

Incentivizing Advanced Load Scheduling in Smart Homes

Ye Xu[†], and David Irwin[†], and Prashant Shenoy*
Department of Electrical and Computer Engineering[†]
School of Computer Science*
University of Massachusetts Amherst

ABSTRACT

In recent years, researchers have proposed numerous advanced load scheduling algorithms for smart homes with the goal of reducing the grid's peak power usage. In parallel, utilities have introduced variable rate pricing plans to incentivize residential consumers to shift more of their power usage to low-price, off-peak periods, also with the goal of reducing the grid's peak power usage. In this paper, we argue that variable rate pricing plans *do not* incentivize consumers to adopt advanced load scheduling algorithms. While beneficial to the grid, these algorithms do not significantly lower a consumer's electric bill. To address the problem, we propose *flat-power pricing*, which directly incentivizes consumers to flatten their own demand profile, rather than shift as much of their power usage as possible to low-cost, off-peak periods. Since most loads have only limited scheduling freedom, load scheduling algorithms often cannot shift much, if any, power usage to low-cost, off-peak periods, which are often many hours in the future. In contrast, flat-power pricing encourages consumers to shift power usage even over short time intervals to flatten demand. We evaluate the benefits of advanced load scheduling algorithms using flat-power pricing, showing that consumers save up to 40% on their electric bill, compared with 11% using an existing time-of-use rate plan.

Categories and Subject Descriptors

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

General Terms

Design, Measurement, Management

Keywords

Energy, Battery, Electricity, Grid

1. INTRODUCTION

Rising electricity prices over the past 20 years combined with a growing awareness of the environmental effects of burning fossil

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

BuildSys'13, November 14 - 15, 2013, Roma, Italy.

Copyright 2013 ACM ACM 978-1-4503-2431-1/13/11 ...\$15.00.

fuels, e.g., air pollution, climate disruption, water contamination, is motivating both energy producers and consumers to better optimize their electricity generation and consumption, respectively.¹ The simplest and most direct way to optimize the electric grid is for consumers to simply use less energy. Unfortunately, despite continuing improvements in energy-efficiency, society's energy demand continues to grow at a rapid pace—estimated to increase 56% by 2040 [5]—driven by both population growth and improving economic conditions in developing countries. Since reducing overall energy consumption presents many complex non-technical challenges, one promising alternative approach to grid optimization has been to focus on reducing its peak power usage.

The magnitude of the grid's peak usage has a disproportionate impact on electricity generation's capital and operational costs, as well as its carbon emissions. For instance, a lower peak usage directly translates to less idle, unused generation capacity and a need for fewer power plants. In addition, the marginal cost to generate each additional watt of electricity increases non-linearly, since utilities usually dispatch the highest-cost "peaking" generators last. In fact, the cost to generate each watt using an oil-based peaking generator can be as much as ten times the cost of using a coal-fired baseload power plant [6]. Finally, since peaking generators tend to be the least efficient, they also produce more carbon emissions per watt. These trends have led utilities to introduce variable rate electricity pricing plans that incentivize residential consumers to shift their power usage to reduce the grid's peak demand. These pricing plans vary the price of electricity throughout the day, such that electricity costs more when grid demand is high, i.e., *peak periods*, than when it is low, i.e., *off-peak periods*. Figure 1 shows an example of how electricity prices vary over a day with real time-of-use (TOU) [14] and real-time pricing (RTP) [16] plans. With RTP plans, rates change every hour of every day based on electricity's real-time price in the wholesale market, while with TOU plans, rates change only a few times each day and each day's rate profile remains constant over long periods, e.g., 3-6 months.

The goal of these new variable rate plans is to incentivize consumers to lower their electric bill by manually altering their behavior, i.e., when they perform certain energy-intensive tasks during the day. For example, if electricity prices are much lower in the evening, consumers might choose to perform energy-intensive tasks, such as doing their laundry or running their dishwasher, at that time, rather than in the middle of the day. In parallel, the benefits of reducing the grid's peak usage have motivated researchers to develop a variety of advanced load scheduling algorithms for buildings and homes that programmatically control when electrical devices (or *loads*) operate to lower a building's electricity bill, e.g.,

¹Research supported by NSF grants CNS-1253063, CNS-1143655, CNS-0916577, CNS-0855128, CNS-0834243, CNS-0845349.




