VPeak: Exploiting Volunteer Energy Resources for Flexible Peak Shaving

Phuthipong Bovornkeeratiroj, John Wamburu, David Irwin, and Prashant Shenoy
{phuthipong,jwamburu,shenoy}@cs.umass.edu,irwin@ecs.umass.edu
University of Massachusetts Amherst

ABSTRACT

Traditionally, utility companies have employed demand response for large loads or deployed centralized energy storage to alleviate the effects of peak demand on the grid. The advent of Internet of Things (IoT) and the proliferation of networked energy devices have opened up new opportunities for coordinated control of smaller residential loads at large scales to achieve similar benefits. In this paper, we present VPeak, an approach that uses residential loads volunteered by their owners for coordinated control by a utility for grid optimizations. Since the use of volunteer resources comes with hard limits on how frequently they can be used by a remote utility, we present machine learning techniques for carefully selecting which days to operate these loads based on expected peak demand. VPeak uses a distributed and heterogeneous pool of volunteer loads to implement flexible peak shaving that can either selectively target hotspots within the distribution network or perform grid-wide peak shaving. Our results show that VPeak is able to shave up to 26% of the total demand when selectively shaving peaks at local hotspots and up to 46.7% of the demand for grid-wide peak shaving.

CCS CONCEPTS

• **Applied computing** → *Forecasting*; • **Hardware** → **Smart grid**.

KEYWORDS

Smart Grid, Peak Prediction, Peak Reduction, Virtual Power Plants, Volunteer Resources, Energy Storage, Electric Vehicle, HVAC

ACM Reference Format:
Phuthipong Bovornkeeratiroj, John Wamburu, David Irwin, and Prashant Shenoy. 2021. VPeak: Exploiting Volunteer Energy Resources for Flexible Peak Shaving. In *The 8th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys '21), November 17–18, 2021, Coimbra, Portugal.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3486611.3486667

1 INTRODUCTION

The advent of networked devices, also known as the Internet of Things (IoT), has resulted in smart and connected devices such as smart thermostats, smart switches, and smart lighting becoming commonplace. This trend has continued, and today's smart buildings consist of new types of electrical loads such as residential energy storage batteries, smart washing machines, and smart electric vehicle chargers that have networked capabilities, allowing remote control via programmatic interfaces.

As IoT-enabled smart electrical devices become more commonplace in smart buildings, they provide increased convenience to their users by allowing remote monitoring and control, often through smartphones. At the same time, they also open up new opportunities for energy optimizations from the grid's perspective. One particular paradigm that is emerging involves home owners *volunteering* their large loads for utility-driven control. In this case, the utility has *limited* ability to directly control a user's volunteered load and can remotely coordinate a number of such devices or loads spread across a large number of customer homes. Such a pool of volunteered energy resources can be viewed as a virtual power plant, and their coordinated control can be used by the grid to perform energy optimizations such as peak demand reduction. Indeed, rudimentary forms of such schemes have already emerged. In Europe, the concept of citizen energy communities (CECs) has emerged to encourage individual citizens to participate in the energy system for energy sharing, cooperative ownership, and control of energy resources [7]. Volunteer-driven approaches are also emerging in the United States. The Nest Smart Thermostat, from Google, for instance, offers a program called *Nest Rush Hour Savings*, where home owners agree to let their electric utility control their HVAC system for a limited number of days in each season *in return for monetary payments* on their energy bill [3]. Similarly, as deployment of residential energy storage batteries such as Tesla Powerwall becomes commonplace in residential solar installations, states such as California and Massachusetts in the U.S. have begun offering monetary incentives to home owners for the ability to *remotely control* the discharging of their batteries on peak demand days [1]. Similarly, residential electric vehicle (EV) chargers with grid control capabilities have been developed [20].

