GridPeaks: Employing Distributed Energy Storage for Grid Peak Reduction

Yasra Chandio*, Aditya Mishra†, Anand Seetharam*
*Department of Computer Science, SUNY Binghamton, USA
†Department of Computer Science and Software Engineering, Seattle University, USA
{ychandio, aseethar}@binghamton.edu,mishraa@seattleu.edu

***Abstract*—Since peak demand dictates the costs and carbon emissions in electricity generation, electric utilities are transitioning to renewable energy to cut peaks and curtail carbon footprint. Although clean and sustainable energy source, intermittent nature of most renewables (e.g., solar, wind) makes it challenging to integrate them with the traditional electric grid. Energy storage could facilitate the integration. Grid-scale energy storage projects have been coming up across the world, but require huge upfront capital costs, and significant time and efforts. An economic and scalable alternative to expensive centralized energy storage is to leverage distributed energy storage across several homes in the grid. Prior research has proposed employing home energy storage for cutting peak demands and electricity bills for customers. However, homes working individually cannot shave aggregate peaks efficaciously since not all home peaks are aligned with the aggregate peaks and homes don't have the aggregate consumption knowledge. To address these limitations, we present GridPeaks, a distributed energy storage system that centrally controls the batteries of the participating homes from a master node deployed at the grid. The master makes battery charging-discharging decisions based on aggregate grid demand. We evaluate GridPeaks on real power consumption data. Our results show that GridPeaks can cut grid-wide peaks by 23% and reduce the generation costs by up to 14%.**

I. Introduction

Eighteen of the 19 warmest years have all occurred since 2001 [4]; 2016 is the warmest year on record. The temperature increase is chiefly driven by increased emissions of greenhouse gases (such as CO2) into the atmosphere. Electricity generation is a significant contributor to the emissions. Around 33% of the total U.S. energy-related CO2 emissions, in 2018, were from electric power sector [5]. Carbon emissions from electricity generation are determined not only by the total energy demand, but also by the magnitude of the maximum energy that needs to be generated per unit time, i.e., the peak demand (measured in kW or MW).

Peak demands dictate both the costs and carbon emissions in electricity generation. Grid's installed generation and transmission capacity needs to be provisioned for the tallest annual peak. Unfortunately, due to peaky nature of demand, lot of the installed capacity sits idle for most time. In the U.S., only about 50% of the capacity is used 100% of the time. Only 5% of the time greater than 90% generation capacity is used [1]. The base load generators that are operational continually often cannot serve peak demands. Utilities need to run extra "peaking" generators to serve the peaks. These "peaking" generators frequently operate on fossil fuels [7] hence have greater carbon emissions compared to base load generators like hydroelectric plants, nuclear plants, solar power [2]. Additionally, due to higher costs of building fuel efficient generators, peaking generators are less efficient than base load generators [7]. The inefficiency adds to carbon footprint.

To curtail electricity's carbon footprint and cut costs, states are transitioning to renewable energy sources. For instance, Hawaii plans to get 100% of its energy from renewable sources by 2045. Although a sustainable energy source, most renewables are intermittent and unpredictable. The intermittence makes them incompatible with the traditional electric grid that operates on the paradigm of "supply follows demand." One possible way of bridging this gap is energy storage. Intermittent energy can be stored in batteries when available. The stored energy can be drawn to serve demand later.

Grid-scale energy storage projects have been coming up across the world. For instance, one of the largest energy storage project, with a gigawatt power rating, is planned for Utah [8]. Such centralized storage installations do facilitate renewable energy integration in the grid and allow peak shaving, but require huge upfront costs, planning, and significant engineering time and efforts.

An economic and scalable alternative to expensive centralized energy storage is to leverage distributed energy storage across several buildings and homes in the grid. Many homes already have battery storage, for instance, homes with electric vehicles (EVs), and photovoltaic deployments; in remote areas and developing nations, where the grid is unstable, battery is used as backup during outages. As the energy storage price continues to drop and more off-the-shelf home storage solutions become available, the storage penetration in residential sector is expected to rise further [6].

Prior research has proposed employing home energy storage for cutting peak demands and electricity bills for customers [10], [11], [13], [14], [20]–[22], [25], [26]. However, this approach has limitations. Using batteries to exploit time-of-use pricing could create a taller rebound peak during low price periods [20], [22]. Employing storage to cut individual home peaks could reduce aggregate peaks (as shown in [22]), but it does not achieve best results since not all home peaks are aligned with the grid-wide aggregate peaks. Furthermore,

978-1-5386-7466-6/18 $31.00 © 2019 IEEE

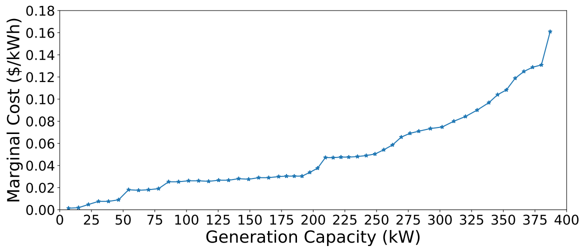

Fig. 1. The marginal cost to generate electricity increases as the demand increase

as homes do not have grid-level consumption knowledge, they cannot shave aggregate peaks efficaciously.

Therefore, to cut peaks and generation costs effectively, overcome limitations of isolated battery deployments, and reduce user electricity costs we present GridPeaks. GridPeaks is a distributed energy storage system where batteries at participating homes are centrally controlled by a master node deployed at the electric grid. Since the master has knowledge of real-time aggregate demand, it uses an online charge-discharge algorithm to control distributed storage so as to optimize the grid demand profile for peak shaving and generation cost savings.