Figure 1: Examples of how electricity rates vary under time-of-use (TOU) and real-time pricing (RTP).

[2, 3, 7, 9, 11, 12, 17, 18]. Instead of requiring consumers to manually alter their behavior, which many consumers may choose not to do regardless of electricity’s price, these scheduling algorithms exploit a limited degree of scheduling freedom available in a subset of loads. This freedom includes the option to transparently *shift*, *slide*, *stretch*, *store*, or *sell* power for some loads without consumer involvement. For example, a scheduling algorithm may partially i) *shift* a thermostatic or timer-driven load’s duty cycle, ii) *slide* a batched load’s start time into the future, iii) *stretch* an elastic load’s operation to reduce its peak usage, iv) *store* power in a battery to alter a load’s profile, or v) *sell* power produced by renewables back to the grid or other buildings, e.g., in a microgrid [22]. Of course, each dimension of scheduling freedom has inherent limitations. For example, a refrigerator may shift its duty cycle only so far before its interior temperature becomes too high and food spoils.

As this prior research shows, widespread adoption of these algorithms would result in significant reductions in the grid’s peak usage, which in-turn would result in a significant decrease in the cost to generate electricity. Unfortunately, in this paper, we argue that today’s variable rate pricing plans do not strongly incentivize consumers to adopt load scheduling algorithms. In particular, we show that optimizing each degree of scheduling freedom, both in isolation or collectively, within reasonable limits does not significantly reduce, e.g., only 11% with TOU pricing, consumer electric bills. We also describe some potential adverse and counterintuitive effects, including grid oscillations and higher peak load, if consumers were to adopt advanced load scheduling algorithms in large numbers under these pricing plans. Ultimately, these plans neither encourage nor support advanced load scheduling, especially at scale. Rather, their goal is to reduce peak demand by incentivizing a small fraction of consumers to manually alter their behavior.

To address the problem, we propose *flat-power pricing* to incentivize consumers to flatten their own demand. Flat-power pricing directly incentivizes load scheduling by charging a higher price for electricity that is above a consumer-specific target set *a priori* by the utility. The approach incentivizes shifting power usage over shorter intervals than variable rate pricing, which results in greater savings from using advanced load scheduling algorithms. In addition, as we discuss, since flat-power pricing encourages consumers (and load scheduling algorithms) to flatten their own demand, it avoids the adverse grid effects caused by the widespread adoption of load scheduling algorithms under today’s variable rate plans. Our hypothesis is that flat-power pricing is better at incentivizing advanced load scheduling algorithms—encouraging their adoption at large scales—than today’s variable rate pricing plans. In evaluating our hypothesis, we make the following contributions.

Combined Load Scheduling. We develop a mixed integer linear programming (MILP) formulation that combines all the degrees of scheduling freedom mentioned above to minimize electricity costs by scheduling multiple types of loads under variable rate pricing. The MILP extends prior work that focuses on optimizing each degree of scheduling freedom in isolation. For a representative home, we show that combined load scheduling is only able to lower the electric bill by 11% and 20% under an existing time-of-use (TOU) and real-time pricing (RTP) plans, respectively.

Flat-power Pricing Model. We introduce our flat-power pricing model, and discuss its benefits relative to existing pricing models, including both variable rate and peak-based pricing, with respect to incentivizing advanced load scheduling algorithms and encouraging consumers to change their behavior.

Evaluation. We alter our combined MILP formulation above to schedule loads to minimize electricity costs under our new flat-power pricing plan. For the same home as above, we show that with flat-power pricing, advanced load scheduling algorithms in a representative home are able to lower the electric bill by up to 40%—significantly more than the savings from existing plans.

2. BACKGROUND

We first introduce the different dimensions of scheduling freedom available to some electric loads, as well as highlight prior research that focuses on optimizing each particular degree of freedom in isolation, either to reduce consumer costs or the grid’s peak usage. In addition, we discuss the limitations of each metric in exploiting variable rate pricing plans to lower a consumer’s electric bill. The primary degrees of freedom we identify are the ability to transparently shift, slide, stretch, store, 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 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* [2, 18].

