Enabling Distributed Energy Storage by Incentivizing Small Load Shifts

DAVID IRWIN, SRINIVASAN IYENGAR, and STEPHEN LEE,
University of Massachusetts Amherst
ADITYA MISHRA, Seattle University
PRASHANT SHENOY and YE XU, University of Massachusetts Amherst

Reducing peak demands and achieving a high penetration of renewable energy sources are important goals in achieving a smarter grid. To reduce peak demand, utilities are introducing variable rate electricity prices to incentivize consumers to manually shift their demand to low-price periods. Consumers may also use energy storage to automatically shift their demand by storing energy during low-price periods for use during high-price periods. Unfortunately, variable rate pricing provides only a weak incentive for distributed energy storage and does not promote its adoption at large scales. In this article, we present the storage adoption dilemma to capture the problems with incentivizing energy storage using variable rate prices.

To address the problem, we propose a simple pricing scheme, called *flat-power pricing*, which incentivizes consumers to shift small amounts of load to flatten their demand rather than shift as much of their power usage as possible to low-price, off-peak periods. We show that compared to variable rate pricing, flat-power pricing (i) reduces consumers' upfront capital costs, as it requires significantly less storage capacity per consumer; (ii) increases energy storage's return on investment, as it mitigates free riding and maintains the incentive to use energy storage at large scales; and (iii) uses aggregate storage capacity within 31% of an optimal centralized approach. In addition, unlike variable rate pricing, we also show that flat-power pricing incentivizes the scheduling of elastic background loads, such as air conditioners and heaters, to reduce peak demand. We evaluate our approach using real smart meter data from 14,000 homes in a small town.

CCS Concepts: • **Hardware** → **Smart grid**; *Energy generation and storage*; *Energy distribution;*

Additional Key Words and Phrases: Energy storage, load scheduling, smart grid, load shifting

ACM Reference Format:
David Irwin, Srinivasan Iyengar, Stephen Lee, Aditya Mishra, Prashant Shenoy, and Ye Xu. 2017. Enabling distributed energy storage by incentivizing small load shifts. ACM Trans. Cyber-Phys. Syst. 1, 2, Article 10 (February 2017), 30 pages.
DOI: http://dx.doi.org/10.1145/3015663

1. INTRODUCTION

As is now well known, a significant fraction of the electric grid's capital and operational expenses (CapEx and OpEx, respectively) result from satisfying its peak power demands. For example, recent work estimates that 10% to 18% of North American CapEx, in terms of energy generation capacity, is idle and wasted over 99% of the

---

This article extends material from a prior conference [Mishra et al. 2013b] and workshop paper [Xu et al. 2013].

Our research was supported by NSF grants IIP-1534080, CNS-1405826, CNS-1253063, CNS-1505422, and the Massachusetts Department of Energy Resources.

Authors' addresses: D. Irwin, S. Iyengar, S. Lee, P. Shenoy, and Y. Xu, 151 Holdsworth Way, Amherst, Massachusetts 01003; A. Mishra, 901 12th Avenue, P.O. Box 222000, Seattle, WA 98122.

DOI: http://dx.doi.org/10.1145/3015663

year [Faruqui et al. 2009]. Similarly, peak demand also influences OpEx by (i) requiring utilities to operate high-cost and inefficient "peaking" generators to meet demand [EIA 2012]; (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 fraction of demand from peak to off-peak periods. To shift demand, prior work simply proposes storing energy during off-peak periods, which increases the off-peak demand, for use during peak periods, which then decreases the peak demand [Carpenter et al. 2012; Daryanian et al. 1989; Govindan et al. 2011; Koutsopoulos et al. 2011; Urgaonkar et al. 2011; van de Ven et al. 2011; Vytelingum et al. 2010]. The recently introduced Tesla Powerwall home battery now advertises such load shifting as an explicit feature [Howard 2015].

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 [Koutsopoulos et al. 2011; Mishra et al. 2015], 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, as the increasing adoption of variable rate pricing plans by utilities [CNT 2011; Ontario Energy Board 2012; Smart Grid Clearinghouse 2011] provide an incentive to install energy storage [Carpenter et al. 2012; Daryanian et al. 1989; Govindan et al. 2011; Urgaonkar et al. 2011; van de Ven et al. 2011; Vytelingum et al. 2010]. Endowing buildings with energy storage also has additional benefits, including (i) providing power during outages, (ii) conditioning power to increase its quality, (iii) dampening fluctuations from local renewable sources, (iv) distributing energy storage's capital costs across consumers rather than requiring utilities to bear them, and (v) decreasing transmission losses by storing energy near its point of consumption.

Since variable rate pricing plans charge higher rates during periods of peak demand during the day, consumers who store energy during off-peak periods at night (when prices are low) for use during peak daytime periods (when prices are high) are able to lower their electricity bill. In addition, consumers with excess local solar generation may store their own local solar generation for later use, which may be preferable to net metering excess solar power to the grid if utilities pay less than the retail rate for net metered solar power. Local energy storage will increasingly be required as renewable penetrations increase, as the grid generally cannot handle large, e.g., $>10\%$, renewable penetrations without energy storage [Thomson 2015]. Many energy storage technologies exist, including pumped water storage, flywheels, and compressed air; however, batteries remain the most viable option for stationary energy storage systems in buildings. Figure 1 shows an example of how electricity prices vary over a day with time-of-use (TOU) [Ontario Energy Board 2012] and real-time pricing (RTP) [Smart Grid Clearinghouse 2011] plans. With RTP plans, rates change every hour of every day based on electricity's real-time price in the wholesale market, whereas with TOU plans, rates change only a few times each day and each day's rate profile remains constant over long periods, e.g., 3 to 6 months.

Prior research analyzes the potential savings for residential [Carpenter et al. 2012; Daryanian et al. 1989; van de Ven et al. 2011; Vytelingum et al. 2010] and industrial [Govindan et al. 2011; Urgaonkar et al. 2011] consumers to install batteries and perform energy arbitrage in the presence of variable rate pricing plans. The focus is largely on cost-benefit analyses of existing pricing plans, which vary electricity's price

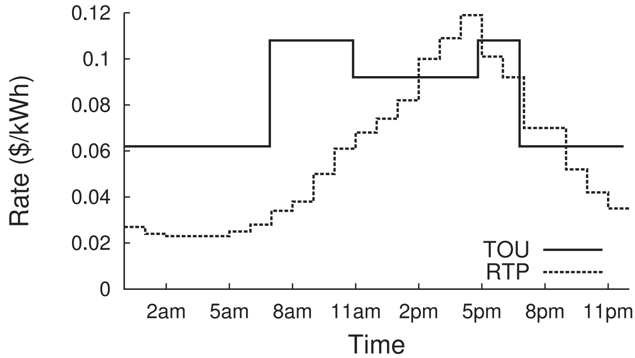

Fig. 1. Examples of existing TOU (from Ontario) and RTP (from Illinois).

per kilowatt-hour (kWh). Unfortunately, as we discuss, these plans provide a weak incentive for distributed energy storage and do not promote its large-scale adoption.

Large upfront capital costs. Since today's pricing plans exhibit low prices during off-peak nighttime periods and high prices during peak daytime periods, they incentivize consumers to shift large amounts of load to the off-peak period. Of course, the cost of batteries limits the amount of storage capacity available to shift demand. In prior work, Mishra et al. [2012] show that for a residential home with near the average U.S. electricity usage, nearly 24kWh of capacity[1] maximizes the return on investment (ROI) when taking into account battery costs. Given typical battery lifetimes, the estimated annual amortized cost to maintain 24kWh of energy storage is then nearly $1416 [Mishra et al. 2012]. Since the annual electricity bill for an average U.S. home approximately $1419 [Cauchon 2011], battery costs prevent (at current price levels) a positive ROI using this much energy storage.

Uncertain benefits. To prevent rebound peaks under large-scale penetrations of energy storage, utilities must alter electricity rates over time as peak and off-peak demand changes. Unfortunately, consumers only benefit from energy storage by exploiting the difference between peak and off-peak prices. With variable rate pricing 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 will narrow, thereby reducing energy storage's benefits [Vytelingum et al. 2010]. In the extreme, if grid demand is near flat, then the price of electricity will be similar at all times [Carpenter et al. 2012; Mishra et al. 2012; Vytelingum et al. 2010]. Once the peak/off-peak price differential is not large enough to compensate for battery conversion losses, there is no benefit to using energy storage. Our prior work estimates that there would be no difference between peak and off-peak prices once just 22% of consumers install 24kWh of energy storage [Mishra et al. 2012], which is consistent with related work [Carpenter et al. 2012]. Vytelingum et al. [2010] also show that a Nash equilibrium—where there is no longer an incentive to install energy storage—is reached well before battery penetration levels increase to the point where peak and off-peak prices converge. Thus, even if energy storage prices decline, consumers will be unlikely to invest in energy storage en masse 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 10× for

[1]Operated at a maximum of 45% depth of discharge.

average consumers in the United States—than the annual savings on an electric bill using current battery costs, electricity rates, and pricing plans [Mishra et al. 2012]. However, prior work does not consider the grid-wide reductions in generation costs from lowering the grid's aggregate peak demands. Unfortunately, with existing pricing plans, these cost savings are distributed (or socialized) across all consumers, as they manifest themselves as cheaper electricity rates. Socializing the savings from distributed energy storage weakens the incentive to adopt it. Strengthening the incentive requires ensuring that consumers investing in energy storage reap its full benefits, especially given its large capital costs.

Ideally, energy storage distributed at buildings throughout the grid would behave like centrally controlled energy storage of equal capacity. In other words, 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 largely flat and constant at the average demand. However, ensuring that the behavior of such a self-organizing distributed system emulates that of an equivalently sized centralized system is challenging. In this case, determining when and how many buildings should charge their batteries requires explicit feedback from the grid and coordination among all buildings, which does not scale. In addition, determining how to fairly distribute the use of energy storage across buildings in the grid poses another challenge.

To address the problem, we propose *flat-power pricing* to incentivize consumers to shift small amounts of load to flatten their demand. Flat-power pricing charges a higher price for electricity that is above a consumer-specific target set a priori by the utility. Since the approach incentivizes shifting power usage over shorter intervals than variable rate pricing, it requires less energy storage capacity per home and also incentivizes scheduling of elastic background loads that can only defer power usage for short periods. As we discuss, optimal charging behavior under flat-power pricing differs significantly from existing peak-based pricing plans that impose a surcharge on consumers based on their absolute peak demand over a billing period, e.g., the peak 15 minutes of energy usage over a month-long billing cycle. In particular, these existing peak-based pricing models strongly incentivize consumers to focus on reducing their absolute peak, whereas flat-power pricing equally incentivizes consumers to reduce any peak above their average. By focusing on the absolute peak, optimal charging algorithms for existing peak-based pricing models may not use energy storage effectively to reduce other peaks, as there is no incentive to do so. However, reducing the nonabsolute peaks is important, as a consumer's absolute peak may not necessarily align with the time of the grid's peak demand. Thus, flat-power pricing equally incentivizes consumers to reduce any peak in their demand, enabling them and the grid to benefit from shifting small amounts of load. Our hypothesis is that by encouraging small load shifts, flat-power pricing incentivizes the use of distributed energy storage, as well as the scheduling of elastic background loads, at large scales. In evaluating our hypothesis, this article makes the following contributions.

Scalable incentives design. We describe the *energy storage adoption dilemma* that arises as the use of energy storage scales. We show that existing charging algorithms and pricing plans cannot simultaneously minimize an electric bill and ensure grid stability at large scales. In particular, preventing rebound peaks, which reverse the peak and off-peak periods, requires some (explicit or implicit) feedback from the grid to signal distributed energy storage to rate-limit charging as demand rises.