Although distributed energy storage, operating independently, across several buildings in the grid can cut peaks and save electricity costs for the customers, it produces sub-optimal results. Our hypothesis is that letting the grid control charging-discharging of the distributed storage at buildings can significantly boost aggregate peak reduction, and generation cost savings. In evaluating our hypothesis, we make the following contributions: 1) We detail GridPeaks' architecture deployed at the participating homes. 2) We present a framework that allows the grid to control charging-discharging of the distributed energy storage from a centralized location. 3) Novel revenue sharing model. If traditional peak based pricing plans are used at homes, centralized control of storage could increase the bill. Therefore, to incentivize participation in GridPeaks, we propose a novel revenue sharing system that employs peak based variable electricity pricing at homes, but allows the grid to share its generation cost savings (from GridPeaks) with the participating homes in proportion to their storage capacity contribution. 4) We evaluate GridPeaks in simulation using power consumption data from real homes and existing variable electricity pricing plans.

II. Background

A. Electricity generation costs

Figure 1 is derived from the actual U.S. electricity generation costs data presented in [12]. The figure illustrates that the generation costs grow super-linearly as the consumption peak increases. Considering the transmission and distribution losses are proportional to the square of the of the current, increased peak demands worsen generation costs since lost electricity is wasted generation. Approximately 10 to 20% of generation costs in the U.S. are incurred servicing just the top 100 hours of peak demand per year [23]. We use the data presented in [12] to derive the generation cost function in Figure 1 for GridPeaks evaluation. The function is derived by scaling the generation cost data in the figure to match the aggregate peak demand in our micro-grid traces.

B. PeakCharge

GridPeaks leverages centrally controlled distributed battery system to shave peaks and cut generation costs. For performance evaluation, we compare GridPeaks with its distributed alternative—PeakCharge [22]. PeakCharge proposes a large scale distributed energy storage system in the grid that is individually managed by homes. In PeakCharge, every home uses an online battery charging-discharging algorithm to cut user electricity bills in presence of peak based pricing plans. The idea being, if individual home peaks are shaved, the aggregate grid-peak would be shaved consequently. To minimize peaks, PeakCharge tries to keep the home's consumption close to its target daily average consumption. If the home's demand is below the target average, PeakCharge charges the battery from the grid at a rate so as to bring up the net consumption from the grid to the target average. Similarly, if the consumption is above the target, PeakCharge draws a fraction of the demand from the battery to bring down the net draw to the target average. More details about PeakCharge are in [22].

C. Variable and peak based electricity pricing

To cut peaks and generation costs, several utilities are transitioning to market-based variable electricity pricing plans such as time-of-use (ToU) pricing. The variable pricing incentivizes customers to lower peak demand by having a higher electricity price during high demand periods (such as dinner times) and a lower price during low demand periods (e.g., late nights). Such pricing plans are in place at several locations including Ontario [24], Illinois [15]. Additionally, to curtail peaks from commercial customers, utilities impose a peak penalty on their highest consumption peak, e.g., [3]. Hence, commercial customers pay for both energy consumption ($/kWh) and every kW of their tallest monthly peak ($/kW).

D. GridPeaks' novel electricity pricing

For a fair comparison with PeakCharge, we evaluate GridPeaks using a hybrid pricing plan needed by PeakCharge, as done in [22]. The hybrid pricing bills customers for both the energy consumption ($/kWh) and the tallest peak ($/kW). For energy consumption, we use the ToU prices from Ontario Energy Board [24]. We adopt a peak surcharge of $0.60/kW based on [22].

However, the hybrid pricing plan, recommended by PeakCharge, is not ideal for GridPeaks since GridPeaks is focused on shaving the aggregate peaks and not all individual home peaks are aligned with the aggregate peaks. Therefore, to incentivize participation in GridPeaks, we propose a novel revenue sharing system that employs peak based variable electricity pricing at homes, but allows the grid to share its generation cost savings (from GridPeaks) with the participating homes in proportion to their storage capacity contribution. We propose a savings split of 20%-80% between the grid and

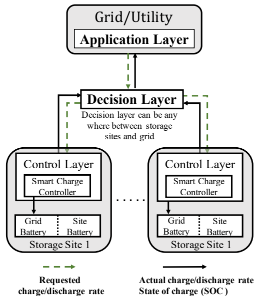

Fig. 2. GridPeaks system architecture

the homes (respectively) since the homes own and maintain the batteries. The proposed split is similar to the split used by transportation network companies, like Lyft and Uber, that do not own the vehicles but provide the enabling technologies.

E. GridPeaks architecture

Figure 2 depicts the GridPeak's system architecture as deployed at the participating homes. The architecture is armed with a charge controller and a discharge controller (such as in [27]) that allow GridPeaks to programmatically control the charge and discharge rate of the battery, respectively. The gateway server continually monitors the home's aggregate consumption via an in-panel energy monitor, and the battery's state of charge. The gateway server periodically sends the monitored data to the central master node. Based on the data received from the participating homes and the aggregate consumption data from the grid generation units, the master node computes and sends the desired charge/discharge rate to each participating home. Section IV-A presents the details of the charging, discharging algorithms executed at the master node. The home gateway server uses the charge-discharge directives from the master node to accordingly charge/discharge the home battery.

GridPeaks controls the battery charge-discharge at the participating homes from a central master node deployed at the grid. The homes are connected to the master via a high speed network link such as high speed Internet connection. The network is used to relay data and the charge-discharge directives. The master node also receives the real-time aggregate demand information from the grid generation units. As long as the master receives the required data in timely fashion (over the network), its exact deployment location is flexible. One possible location could be the bulk power substation.

