, the peak-aware algorithm is not able to perfectly flatten demand due to inaccuracy in choosing the target average, which may result in high peaks if a home's battery is empty during a period of high demand. While predicting average demand over a long period is more accurate than predicting hourly demand over the same period, it is not perfect due to changing consumer behavior. Despite the inaccuracy, the approach comes within 18% of the optimal centralized approach.

5.2 Consumer-scale Effects

With distributed energy storage in use large scales, the peak reductions and cost savings for individual consumers mirror the reductions and savings in the grid. To show this, we use the TOU pricing scheme, which, as shown above, performs similarly to DART. We take a representative home with near the average demand for a home in our trace and look at its individual peak reduction and cost savings (as both a percentage and in dollars) as function of the home's usable energy storage capacity for our peak-aware algorithm. Figure 13 shows the results. As in the grid, the peak-aware algorithm reaches its maximum peak reduction

(Figure 13(a)) when usable storage capacity is only 20% of average demand (rather than 50% with a greedy approach like SmartCharge [15]). Further, for less than 2X the energy storage capacity, the percentage cost savings in the electric bill (Figure 13(b)) is in the same 10-15% range as we found in SmartCharge. Thus, the ROI for the consumer is much higher, since consumers achieve similar savings using much less energy storage capacity, which overwhelmingly dominates the cost of the system. Finally, Figure 13 shows the average per day dollar savings, which mirrors the trend in the overall percentage savings.

The results above clearly demonstrate the benefits, as the use of energy storage scales, of incentivizing consumers to flatten their own demand using a peak demand surcharge. For the similar cost savings per consumer using energy storage, the approach requires much less storage capacity, resulting in a higher ROI. Further, this ROI does not diminish as more consumers use energy storage, since the utility is able to control each consumer's incentive to flatten demand using the peak demand surcharge independent of the electricity rates. This also mitigates free riding, since consumers without energy storage are less able to control their peak demand and benefit from reducing their peak. By incentivizing all consumers to use energy storage, the approach encourages distributing the aggregate energy storage necessary to flatten grid demand across a wide set of consumers, which all share in the savings. In prior work, with only variable electricity rates, the consumers not using energy storage also benefit from lower overall generation costs (and electricity rates) as storage capacity increases, which, in turn, diminishes the savings for the consumers that use energy storage.




Figure 13: Percentage peak reduction (a), percentage cost savings (b), and dollar cost savings (c) for an individual home using our peak-aware algorithm as the home’s energy storage capacity varies using our peak-aware algorithm under a peak-demand surcharge.




Figure 12: Increasing the peak demand surcharge prevents rebound peaks in the grid by incentivizing consumers to flatten their demand.

6. RELATED WORK

As mentioned in Section 1, numerous researchers have studied the use of energy storage at homes and buildings to shift demand and cut electricity bills under emerging variable rate electricity pricing plans. Daryanian et al. [7] was the first to propose this form of energy arbitrage. This work, as well as work by van de ven et al. [23], study the problem from a theoretical standpoint, e.g., assuming certain demand distributions, without evaluating their solutions on real data. More recently, our own work on SmartCharge [15], as well as work by Carpenter et al. [4], study a similar problem in a realistic setting taking into account battery inefficiencies, stochastic demand in residential settings, and existing variable rate pricing plans in Ontario and Illinois. Both papers mention the problems with scaling distributed energy storage to many consumers, but neither i) explores the full implications of large scale adoption, including the decreasing ROI for consumers with storage as adoption scales nor ii) proposes or evaluates a solution to the problem. Johnson et al. [12] formulate the peak shaving problem based on a pricing plan where customers are billed for their peak usage. The authors present an optimal offline algorithm, and a competitive online algorithm for solving the problem. However, they do not focus on or evaluate the proposed algorithms at large scales, as the use of energy storage increases.

While the data sets in these papers are different, they both show that $\sim 20\%$ of homes using energy storage maximizes the grid’s peak reduction. After this point, rebound peaks and simultaneous battery charging begin to reduce energy storage’s benefits, ultimately leading to a higher peak usage than without energy storage if prices do not react to demand. In earlier work, Vytelingum et al. [24] shows formally that under variable electricity rate pricing plans there is a Nash equilibrium that maximizes social welfare, e.g., cost savings, once only 38% of U.K. households use energy storage (based

on a U.K. data set). Although slightly higher than the $\sim 20\%$ of homes found above, the paper’s trend is the same: beyond a certain point with existing variable rate electricity prices the benefits of consumers installing energy storage begin to decrease. We argue that, due to the high cost of batteries, when designing incentivizes for distributed energy storage, the goal should be to encourage the distribution of aggregate capacity as widely as possible among consumers.