Prior work focuses on scheduling these shiftable on-off loads either to i) align as much load as possible with renewables [18], ii) defer loads during times of grid constraint [7], or iii) flatten a building’s demand profile [2, 9]. 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 exploit the TOU/RTP pricing models from Figure 1, 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 non-preemptible, batch task for a fixed, well-known period of time [18]. These loads are first initialized by a user, and then run for a pre-determined 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 an 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 user's 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 above, users are often unwilling to delay start times many hours into the future, e.g., from daytime to nighttime, which severely limits scheduling freedom.

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 [17]. 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 cited in prior work include washing machines, dryers, dishwashers, ovens, stoves, refrigerators, freezers, air conditioners, electric water heaters, and electric space heaters. Of course, schedulers cannot just arbitrarily stretch a load, since 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, while low heat for a long duration is better for delicate fabrics. Thus, prior work [17] places a relatively low upper limit ($\sim 10\%$) on stretching a load's running time.

Storing large amounts of power during low price periods and then using this stored power during high price periods is the most effective way to transparently schedule loads and lower a consumer's electricity bill under today's variable rate pricing plans. Prior work explores many different aspects of scheduling energy storage, largely in the form of lead-acid batteries, to exploit this opportunity for arbitrage [3, 4, 10, 11, 20, 21]. Unfortunately, energy storage, especially at small scales, is expensive. As prior work shows, either electricity prices must rise, or battery prices must fall, by an order of magnitude before energy storage's return-on-investment (ROI) would near a break-even point [3, 11]. One reason for the high ROI is that TOU/RTP pricing plans require consumers to install a large amount of storage capacity to shift as much load as possible from high-price daytime periods to low-price nighttime periods.

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 [12, 18, 23]. Aligning a home's load with renewable generation also decreases grid transmission losses, since it increases the amount of power consumed at the point of generation. The scheduling limitations above also impact the freedom to align load with renewable generation.

3. COMBINED LOAD SCHEDULING

The prior research described in §2 focuses on optimizing different dimensions of scheduling freedom in isolation. In this paper, we want to examine the performance of jointly optimizing scheduling for all types of loads, e.g., shiftable, stretchable, slidable, etc., to quantify the maximum potential benefits from advanced load scheduling. As in prior work, we formulate our combined schedul-

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 kWh
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

Table 1: Parameter definitions for linear program.

ing problem as a mixed integer linear program (MILP). In particular, we extend the MILP used in Parasol [8]: a solar-powered micro-data center that schedules batch computing jobs with well-known running times and deadlines. As with our MILP, Parasol's objective is to minimize electricity costs under TOU/RTP rate plans and maximize the use of solar power, albeit for a data center instead of a residential home. While Parasol accounts for energy storage, net metering, and slide loads (which are akin to batch computing jobs), we extend our MILP to also schedule shiftable and stretchable loads. Of course, our work differs from Parasol, since a home is substantially different than a data center in terms of its users, workloads, loads, and scheduling freedom.

3.1 Joint Optimization

Our MILP uses the parameters listed in Table 1. 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 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.

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

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

$$\sum_{i=0}^T B_{green} + \sum_{i=0}^T B_{grid} - \frac{\sum_{i=0}^T L_{battery}}{e} \leq C \quad (3)$$

$$\forall i \in T | L_{battery}^i > 0, N_{green} = 0 \quad (4)$$

$$\forall i \in \{T | L_{battery}^i > 0\}, B_{green}^i + B_{grid}^i = 0 \quad (5)$$