Flat-power pricing model. To resolve the dilemma, we introduce flat-power pricing and outline its benefits relative to existing pricing models, including both variable rate and peak-based pricing, with respect to incentivizing distributed energy storage adoption, as well as the scheduling of elastic background loads. We present simple scheduling algorithms for battery charging and discharging, along with different types of elastic background loads, that optimize for flat-power pricing.

Grid- and consumer-scale evaluation. We evaluate both the grid- and consumer-scale effects of flat-power pricing. For our grid-scale evaluation, we leverage nearly a year's worth of smart meter data from 14,000 homes in a small town (at 5-minute granularity) to demonstrate the impact of distributed energy storage on the grid under flat-power pricing. For our consumer-scale evaluation, we use extensive data from a real home to demonstrate the impact of scheduling of background loads under flat-power pricing. Our analysis shows that when compared to existing pricing plans, flat-power pricing results in (i) less upfront capital costs, as it requires significantly less energy storage capacity per consumer; (ii) reduces grid generation costs, as it mitigates free riding and maintains the incentive to adopt energy storage at large scales; iii) only requires aggregate storage capacity within 31% of optimal; and iv) encourages the scheduling of elastic background loads.

2. BACKGROUND

Our work leverages the use of battery-based energy storage and elastic background load scheduling to reduce electricity costs. We assume a load scheduling system that is capable of programmatically controlling battery charging and discharging, as well as the scheduling of background loads, based on electricity rates. Existing pricing plans vary electricity rates throughout the day based on demand such that rates during high-demand periods, e.g., during the day, are significantly higher than rates during low-demand periods, e.g., at night. To be cost effective, these systems must (i) limit energy storage capacity due to battery costs, which, amortized over their lifetime, are currently \$100 to \$200 per year per kilowatt-hour of usable capacity for the VRLA/AGM lead-acid variety widely used in stationary energy storage systems, and (ii) account for the nearly 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 approximately 12kWh of usable capacity.

2.1. System Architecture

Previously proposed architectures for leveraging energy storage [Mishra et al. 2012] 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 expectations of future prices and demand. Of course, when batteries supply power, the building's load dictates the rate of discharge based on 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 that indicates a blackout or brownout. The coarse switching architecture works well in prior systems, as 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

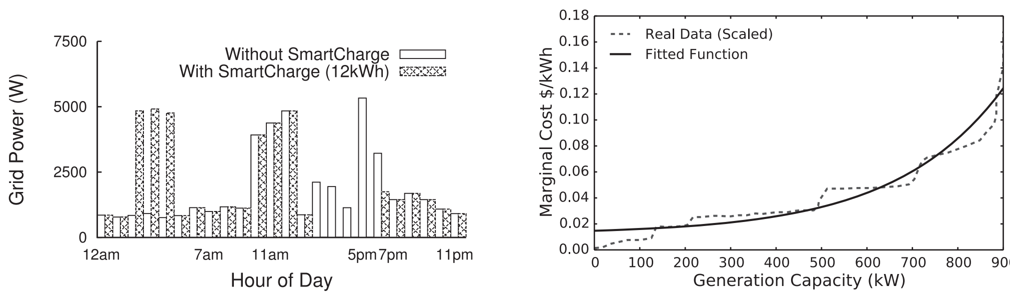

(a) Switch-based Architecture (b) Dispatch Stack Model

Fig. 2. Switch-based architectures (a) do not lower an individual building's peak demand [Mishra et al. 2012]. The model we use (b) 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 Federal Energy Regulatory Commission report [Federal Energy Regulatory Commission 2009].

necessary for our approach, which encourages buildings to flatten demand by shifting small amounts of load rather than simply shift large amounts of demand from daytime to nighttime. As Figure 2(a) demonstrates, for individual buildings, the simple switching architecture does not significantly reduce an individual building's peak demand (in this case, the average demand each hour) [Mishra et al. 2012]. The figure illustrates how, due to off-peak battery charging, the switch-based SmartCharge system simply shifts the original peak demand to a lower-price nighttime period to minimize electricity costs. Importantly, such existing switch-based architectures do not significantly lower a building's peak demand—they only move it to the lowest-priced period.

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 simply switching some subset of circuits to the battery to achieve a specific rate of discharge. An alternative approach is to use a utility-interactive battery inverter to connect a battery bank in parallel to the grid and programmatically control its rate of discharge. Currently, controllers capable of programmatically setting the rate of discharge at fine granularities are available for testing equipment [Zivan 2012].

However, utility-interactive battery inverters—currently used by photovoltaic (PV) systems to operate in both grid-tied and off-grid modes—can also be co-opted for controlling the rate of discharge. In grid-tied mode, these inverters typically maintain the battery at a "set point" voltage such that if the voltage rises above the set point (indicating a higher state of charge), the inverter draws power from the battery bank, and if the voltage falls below the set point, the inverter charges the battery bank. In grid-tied mode, both the PV system and the battery bank connect to the inverter and then converts their DC power output into AC power at 60Hz synchronized with the grid's 60Hz power. The current advantage of utility-interactive battery inverters is they can operate in off-grid mode during an outage such that the inverter is capable of charging the battery from the PV system and supplying AC power to loads. A specific example of a utility-interactive battery inverter that includes the preceding features is the Conext SW-4024.[2] Since utility-interactive inverters are designed for residential

[2]See http://solar.schneider-electric.com/product/conext-sw-na/.

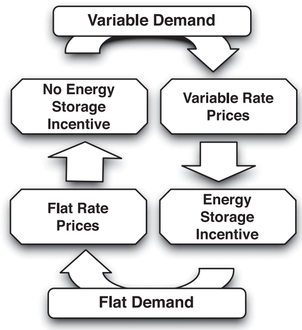

Fig. 3. Idealized depiction of the storage adoption dilemma.

PV systems, they are capable of operating at residential-scale loads. For example, the Conext SW-4024 is rated for 8kW peak load, which is above the peak load for most homes.

Although current utility-interactive battery inverters are configured to maintain a set point voltage for backup power, their basic mechanisms can also be configured to programmatically control the rate of battery discharge by controlling the set point over time. Such programmatic control is expected to become more widespread, as the Tesla Powerwall and other energy storage devices are designed for similar peak shaving features. As a result, the Public Utilities Commission of California recently introduced regulations and policies for net energy metering from energy storage devices.[3] Such programmatic control also has uses beyond net metering and is a prerequisite for recent work on privacy-preserving smart grid techniques using batteries [McLaughlin et al. 2011; Yang et al. 2012]. For our work, we assume the use of a utility-interactive inverter with the ability to programmatically adjust its set point voltage to control a battery bank's charging and discharging rate.

Finally, both this and prior work derive from the fact that the marginal cost for a utility to generate each additional watt of power increases nonlinearly 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 2(b) shows the demand-cost function we use to compute generation costs based on demand and demonstrates the nonlinear relationship between cost and demand. To derive our function, we scaled real demand-cost data for the southeastern United States from a 2008 report by the Federal Energy Regulatory Commission to match the peak demand in our trace data, discussed in Section 5, while also ensuring a median electricity cost of 10¢/kWh, which is near the average cost of electricity in the United States [Federal Energy Regulatory Commission 2008]. We then fit an exponential function to this data to estimate the effect on costs in our evaluation.

2.2. The Storage Adoption Dilemma

Figure 3 depicts the *storage adoption dilemma*, a variant of the classic prisoner's dilemma, 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 2(b). Variable pricing,

[3]See http://docs.cpuc.ca.gov/PublishedDocs/Published/G000/M091/K247/91247687.pdf.

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, the difference between the peak and off-peak electricity prices narrows to reflect the new demand distribution. Unfortunately, such flatter prices eliminate the incentive to adopt and use energy storage, which causes demand to vary again and the cycle to repeat. Of course, our depiction is idealistic. In practice, the cycle may stall due to weak incentives for energy storage, as outlined in Section 1, and completing each step would take a long time, potentially requiring significant regulatory changes and large capital investments by consumers. Prior work [Vytelingum et al. 2010] has also shown (under certain assumptions) that a Nash equilibrium would be reached that would eliminate the incentive for additional consumers to install energy storage (under their assumptions, once 38% of consumers installed 4kWh of energy storage), thereby breaking the cycle. This equilibrium is well below both the optimal amount of storage required by a centralized system and 100% consumer adoption, and it exhibits all of the problems cited in Section 1: large upfront capital costs, uncertain returns, socialized benefits, and free riders. Thus, although the equilibrium may be acceptable to a utility, it does not preserve the incentive for consumers to install and maintain energy storage. In addition, compared to prior work, our flat-power pricing model avoids the game-theoretic aspects of current variable rate pricing models, as under flat-power pricing the amount consumers pay is based only on their own usage and is independent of other consumers' usage.

2.3. A Centralized Approach

Before examining a distributed approach to energy storage, we first define and consider an optimal centralized battery charging scheme. Ideally, to minimize generation costs based on the demand-cost function from Figure 2(b), the optimal approach would shift aggregate grid demand such that it was the same—equal to average demand—at all times. If we assume a centrally controlled battery array, then a simple algorithm for controlling a centralized system would simply charge and discharge batteries whenever grid demand is below or above the average demand, respectively, such that demand is always equal to the average. With this algorithm, the minimum energy capacity necessary to optimally flatten demand is equal to the maximum capacity ever required 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 several benefits.

As an example, Figure 4 depicts a grid demand profile with average power every 5 minutes from one day across nearly 14,000 homes, as well as the day's average daytime and nighttime demand. In this case, the maximum capacity required to charge or discharge the battery occurs between 12am and 6:30am (equivalent to the area between the instantaneous demand and the average demand over that time period). The optimal approach requires 43.3MWh of energy storage capacity to charge batteries during this nighttime period and would reduce generation costs by 24.2% based on the demand-cost function in Figure 2(b). If this storage capacity were distributed evenly among all homes, then each home would need only 3.09kWh of usable energy storage. This capacity is nearly 4× less than the 12kWh of usable (or 24kW of rated) capacity that each home requires to maximize energy storage's ROI estimated by prior work [Mishra et al. 2012], which uses existing variable rate pricing plans with an optimal charging algorithm that assumes future knowledge of demand and prices. Since battery costs scale linearly with capacity, maintaining 4× less capacity decreases costs by 4× (from $1416 amortized per year to maintain 12kWh to $354 per year to maintain 3.09kWh). The example demonstrates how minimizing the aggregate battery capacity and

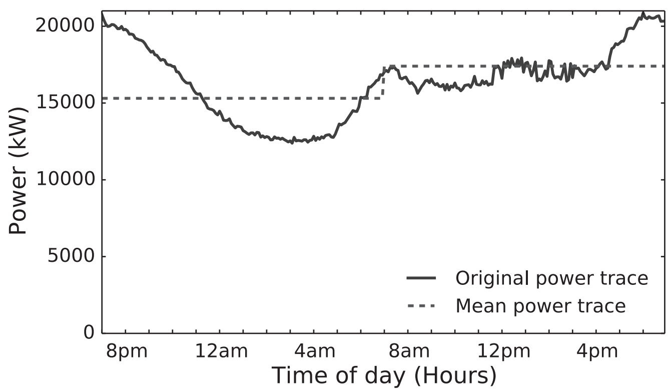

Fig. 4. Instantaneous day and night average grid demand for 14,000 homes in our town smart meter data over a single day.

distributing it as widely as possible among consumers reduces the cost per consumer of energy storage.