III. Problem Statement

Although energy storage at individual homes can cut peaks and electricity bills, the choice of when and how much to charge-discharge the battery presents interesting tradeoffs. Allowing the homes to independently make battery charging-discharging decisions could minimize each home's peak demand and reduce its electricity bill, but may not shave the aggregate peak—since individual home peaks do not always sync with the grid-wide peak. If we instead allow the grid to control charging-discharging of the batteries deployed at homes, it could cut the aggregate peaks effectively. However, as not all the peaks occurring at an individual home align with the grid-peak, individual homes may not obtain the best peak reduction. If the customers are billed for their tallest peaks, a sub-optimal peak reduction could increase their bill, hence is undesirable.

We define the problem of employing distributed energy storage for grid-peak reduction as follows. Given the aggregate power consumption of all the homes in the grid, and electricity pricing plan the problem is to design an online distributed energy storage charging-discharging algorithm to cut the participating home's electricity bills and shave grid-wide aggregate peak demand. The algorithm should determine: 1) the homes whose batteries should be charged or discharged at a given time; 2) for battery discharge, the fraction of the home's demand to be consumed from the grid and the fraction to be drawn from the energy storage; 3) for battery charge, the amount of energy to be fed into the battery so as to shave the aggregate demand peaks and cut generation costs while reducing electricity bill for the homes.

IV. GridPeaks Framework and Algorithm

GridPeaks hides the details of distributed energy storage and instead presents an abstraction of unified central battery to the grid. The utility grid can leverage GridPeaks to any desired end including (but not limited to) peak reduction, generation cost saving, frequency regulation, storing excess renewable energy, satisfying excess demand in lieu of dispatching "peaking" generators. The flexibility inherent in battery storage presents diverse possibilities. However, in this paper, we focus on applying GridPeaks for peak shaving and generation cost savings. One way of achieving those goals is to minimize load variations on utility generators by keeping the net demand close to a target daily average value. In this paper, we adopt the aforementioned approach to the desired goals. In Section IV-A, we present the GridPeaks framework along with the details of translating charging-discharging decisions on a centralized battery abstraction to the actual distributed storage system. Next, Section IV-B talks about applying the framework for curtailing peaks and generation costs.

A. Storage Sharing Framework

In this section, we describe the GridPeaks framework that allows the utility provider to seamlessly control the distributed energy storage. Let us assume that there are n homes and the battery capacity at home h is B_h. The total battery capacity available to the grid, B_G, across n sites is $B_G = \sum_{h=1}^{n} B_h$. We assume a controller, that is deployed at the central master node, helps the utility provider to effectively manage the distributed energy storage. Our framework is inspired from prior work [18], in which the authors design a virtualization layer to enable multiple homes to share underlying physical resources (i.e., battery and solar array). The objective of our framework is to hide all the complexities involved in controlling the distributed set of batteries and present the utility an intuitive interface that resembles that of a single battery. Therefore, from an operational perspective, the controller views the distributed set of batteries as single battery of total capacity B_G and takes aggregate charging/discharging decisions based on that.

The controller keeps track of the the individual battery capacities, their state of charge (i.e., *SOC*), and how much to charge/discharge individual batteries. The battery capacity available to the grid at each site is fetched at the start and is updated only when the battery capacity at the site changes. The controller periodically communicates with the smart home gateway servers at the participating homes to set the battery charge-discharge rates and fetch the state of charge of the individual batteries at each site. We denote the current state of charge for the aggregate battery as $SOC_G(t)$, which is essentially computed by adding the available battery capacities at the participating homes. Let us assume that $C_G(t)$ and $D_G(t)$ denote the desired aggregate charge and discharge rates, respectively, at time t on the aggregate battery B_G. The $C_G(t)$ and $D_G(t)$ are determined by our grid-level peak shaving algorithm, described in Section IV-B.The controller takes into account the state of charge $SOC_h(t)$ for home h at time t and the allowed depth of discharge DOD_h for home h, to determines how the grid-level aggregate charging/discharging decisions are handled at the individual homes. While doing so, the controller also determines if the aggregate charge/discharge rates can be achieved by the distributed storage and returns the actual charge rate $C_G^{'}(t)$ and discharge rate $D_G^{'}(t)$, which may or may not be the same as $C_G(t)$, $D_G(t)$ depending on the combined state of charge and remaining capacity of the distributed storage.

Table I provides an overview of the variables. We next describe the battery charging and discharging decisions undertaken at the individual battery level.

Charging: Algorithm 1 describes the algorithm for charging the batteries. The charge operation takes the aggregated charge requirement at time t, $C_G(t)$, and the set of all participating homes (**H**) as input (in line 1) and determines the charging actions to be undertaken at the individual home level. The magnitude of $C_G(t)$ is determined by the grid-level peakshaving algorithm, as described in Section IV-B. Line 2 gets the set of homes with full battery. These homes need to be excluded from further charging considerations since they have no more room in their battery. Line 3 computes the total energy that the grid would store in next time step "dt" if it could charge at the rate of $C_G(t)$. It is obtained by multiplying the charging rate with the interval length. The while loop in line 4 continues until the desired aggregate charge rate is distributed among participating homes or all the homes have run out of room to charge. Initial value of $C_h(t)$ for all homes is assumed zero. In line 5, E_C/dt gives us the charge rate required to provide energy E_C in time dt. We divide this required rate by the number of homes in the set $H - HF$, i.e., $|H - HF|$, which means that the charge rate is equally divided among all homes whose batteries are not full. Whether a given home would be able to sustain the computed charge rate depends on its available battery capacity. Therefore, in line 6, we compute the remaining battery capacity at home h. The nested for loop in line 7 iterates through every home in set $H - HF$. Line 8 determines whether the remaining capacity in the battery is less than required by the computed charge rate in line 5. If so, in line 9, we compute the charge rate the home can sustain at its current remaining capacity for the next time step and modify the assigned charge rate accordingly.In line 10, we add this home to the set of homes with full battery (**HF**), since charging at the modified rate (computed in line 9) the home would be full at the end of the next time step. Line 11 updates E_C by subtracting the energy that will be fed into home h.