$$\forall i \in T, L_{green}^i + B_{green}^i + N_{green}^i \leq g_i \quad (6)$$

$$\forall i \in T, L_{grid}^i + B_{grid}^i = P_{grid}^i \quad (7)$$

$$\forall i \in \{T | L_{grid}^i > 0\}, N_{green}^i = 0 \quad (8)$$

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

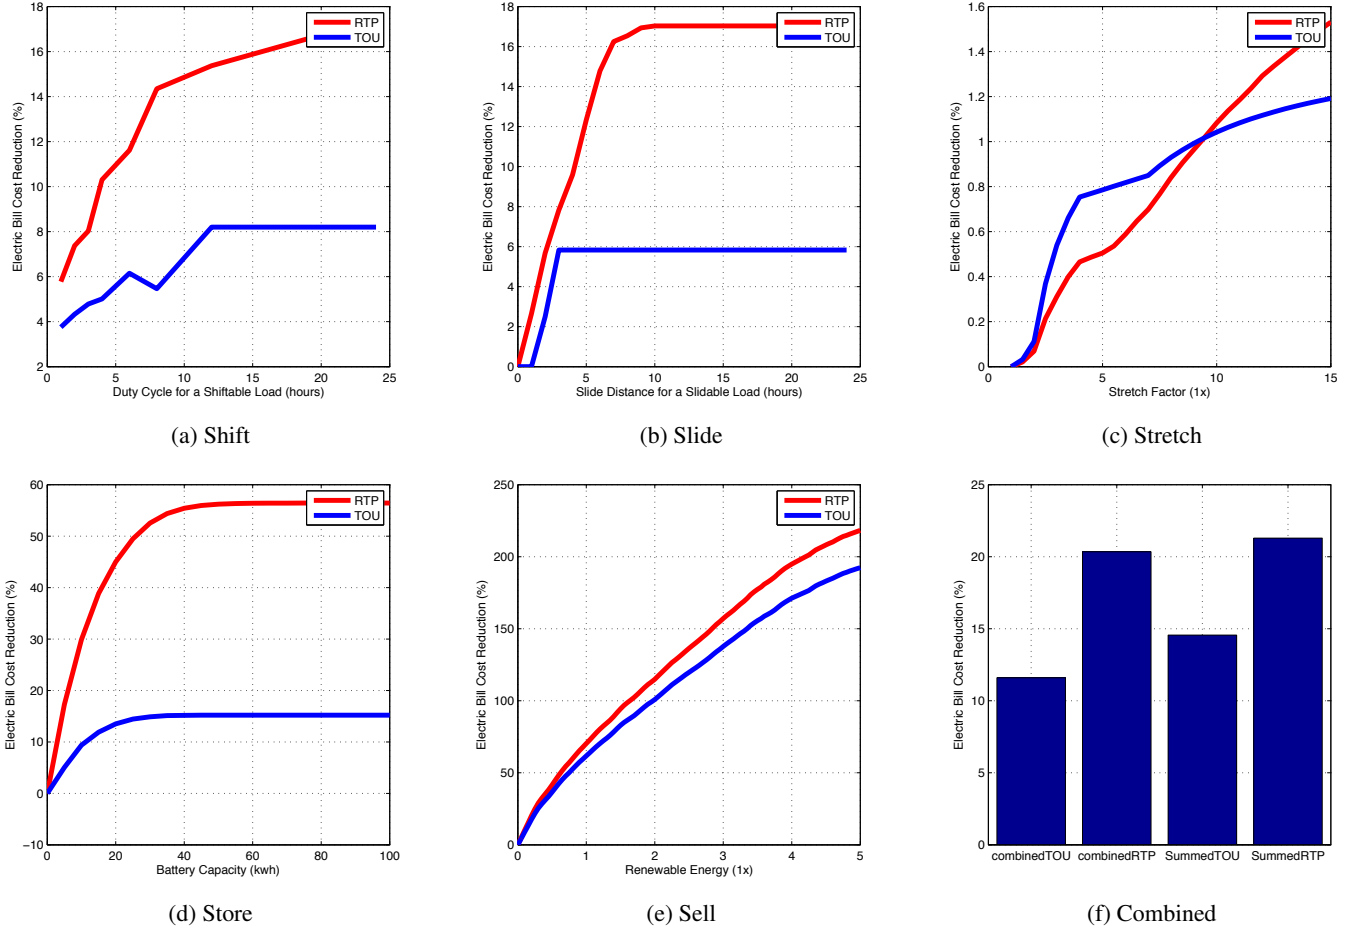


Figure 2: Cost savings from optimizing each degree of scheduling freedom (shifting, sliding, stretching, storing, and selling) for electric loads in isolation ((a)-(e)), and in combination (f) within reasonable limits.

$$\forall L_{shift}^i, \forall i \in \frac{T}{L_{shift}^i[period]}, \sum L_{shift}^i[power] = L_{shift}[power] \quad (10)$$

Briefly, the constraints ensure the following invariants: (1) the battery’s charging rate is not more than its capacity divided by 4, i.e., a $C/4$ charge rate; (2) the energy charged to the battery never exceeds the energy discharged from it; (3) the energy stored in the battery never exceeds its capacity; (4) 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 non-linear mutual exclusion constraints. We convert these to linear constraints by introducing a binary variable $b \in \{0, 1\}$ and replacing each non-linear mutual exclusion constraint with five linear constraints that enforce the 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, since 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 each day, we run the MILP for each possible stretched duration.