3. FLAT-POWER PRICING

The storage adoption dilemma discourages consumers from adopting distributed energy storage at large scales (>20% of consumers). Unfortunately, variable rate electricity prices incentivize consumers with energy storage to use battery charging algorithms that incentivize charging batteries during the lowest-price periods to minimize electricity costs. At large scales, the use of these variable rate charging algorithms by distributed consumers may exhibit one of multiple possible outcomes. In particular, these algorithms will result in either (i) large rebound peaks (if prices do not react to changing demand), (ii) grid instability (if consumers react quickly to changing demand causing price oscillations), or (iii) no benefit to the consumer (if consumers react slowly to changing demand causing prices to slowly equalize over time). 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 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 utilities should encourage consumers to minimize their peak demand by incentivizing them to flatten their demand rather than simply shift as much demand to the lowest-price period. Of course, the more consumers attempt to flatten their own demand, the more the grid's aggregate demand will flatten. To this end, we propose a new type of electricity pricing plan for residential consumers, which we call *flat-power pricing*. Note that although electricity's price is set by supply and demand in the wholesale market, utilities have the freedom to offer consumers different pricing models. For instance, neither flat-rate pricing, which charges the same price per kilowatt-hour at all times, nor TOU pricing conform to wholesale prices. These pricing plans, instead, simplify consumer pricing relative to the wholesale market.

Our flat-power pricing model charges consumers $\$\beta$ per kilowatt-hour for power usage less than a consumer-specific target and then charges $\$(1 + \alpha)\beta$ for any usage above that target. Figure 5 depicts how flat-power pricing works, where the consumer-specific target is near 1kW, the energy usage in blue costs $\$\beta$ per kilowatt-hour, and

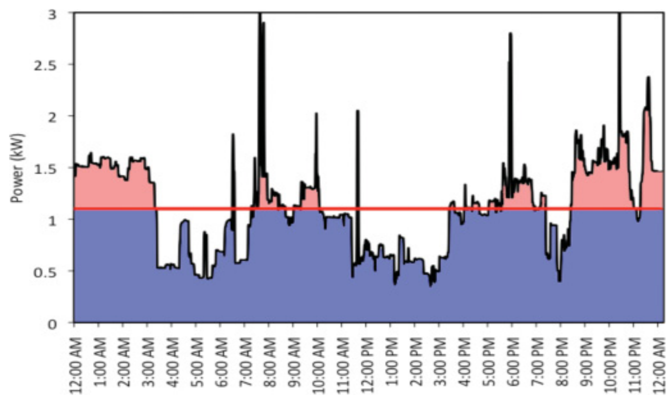

Fig. 5. Flat-power pricing charges consumers $\$\beta$ per kilowatt-hour for power usage less than a consumer-specific target (in blue) and $\$(1+\alpha)\beta$ for usage above the target (in red). The magnitude of α determines the incentive to flatten demand.

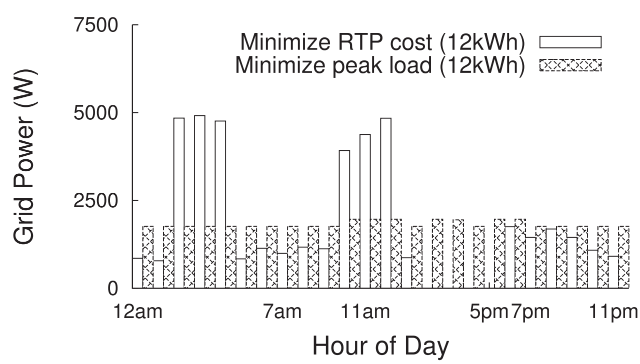

Fig. 6. Although 12kWh of energy storage is capable of shifting only a fraction of demand to the low-price period, it is more than enough to completely flatten the demand from Figure 2.

the energy usage in red costs $\$(1+\alpha)\beta$. Unlike variable rate pricing, flat-power pricing directly incentivizes consumers to flatten their own demand, where the α parameter controls the magnitude of the incentive, rather than shift as much demand as possible to low-price nighttime periods. Importantly, the approach aligns consumer incentives with the grid's incentive to reduce its peak and flatten demand without requiring the grid to actively coordinate distributed battery charging and discharging, e.g., by explicitly what fraction of homes are charging and discharging over time. In addition, flattening a consumer's demand takes significantly less energy storage capacity than shifting all of it to the lowest-price period. For example, Figure 6 shows that although 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 2.

Thus, flat-power pricing encourages distributing aggregate storage capacity widely across consumers, requires less storage capacity per consumer, and results in lower upfront capital costs. 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). The approach also prevents rebound peaks and grid instability, as consumers are flattening, rather than shifting, their demand, and maintains the incentive to use energy storage as capacity scales, as consumers always benefit from not incurring a penalty, based on α, regardless of other consumers' behavior. Finally, utilities can mitigate free riding by altering α's value, in addition to β, as generation costs change, as only the set of consumers with energy storage are able to automatically optimize for peak demand. A higher α and lower β penalizes consumers with energy storage less than consumers without it. We elaborate on these advantages in the following.

3.1. Advantages

Scalability. Flat-power pricing increases the incentive for consumers to install and use small amounts of local energy storage. Although prior research has demonstrated the benefits of energy storage to the grid, i.e., to reduce its generation costs and carbon footprint, handle intermittent renewable energy sources, and increase reliability, it has not addressed how to incentivize consumers to adopt it. As we show in Section 5, flat-power pricing (for suitable values of α) provides consumers a strong monetary incentive to use small amounts of energy storage (or scheduling of elastic background loads) to flatten their grid demand. In contrast, existing variable rate pricing plans only incentivize shifting large amounts of load from low-price nighttime periods to high-price daytime periods, which necessitates large amounts of energy storage. In addition, unlike with variable rate plans, flat-power pricing requires only limited coordination with a utility, as we expect α and the consumer-specific target to rarely change. In contrast, consumers who opt to use variable rate plans must react appropriately to constantly changing prices and price forecasts, often daily via the Web, e.g., https://www.powersmartpricing.org/prices/.

Stability. One issue with existing variable rate plans is that, as we discuss, if consumers were to adopt load scheduling algorithms at scale, all consumers would chase low prices in tandem, potentially resulting in grid instability. In the extreme, if all consumers shift all of their power usage to the lowest-price period, the result would be a rebound peak much greater than the original peak. In some sense, variable rate pricing plans implicitly rely on the assumption that the price elasticity of electricity demand is low: only a small fraction of consumers are likely willing to change their behavior to lower their electricity bill. Oddly, this assumption actually relies directly on there being a weak incentive to install energy storage. However, the presence of distributed energy storage at scale would invalidate the assumption by making demand highly elastic with price. Thus, even slight price variations could have dramatic effects on the grid, akin to flash crashes. In contrast, flat-power pricing is more conducive to algorithmic optimization, as consumers focus on flattening their own demand rather than reacting (potentially in tandem) to changing grid conditions. If all consumers flatten their own demand, it achieves the same goal as variable rate plans: a flatter grid demand with a lower peak.

Utilities already use a variety of pricing plans that incorporate a peak-based component, typically for large industrial consumers [Braithwait et al. 2007]. We discuss less common pricing models that incorporate a peak-based component in Section 6, many of which attempt to address similar problems as flat-power pricing. In the following, we compare the incentives from the most common variant of peak-based pricing with flat-power pricing, which charges consumers a demand charge based on their absolute peak usage in kilowatts over a small time window, e.g., 30 minutes, in the billing period. This is similar to flat-power pricing in that it encourages consumers to reduce their absolute peak demand. As a specific example, our local utility includes a demand component that charges \$12.52/kW of peak power based on the maximum average power consumed over a 30-minute interval within the billing period.[4] Note that this peak demand charge differs from the variable energy usage charge, which is on the order of $0.10/kWh of energy used over the billing period.

Although this peak-based component strongly incentivizes consumers to reduce their absolute peak, it does not incentivize consumers to reduce their nonabsolute peaks. In the extreme, if a consumer incurs a large peak usage on the first day (or hour) of the billing cycle, he or she has no incentive to reduce their peaks for the remainder of the cycle. Thus, optimal charging algorithms for peak-based pricing focus on predicting when the peak demand will occur, as well as saving energy to ensure that it is available to reduce demand during these peak periods. However, since the focus is on the absolute peak, charging algorithms may not use energy storage effectively to reduce other peaks. Yet reducing the nonabsolute peaks is also important, as the absolute peak may not align with the grid's peak demand. In addition, charging algorithms for peak-based pricing are highly sensitive to predictions of when the absolute peak will occur each billing period. In prior work, we found that accurately predicting when a 30-minute peak energy usage period will occur over a month-long billing period is challenging [Mishra et al. 2013b], hindering any algorithm that must optimize for the absolute peak. Accurate predictions are challenging because consumers often cannot control their absolute peak, especially if caused by a single appliance, such as an electric dryer.

Our flat-power pricing model does not suffer from these drawbacks, as consumers need only optimize for the short-term goal of maintaining power usage at or below the target level α rather than precisely planning their usage over a month-long billing period. As a result, optimizing for flat-power pricing does not require accurately predicting demand over long time horizons and does not penalize as much as peak-based pricing if short-term predictions of future demand are inaccurate. For example, if an incorrect prediction causes a battery charging algorithm to generate a power peak that exceeds the target level α for a brief period, the consumer is only charged for the period their power usage was over the peak, whereas with peak-based pricing the power peak could dictate the consumer's peak demand surcharge for the entire month-long billing period. In addition, predicting a home's average demand over a long period, e.g., an entire day, which is necessary to optimize for flat-power pricing, is generally more accurate than predicting it over shorter time periods, e.g., every 15 minutes throughout the day, which is necessary to optimize for the absolute peak demand.

Renewable integration. Flat-power pricing also better accommodates high penetrations of local intermittent renewable energy sources, such as wind and solar, compared to variable rate plans. Today, consumers who leverage local renewables use "net metering" to release any excess energy into the grid. However, net metering is not feasible at large scales, as the intermittent and uncontrollable nature of renewable generation would make it difficult (or impossible) for utilities to balance the grid's supply and demand. This is one reason most states place strict caps on the aggregate number of consumers who may net meter. For example, in Georgia, only 0.2% of subscribers may net meter; additional subscribers must either waste excess renewable energy or store it in a battery. Of course, these caps are currently not an issue, as renewable adoption is still in its infancy.

However, rising long-term electricity prices combined with the increase in efficiency of solar panels and dropping solar panel prices has led to a dramatic rise in the number of local renewable installations. To scale renewable generation, consumers will need to

[4] See http://www.hged.com/customers/payment-billing-rates/rates/electric-rates/183E%20-%20Large%20General%20Service.pdf.

dampen the renewable variations that they expose to the grid. By incentivizing consumers to flatten their grid demand, flat-power pricing encourages this dampening. Thus, in addition to encouraging consumers to shift their usage, it also encourages consumers to control battery charging and discharging to absorb renewable fluctuations, as it penalizes consumers for these variations.

Simplicity. Finally, in addition to being insensitive to predictions, flat-power pricing is also simple for consumers to understand. The primary objective of variable rate pricing plans is to incentivize consumers to manually alter their behavior by varying electricity prices. In practice, however, consumers may find it difficult to determine how to appropriately alter their behavior, given that prices may rise and fall each hour or day based on electricity's real-time balance of supply and demand. Existing peak-based pricing plans that include an additional surcharge based on a consumer's peak demand add yet another layer of complexity. In contrast, flat-power pricing only requires consumers to know and respond to their consumer-specific target. In this case, a simple energy monitor displayed in a prominent place would indicate whether or not consumers are above their target and thus are incurring higher prices. Therefore, flat-power pricing is more conducive to manual optimization by consumers compared to variable-rate and peak-based pricing models.