Utility control of volunteered residential energy resources can be viewed as a form of demand response for peak load shaving. However, the volunteer nature of energy resources introduces new research challenges. First, there is a decoupling of *ownership* and *control* of volunteer resources, with ownership remaining with the home owner and control being split between the user and the utility. Since volunteered resources belong to end users (i.e., home owners) and not to the utility, inconvenience to users should be avoided or minimized. If users get inconvenienced due to "aggressive" utility control of their devices (e.g., if excessive deferral of heating or cooling causes discomfort, or if aggressive deferral of EV charging yields cars with little or no charge), they will be reluctant to participate in such schemes despite the monetary incentives.

BuildSys '21, November 17–18, 2021, Coimbra, Portugal

ACM ISBN 978-1-4503-9114-6/21/11...$15.00
https://doi.org/10.1145/3486611.3486667

Second, utility control of volunteered resources often comes with hard limits on the number of times the resources can be used by the utility each month or over a season. Thus, their capabilities must be judiciously leveraged only when absolutely necessary; incorrect or overaggressive use may hamper the utility's ability to perform optimization later on in a month or a season when a peak occurs.

Third, in the future, the pool of volunteered resources from smart buildings will be heterogeneous—comprising disparate devices such as the HVAC system, EVs, energy storage batteries, and more. Since each device has different inherent characteristics and has a different impact on the user, their characteristics should be considered to determine how to intelligently use them to optimize energy reduction and user-perceived inconvenience. These challenges require the design of new approaches for control and management of volunteered energy resources for VPP-enabled energy optimization. Current approaches such as Nest's Rush Hour savings or volunteer battery programs are homogeneous, with only one type of resource being controlled in each program. Applying such approaches to a heterogeneous pool of resources has not been addressed by prior work.

To address these challenges, we present VPeak, a new approach that uses volunteer resources for flexible peak shaving. In designing and evaluating VPeak, our paper makes the following contributions.

- We introduce the notion of *volunteer energy resources* and formulate the problem of coordinated control of a heterogeneous pool of volunteer energy resources and its use for grid energy optimizations despite limited control ability by the utility.
- We present a machine learning-based algorithm for peak day selection that judiciously selects on which days to exercise control of volunteer resources. A characteristic of our algorithm is its ability to handle prediction errors that are inevitable when predicting future peak demand.
- We present resource selection algorithms that intelligently select energy resources for coordinated utility control of volunteered resources. A key contribution of our approach is its ability to perform *selective shaving* of local hot-spots that occur in the distribution grid as well as *grid-wide shaving* of peaks. Our algorithm employs a flexible optimization method that iteratively uses an ordered set of resources to perform peak demand reduction.
- We evaluate the efficacy of our algorithms using real-world energy traces from 15,242 homes. Our results show that VPeak's peak day classification outperforms the state of the art approaches by up to 20%. Further, VPeak is able to shave up to 26% of the total demand when selectively shaving peaks at local hotspots and up to 46.7% of the demand for grid-wide peak shaving.

2 BACKGROUND

In this section, we provide background on the use of volunteered energy resources for grid energy optimization.

2.1 The Case for Volunteer Energy Resources

The use of volunteer energy resources in the grid is analogous to the concept of *volunteer computing* in the field of distributed computing. In volunteer computing, users volunteer *idle* computing cycles on their PC for use by a third-party [4]. The network of volunteered computers forms a large distributed cluster over the Internet that is used to run large compute jobs whenever any volunteered PCs are idle. Volunteers donate the CPU cycles on their machine (either for free or for a small monetary incentive) for projects such as SETI@Home or cancer research [22]. We discuss how such a volunteer model can be adopted by the energy system and the grid.

Traditionally, grid utility companies have used techniques such as demand response to reduce demand during peak load periods. Conventional demand response has typically been limited to large commercial and industrial customers who can yield a "significant" reduction in the overall grid demand by reducing their own energy consumption upon being instructed to do so [9, 23]. Demand response customers receive monetary compensation for participating in demand response energy reduction events. Residential customers have typically not been allowed to participate in such demand response schemes since they were assumed to be too "small" individually to provide a meaningful reduction in grid's aggregate demand. Further, traditional demand response approaches require manual action by customers to reduce their usage, which was assumed to be too cumbersome to implement across large numbers of "small" residential customers. While automated demand response approaches for residential settings have been studied by researchers [5, 17], they have not seen widespread adoption by utility companies.