3.2 Scheduling’s Benefits

The purpose of our MILP 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 both the TOU and RTP pricing plans depicted in Figure 1. In this case, we use data from Home-A in the Smart* dataset [1].² Since we use the same RTP prices from Figure 1 each day, our RTP

²From July 1st, 2012 to August 30th, 2012.

plan has significantly more opportunity to lower costs by scheduling loads than a real RTP plan. Thus, our RTP plan represents a rough upper-bound on the size of the price differential between the lowest-cost and highest-cost periods: a real RTP plan would achieve 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 [11, 15]. Thus, our results also represent an upper-bound on the cost savings due to scheduling.

Figure 2 shows the extent to which optimizing each degree of scheduling freedom lowers Home-A’s electricity bill. 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 (a)-(e), 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, the battery capacity for energy storage, 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 the extreme RTP pricing. For example, as the length of the duty cycle for shiftable loads increases, the scheduler has more freedom to shift power usage long periods of time without violating the constraint that energy usage within a duty cycle must be constant. In practice, though, 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%. 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 [19]. Similarly, while 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 [17]. While 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 example, in recent work, we estimate that each kWh of usable energy storage costs \$118 per year to install and maintain [11].

Finally, Figure 2(f) 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 two-hour duty cycle for the shiftable loads, four hours for the delay time of slidable loads, and $2x$ for the stretch factor of stretchable loads, with no battery capacity or renewable energy. We believe 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 excessive inconvenience to users. We do not include any battery capacity or renewable energy due to their high upfront capital costs. The penetration rate of these technologies in current homes is low. In

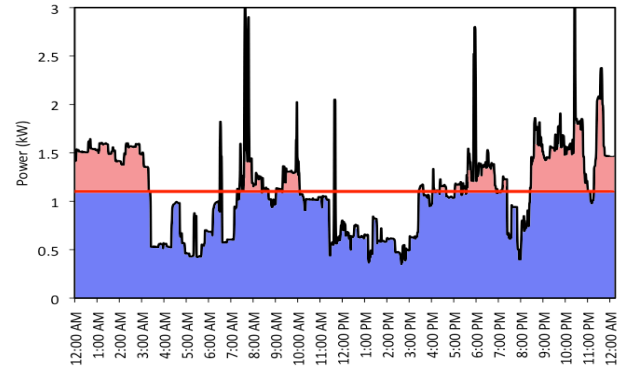


Figure 3: Flat-power pricing charges consumers β per kWh for power usage less than a consumer-specific target (in blue) and $\$(1 + \alpha)\beta$ for usage above the target (in red).

addition, many homes are not well-suited for supporting on-site renewable energy sources. For example, solar panels work best on a south-facing roof angled at 33° without any nearby tree cover.

Our results in Figure 2(f) shows that the benefit from jointly scheduling based on the constraints above is actually less than the sum of the benefits from scheduling each dimension independently, since 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 TOU case. Since many of our assumptions are conservative, the benefits are likely to be much less in practice. Note that, if we add in 10kWh of energy storage capacity and 1x of solar generation, the savings do improve to 40% and 25% in the TOU and RTP cases, respectively, without accounting for capital costs. This illustrates that TOU/RTP pricing plans incentivize capital-intensive battery-based energy storage and renewable energy sources over other forms of advanced load scheduling.

4. FLAT-POWER PRICING

The previous section’s results demonstrate that existing TOU/RTP pricing plans provide consumers little incentive to adopt advanced load scheduling. To address the problem, we propose a new type of electricity pricing for residential consumers, which we call *flat-power pricing*. Note that, while electricity’s price is set by supply and demand in the wholesale market, utilities have the freedom to offer consumers different prices. For instance, neither flat-rate pricing, which charges the same price for kWh at all times, nor TOU pricing conform to wholesale prices. Instead, these pricing plans simplify consumer pricing relative to the wholesale market.

Our flat-power pricing model charges consumers β per kWh for power usage less than a consumer-specific target, and then charges $\$(1 + \alpha)\beta$ for any usage above that target. Figure 3 depicts how flat-power pricing works, where the consumer-specific target is near 1kW, the energy usage in blue costs β per kWh, and the energy usage in red costs $\$(1 + \alpha)\beta$. Unlike variable rate pricing, flat-power pricing directly incentivizes consumer’s 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. Encouraging consumers to flatten their own demand using flat-power pricing has a number of advantages in promoting advanced load scheduling, as described below.