3.2. Drawbacks

Whereas flat-power pricing has the numerous advantages described earlier, it also has some potential drawbacks. For example, encouraging consumers to flatten their own demand, in aggregate, may require consumers to install more energy storage capacity than necessary 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 to 12pm and 2kW from 12pm to 12am, whereas the second home uses 2kW from 12am to 12pm and 1kW from 12pm to 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 smart meter data from 14,000 homes in Section 5.

Another potential drawback is that flat-power pricing alters battery charging patterns, which affect lifetime. Battery cycle lifetime is typically rated for a fixed number of charge-discharge cycles to specified DOD, e.g., where each cycle starts at 100% charged and then goes to a fixed DOD, e.g., 50% DOD. Whereas prior switch-based architectures only charge and discharge batteries once per day and conform to these ratings, flat-power pricing requires variable charging and discharging throughout the day to different, often shallower, DODs. Thus, we evaluate battery lifetime under flat-power pricing in Section 5 based on a fixed amp-hour throughput model of battery lifetime, which is a standard model for evaluating battery cycle lifetime under dynamic charging and discharging patterns [Bindner et al. 2005]. The amp-hour throughput lifetime model assumes the total amount of energy charged/discharged to a battery over its lifetime is fixed. Thus, the model handles dynamic and nonuniform charging and discharging patterns by deducting the total amount of energy charged/discharged on a cycle from the battery's fixed lifetime amount of energy charge/discharge. The amp-hour throughput model of battery lifetime is conservative in our evaluation, as it weights any battery charge or discharge equally regardless of the current DOD. However, in reality, charging and discharging at higher DODs have a much larger effect on battery lifetime than at smaller DODs. With flat-power pricing, there tend to be many small charging cycles to small DODs compared to existing variable rate pricing that necessitates a small number of charge/discharge cycles to high DODs.

3.3. Discussion

Flat-power pricing requires setting an appropriate β, α, and consumer-specific target. We expect β to be set similarly to flat-rate pricing plans, which charge the same rate for electricity at all times. For the consumer-specific target, one way to choose it is based on each home's historical average per-day power usage. Using each home's average power is effective, as an optimal load scheduling algorithm would result in the home consuming its average power at all times. Of course, average power changes over the course of the year and week, e.g., weekends versus weekdays, so the target would also need to change. However, using the average power to set the target raises concerns about fairness and manipulation. For example, inefficient homes that have historically consumed more energy overall would pay lower prices for higher levels of use than efficient homes. In addition, homes might attempt to manipulate their target by artificially increasing their consumption to increase their target. Once their target increases, they could then resume a normal level of consumption at a lower price. As a result, we advocate setting targets based on the average power usage of homes with similar size and heating systems, e.g., electric versus gas, both of which are available in public records. To promote energy efficiency, many utilities already maintain such records to provide consumers a ranking of their electricity usage compared to their peers.

One other potential point of concern is that flat-power pricing encourages overoptimizing the use of energy storage. Since the grid aggregates electricity demand over large numbers of consumers, shifting demand over short time periods should not affect the aggregate load profile, as this demand is already highly multiplexed across homes within the grid. Although true today, the introduction of a high penetrations of intermittent solar and wind energy sources into the grid is likely to significantly alter this dynamic and increase the benefits of shifting load over shorter time intervals. We evaluate the impact under increasing levels of renewable penetration in Section 5. In addition, future microgrids, which promise to increase grid reliability by leveraging local energy sources, will benefit much less from aggregation effects [Zhu et al. 2013]. Finally, the potential benefits of incentivizing all consumers to adopt small amounts of energy storage, rather than none of them (which is effectively the case today), likely outweighs the cost of overoptimization.

4. SCHEDULING ALGORITHMS

In this section, we outline algorithms for battery charging and discharging, as well as scheduling of elastic background loads, to optimize electricity costs for flat-power pricing. In both cases, we discuss how these algorithms compare to those that optimize for variable rate pricing.

4.1. Battery Charging and Discharging

To control battery-based energy storage under variable electricity prices, prior work runs a linear programming (LP) at the beginning of each day and takes as input predictions of the next day's demand and electricity price profiles [Mishra et al. 2012; Carpenter et al. 2012]. Assuming that the predictions of next-day demand and electricity price profiles are accurate, the LP-based solution is optimal. The LP-based solution under variable rate pricing is not highly sensitive to prediction accuracy and is generally near the optimal even when predictions are not accurate. In contrast, as discussed earlier, an LP-based solution that incorporates an additional surcharge based on the absolute peak is highly sensitive to the prediction of when the peak will occur and generates highly nonoptimal results even for peaks predictions that are slightly shifted in time, e.g., by an hour.

In general, all battery control strategies that exploit arbitrage under existing variable rate pricing models follow a greedy approach that charges batteries during lower-price periods and discharges them during higher-price periods. Note that the particular control strategy determines how much to charge and discharge the battery based on its assumptions of battery characteristics and predictions of future demand and prices. In contrast, under flat-power pricing, rather than charge the battery to full capacity during the lowest-price period at night, our charging/discharging algorithm attempts to flatten a building's demand to prevent paying the penalty based on α. In particular, we select a target average power for flattening and then simply charge and discharge batteries whenever demand is below or above the target, respectively, to ensure that demand is near the target. Note that rather than run our algorithm once per day using predictions of next-day demand, as with the variable rate algorithms earlier, the algorithm naturally operates in an online manner, adjusting the charging and discharging of the battery in real time based on changing demand in relation to the chosen target. This algorithm works well as long as (i) the target flattening threshold 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 demands when necessary; if the target is too large, then it will store more energy than necessary throughout the day. However, importantly, although 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 presented previously. 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 [Sharma et al. 2010], e.g., the average power of a home each day varies less than its average power every 10 minutes. As a result, accurate predictions of average power over long periods, e.g., a day or month, do not require sophisticated methods [Sharma et al. 2010, 2011]. Thus, we simply use the average demand over the previous interval.

4.2. Elastic Background Load Scheduling

Flat-power pricing enables consumers to benefit from small, rather than large, shifts in their load. In addition to using a battery to shift load to flatten demand, consumers may also schedule elastic background loads, such as refrigerators, freezers, heaters, dehumidifiers, humidifiers, dryers, and dishwashers. As with a battery, each of these loads expose different degrees and types of scheduling freedom. However, this freedom is limited, as generally these loads cannot be deferred for long periods of time. For example, a refrigerator's cycle may have enough slack to delay it for a half hour, but it cannot be delayed for multiple hours without causing its temperature to rise too high. As a result, scheduling elastic background loads is not beneficial under variable rate pricing plans. However, flat-power pricing introduces an incentive to schedule these loads to flatten demand by shifting small amounts of power.

Whereas prior research focuses on optimizing different dimensions of load scheduling freedom in isolation, we examine the performance of jointly optimizing scheduling for all types of loads to quantify the maximum benefits from scheduling loads under flat-power pricing. As in prior work, we formulate a combined scheduling problem as a mixed integer linear program (MILP). In particular, we extend the MILP used in Parasol [Goiri et al. 2013]: a solar-powered microdata center that schedules batch computing jobs with well-known running times and deadlines. Similar to our MILP, Parasol's objective is to minimize electricity costs under variable rate pricing plans and maximize the use of solar power, albeit for a data center instead of a residential home. Of course, our work differs from Parasol, as a building is different from a data center in terms of its users, workloads, loads, and scheduling freedom.

4.2.1. Degrees of Freedom. The primary degrees of scheduling freedom for loads we identify are the ability to transparently shift, slide, stretch, or sell power. Note that some loads may exhibit multiple degrees of scheduling freedom.

Shifting power refers to scheduling on-off loads that operate over one or more duty cycles such that within each duty cycle the load is on for some period of time and off for some period of time. The length of the duty cycle and the amount of time a load is on within each cycle may be either static, e.g., if driven by a timer, or dynamic, e.g., if driven by a thermostat. In either case, the load must be on for a fixed amount of time during a cycle to satisfy its objective, such as keeping the temperature of an enclosed space within a specified range or guardband. However, as long as the scheduler ensures that the load is on for this time each cycle, it is free to determine when the load is on during the cycle without violating its objective. This freedom is often referred to as slack [Barker et al. 2012b; Taneja et al. 2010].

Prior work focuses on scheduling these shiftable on-off loads either to (i) align as much load as possible with renewables [Taneja et al. 2010], (ii) defer loads during times of grid constraint [Ganu et al. 2012], or (iii) flatten a building's demand profile [Barker et al. 2012b; Karmakar et al. 2013]. However, prior work does not address the lack of monetary incentive for consumers to shift loads. Since the duty cycle length for common shiftable loads, such as an air conditioner or refrigerator, are at most a couple of hours, and often much less, schedulers are only able to shift these loads within narrow time periods without violating their objectives. As a result, shiftable loads cannot generally exploit variable rate pricing models, which require consumers to shift load to off-peak periods that may be many hours from the peak periods.

Sliding power refers to deferring the start time of a batched load, such as a dishwasher, washing machine, or dryer, that performs a usually nonpreemptible batch task for a fixed, well-known period of time [Taneja et al. 2010]. These loads are first initialized by a user and then are run for a predetermined amount of time without user intervention before being emptied and reset by the user. In some cases, these loads are pipelined. For example, clothes are usually washed and then dried in sequence. In addition, some loads may also be preemptible, in which case, once started, they also act as shiftable loads that operate over a single duty cycle. As long as a user has initialized a load, e.g., filled it with clothes or dishes, a scheduler has the freedom to indefinitely delay its start time. Of course, delaying the start time also delays the end time, which may in turn cause delays in the pipeline. The primary constraint for transparently scheduling slide loads is the amount of inconvenience a user is willing to tolerate. Since users directly control when to initialize and start slide loads, incentivizing them to change when they operate these loads is an important goal of existing variable rate pricing plans. Unfortunately, as earlier, users are often unwilling to delay start times many hours into the future, e.g., from daytime to nighttime, which limits scheduling freedom.

Table I. Parameter Definitions for Linear Program

Parameter	*Definition*
T	Time in T discrete intervals 1 to T
p_i	Home power demand in interval i
$\vec{L_{slide}}$	Vector of tuples for slide load schedules
$\vec{L_{shift}}$	Vector of tuples for shiftable load schedules
g_i	Renewable generation (kWh) in interval i
α	Percentage of cost paid for net metering
C	Battery capacity in kilowatt-hours
e	Battery efficiency, ≤ 1
m_i	Electricity cost in interval i
L_{green}	Demand satisfied by renewables
L_{grid}	Demand satisfied by grid
$L_{battery}$	Demand satisfied by battery
B_{green}	Renewable energy charged to battery
B_{grid}	Grid power used to charge battery
N_{green}	Renewable energy net metered to grid
P_{grid}	Power used from the grid

Stretching power refers to extending an elastic load's running time while lowering its average power usage to keep its energy consumption constant for a particular task [Srikantha et al. 2012]. Typically, elastic loads utilize resistive heating elements or variable drive motors, which enable a scheduler to precisely adjust their temperature or speed, respectively, to dictate a specified running time. Examples of elastic loads include washing machines, dryers, dishwashers, ovens, stoves, refrigerators, freezers, air conditioners, electric water heaters, and electric space heaters. Of course, schedulers cannot arbitrarily stretch a load, as the average power usage and duration of a task affects its operation. For example, running dryers at high heat for a short duration works well for heavy fabrics, whereas low heat for a long duration is better for delicate fabrics. Thus, prior work [Srikantha et al. 2012] places a low upper limit (approximately 10%) on stretching a load's running time.