The emergence of low-cost embedded networking technologies and the Internet of Things (IoT) has resulted in many residential loads gaining networked capabilities. These IoT devices and loads expose networked interfaces (APIs), often through cloud servers, for remote monitoring and control by their users (e.g. through smartphone applications). The networked capabilities, while originally designed for user convenience, also create new opportunities for grid energy optimizations.

What is Volunteer Energy Management? In volunteer energy management, residential users can volunteer one or more of their networked loads to their electric utility company. The utility company gains the ability to remotely control any such volunteered load for performing grid optimizations. Users may choose to voluntarily participate in such schemes either for the social good or for the monetary incentives offered by such approaches.

Why is it useful? These emergence of IoT-based loads allows utility companies to offer energy optimization programs (i.e., demand response) to their *residential* customers. Specifically, these networked capabilities can enable coordinated control at scale. A utility can now control and coordinate thousands or tens of thousands of such loads and use them to perform significant energy optimizations in the aggregate. Consequently, even though a single residential customer may be a "small" energy consumer, when aggregated at larger scales, these consumers and their loads still provide opportunities for significant energy reductions. Moreover, automation of such coordination via networked APIs eliminates the need to manually coordinate optimization requests across large number of residential customers.

What types of loads can be volunteered? Many types of large residential loads with networked control capabilities are amenable for volunteer energy management. These include HVAC systems (via smart thermostats), residential energy storage batteries, EV chargers for electric cars, and smart appliances such as washing machines.

Optimization	Resources	Ownership	Control (who, when)
Traditional DR	Large loads	Commercial customer	Consumers, DR signal
Grid peak shaving	Grid Storage	Utility	Utility, during peak
Customer shaving	Large loads	Commercial customer	Customer, Price-driven
Volunteer resources	Networked loads	Residential customer	Utility, during peak

Table 1: Decoupled ownership and control of volunteer energy resources

	Control	Season
Battery	Discharge	All
HVAC	Change setpoint	summer/winter
EV	Reduce charge rate	All

Table 2: Types of resource control decisions

Through the ability to defer HVAC system loads or appliances, discharge batteries, or reduce EV charging rates across large number of users in a coordinated fashion, the utility can implement automated methods for peak demand reduction in residential settings. As noted in Section 1, there are already volunteer programs that allow residential customers to enroll their HVAC systems (when managed by smart thermostats such as Nest) or residential battery systems for remote utility control [1, 3]. Currently, each such program is offered separately for a single type of residential load (e.g. smart thermostat-based HVAC). Our work envisions a future where each user is offered a unified volunteer program to offer any subset of their networked loads for utility control, and the utility uses this heterogeneous collection of volunteered resources as a Virtual Power Plant to perform coordinated and automated energy optimizations at grid scale. Our current work considers a heterogeneous mix of three residential loads: HVAC, EV charging, and energy storage batteries, which offer the greatest promise for energy savings. However, our approach is general and can be extended to a more diverse mix of volunteer loads in the future.

2.2 Differences from Traditional Approaches

We now discuss key differences between prior approaches and the use of volunteer energy resources (See Table 1).

In traditional demand response, the end customers (typically commercial ones) make control decision (often manually, but sometimes in automated fashion) of which loads to turn off or defer. They receive monetary incentives to participate in demand response schemes and are contractually obligated to reduce their usage whenever requested by the utility. Such methods have been used for decades for large energy customers [9, 14, 23]. Peak load reduction or peak shaving is a related form of energy optimization. Such techniques come in different forms. In utility-driven peak load reduction, the utility deploys resources such as energy storage batteries in its distribution grid and operates them during peak demand hours [12, 19]. An alternative approach is customer-driven peak load reduction. In this approach, the customer is incentivized through pricing signals to reduce usage during peak hours. For example, time-of-use (ToU) pricing or demand charges that add peak surcharges can incentivize users to reduce peak usage.