While the work above focuses on residential settings, prior work has also looked at similar problems from the perspective of industrial consumers, particularly data centers [11, 22], but has not examined the impact of storage at scale. Prior work also highlights the effect of variable rate pricing on grid stability [17, 18], showing that real-time pricing has the potential to create an unstable closed feedback loop. We show this experimentally in Figure 5 in the presence of large-scale energy storage. Finally, we know of no work that proposes and evaluates using a peak demand surcharge to maintain a stable grid and prevent rebound peaks by incentivizing consumers to flatten their own demand.

7. CONCLUSION

This paper examines the effects of using energy storage distributed at buildings and homes throughout the grid to flatten grid demand. In particular, we show that as more consumers adopt energy storage, a number of problems arise that impact grid stability and generation costs. As a result, we propose to augment traditional variable rate electricity pricing plans with a substantial peak demand surcharge, which incentivizes consumers to flatten their demand rather than shift it all to a low-price period. Utilities already use peak demand surcharges for large industrial consumers; we argue that, to incentivize distributed energy storage, they may want to broaden their use to other consumers.

We then design PeakCharge, which includes an online algorithm to minimize electricity costs in the presence of variable rates and the peak demand surcharge. Using a closed-loop simulator, we show that our algorithm is effective at both i) maintaining the incentives for consumers to use energy storage at large scales and ii) ensuring grid stability. Further, our results indicate the aggregate energy storage capacity to flatten grid demand by incentivizing consumers to flatten their own demand is within 18% of the minimum, optimal capacity to flatten grid demand in a centralized system. Since flattening a home’s demand requires over 2X less storage capacity per consumer to maximize consumer savings, it significantly reduces the ROI of energy storage, which is dominated by battery costs.

References

- [1] Electric Generator Dispatch Depends on System Demand and the Relative Cost of Operation. In *U.S. Energy Information Administration: Today in Energy*, August 17th 2012.
- [2] S. Barker, A. Mishra, D. Irwin, E. Cecchet, P. Shenoy, and J. Albrecht. Smart*: An Open Data Set and Tools for Enabling Research in Sustainable Homes. In *SustKDD*, August 2012.
- [3] S. Barker, A. Mishra, D. Irwin, P. Shenoy, and J. Albrecht. Smartcap: Flattening peak electricity demand in smart homes. In *PerCom*, March 2012.
- [4] 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.
- [5] D. Cauchon. Usatoday, household Electricity Bills Skyrocket. <http://www.usatoday.com/money/industries/energy/story/2011-12-13/electric-bills/51840042/1>, December 13th 2011.
- [6] CNT energy. Dynamic Pricing and Smart Grid. <http://www.cntenergy.org/pricing/>, 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] A. Faruqui, R. Hledik, and J. Tsoukalis. The Power of Dynamic Pricing. *Electricity Journal*, 22(3):42–56, February 2009.
- [9] State of the Markets Report 2008. Technical report, Federal Energy Regulatory Commission, August 2009.
- [10] T. Ganu, D. Seetharam, V. Arya, R. Kunnath, J. Hazra, S. Husain, L. DeSilva, and S. Kalyanaraman. nPlug: A Smart Plug for Alleviating Peak Loads. In *e-Energy*, May 2012.
- [11] S. Govindan, A. Sivasubramaniam, and B. Urgaonkar. Benefits and Limitations of Tapping into Stored Energy for Datacenters. In *ISCA*, June 2011.
- [12] M. P. Johnson, A. Bar-Noy, O. Liu, and Y. Feng. Energy Peak Shaving with Local Storage. *Sustainable Computing: Informatics and Systems*, 1(3), 2011.
- [13] I. Koutsopoulos, V. Hatzi, and L. Tassiulas. Optimal Energy Storage Control Policies for the Smart Power Grid. In *SmartGridComm*, September 2011.
- [14] S. McLaughlin, P. McDaniel, and W. Aiello. Protecting Consumer Privacy from Electric Load Monitoring. In *CCS*, October 2011.
- [15] A. Mishra, D. Irwin, P. Shenoy, J. Kurose, and T. Zhu. SmartCharge: Cutting the Electricity Bill in Smart Homes with Energy Storage. In *e-Energy*, May 2012.
- [16] Ontario Energy Board: Electricity Prices. <http://www.ontarioenergyboard.ca/OEB/Consumers/Electricity/Electricity+Prices>, 2012.
- [17] M. Roozbehani, M. A. Dahleh, and S. K. Mitter. On the stability of wholesale electricity markets under real-time pricing. In *CDC*, December 2010.
- [18] M. Roozbehani, M. Rinehart, and M. A. Dahleh. Real-time pricing with ex-post adjustments in competitive electricity markets: Volatility and efficiency analysis. In *Allerton*, September 2011.
- [19] N. Sharma, J. Gummeson, D. Irwin, and P. Shenoy. Cloudy computing: Leveraging weather forecasts in energy harvesting sensor systems. In *SECON*, June 2010.
- [20] N. Sharma, P. Sharma, D. Irwin, and P. Shenoy. Predicting Solar Generation from Weather Forecasts Using Machine Learning. In *SmartGridComm*, October 2011.
- [21] Smart Grid Information Clearinghouse: Legislation and Regulation. <http://www.sgiclearinghouse.org/Legislation?q=node/1705>, 2011.
- [22] R. Urgaonkar, B. Urgaonkar, M. Neely, and A. Sivasubramaniam. Optimal Power Cost Management Using Stored Energy in Data Centers. In *SIGMETRICS*, March 2011.