Selling power, e.g., via net metering, is an attractive option for homes that generate their own power using on-site renewables, such as solar power. However, in many cases, the price utilities pay consumers for power is less than the price they charge them for it. In addition, current laws often place strict caps on the amount of power a utility must buy back from a consumer. These dynamics alter the scheduling problem by incentivizing consumers to not only transfer load to low-price periods but also to align as much of it as possible with renewable generation [Mishra et al. 2013a; Taneja et al. 2010; Zhu et al. 2011]. Aligning a home's load with renewable generation also decreases grid transmission losses, as it increases the amount of power consumed at the point of generation. The preceding scheduling limitations also impact the freedom to align load with renewable generation.

4.2.2. Joint Optimization. We define an MILP that jointly optimizes the degrees of scheduling freedom above using the parameters listed in Table I. We model both $L_{shift}[\vec{power}]$ and $L_{slide}[\vec{power}]$ as vectors of three-tuples that specify each load's start time, running time, and power usage. We assume that slide loads also have a completion deadline. We then divide each day into T discrete intervals of length l from 1 to T with the objective of minimizing $\sum_{i=0}^{T}(m_i * P_{grid}^i - \alpha * m_i * N_{green}^i)$ each day, i.e., the net bill after any net metering, given the following constraints.

$$\forall i \in T, B_{green} + B_{grid} \leq \frac{C}{4} \tag{1}$$

$$\sum_{i=0}^{T} L_{battery} - e * \sum_{i=0}^{T} (B_{green} + B_{grid}) \leq 0 \tag{2}$$

$$\sum_{i=0}^{T} B^{i}_{green} + \sum_{i=0}^{T} B^{i}_{grid} - \frac{\sum_{i=0}^{T} L_{battery}}{e} \leq C \tag{3}$$

$$\forall i \in T \,|\, L^{i}_{battery} > 0, N_{green} = 0 \tag{4}$$

$$\forall i \in \{T \,|\, L^{i}_{battery} > 0\}, B^{i}_{green} + B^{i}_{grid} = 0 \tag{5}$$

$$\forall i \in T, L^{i}_{green} + B^{i}_{green} + N^{i}_{green} \leq g_i \tag{6}$$

$$\forall i \in T, L^{i}_{grid} + B^{i}_{grid} = P^{i}_{grid} \tag{7}$$

$$\forall i \in \{T \,|\, L^{i}_{grid} > 0\}, N^{i}_{green} = 0 \tag{8}$$

$$\begin{aligned} \forall i \in T, L^{i}_{battery} + L^{i}_{grid} + L_{green} \\ = L_i - L^{i}_{slide}[\vec{power}] - L^{i}_{shift}[\vec{power}] \end{aligned} \tag{9}$$

$$\begin{aligned} \forall \vec{L^{i}_{shift}}, \forall i \in \frac{T}{L^{i}_{shift}[period]}, \sum L^{i}_{shift}[power] \\ = L_{shift}[power] \end{aligned} \tag{10}$$

Briefly, the constraints ensure the following invariants. Constraint (1) dictates that the battery's charging rate is not more than its capacity divided by 4, i.e., a C/4 charge rate. We use a maximum charging rate of C/4, as recommendations for maximum charging rates for sealed lead-acid batteries range between C/5 and C/3 charging rates [Bullock and Salkind 2011]. However, we note that maximum charging rates are generally not required for either flat-power pricing or existing variable rate pricing models. Since flat-power pricing incentivizes consumers to operate at their average demand, it naturally reduces the amount a battery must discharge over a small time window. In addition, maximum charging rates are generally not necessary in variable rate pricing models, as they typically offer long low-rate periods for charging at night. Constraint (2) ensures that the energy charged to the battery never exceeds the energy discharged from it. Similarly, constraint (3) is a condition to ensure that the energy stored in the battery never exceeds its capacity, i.e., the maximum amount of energy it is capable of storing. Finally, constraint (4) dictates that net metering and battery charging do not occur simultaneously, (5) battery charging and discharging do not occur simultaneously, (6) renewables can charge the battery, (7) renewables can be net metered, (8) consuming grid power and net metering cannot occur simultaneously, (9) every load is powered by only one energy source, and (10) the amount of power shiftable loads consume per period is constant.

Constraints (4), (5), and (8) are nonlinear mutual exclusion constraints. We convert these to linear constraints by introducing a binary variable $b \in \{0, 1\}$ and replacing each nonlinear mutual exclusion constraint with five linear constraints that enforce the

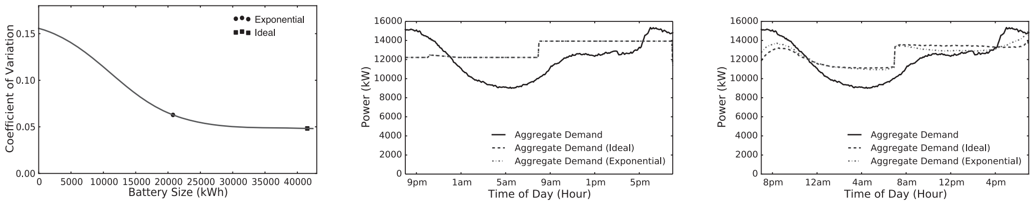

(a) Battery size (b) Centralized (c) Distributed

Fig. 7. (a) The coefficient of variation of grid power usage over 11 months as a fraction of centralized battery capacity under flat-power pricing. (b) Aggregate demand without energy storage and with centralized energy storage using ideal and exponential capacity under flat-power pricing. (c) Aggregate demand without energy storage and with each home using the ideal amount under flat-power pricing. Note that aggregate demand (ideal) and aggregate demand (exponential) nearly precisely overlap each such that it is difficult to distinguish them.

same invariant. In this case, we replace any constraint of the form $\forall i \in \{T|x > 0\}, y = 0$ with $x - \infty * b \leq 0$, $-\infty * x + b \leq 0$, $y + \infty * b \leq \infty$, $x \geq 0$, and $y \geq 0$. In addition, since neither slidable loads nor stretchable loads map well to linear constraints, we use brute-force methods to determine an optimal schedule. For each slide load each day, we simply run the MILP multiple times for each possible start time of the slide load and then use the minimum cost schedule. Of course, as the number of slide loads increases, we must run the MILP for each possible combination of start times, which increases exponentially as the number of slide loads increases. However, the approach is computationally tractable in practice, as the number of slide loads is typically small, e.g., usually three or less, and they do not run everyday. Likewise, for each stretchable load, we run the MILP for each possible stretched duration.

5. EVALUATION

We evaluate the ability of flat-power pricing to incentivize both distributed energy storage and elastic load scheduling. We examine both grid- and consumer-scale effects. For our grid-scale experiments, we use smart meter data from 14,000 homes in a small town over 11 months, whereas for our consumer-scale experiments, we use fine-grain per-load energy data from Home-A in the Smart* dataset [Barker et al. 2012a].[5] To estimate the effect of flat-power pricing on generation costs, we use the demand-cost function in Figure 2(b), which has been scaled based on the maximum generation capacity in our traces.

5.1. Distributed Energy Storage

We first evaluate the aggregate battery capacity necessary with flat-power pricing, both when it is centralized and when it is distributed across all homes. For this experiment, we assume that we set a new target threshold per home every 12 hours such that the home attempts to flatten demand over the 12-hour period. Figure 7(a) plots the coefficient of variation—the ratio of the standard deviation σ to the mean μ—of the grid's power usage over the 11-month period with different battery capacities when operated in a centralized manner. The coefficient of variation is a normalized measure of variation in data such that the lower the coefficient of variation, the flatter the grid's demand profile. We label with a square the point on the x-axis where an increasing battery capacity does not change the grid's demand profile as *ideal*, as this is the flattest possible profile with any battery size. The graph shows that as the battery capacity decreases on the x-axis, the coefficient of variation remains steady until a knee point

[5] From July 1, 2012, to August 30, 2012.

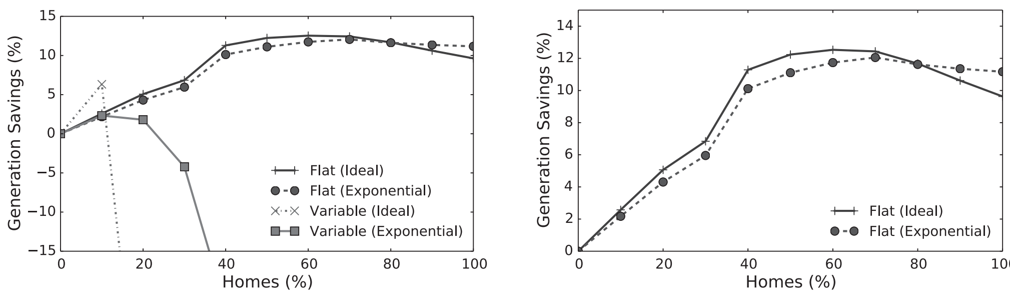

(a) Centralized (b) Distributed

Fig. 8. (a) Comparison in generation savings when the fraction of homes on the x-axis use energy storage under a variable rate versus flat-power pricing. (b) A zoomed-in view of (a) that focuses on the savings under flat-power pricing.

where it begins to increase. We label this knee point with a circle as *exponential*, since this is the point where the coefficient of variation begins to increase exponentially with decreasing battery size. The graph demonstrates that a significant fraction of the possible flattening can be achieved with significantly less than the ideal-sized battery capacity.

In our city-scale dataset, over the course of the 11-month period, the ideal centralized battery capacity is 43.3MWh, whereas the distributed capacity necessary to maintain the knee point above is only 62.9MWh, which is within 31% of the ideal centralized capacity. In addition, in the distributed case, the median battery size is only 3.44kWh per consumer on average, which is significantly less than the 12kWh necessary to maximize the benefits under variable rate pricing plans [Mishra et al. 2012]. Figure 7(b) and (c) illustrate the effect on demand of centralized versus distributed use of energy storage over a single day. The figures show that the centralized and distributed use of energy storage yields essentially the same demand pattern, as the profile in Figure 7(b) and (c) is nearly the same. However, in the distributed case, the demand pattern stems from thousands of consumers independently charging and discharging their batteries based on the incentives introduced by flat-power pricing without any centralized coordination. In both cases, note that aggregate demand (ideal) and aggregate demand (exponential) nearly precisely overlap each such that it is difficult to distinguish them. This demonstrates the diminishing returns of increasing aggregate battery capacity (whether centralized or decentralized) beyond the knee point. The graph also shows that distributed energy storage under flat-power pricing is similar to that of a centralized energy storage system.

We next examine the effect of rebound peaks when consumers use energy storage at large scales. Figure 8(a) shows the 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 flat-power pricing and variable rate prices. For these experiments, we set the battery capacity equal to both its ideal value, i.e., where a greater capacity will have no effect on the experiment, and a lesser value that performs similarly, e.g., equivalent to the "knee" of the battery-cost curve. The figure shows that as the number of homes using energy storage scales up, using the ideal charging algorithm (the SmartCharge algorithm from prior work [Mishra et al. 2012]), increases generation costs, i.e., the savings are negative, due to simultaneous battery charging and large rebound peaks after roughly 20% of homes use energy storage with variable rate pricing. In contrast, flat-power pricing steadily decreases the grid's generation costs as more homes use energy storage, signaling that homes and the grid are successfully flattening demand.