Figure 4: Example of how power generation from a 10kW rooftop home solar installation can vary widely minute-to-minute.

4.1 Advantages

Scalability. First and foremost the approach increases the incentive for consumers to adopt advanced load scheduling. While prior research has demonstrated the benefits of such scheduling to the grid, i.e., to reduce its generation costs and carbon emissions, it has not addressed how to incentivize consumers to adopt it. As we show in §5, flat-power pricing (for suitable values of α) provides consumers a strong monetary incentive to schedule loads even over the relatively short time intervals possible with shifting, sliding, and stretching. In contrast, §3 shows that scheduling these loads provides consumers little savings under existing TOU/RTP plans. In addition, unlike RTP plans, flat-power pricing requires only limited coordination with a utility, since we expect α and the consumer-specific target to rarely change. In contrast, consumers that participate in RTP plans must react appropriately to constantly changing prices and price forecasts, often daily via the web, e.g., <http://www.powersmartpricing.org/chart>.

Stability. One issue with existing TOU/RTP plans is that, 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, TOU/RTP 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 on there being a weak incentive to schedule loads. However, the presence of advanced load scheduling 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, e.g., akin to flash crashes. Thus, flat-power pricing is more conducive to algorithmic optimization, since 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 stated goal as TOU/RTP plans: a flat grid demand.

Utilities already use peak-based pricing plans for some, typically large, industrial consumers. Peak-based pricing, which charges consumers based on their absolute peak usage over a billing period, is similar to flat-power pricing in that it encourages consumers to flatten their own demand. However, peak-based pricing is not as conducive to algorithmic optimization, since it requires scheduling algorithms to know with high accuracy *when* the peak will occur each billing period. As part of prior work, we have found that accurately predicting when a peak will occur is challenging [13],

hindering load scheduling algorithms that must optimize for the absolute peak. Consumers are typically charged based on their peak usage over an entire month-long billing period, requiring accurate predictions of power usage over a small time window, e.g., fifteen minutes, as far as one month in advance. Further, consumers often cannot control their absolute peak, especially if it is caused by a single appliance, such as an electric dryer. In these cases, peak-based pricing would have no effect on the incentive to schedule load, since consumers would not be able to alter the peak.

Our flat-power pricing model does not suffer from these drawbacks. Instead, consumers and load scheduling algorithms need only optimize for the short-term goal of maintaining power usage below the target level α , rather than precisely plan their usage over month-long billing periods. As a result, optimizing for flat-power pricing does not require accurately predicting demand over long time horizons, and does not penalize algorithms as much as peak-based pricing if short-term predictions of future demand turn out to be inaccurate. For example, if an incorrect prediction causes a load scheduling 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.

While we do not evaluate it here, our flat-power pricing model likely also applies to large industrial consumers. As described above, flat-power pricing offers multiple advantages over surcharges based on absolute peak demand, which may also incentivize a broader range of industrial consumers, e.g., those that have little to no control over their absolute peaks, to flatten their load profile. **Simplicity.** In addition to being conducive to algorithmic optimization, flat-power pricing is also simple for consumers to understand. The primary objective of TOU/RTP plans is to incentivize consumers to manually alter their behavior by varying electricity prices. In practice, though, 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 balance of supply and demand. In contrast, flat-power pricing only requires consumers to know and respond to their consumer-specific target. In this case, a simple energy meter, such as a TED monitor, displayed in a prominent place would indicate whether or not consumers are above their target, and thus incurring higher prices.

4.2 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. We evaluate consumer savings for different values of α in §5. For the consumer-specific target, one way to choose it is based on each home's historical average power usage, e.g., over the last 5 years. Using each home's average power seems reasonable, since an optimal load scheduling algorithm would result in the home consuming its average power at all times. However, using average power to set the target does raise 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




Figure 5: Cost savings from optimizing each degree of scheduling freedom (shifting, sliding, stretching, storing, selling and in combination) under flat-power pricing within reasonable limits.

provide consumers a ranking of their electricity usage compared to their peers in their monthly bill.