Variable	Description	Unit
n	Number of sites	-
dt	Time step	hour
$\mathbf{H}$	Set of sites	-
$\mathbf{HF}$	Set of sites with full battery	-
$\mathbf{HE}$	Set of sites with empty battery	-
k	Fraction of battery shared with the grid	-
$P_D(t)$	Power demand at time t	kW
$P_D^{'}(t)$	Power demand at time t after GridPeaks	kW
P_D^{limit}	Demand limit, typically average demand	kW
B_h	Battery for site h	kWh
B_h^G	Battery for grid from site h	kWh
B^G	Battery available to the grid	kWh
l	B_h/l defines the maximum charge rate	-
$SOC_h(t)$	State of charge for battery of site h at time t	-
DOD_h	Depth of discharge for battery of site h	-
$SOC_G(t)$	State of charge for grid battery at site t	-
$C_h(t)$	Charging rate of battery for site h at time t	kW
$D_h(t)$	Discharging rate of battery for site h at time t	kW
$C_G(t)$	Charging rate required by grid at time t	kW
$D_G(t)$	Discharge rate required by grid at time t	kW
$C_G^{'}(t)$	Charging rate achieved by battery at time t	kW
$D_G^{'}(t)$	Discharge rate achieved by battery at time t	kW
$C_G^{max}(t)$	Maximum charging rate available to grid	kW
$D_G^{max}(t)$	Maximum discharging rate available to grid	kW
$E_D(t)$	Energy discharged between time t and $t + dt$	kWh
$E_C(t)$	Energy charged between time t and $t + dt$	kWh
$E_h^{full}(t)$	Energy required to fully charge the battery at site h at time t	kWh
$E_h^{empty}(t)$	Energy available till empty battery at site h at time t	kWh
$E_h(t)$	Energy stored in battery at site h at time t	kWh

TABLE I
VARIABLE DEFINITIONS

The charge distribution process (inside the while loop in line 4) keeps repeating until either we have assigned all the desired energy to homes, or several batteries are full such that the algorithm detects there isn't enough spare capacity to charge the desired amount. In the latter case, the desired charge rate by the grid ($C_G(t)$) cannot be achieved. When the loops terminate, in line 12, the algorithm computes $C_G^{'}(t)$, the sum of the charge rates assigned to all the batteries. If there was enough spare capacity, $C_G^{'}(t) = C_G(t)$. Otherwise $C_G^{'}(t) < C_G(t)$.

Discharging: Algorithm 2 describes the algorithm for discharging the batteries. The discharge algorithm takes the desired aggregate discharge rate at time t, $D_G(t)$, and the set of all participating homes (H) as the input. The magnitude of $D_G(t)$ is determined by the grid-level peakshaving algorithm, as described in Section IV-B.

Algorithm 1 Grid Battery Charging

1: **function** CHARGE($C_G(t), \mathbf{H}$)
2: $\mathbf{HF} = \{h : h \in \mathbf{H} \wedge SOC_h(t) = 1\}$
3: $E_C = C_G(t) \times dt$
4: **while** ($E_C > 0 \wedge \{\mathbf{H} - \mathbf{HF}\} \neq \emptyset$) **do**
5: $C_h(t) = \frac{E_C}{|\mathbf{H}-\mathbf{HF}|\times dt} + C_h(t), \quad \forall h \in \{\mathbf{H} - \mathbf{HF}\}$
6: $E_h^{full} = B_h^G \times (1 - SOC_h(t)), \quad \forall h \in \{\mathbf{H} - \mathbf{HF}\}$
7: **for** ($h \in \{\mathbf{H} - \mathbf{HF}\}$) **do**
8: **if** ($C_h(t) \times dt \geq E_h^{full}$) **then**
9: $C_h(t) = E_h^{full}/dt$
10: $\mathbf{HF} = \mathbf{HF} + \{h\}$
11: $E_C = E_C - C_h(t) \times dt$
12: $C'_G(t) = \sum_{h=1}^{n} C_h(t)$
13: return $C'_G(t)$

Algorithm 2 Grid Battery Discharging

1: **function** DISCHARGE($D_G(t), \mathbf{H}$)
2: $\mathbf{HE} = \{h : h \in \mathbf{H} \wedge SOC_h = DOD_h\}$
3: $E_D = D_G(t) \times dt$
4: **while** ($E_D > 0 \wedge \{\mathbf{H} - \mathbf{HE}\} \neq \emptyset$) **do**
5: $D_h(t) = \frac{E_D}{(|\mathbf{H}-\mathbf{HE}|\times dt)} + D_h(t), \ \forall h \in \{\mathbf{H} - \mathbf{HE}\}$
6: $E_h^{empty} = B_h^G \times (SOC_h(t) - DOD_h)$
7: **for** ($h \in \{\mathbf{H} - \mathbf{HE}\}$) **do**
8: **if** ($D_h(t) \times dt \geq E_h^{empty}$) **then**
9: $D_h(t) = E_h^{empty}/dt$
10: $\mathbf{HE} = \mathbf{HE} + \{h\}$
11: $E_D = E_D - D_h(t) \times dt$
12: $D'_G(t) = \sum_{h=1}^{n} D_h(t)$
13: return $D'_G(t)$