In all of the above approaches, *the ownership and control of energy resources lie with the same entity*—either the utility or the end customer—which allows complete control or flexibility in how they are used (See Table 1). In our work, the ownership and control of energy resources are *decoupled*. The energy loads being volunteered are owned by the residential user. While they still retain control over their devices, they cede limited control of the resources to the utility. They are incentivized for *monetary or social* reasons to volunteer resources, and the utility is guaranteed control over these resources but for a limited number of days and limited duration for each event. The decoupling of ownership and control of volunteered energy resources for grid energy optimization raises new challenges not addressed by prior work in demand response or peak load reduction.

First, the utility is allowed to use volunteered resources only for a small number of days every month or season. Consequently, it must carefully predict and choose the n days that are the peak demand days of the month or season and ensure that no resource is used more than k times across those n days (k is the limit on the number of times a volunteer resource can be used). If the choice is made too aggressively, it may use up its quota of n days before the month ends and not have any room to use such resources if a peak day occurs later in the month. Conversely, if it is too conservative, it may miss a peak day while "saving" some of its n slots for the future.

Second, since resources are volunteered, the utility must use them judiciously—any control decisions that cause inconvenience to users will make them reluctant to participate in such volunteer-based approaches. For instance, deferring of EV charging decisions should still ensure the car is charged fully by the next morning before the user leaves for work. Similarly, deferring HVAC usage on a very hot day should not cause excessive discomfort (e.g. which can be avoided using higher setpoints or by pre-cooling [25]).

Third, since volunteered energy resources come in different types (HVAC, EVs, batteries, smart appliances), the user may have different preferences in how they should be controlled by a utility or may have priorities on which ones they are willing to defer first. Further, the type of control decision to reduce the energy usage from each type of energy load is itself different (see Table 2). Hence, utility control decisions should consider user preferences as well as the heterogeneous nature of these loads.

2.3 Problem Statement

Consider a distribution grid comprising h homes served by a utility. Each home is provided an option of volunteering its remotely controllable loads for energy optimizations under utility control. While in theory, any load with remote control capabilities can be volunteered, in this work, we consider the three significant loads, namely the HVAC system, EV chargers, and energy storage batteries, since these "dominant" loads can provide the most significant demand reduction for any individual home. Not all three loads may be present in a home (e.g., a home may lack a battery or EV), and even if present, a home owner may not volunteer them for utility control. We assume that each home volunteers any arbitrary subset of these three loads, when present, for utility control. Each volunteered load can be used by the utility no more than k times each month and only during the peak demand hours of a day. (e.g. during the peak evening hours). So long as no volunteered load is used more than k times a month, it is possible for the utility to implement any policy that chooses varying subsets on different peak days based on its needs.

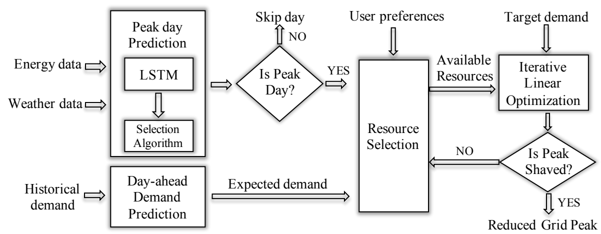

Figure 1: Overview of our VPeak approach.

Our work considers two types of peak reduction methods: *selective* and *grid-wide.* Due to the highly distributed nature of volunteer resources, the utility can implement a precise and selective peak shaving approach where the peak load seen by only the most heavily loaded transformers is shaved, which can be done by activating volunteer resources in homes connected to those transformers. Alternatively, the utility can also choose to shave the aggregate peak seen by the entire distribution grid. In either case, the utility's goal is to reduce the aggregate energy demand across on the top-n peak demand days of each month. To minimize the uncertainty and inconvenience of users, we assume that n is typically small (e.g. n=3 or 5 days each month). We assume that volunteered loads expose network APIs that allow a utility to make control decision as shown in Table 2 to reduce demand.