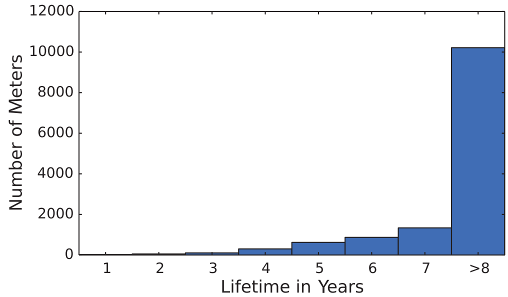

Fig. 9. Histogram of battery lifetime based on an amp-hour throughput model under flat-power pricing.

Notice that the maximum savings from both the variable rate pricing and flat-power pricing are similar: this reflects that in both cases, the aggregate energy storage is sufficient to flatten demand. The difference with flat-power pricing is that it distributes this capacity across 100% of consumers, whereas variable rate pricing distributes it across a much smaller set. Figure 8(b) zooms into the grid's generation savings under flat-power pricing from Figure 8(a) and shows that flat-power pricing reduces generation costs over the course of the year by up to 13%. Note that at high adoption rates under flat-power pricing, the ideal battery capacity that results in the flattest possible profile has less savings than the nonideal exponential battery capacity based on the knee of the capacity-variation trade-off. This occurs because the ideal battery capacity is significantly larger than the exponential capacity. The increased battery capacity at high adoption rates results in more energy conversion losses, which increases the overall demand on the grid and reduces the savings from using batteries.

We next evaluate the lifetime of batteries under our charging algorithm for flat-power pricing. We evaluate battery lifetime under an amp-hour throughput model, which is the standard model for evaluating battery lifetime under dynamic and nonuniform charging and discharging patterns [Bindner et al. 2005]. Recall from Section 3.2 that the amp-hour throughput lifetime model assumes that the total amount of energy charged/discharged to a battery over its lifetime is fixed and simply deducts the amount of energy stored/released by a battery on each cycle (based on the depth of the cycle's discharge or charge). Thus, since large charges and discharges cause more current to pass through the battery, they also have a larger negative affect on its lifetime. Figure 9 shows that most batteries last more than 8 years under flat-power pricing when using our charging algorithm to flatten demand. If we assume that the maximum battery lifetime is 8 years, independent of the charging cycles, then the average lifetime under flat power pricing is 7.47 years. By comparison, under the same lifetime model, using an existing variable rate charging algorithm [Mishra et al. 2012] that typically charges and discharges batteries once per day (to a 50% DOD), as with a variable rate pricing plan, the average battery lifetime would be 6.8 years. Thus, flat-power pricing increase battery lifetimes compared to existing variable rate pricing models. The increase is due to the much shallower discharge depths required by flat-power pricing.

Finally, Figure 10 shows the effect of flat-power pricing on grid demand with increasing levels of intermittent solar energy. Supporting high levels of solar penetration requires energy storage to absorb large fluctuations in renewable output. For example, Figure 10(a) shows that even on a relatively sunny day, the energy generated by this 10kW home solar installation rises and falls dramatically based on passing clouds. Mechanical generators cannot offset such large fluctuations to balance supply and demand. Figure 10(b) then shows the ideal battery capacity under flat-power pricing as the percentage of consumers with 10kW solar systems scales up. As is seen in the

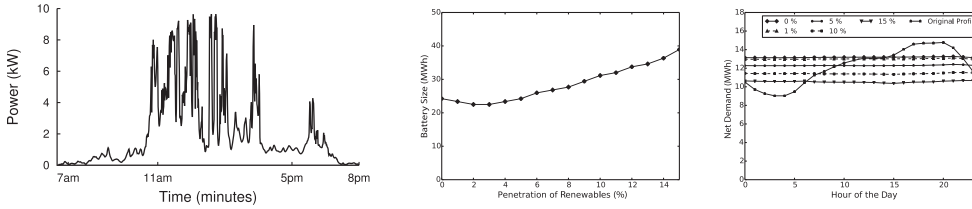

Fig. 10. Example of how power generation from a 10kW rooftop home solar installation can vary widely minute to minute (a). Ideal battery capacity with flat-power pricing for increasing levels of solar penetration (b) and associated load curves under increasing levels of solar penetration (c).

figure, this capacity drops initially, as increased solar energy decreases the difference in daytime and nighttime consumption but then climbs steadily as more solar energy introduces a more stochastic grid profile. The increase in required energy storage capacity under increasing levels of solar penetration increases the importance of incentivizing users to install distributed energy storage. Finally, Figure 10(c) shows how the grid's demand profile still remains nearly flat despite increasing levels of solar penetration. Note that this graph only shows the grid's demand profile and does not include the additional solar generation.

5.2. Elastic Load Scheduling

The purpose of our MILP from Section 4.2.2 is to quantify the cost savings from optimizing each dimension of scheduling freedom, both in isolation and in combination over a 60-day period in a representative home for TOU, RTP, and flat-power pricing plans. The TOU and RTP plans are depicted in Figure 1. Since we use the same RTP prices from Figure 1 each day, our RTP plan has significantly more opportunity to lower costs by scheduling loads than a real RTP plan. Thus, our RTP plan represents a conservative upper bound on the size of the price differential between the lowest- and highest-cost periods: a real RTP plan would result in much less savings. In our experiments, we run our MILP at the beginning of each day with $T = 24$, assuming that we know the home's power demand p_i, renewable generation g_i, and electricity cost each interval. In practice, the scheduler would require predictions for these parameters [Mishra et al. 2012; Sharma et al. 2010]. Therefore, our results also represent an upper bound on the cost savings due to scheduling.

5.2.1. Variable Rate Pricing Plans. Figure 11(a) and (c) and Figure 12(a) and (c) show the extent to which optimizing each degree of scheduling freedom lowers Home-A's electricity bill under existing variable rate pricing plans. In this case, the primary shiftable load in the home is an air conditioner that on average consumes 13kWh per day such that it is on for 33% of its duty cycle. The primary slidable and stretchable loads include a dryer, washing machine, and dishwasher. For each graph, the y-axis shows the percentage cost savings from scheduling each type of load in isolation, and the x-axis represents the degree of scheduling freedom for each load—in this case, the duty cycle length for shiftable loads, the maximum delay for slidable loads, the stretch factor (as a multiple of the original running time) for stretchable loads, and the amount of renewable energy available to sell through net metering. For renewables, the x-axis represents a multiplicative factor applied to a real solar trace sized such that it provides 50% of the home's energy. For example, the two on the x-axis represents a trace that provides 100% of the home's energy.

As expected, as the scheduling freedom increases, so do the savings. Unfortunately, practical values in each case are generally low, with each offering less than a 10% reduction in costs even with our extreme RTP pricing model and generous assumptions,

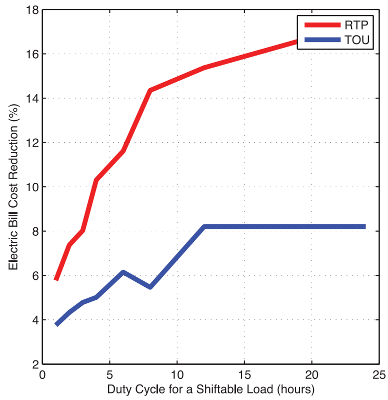

(a) Shift (Variable rate pricing)

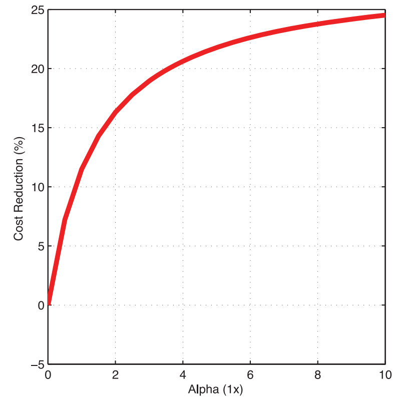

(b) Shift (Flat-power pricing)

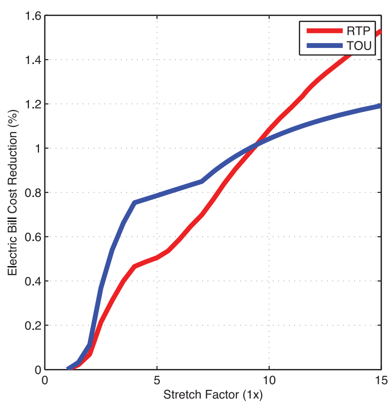

(c) Stretch (Variable rate pricing)

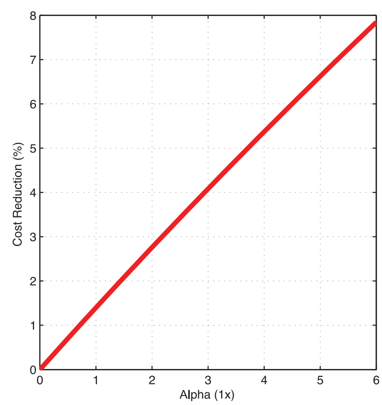

(d) Stretch (Flat-power pricing)

Fig. 11. Cost savings from optimizing each degree of scheduling freedom (shifting and sliding) for electric loads using variable rate pricing and flat-power pricing within reasonable limits.

e.g., future knowledge of demand and prices, and a high price differential. For example, as the length of the duty cycle for shiftable loads increases, the scheduler has more freedom to shift power usage for long periods of time without violating the constraint that energy usage within a duty cycle must be constant. In practice, however, common shiftable loads, such as refrigerators, freezers, heaters, and air conditioners, have duty cycles of only a few hours or less, which results in savings of less than 10%. This low savings is due to the fact that RTP and TOU plans require loads to shift their usage over long periods (from high-price daytime periods to low-price nighttime periods).

Significantly increasing the duty cycle for these loads is not possible without redesigning them to incorporate some form of energy storage, e.g., such as thermal storage [Taneja et al. 2013]. Similarly, although slide loads significantly reduce costs if deferred multiple hours into the future (to low cost nighttime periods), such long delays impose a significant burden on users. Likewise, stretching loads only provides significant savings for unrealistically large stretch factors, e.g., $> 3x$, that are much greater than the 10% assumed in prior work [Srikantha et al. 2012]. Although energy storage and renewable energy are capable of lowering electric bills, they both require a large capital investment that generally negates any savings on an electric bill. For

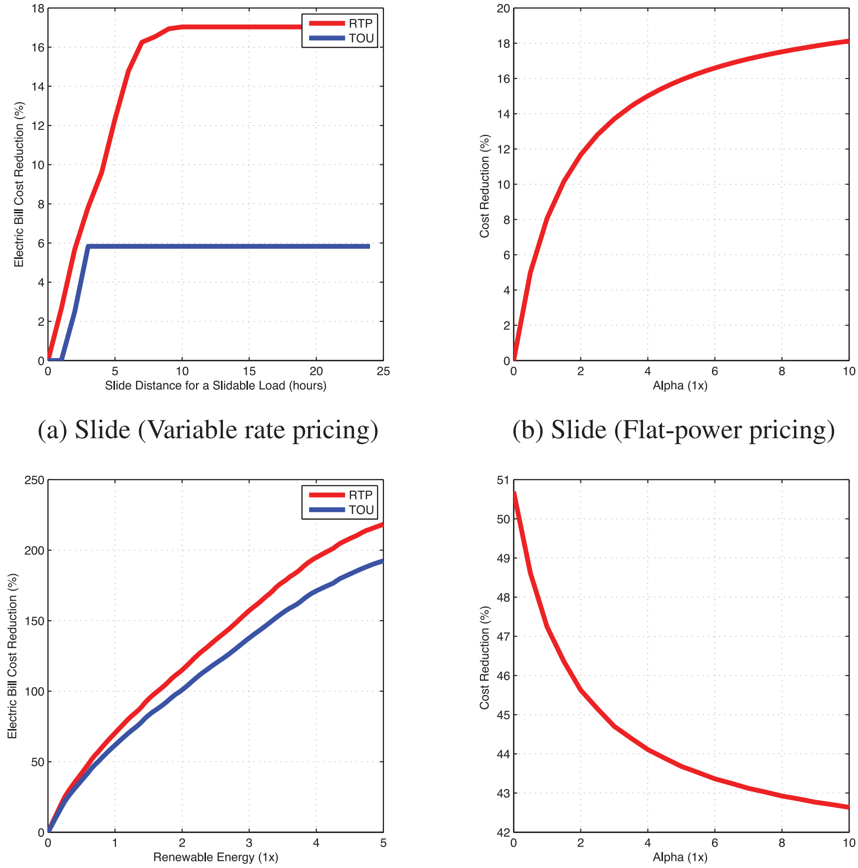

(a) Slide (Variable rate pricing) (b) Slide (Flat-power pricing)

(c) Sell (Variable rate pricing) (d) Sell (Flat-power pricing)

Fig. 12. Cost savings from optimizing each degree of scheduling freedom (stretching and selling) under variable rate pricing and flat-power pricing within reasonable limits.

example, prior work estimates that each kilowatt-hour of usable energy storage costs $118 per year to maintain [Mishra et al. 2012].