One other potential point of concern is that flat-power pricing encourages over-optimizing load scheduling. Since the grid aggregates electricity demand over large numbers of consumers, scheduling loads over short time periods will not affect the aggregate load profile, since these loads are already highly multiplexed across homes within the grid. While true today, the introduction of a high percentage of intermittent solar and wind energy sources is likely to alter this dynamic and increase the benefits of scheduling over shorter time intervals. For example, Figure 4 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. In addition, future microgrids, which promise to increase grid reliability by leveraging local energy sources, will benefit much less from aggregation effects [22]. Finally, the potential benefits of incentivizing all consumers to adopt advanced load scheduling, rather than *none* of them (which is effectively the case today), likely outweighs the cost of over-optimization. However, we plan to explore and quantify this cost-benefit tradeoff in-depth as part of future work. As part of this exploration, we also plan to consider the affect of load scheduling and our incentives on a utility’s revenue. Since utilities are often profit-driven, any savings incurred by customers must be offset by a greater reduction in the utility’s operating costs to incentivize the utility to change

their pricing plan. The economics of power generation imply that a flatter demand profile will significantly reduce operational costs, but do not evaluate the extent of the reduction in this paper (and whether it exceeds the 40% savings by the consumer in §5).

5. EVALUATION

Figure 5(a)-(e) shows the savings from scheduling each type of load in isolation 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 2(f): a two-hour duty cycle for the shiftable loads, four hours for the delay time, and $2x$ for the stretch factor. For storage and selling, we use a 10kWh battery capacity and our solar energy harvesting trace scaled at $1x$. 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 5(a) shows that the savings from shifting loads is 18% for $\alpha = 3$, while in Figure 2(a) it is 4-7%. Likewise, Figure 5(b) shows that the savings from sliding loads is 13% for $\alpha = 3$, while in Figure 2(b) it is $\sim 6\%$. Stretching and storing power also show greater benefits. Flat-power pricing is particularly attractive for storage, since it takes much less storage capacity to flatten a home’s demand than to shift all of its demand to the low-price period. We explore this aspect of scheduling storage in recent work [13], although using a peak-based pricing plan based on a

home's absolute peak. As we mention, one challenge with peak-based pricing is the difficulty in predicting when the absolute peak will occur. For scheduling storage, flat-power pricing eliminates this challenge, while preserving the incentive to flatten demand, making it more conducive to algorithmic optimization. Interestingly, solar harvesting offers roughly the same benefit as with the TOU/RTP plans ($\sim 50\%$) at 1x scale (from Figure 5(a)). The reason is that with TOU/RTP plans, 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 energy generation.

Finally, Figure 5(f) shows the savings from the same combined load scheduling as in Figure 2(f) for increasing values of α . Without any energy storage or renewable energy, 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 addition, if we add in 10kWh of energy storage and 1x of solar generation, the savings are 77% for $\alpha = 3$, compared with 40% and 25% when adding energy storage and 1x 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. CONCLUSION

Researchers have proposed numerous advanced scheduling algorithms for electrical loads that, in combination, would reduce the grid's peak usage. However, as we show, there is little monetary incentive for consumers to adopt these scheduling algorithms. While utilities are introducing new variable rate pricing plans to incentivize consumers to shift their power usage to low-price, off-peak periods, most loads do not have enough scheduling freedom to shift their usage many hours into the future. Thus, we propose flat-power pricing to directly incentivize consumers to flatten their own demand, rather than shift it to low-price periods. To evaluate flat-power pricing, we formulate a MILP that optimally schedules loads with different degrees of scheduling freedom—the ability to shift, slide, stretch, store, and sell—to minimize costs under both TOU/RTP pricing and flat-power pricing. In each case, we find that flat-power pricing offers greater savings than a TOU/RTP plan for each degree of scheduling freedom both in isolation and in combination. In addition, we argue that flat-power pricing, unlike variable rate pricing, is also scalable, stable, and simple. Overall, our results suggest that flat-power pricing may be a better alternative than variable rate pricing for incentivizing consumers to adopt advanced load scheduling to reduce the grid's peak usage.