Line 2 gets the set of homes with empty battery. These homes need to be excluded from further consideration since they have no stored charge. Line 3 computes the total energy the grid would discharge if it could discharge at the desired rate of $D_G(t)$ during the next time step "dt." It is obtained by multiplying the discharge rate with the interval length. The while loop in line 4 runs until the desired discharge rate has been distributed among the participating homes or all the homes have run out of remaining charge. Initial value of $D_h(t)$ for all homes is zero. In line 5, E_D/dt gives us the discharge rate required to discharge total energy E_D in time "dt." We divide this required rate by the number of homes in the set $H - HE$, i.e., $|H - HE|$, which means that the discharge rate is equally divided among all the homes with non-empty batteries. Whether a given home would be able to sustain the computed discharge rate depends on the remaining charge in its battery. Therefore, in line 6, we computed the remaining charge stored at home h. The nested for loop in line 7 iterates through every home in the set $H - HE$. Line 8 determines whether the remaining charge in the battery is is less than required by the computed discharge rate in line 5. If so, in line 9, we compute the discharge rate that the home can sustain at its current remaining charge for the next time step and modify the remaining discharge rate accordingly. In line 10, we add this home to the set of homes with empty battery (HE), since discharging at the modified rate (computed in line 9) the home would be empty at the end of the next time step. Line 11 updates E_D by subtracting the energy that would be drawn out of home h.

Algorithm 3 Battery State of Charge

1: **function** SOC()
2: $E_h(t) = B_h^G \times SOC_h(t), \qquad \forall h \in \{\mathbf{H}\}$
3: $SOC_G(t) = E_h(t)/B^G$
4: return $SOC_G(t)$

The discharge rate distribution process inside the while loop (line 4) iterates until either all the desired discharge energy ($D_G(t)$) has been assigned to the participating homes, or several batteries are empty such that the algorithm detects there isn't enough available capacity to discharge the desired amount. In the latter case, the desired discharge rate $D_G(t)$ cannot be achieved. After the loop terminates (in line 12) the algorithm computes $D'_G(t)$, the sum of discharge rates assigned to all the batteries. If there was enough available charge in the batteries, $D'_G(t) = D_G(t)$. Else, $D'_G(t) < D_G(t)$.

Battery State of Charge: The framework also provides a state of charge function, presented in Algorithm 3, that computes and returns the aggregate state of charge (SOC) across the distributed batteries. The function first checks the state of charge for every participating battery and computes the total energy stored in those batteries. The SOC for the aggregate grid-wide storage is computed by dividing the total stored energy by the total capacity available to the grid. The SOC for the battery is in range 0 to 1.

B. GridPeaks Algorithm

In this section, we describe how an electric utility can leverage the GridPeaks framework for peak reduction. Algorithm 4 provides an overview of the algorithm. The utility can use the GridPeaks abstraction to achieve any desired target net demand on its generation units. Here we assume, to minimize load variations on its generators and cut generation expenses, the utility strives to keep the net demand close to a target daily average value. If the current aggregate demand is less than the target average, the utility requests GridPeaks to charge the distributed battery system at a rate $C_G(t)$ that is equal to the difference between the target average and the current demand. The GridPeaks' computed charge rate ($C_G(t)$) is capped by the maximum charge rate for the entire battery system, the algorithm picks the smaller of the two values. The maximum charge rate limit is simply the sum of maximum charge rates of all the batteries in the distributed storage system. The maximum charge rate at each site h is computed as B_h/l, where l is an integer that is defined by the battery specification and generally takes value between 4 and 8. After computing $C_G(t)$, the algorithm calls the $CHARGE$ function (defined in Algorithm 1), which returns the actual charge rate $C'_G(t)$ that can be achieved at the time. The final net aggregate demand on utility generators, $P_D^{'}(t)$, is simply the sum of original aggregate demand $P_D(t)$ and the achieved charge rate $C_G^{'}(t)$.

On the other hand, if the current aggregate demand is greater than the target average demand, GridPeaks discharges the distributed storage system at a rate $D_G(t)$, which is equal to the difference between current total demand and target average demand. Again, the algorithm caps the computed $D_G(t)$ value with the maximum discharge rate for the battery. The maximum discharge rate limit is simply the sum of maximum discharge rates of all the participating batteries. The maximum discharge rate at each site h is typically set at a rate that will discharge the battery in an hour; for instance a 12 kWh battery can be discharged in an hour by drawing power at 12 kW–the max discharge rate. After computing $D_G(t)$, the algorithm calls the $DISCHARGE$ function (defined in Algorithm 2), which returns the actual discharge rate $D_G^{'}(t)$ that can be achieved at the time. The final net aggregate demand on utility generators $P_D^{'}(t)$, in this case, is simply the sum of the original demand $P_D(t)$ and the achieved discharge rate $D_G^{'}(t)$.

Algorithm 4 GridPeaks

Require: $P_D^{limit}, P_D(t), \mathbf{H}$

if $(P_D(t) < P_D^{limit})$ **then**

  $C_G(t) = min(P_D^{limit} - P_D(t), C_G^{max})$

  $C_G^{'}(t) = CHARGE(C_G(t), \mathbf{H})$

  $P_D^{'}(t) = P_D(t) + C_G^{'}(t)$

else

  $D_G(t) = min(P_D(t) - P_D^{limit}, D_G^{max})$

  $D_G^{'}(t) = DISCHARGE(D_G(t), \mathbf{H})$

  $P_D^{'}(t) = P_D(t) - D_G^{'}(t)$

V. Evaluation

In this section, we evaluate GridPeaks using real-world power consumption data from 114 homes over a week (February 8 to February 14, 2016) [9]. We compare GridPeaks with PeakCharge [22], the closest distributed energy storage system. Our results show that GridPeaks outperforms PeakCharge in aggregate peak reduction and generation cost savings.