Finally, Figure 13(a) shows the combined benefits from jointly optimizing the scheduling for each type of load for "reasonable" values of the scheduling freedom and then compares it to the sum of benefits from optimizing each load in isolation. In this case, we chose a 2-hour duty cycle for the shiftable loads, 4 hours for the delay time of slidable loads, and $2\times$ for the stretch factor of stretchable loads, with no battery capacity or renewable energy. We believe that these represent a reasonable upper limit on the scheduling constraints for a large class of today's homes. Higher limits would impose additional requirements on the loads or pose an excessive inconvenience to users. We do not include any battery capacity or renewable energy due to their high upfront capital costs, as we evaluated them independently in the previous section.

Our results in Figure 13(a) shows that the benefit from jointly scheduling based on the preceding constraints is actually less than the sum of the benefits from scheduling each dimension independently, as optimizing one degree of freedom often prevents optimizations from another. Overall, the results show that advanced load scheduling offers at most 20% savings in the extreme RTP case and 11% in the more common

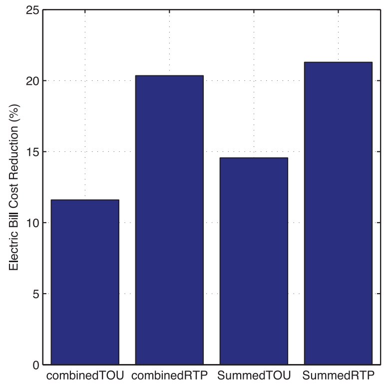

(a) TOU/RTP Pricing

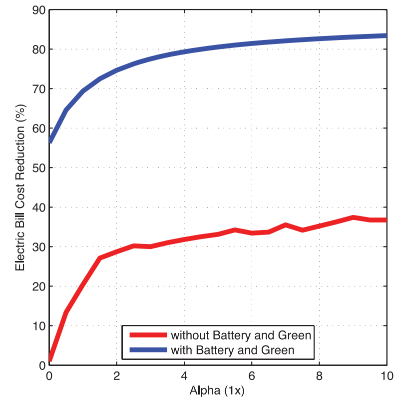

(b) Flat-power Pricing

Fig. 13. (a) Cost savings when optimizing the scheduling of shift, stretch, and slide loads both independently and jointly under TOU and RTP pricing models without energy storage or renewable energy. (b) Cost savings when optimizing the same loads under a flat-power pricing model for different values of α both with and without energy storage and renewable energy.

TOU case. Since many of our assumptions are conservative, the benefits are likely to be much less in practice.

5.2.2. Flat-Power Pricing. By contrast, Figure 11(b) and (d) and Figure 12(b) and (d) show the savings from scheduling each type of load in isolation under our flat-power pricing plan for increasing values of α. For these experiments, we chose the consumer-specific target to be equal to the home's average power. In this case, for each load, we evaluate it using the "reasonable" values from Figure 13(a): a 2-hour duty cycle for the shiftable loads, 4 hours for the delay time, and $2\times$ for the stretch factor. For storage and selling, we use a 10kWh battery capacity and our solar energy harvesting trace scaled at $1\times$. Our results show that in each case, the cost savings for modest values of α, such as 3, are higher than for the TOU/RTP plans.

For example, Figure 11(b) shows that the savings from shifting loads is 18% for $\alpha = 3$, whereas in Figure 11(a) it is 4% to 7%. Likewise, Figure 12(b) shows that the savings from sliding loads is 13% for $\alpha = 3$, whereas in Figure 12(a) it is approximately 6%. Stretching power also show greater benefits. Interestingly, solar harvesting offers roughly the same benefit as with the TOU/RTP plans (nearly 50%) at $1\times$ scale (from Figure 11(b)). The reason is that with TOU/RTP, solar panels produce energy near the optimal time to minimize costs, i.e., during peak periods. In addition, for increasing values of α, solar energy actually sees decreasing relative savings with flat-power pricing. Since we choose the target to be equal to the home's average power, the increasing amount of solar energy lowers the consumer-specific target, resulting in higher prices due to more electricity usage being above the target. The dynamic demonstrates that the consumer-specific target should be independent of a home's renewable generation.

Finally, Figure 13 shows the savings for combined load scheduling under both variable rate TOU/RTP pricing models (Figure 13(a)) and our flat-power pricing model (Figure 13(b)). As discussed earlier, Figure 13(a) shows the savings under existing pricing models for both independently optimizing each scheduling degree of freedom and from jointly optimizing them under representative TOU and RTP pricing models. Figure 13(b) then shows the savings from optimizing the same degrees of scheduling freedom under flat-power pricing for different values of α. As in Figure 13(a), we do not consider energy storage or renewable energy, as we evaluate their effects independently in the previous section. Without any energy storage or renewable energy, in Figure 13(b) the savings are 30% for an $\alpha = 3$ and up to 40% with higher α values compared to the 11% and 20% savings in the TOU and RTP cases, respectively, in Figure 13(a). In addition, if we add in 10kWh of energy storage and $1\times$ of solar generation, the savings are 77% for $\alpha = 3$ compared to 40% and 25% when adding energy storage and $1\times$ of solar generation in the TOU and RTP cases. The results show that flat-power pricing offers consumers a greater incentive in all cases to schedule loads than existing TOU/RTP plans, enabling utilities to control the incentive by raising or lowering the α parameter.

6. RELATED WORK

This article focuses on combining a pricing model, charging algorithm, and load scheduling to incentivize small load shifts. There is a large body of prior work on each of these aspects. However, prior work generally focuses on only a single aspect. For example, as we discuss in the following, prior work investigates battery charging algorithms in the context of existing variable-rate pricing models. In contrast, this article's primary contribution lies in its combination of these aspects to incentivize small load shifts. We discuss prior work on each aspect in turn.

Pricing models. Section 3.1 compares flat-power pricing to a common hybrid pricing model for industrial consumers that incorporates a peak-based demand charge; however, utilities are experimenting with a variety of other less common hybrid pricing models, as outlined by Braithwait et al. [2007] and Faruqui and Eakin [2002] in surveys on pricing models. For example, Liu et al. [2013] analyze workload scheduling algorithms data centers under coincident peak pricing; this type of pricing is similar to a peak-based demand charge except the charge in dollars per kilowatt-hour is based on power usage when the grid's demand peaks rather than when the consumer's own demand peaks. Optimizing battery charging and load scheduling for coincident peak pricing has problems similar to the peak-based demand charges detailed in Section 3.1. In other words, the scheduling is sensitive to accurate predictions of when grid demand will peak, and the pricing model does not incentivize reducing the nonabsolute peak. Flat-power pricing is most similar to existing consumer baseline load (CBL) pricing models, which charge consumers a fixed price for conforming to a baseline load and penalize them from deviating from the baseline [Faruqui and Eakin 2002]. However, with CBL pricing, the baseline load may follow any profile set by the utility and generally follows a typical profile with high power usage during the day and low power usage at night. In contrast, flat-power pricing incentivizes a flat demand profile by penalizing usage above a target.

Energy storage. 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. [1989] was the first to propose this form of energy arbitrage. This work and work by van de Ven [2011] study the problem from a theoretical standpoint, e.g., assuming certain demand distributions without evaluating their solutions on real data. More recently, work by Mishra et al. on SmartCharge [Mishra et al. 2012] and Carpenter et al. [2012] study a similar problem in a realistic setting when taking into account battery inefficiencies, stochastic demand in residential settings, and existing variable rate pricing plans in Ontario and Illinois. Both works highlight the problems with scaling distributed energy storage to many consumers, but neither (i) explores the full implications of large-scale energy storage adoption, including the decreasing costs for consumers with storage as adoption scales up, nor (ii) proposes or evaluates a solution to the problem.

Johnson et al. [2011] 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 beyond a significant fraction of usage.

Koutsopoulos et al. [2011] explore the problem of managing large-scale energy storage from the perspective of a utility operator. In this case, the utility controls when to charge and discharge battery-based storage to minimize generation costs, assuming that the marginal cost to dispatch generators is similar to Figure 2(b). However, this approach focuses on large centralized energy storage facilities, whereas we target decentralized control of distributed energy storage.

Although the datasets in the preceding works are different, they both show that nearly 20% of homes using energy storage maximize 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 prior work, Vytelingum et al. [2010] show 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 UK households employ energy storage (based on a UK dataset). Although slightly higher than the nearly 20% of homes found earlier, the work by Vytelingum et al. shows the same trend as prior work: 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. In addition, prior work also examines a similar problem when homes also have distributed solar generation, as energy storage is useful in such systems [Mishra et al. 2013a]. Our experiments show that flat-power pricing also incentivizes consumers to use energy storage to flatten demand as the renewable penetration levels increase.

Whereas the preceding work focuses on residential settings, prior work has also looked at similar problems from the perspective of industrial consumers, particularly data centers [Govindan et al. 2011] but has not examined the effects from the use of energy storage at large scales. Prior work also highlights the effect of variable rate pricing on grid stability [Roozbehani et al. 2010], showing that RTP has the potential to create an unstable closed feedback loop. Finally, we know of no work that proposes and evaluates using a pricing plan to maintain a stable grid and prevent rebound peaks by incentivizing consumers to flatten their own demand. Further, as discussed in Section 4.2.1, numerous researchers have studied load scheduling algorithms for different types of loads with different degrees of freedom. However, there is no incentive under variable rate electricity plans to use these load scheduling algorithms, as most elastic loads generally cannot shift their power usage over the long time periods necessary to decrease electricity costs. Our work shows how to jointly optimize for scheduling multiple loads with different degrees of scheduling freedom and quantifies the benefit of scheduling such loads under flat-power pricing.