7. REFERENCES

- [1] S. Barker, A. Mishra, D. Irwin, E. Cecchet, P. Shenoy, and J. Albrecht. Smart*: An Open Data Set and Tools for Enabling Research in Sustainable Homes. In *SustKDD*, August 2012.
- [2] S. Barker, A. Mishra, D. Irwin, P. Shenoy, and J. Albrecht. SmartCap: Flattening Peak Electricity Demand in Smart Homes. In *PerCom*, March 2012.
- [3] T. Carpenter, S. Singla, P. Azimzadeh, and S. Keshav. The Impact of Electricity Pricing Schemes on Storage Adoption in Ontario. In *e-Energy*, May 2012.
- [4] B. Daryanian, R. Bohn, and R. Tabors. Optimal Demand-side Response to Electricity Spot Prices for Storage-type Customers. *TPS*, 4(3), August 1989.
- [5] International Energy Outlook 2013. Technical report, U.S. Energy Information Administration, July 2013.
- [6] State of the Markets Report 2008. Technical report, Federal Energy Regulatory Commission, August 2009.
- [7] T. Ganu, D. Seetharam, V. Arya, R. Kunnath, J. Hazra, S. Husain, L. DeSilva, and S. Kalyanaraman. nPlug: A Smart Plug for Alleviating Peak Loads. In *e-Energy*, May 2012.
- [8] I. Goiri, W. Katsak, K. Le, T. Nguyen, and R. Bianchini. Parasol and GreenSwitch: Managing Datacenters Powered by Renewable Energy. In *ASPLOS*, March 2013.
- [9] G. Karmakar, A. Kabra, and K. Ramakrishnan. Coordinated Scheduling of Thermostatically Controlled Real-time Systems under Peak Power Constraint. In *RTAS*, April 2013.
- [10] I. Koutsopoulos, V. Hatzi, and L. Tassiulas. Optimal Energy Storage Control Policies for the Smart Power Grid. In *SmartGridComm*, September 2011.
- [11] A. Mishra, D. Irwin, P. Shenoy, J. Kurose, and T. Zhu. SmartCharge: Cutting the Electricity Bill in Smart Homes with Energy Storage. In *e-Energy*, May 2012.
- [12] A. Mishra, D. Irwin, P. Shenoy, J. Kurose, and T. Zhu. GreenCharge: Managing Renewable Energy in Smart Homes. *JSAC*, 31(7), 2013.
- [13] A. Mishra, D. Irwin, P. Shenoy, and T. Zhu. Scaling Distributed Energy Storage for Grid Peak Reduction. In *e-Energy*, May 2013.
- [14] Ontario Energy Board: Electricity Prices. <http://www.ontarioenergyboard.ca/OEB/Consumers/Electricity/Electricity+Prices,2012>.
- [15] N. Sharma, J. Gummeson, D. Irwin, and P. Shenoy. Cloudy Computing: Leveraging Weather Forecasts in Energy Harvesting Sensor Systems. In *SECON*, June 2010.
- [16] Smart Grid Information Clearinghouse: Legislation and Regulation. <http://www.sgiclearinghouse.org/Legislation?q=node/1705,2011>.
- [17] P. Srikantha, C. Rosenberg, and S. Keshav. An Analysis of Peak Demand Reductions due to Elasticity of Domestic Appliances. In *e-Energy*, May 2012.
- [18] J. Taneja, P. Dutta, and D. Culler. Towards Cooperative Grids: Sensor/Actuator Networks for Promoting Renewables. In *SmartGridComm*, October 2010.
- [19] J. Taneja, K. Lutz, and D. Culler. Flexible Loads in Future Energy Networks. In *e-Energy (poster)*, May 2013.
- [20] P. van de ven, N. Hegde, L. Massoulie, and T. Salonidis. Optimal Control of Residential Energy Storage Under Price Fluctuations. In *ENERGY*, May 2011.
- [21] P. Vytelingum, T. Voice, S. Ramchurn, A. Rogers, and N. Jennings. Agent-based Micro-storage Management for the Smart Grid. In *AAMAS*, May 2010.
- [22] T. Zhu, Z. Huang, A. Sharma, J. Su, D. Irwin, A. Mishra, D. Menasche, and P. Shenoy. Sharing Renewable Energy in Smart Microgrids. In *ICCPs*, April 2013.
- [23] T. Zhu, A. Mishra, D. Irwin, N. Sharma, P. Shenoy, and D. Towsley. The Case for Efficient Renewable Energy Management in Smart Homes. In *BuildSys*, November 2011.