Experimental parameters used, unless stated otherwise, are as follows. For every home, battery capacity C is 20% of its daily average consumption, which is in line with the findings of [22], where authors showed that a battery with usable capacity greater than 20% of the daily average consumption was sufficient to get maximum peak reduction at individual homes. The maximum battery charging rate in our experiments is C/4, i.e., the battery charges to full capacity in 4 hours, which translates to roughly a C/8 rate for a battery used at 45% depth-of-discharge (DOD). The C/4 maximum charging rate is within the maximum possible charge rate of C/3 for sealed lead-acid batteries [19]. We use a peak demand surcharge of $0.60/kW on the highest 30-minute peak. Peaks are computed at 30 minute granularity since several utilities charge customers for 30 minute peaks, e.g., [3]. For both the PeakCharge and GridPeaks, we use a hybrid pricing plan, similar to the one used in [22], where the customers are charged for both total energy consumption in the billing cycle ($/kWh component) and tallest peak ($/kW component); the energy consumption is billed using ToU pricing similar to the one employed by Ontario Energy Board [24], which offers time-of-use electricity pricing based on three usage periodsoff-peak (6.5/kWh), mid-peak (9.5/kWh), and on-peak (13.2/kWh).

A. Qualitative Results

Figure 3(a) shows the aggregate grid demand on a particular day. Peak demand for this day is 287 kW. Figures 3(b) and 3(c) show the grid-level consumption after employing PeakCharge and GridPeaks (at participating homes) respectively. In both cases, aggregate battery capacity used is 1056 kWh. Figure 3(b) shows that PeakCharge successfully shaves the initial peaks, but is unable to shave the day's tallest peak. However, GridPeaks shaves all the peaks. This happens because PeakCharge is designed to shave individual home peaks without knowledge of the grid-wide consumption. Not all peaks at homes are aligned with the aggregate peak. PeakCharge uses the stored energy in shaving the initial home peaks and runs out of charge for the tallest grid-wide peak. In contrast, GridPeaks has the knowledge of aggregate consumption. It only discharges the energy needed for shaving the grid-wide peak to a target average value, which may not require shaving all the individual home peaks. Hence, GridPeaks uses much less energy for peak reduction and it has the energy required for curtailing all of the day's peaks. In summary, the figures demonstrate that better grid-level peak reduction (and generation cost savings since peaks drive generation costs) can be achieved if the distributed energy storage is centrally controlled, as done in GridPeaks.

B. Generation Cost Savings and Peak Reduction

In Figures 4 and 5, we investigate the aggregate peak reduction and generation cost savings across the grid, respectively, as a function of battery capacity. The capacity at homes goes from 0% to 100% of the home's average daily consumption. We observe that generation cost savings and peak reduction increase sharply and then flatten out when the battery capacity reaches 20%; this happens because a homes total consumption is fixed. Adding more storage after a certain point cannot reduce the peak further. The figures also show that GridPeaks can provide significantly higher peak reduction and generation cost savings in comparison to the PeakCharge. PeakCharge reduces the peaks up to 10%, whereas GridPeaks can achieve up to 23% reduction—a boost of more than 100%. Similarly, PeakCharge results in around 7% generation cost savings, where GridPeaks achieves 14%. Figure 6 shows the actual dollar generation costs (derived from the data in Figure 1) in our experimental microgrid for raw demand, GridPeaks, and PeakCharge.

Figures 7 and 8 show the peak reduction and generation cost savings, respectively, as a function of fraction of homes participating in GridPeaks. Both GridPeaks and PeakCharge, for a storage penetration of less than 5% do not yield any