Parameter	Definition
T	Time in T discrete intervals 1 to T
I	Length of each of the discrete time intervals
s_i	Power charged in interval i
d_i	Power discharged in interval i
e	Battery efficiency, $0 \leq e \leq 1$
p_i	Grid power demand in interval i
c_i	Power cost per kWh in interval i
m_i	Charge for electricity in interval i
C	Battery capacity in kWh
l_i	Aggregate grid demand in interval i .
L	Peak grid demand over all intervals
c'	Peak demand surcharge

Table 1: Parameter definitions for linear program.

- [23] P. van de ven, N. Hegde, L. Massoulie, and T. Salonidis. Optimal Control of Residential Energy Storage Under Price Fluctuations. In *ENERGY*, May 2011.
- [24] P. Vytelingum, T. Voice, S. Ramchurn, A. Rogers, and N. Jennings. Agent-based Micro-storage Management for the smart grid. In *AAMAS*, May 2010.
- [25] W. Yang, N. Li, Y. Qi, W. Qardaji, S. McLaughlin, and P. McDaniel. Minimizing Private Data Disclosures in the Smart Grid. In *CCS*, October 2012.
- [26] Zivan. Battery Discharge at Constant Current: Technical Features and User Manual. <http://www.zivanusa.com/pdf/SBM%20Gb.pdf>, 2012.

APPENDIX

Below is the modification of the LP from SmartCharge to minimize an electricity bill using energy storage in the presence of a peak demand surcharge, given future knowledge of next-day prices and next-day demand each hour. Table 1 defines the optimization’s parameters. The formal objective is to minimize $\sum_{i=1}^T m_i$ each day, given constraints below. The first five constraints are present in SmartCharge’s original LP: constraints (1) and (2) ensure positive energy is charged and discharged from the battery, constraint (3) bounds the battery’s charging rate, constraint (4) preserves conservation of energy (including energy conversion efficiency), and constraint (5) bounds the battery’s capacity. The final three constraints (in bold) are necessary to optimize for a peak demand surcharge: constraint (6) computes the bill based on variable rate prices and the peak demand surcharge, constraint (7) represents grid’s demand in the i th interval, and constraint (8) is the size of the peak demand surcharge.

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

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

$$d_i \geq 0, \forall i \in [1, T] \quad (4)$$

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

$$\left(\sum_{t=0}^i s_t - \sum_{t=0}^i d_t/e \right) * I \leq C, \forall i \in [1, T] \quad (6)$$

$$\mathbf{m}_i = \left(\mathbf{p}_i + \mathbf{s}_i - \mathbf{d}_i \right) * \mathbf{I} * \mathbf{c}_i + \mathbf{L} * \mathbf{c}', \forall i \in [1, T] \quad (7)$$

$$l_i = \mathbf{p}_i + \mathbf{s}_i - \mathbf{d}_i, \forall i \in [1, T] \quad (8)$$

$$l_i \leq \mathbf{L}, \forall i \in [1, T] \quad (9)$$

Acknowledgements. We thank our shepherd, Chi-Kin Chau, and reviewers for their comments. This research was supported by NSF grants CNS-1143655, CNS-0916577, CNS-0855128, CNS-0834243, CNS-0845349, CNS-1217791.