Load scheduling. As we discuss in Section 4.2.1, there is a large body of work on scheduling elastic background loads in lieu of using energy storage. In general, this work focuses on device-specific algorithms that differ based on the degrees of scheduling freedom available to the device. For example, devices may be able to shift [Barker et al. 2012b; Ganu et al. 2012; Karmakar et al. 2013], slide [Taneja et al. 2010], or stretch [Srikantha et al. 2012] power. However, prior work develops load scheduling algorithms for existing variable rate pricing models. In this article, we quantify the benefits of optimizing these degrees of freedom under flat-power pricing. We show that many devices only have the freedom to perform small load shifts, which, even though they benefit the grid, are not incentivized under existing variable rate pricing models. In contrast, flat-power pricing introduces an incentive for devices to perform such small load shifts.

Summary. In summary, our primary contribution is a flat-power pricing model that encourages small load shifts by consumers, which better incentivizes distributed energy storage and scheduling of elastic background loads compared to existing pricing models. Although Xu et al. [2013] explore the flat-power pricing model, the work does examine the model at grid scale in the context of distributed energy storage and increasing levels of solar penetration. In particular, this work evaluates the pricing model on 5-minute interval data from 14,000 utility smart meters deployed in a real town. Although Mishra et al. [2013b] also focus on improving the incentives for distributed energy storage, they apply existing peak-based pricing that is sensitive to the accuracy of predictions of peak energy usage and evaluate the approach on a small-scale synthetic dataset. Since accurately predicting the time of peak energy is challenging, the approach does not work well in practice.

7. CONCLUSION

This article proposes a simple pricing scheme called *flat-power pricing* to incentivize consumers to shift small amounts of load to flatten their demand rather than shift as much of their power usage as possible to low-price, off-peak periods. We show that compared to variable rate pricing, flat-power pricing better incentivizes distributed energy storage and the use of advanced load scheduling algorithms. For example, our results show that the use of distributed energy storage under flat-power pricing requires aggregate storage capacity within 31% of an optimal centralized system. We also show that advanced load scheduling algorithms under flat-power pricing are able to save consumers up to 40% on their electric bill compared to 11% using an existing TOU rate plan.

REFERENCES

S. Barker, A. Mishra, D. Irwin, E. Cecchet, P. Shenoy, and J. Albrecht. 2012a. Smart*: An open data set and tools for enabling research in sustainable homes. In *Proceedings of the 2012 SustKDD Conference (SustKDD'12)*.

Sean Barker, Aditya Mishra, David Irwin, Prashant Shenoy, and Jeannie Albrecht. 2012b. Smartcap: Flattening peak electricity demand in smart homes. In *Proceedings of the 2012 PerCom Conference (PerCom'12)*.

H. Bindner, T. Cronin, P. Lundsager, J. Manwell, U. Abdulwahid, and I. Baring-Gould. 2005. *Lifetime Modelling of Lead Acid Batteries*. Technical Report. Riso National Laboratory, Roskilde, Denmark.

S. Braithwait, D. Hansen, and M. O'Sheasy. 2007. *Retail Electricity Pricing and Rate Design in Evolving Markets*. Technical Report. Edison Electric Institute, Washington, DC.

Kathryn R. Bullock and Alvin J. Salkind. 2011. Valve regulated lead-acid batteries. In *Linden's Handbook of Batteries* (4th ed.), D. Linden (Ed.). McGraw-Hill, New York, NY, 17.1–17.39. T. Carpenter, S. Singla, P. Azimzadeh, and S. Keshav. 2012. The impact of electricity pricing schemes on storage adoption in Ontario. In *Proceedings of the 2012 e-Energy Conference (e-Energy'12)*.

D. Cauchon. 2011. Household electricity bills skyrocket. *USAToday*. Retrieved January 9, 2017, from http://www.usatoday.com/money/industries/energy/story/2011-12-13/electric-bills/51840042/1.

CNT. 2011. Dynamic pricing and smart grid. *CNT Energy*. Retrieved January 9, 2017, from https://hourlypricing.comed.com/live-prices/.

B. Daryanian, R. Bohn, and R. Tabors. 1989. Optimal demand-side response to electricity spot prices for storage-type customers. *IEEE Transactions on Power Systems* 4, 3, 897–903.

EIA. 2012. Electric Generator Dispatch Depends on System Demand and the Relative Cost of Operation. Retrieved January 9, 2017, from http://www.eia.gov/todayinenergy/detail.cfm?id=7590.

A. Faruqui and K. Eakin. 2002. *Electricity Pricing in Transition*. Springer.

A. Faruqui, R. Hledik, and J. Tsoukalis. 2009. The power of dynamic pricing. *Electricity Journal* 22, 3, 42–56.

Federal Energy Regulatory Commission. 2009. *State of the Markets Report 2008*. Technical Report. Federal Energy Regulatory Commission.

T. Ganu, D. Seetharam, V. Arya, R. Kunnath, J. Hazra, S. Husain, L. DeSilva, and S. Kalyanaraman. 2012. nPlug: A smart plug for alleviating peak loads. In *Proceedings of the 2012 e-Energy Conference (e-Energy'12)*.

I. Goiri, W. Katsak, K. Le, T. Nguyen, and R. Bianchini. 2013. Parasol and GreenSwitch: Managing datacenters powered by renewable energy. In *Proceedings of the 2013 ASPLOS Conference (ASPLOS'13)*.

S. Govindan, A. Sivasubramaniam, and B. Urgaonkar. 2011. Benefits and limitations of tapping into stored energy for datacenters. In *Proceedings of the 2011 ASCA Conference (ASCA'11)*.

B. Howard. 2015. What the Tesla Powerwall Home Battery Means: Inexpensive Time-Shifting for Solar Energy. Retrieved January 9, 2017, from http://www.extremetech.com/extreme/204702-what-the-tesla-powerwall-home-battery-means-inexpensive-time-shifting-for-solar-energy.

Matthew P. Johnson, Amotz Bar-Noy, Ou Liu, and Yi Feng. 2011. Energy peak shaving with local storage. *Sustainable Computing: Informatics and Systems* 1, 3, 177–188. DOI:http://dx.doi.org/10.1016/j.suscom.2011.05.001

G. Karmakar, A. Kabra, and K. Ramakrithnan. 2013. Coordinated scheduling of thermostatically controlled real-time systems under peak power constraint. In *Proceedings of the 2013 RTAS Conference (RTAS'13)*.

I. Koutsopoulos, V. Hatzi, and L. Tassiulas. 2011. Optimal energy storage control policies for the smart power grid. In *Proceedings of the 2011 SmartGridComm Conference (SmartGridComm'11)*.

Z. Liu, A. Wierman, Y. Chen, B. Razon, and N. Chen. 2013. Data center demand response: Avoiding the coincident peak via workload shifting and local generation. *Performance Evaluation* 70, 10.

S. McLaughlin, P. McDaniel, and W. Aiello. 2011. Protecting consumer privacy from electric load monitoring. In *Proceedings of the 2011 CCS Conference (CCS'11)*.

Aditya Mishra, David Irwin, Prashant Shenoy, Jim Kurose, and Ting Zhu. 2012. SmartCharge: Cutting the electricity bill in smart homes with energy storage. In *Proceedings of the 2012 e-Energy Conference (e-Energy'12)*.

Aditya Mishra, David Irwin, Preshant Shenoy, Jim Kurose, and Ting Zhu. 2013a. GreenCharge: Managing renewable energy in smart buildings. *IEEE Journal on Selected Areas in Communications*, 7, 1281–1293.

Aditya Mishra, David Irwin, Prashant Shenoy, and Ting Zhu. 2013b. Scaling distributed energy storage for grid peak reduction. In *Proceedings of the 2013 e-Energy Conference (e-Energy'13)*.

A. Mishra, R. Sitaraman, D. Irwin, T. Zhu, P. Shenoy, B. Dalvi, and S. Lee. 2015. Integrating energy storage in electricity distribution networks. In *Proceedings of the 2015 e-Energy Conference (e-Energy'15)*.

Ontario Energy Board. 2012. Electricity Prices. Retrieved January 9, 2016, from http://www.ontarioenergyboard.ca/OEB/Consumers/Electricity/Electricity+Prices.

M. Roozbehani, M. A. Dahleh, and S. K. Mitter. 2010. On the stability of wholesale electricity markets under real-time pricing. In *Proceedings of the 2010 CDC Conference (CDC'10)*.

N. Sharma, J. Gummeson, D. Irwin, and P. Shenoy. 2010. Cloudy computing: Leveraging weather forecasts in energy harvesting sensor systems. In *Proceedings of the 2010 SECON Conference (SECON'10)*.

N. Sharma, P. Sharma, D. Irwin, and P. Shenoy. 2011. Predicting solar generation from weather forecasts using machine learning. In *Proceedings of the 2011 SmartGridComm Conference (SmartGridComm'11)*.

Smart Grid Information Clearinghouse. 2011. Legislation and Regulation. Retrieved January 9, 2017, from http://www.sgiclearinghouse.org/Legislation?q=node/1705.

P. Srikantha, C. Rosenberg, and S. Keshav. 2012. An analysis of peak demand reductions due to elasticity of domestic appliances. In *Proceedings of the 2012 e-Energy Conference (e-Energy'12)*.

J. Taneja, P. Dutta, and D. Culler. 2010. Towards cooperative grids: Sensor/actuator networks for promoting renewables. In *Proceedings of the 2010 SmartGridComm Conference (SmartGridComm'10)*.

J. Taneja, K. Lutz, and D. Culler. 2013. Flexible loads in future energy networks. In *Proceedings of the 2013 e-Energy Conference (e-Energy'13)*.

P. Thomson. 2015. When the Grid Says 'No' to Wind and Solar Power, This Company's Technology Helps It Say 'Yes' Again. Retrieved January 9, 2017, from http://www.pri.org/stories/2015-04-07/when-grid-says-no-wind-and-solar-power-companys-technology-helps-it-say-yes-again.

R. Urgaonkar, B. Urgaonkar, M. Neely, and A. Sivasubramaniam. 2011. Optimal power cost management using stored energy in data centers. In *Proceedings of the 2011 SIGMETRICS Conference (SIGMETRICS'11)*.

P. van de Ven, N. Hegde, L. Massoulie, and T. Salonidis. 2011. Optimal control of residential energy storage under price fluctuations. In *Proceedings of the 2011 ENERGY Conference (ENERGY'11)*.

P. Vytelingum, T. Voice, S. Ramchurn, A. Rogers, and N. Jennings. 2010. Agent-based micro-storage management for the smart grid. In *Proceedings of the 2010 AAMAS Conference (AAMAS'10)*.

Ye Xu, David Irwin, and Prashant Shenoy. 2013. Incentivizing advanced load scheduling in smart homes. In *Proceedings of the 2013 BuildSys Conference (BuildSys'13)*.

W. Yang, N. Li, Y. Qi, W. Qardaji, S. McLaughlin, and P. McDaniel. 2012. Minimizing private data disclosures in the smart grid. In *Proceedings of the 2012 CCS Conference (CCS'12)*.

T. Zhu, Z. Huang, A. Sharma, J. Su, D. Irwin, A. Mishra, D. Menasche, and P. Shenoy. 2013. Sharing renewable energy in smart microgrids. In *Proceedings of the 2013 ICCPS Conference (ICCPS'13)*.

Ting Zhu, Aditya Mishra, David Irwin, Navin Sharma, Prashant Shenoy, and Don Towsley. 2011. The case for efficient renewable energy management in smart homes. In *Proceedings of the 2011 BuildSys Conference (BuildSys'11)*.

Zivan. 2012. *Battery Discharger at Constant Current: Technical Features and User Manual*. Retrieved January 9, 2016, from http://www.zivanusa.com/pdf/SBM.

Received November 2015; revised November 2016; accepted November 2016