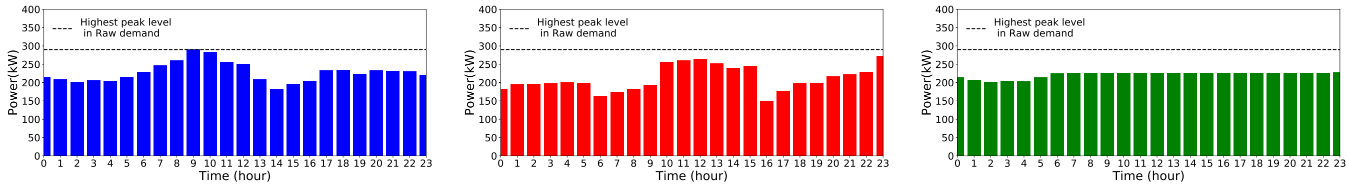

(a) Raw Demand (b) After PeakCharge (c) After GridPeaks

Fig. 3. Average hourly power consumption of a representative day

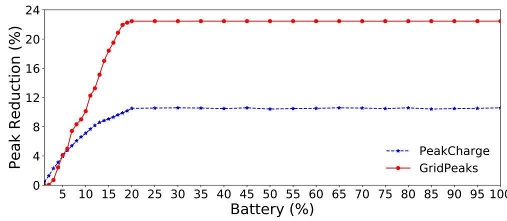

Fig. 4. Peak reduction vs. battery capacity

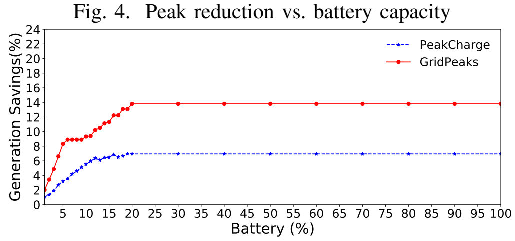

Fig. 5. Generation cost savings vs. battery capacity

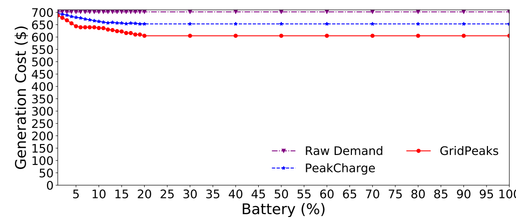

Fig. 6. Generation dollar saving vs Battery Percentage

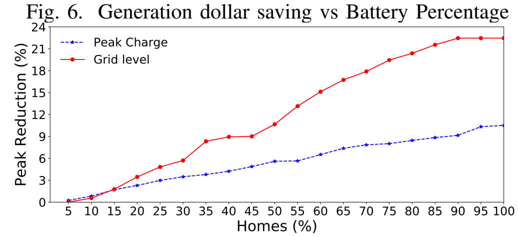

Fig. 7. Peak reduction vs. percentage of homes with battery.

peak reduction or generation savings. This is primarily due to the small storage capacity across the grid; since both algorithms operate in online fashion, the stored energy is used up in shaving smaller peaks earlier in the day. However, as the adoption increases, both PeakCharge and GridPeaks steadily cut the aggregate peaks and generation costs. Here too, GridPeaks outperforms PeakCharge. For instance, GridPeaks achieves generation cost savings of up to 14%, as opposed to 7% from PeakCharge. GridPeaks' better performance can be attributed to its knowledge of aggregate demand and targeting only aggregate peaks. It does not shave every individual peak at homes hence saves energy for cutting more peaks over the day and can achieve better reduction on tall peaks. Since generation costs are driven by peak demand, GridPeaks does well here too.

C. Benefit to Customers

Figure 9 shows the peak reduction and cost savings in energy bills for individual homes. Figure 9(a) shows that PeakCharge achieves significantly greater peak reduction than GridPeaks at individual homes. This is expected since PeakCharge is optimized to shave all possible home-level peaks. However, GridPeaks optimizes aggregate grid-level peak reduction. It doesn't necessarily shave home peaks that are not aligned with grid-level peaks. Figure 9(b) shows the dollar savings per day for the customers. The figure shows PeakCharge can save up to \$0.7/day, whereas GridPeaks saves up to \$0.20/day. On including a 20%-80% generation savings split between the grid and homes, GridPeaks' savings go up to \$0.55/day. For fair comparison, we have employed a peak-based pricing plan for evaluation at the homes, which is optimized for PeakCharge ([22]). Therefore, PeakCharge's better performance is not surprising. Besides, cost savings from GridPeaks are underestimated in these results. We only evaluated generation cost savings here. However, the utility grid can additionally leverage GridPeaks for frequency regulation, transmission cost savings, etc. Given significant unexplored cost saving potential of GridPeaks, we believe, overall GridPeaks will outperform PeakCharge, especially under an appropriate pricing plan (not optimized for PeakCharge).

VI. Related Work

Most of the research on grid peak reduction has focused on consumer-centric approaches that use consumer energy storage in response to the variable electricity pricing to either save on the bill or reduce home-level peak. In [20], authors present a machine learning based approach that forecasts demand and controls the charging of energy storage to reduce the electricity bill. In [22], authors purpose a peak-based surcharge and present a peak-aware charging algorithm to optimize the use of energy storage in the presence of a peak demand surcharge. In [11], authors present an approach that employs energy storage to intelligently utilize the renewable energy by moving a fraction of harvested energy to peak intervals. We argue that the using energy storage to reduce home-level peaks is myopic and present an approach that centrally controls the distributed storage to shave the grid-level peaks. Our approach beats the aggregate peak reduction presented by all the prior work. Another body of work focuses on devising optimal online strategies for controlling energy storage for scenarios where differential and peak-based pricing is present. In [17], authors present optimization framework for finding practical operating strategy for peak-demand pricing when there is price volatility present. In [16], authors developed an adaptive controller framework for battery systems based on neural

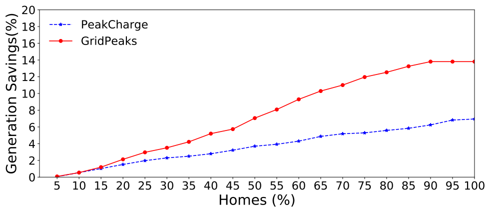

Fig. 8. Generation cost savings vs. percentage of homes with battery.

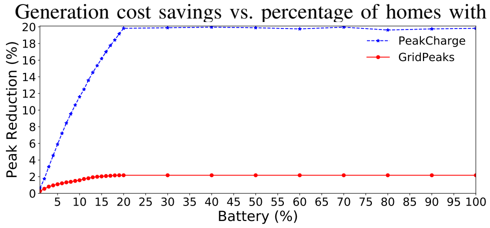

(a) Percentage peak reduction

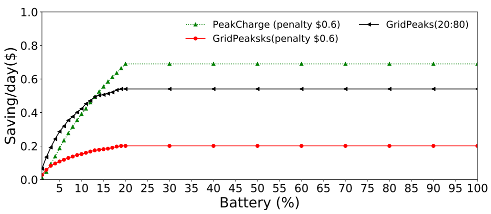

(b) Dollar Savings

Fig. 9. Customer gains vs. variation in battery capacity

networks, model predictive control, and system simulation. They demonstrated that their controller approaches MPC-level performance at a fraction of the online computational cost of MPC on a PV-battery system with time-of-use pricing. While this ensemble of work computes the cost savings, it does not quantify the peak reduction at home or grid level. On contrary, our work quantifies the impact giving the control of storage to the grid when real-world prices are present.

VII. Conclusion

We presented GridPeaks. In GridPeaks, a central master node runs an online algorithm to determine charge-discharge rate of batteries deployed at the participating homes so as to shave the demand peaks and cut generation costs. GridPeaks uses a novel revenue sharing model that allows the grid to reward participating customers by proportionately disbursing generation cost savings. Our experiments show that GridPeaks can cut grid-wide peaks by 23% and reduce generation costs up to 14%. Individual homes can save more than $0.5/day. GridPeaks lends itself to several vital grid applications such as frequency regulation, storing excess renewable energy, cutting transmission costs, etc. In this paper we only explored the monetary benefits from generation cost savings. Employing GridPeaks for other applications could further boost savings for both the grid and the participating customers. We would like to explore these possibilities in the future work.

References

[1] Alternative Fuels Data Center: Electricity Production and Distribution. https://afdc.energy.gov/fuels/electricity_production.html.

[2] Carbon Footprint of Electricity Generation. https://www.parliament.uk/documents/post/postpn268.pdf, October 2006.

[4] Global Temperature — Vital Signs Climate Change: Vital Signs of the Planet. https://climate.nasa.gov/vital-signs/global-temperature/, June 2019.

[3] Large Industrial Contract Schedule. http://www.hged.com/customers/rates/electric-rates/199E%20Lg%20Industrial%20Contr%20Sched.pdf, 2017.

[5] How much of U.S. carbon dioxide emissions are associated with electricity generation? - FAQ - U.S. Energy Information Administration (EIA). https://www.eia.gov/tools/faqs/faq.php?id=77&t=11, May 2019.

[6] Ion age: why the future will be battery powered. https://www.theguardian.com/news/2019/jan/14/on-the-charge-why-batteries-are-the-future-of-clean-energy, 2019.

[7] Peaking power plant. https://en.wikipedia.org/wiki/Peaking_power_plant, Jan 2019.

[8] 'world's Biggest' Energy Storage Project Planned For Utah. https://www.forbes.com/sites/johnparnell/2019/05/30/worlds-biggest-energy-storage-project-planned-for-utah/, May 2019.

[9] 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 *ACM SustKDD*, 2012.

[10] T. Carpenter, S. Singla, P. Azimzadeh, and S. Keshav. The impact of electricity pricing schemes on storage adoption in ontario. In *2012 Third International Conference on Future Systems: Where Energy, Computing and Communication Meet (e-Energy)*, pages 1–10. IEEE, 2012.

[11] R. L. de Pontes, A. Mishra, A. Seetharam, M. Shekhar, and A. Ramesh. Greenpeaks: Employing renewables to effectively cut load in electric grids. In *2018 IEEE International Conference on Smart Computing(SMARTCOMP)*, pages 123–130. IEEE, 2018.

[12] State of the Markets Report 2008. Technical report, Federal Energy Regulatory Commission, August 2009.

[13] S. Govindan, A. Sivasubramaniam, and B. Urgaonkar. Benefits and limitations of tapping into stored energy for datacenters. In *2011 38th Annual International Symposium on Computer Architecture (ISCA)*, pages 341–351. IEEE, 2011.

[14] N. Hegde, L. Massoulié, T. Salonidis, et al. Optimal control of residential energy storage under price fluctuations. In *Energy*, 2011.

[15] Residential Real Time Pricing Programs. https://www.pluginillinois.org/realtime.aspx, 2019.

[16] F. Kazhamiaka, S. Keshav, and C. Rosenberg. Adaptive battery control with neural networks. In *Proceedings of the Tenth ACM International Conference on Future Energy Systems*, e-Energy '19, pages 536–543, New York, NY, USA, 2019. ACM.

[17] F. Kazhamiaka, C. Rosenberg, and S. Keshav. Practical strategies for storage operation in energy systems: Design and evaluation. *IEEE Transactions on Sustainable Energy*, 7(4):1602–1610, Oct 2016.

[18] S. Lee, P. Shenoy, K. Ramamritham, and D. Irwin. vSolar: Virtualizing Community Solar and Storage for Energy Sharing. In *e-Energy*, 2018.

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

[20] A. Mishra, D. Irwin, P. Shenoy, J. Kurose, and T. Zhu. Smartcharge: Cutting the Electricity Bill in Smart Homes with Energy Storage. In *ACM e-Energy*, 2012.

[21] A. Mishra, D. Irwin, P. Shenoy, J. Kurose, and T. Zhu. GreenCharge: Managing Renewable Energy in Smart Buildings. *Selected Areas in Communications, IEEE Journal on*, 2013.

[22] A. Mishra, D. Irwin, P. Shenoy, and T. Zhu. Scaling distributed energy storage for grid peak reduction. In *Proceedings of the Fourth International Conference on Future Energy Systems*, e-Energy '13, pages 3–14, New York, NY, USA, 2013. ACM.

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

[24] Ontario Energy Board: Electricity Prices. http://www.ontarioenergyboard.ca/OEB/Consumers, 2019.

[25] R. Urgaonkar, B. Urgaonkar, M. J. Neely, and A. Sivasubramaniam. Optimal power cost management using stored energy in data centers. In *Proceedings of the ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems*, pages 221–232.ACM, 2011.

[26] P. Vytelingum, T. D. Voice, S. D. Ramchurn, A. Rogers, and N. R. Jennings. Agent-based micro-storage management for the smart grid. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1-Volume 1*, pages 39–46. International Foundation for Autonomous Agents and Multiagent Systems, 2010.

[27] Zivan. Battery Discharge at Constant Current: Technical Features and User Manual. http://www.zivanusa.com/pdf/SBM%20Gb.pdf, 2012.