3 VPEAK DESIGN OVERVIEW

In this section, we present an overview of our Volunteer Peak shaving (VPeak) approach, which is depicted in Figure 1. As shown, our VPeak approach for utility control of volunteer energy resources for peak shaving consists of four components: (1) Peak day selection algorithm, (2) Peak hour prediction, (3) Intelligent resource selection, and (4) Iterative optimization to iteratively reduce peak demand.

VPeak's peak day selection algorithm is responsible for determining whether the next day is estimated to be a peak demand day where some, or all, of the volunteer resources should be activated. Our approach for peak day selection, discussed in §4, is based on machine learning and is designed to mitigate the impact of prediction errors that are inevitable when selecting future days as peak days. In the event the following day is chosen for peak demand shaving, VPeak's next component determines the hours of the day when the peak demand occurs and the hourly demand during that period (§4.3). Our hourly peak demand predictor is also based on machine learning and used a modified version of the approach in [24]. Next, VPeak's resource selection algorithm (§5) determines which volunteer resources to use and in what order to reduce the peak demand, until the desired level of peak reduction is achieved. VPeak supports two types of demand reduction—selective shaving, where demand reduction is performed only at selected distribution transformers that see local hotspots, and grid-wide shaving, where demand reduction is performed at all transformers. We note that prior work has only studied the shaving of grid-wide peaks, whereas VPeak is also capable of performing *highly targeted* peak reductions In either case, VPeak's resource selection algorithm offers multiple policies that determine the order in which heterogeneous volunteer resources should be used until the desired reduction is achieved. These policies use an iterative optimization algorithm (§5.2) to optimize the reduction in demand while respecting various constraints. We have implemented a prototype of VPeak based on the techniques in §4 and 5, which is available on GitHub: https://github.com/umassos/buildsys21-vpeak.

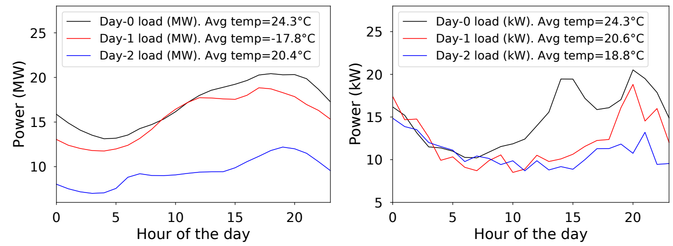

(a) **(b)**

Figure 2: (2a) Power load experienced by the whole grid on three different days, and (2b) power load experienced by a transformer on three different days.

4 PEAK DAY SELECTION

In this section, we begin with an intuition behind our peak day selection approach, and then present the details of our machine learning-based algorithm.

4.1 Basic Approach

The utility's goal is to shave the peak demand on up to n days each month, either selectively in overloaded parts of the distribution grid (e.g., highly loaded transformers), or broadly for the distribution grid as a whole. Typically the decision of whether the following day will be selected as a peak day is done at the end of the current day based on an estimate of how "tall" the next day's peak is likely to be.

The magnitude of the daily peak, whether at an individual transformer or the entire grid, depends on many factors. The primary factor governing the peak demand is weather, and specifically, the next day temperature. To illustrate the correlation between temperature and peak demand, consider the hourly demand seen by a cluster of 14 homes fed by a single transformer as well as those seen by a larger group of 15,000 homes in a small city in the Northeast United States shown in Figure 2. Over the course of a single day, the demand exhibits a well known evening peak with lower demand in night hours. While this overall behavior repeats each day, the *height* of the observed evening peak varies from one day to another. Figures 2a and 2b show the peaks on several different days of the same month, and as can be seen, the peak demand is higher when the temperature is higher. Since the magnitude of peak demand will vary from one day to another, the utility needs an algorithm that, at the end of each day, performs day-ahead calculation to determine if the next day will be one of the n highest peak days of the month.

While day-ahead demand prediction (which involves predicting a time series of next day demand) is widely studied in the literature [11, 13, 21, 26], day-ahead peak prediction is less well studied. Peak prediction to select the top-n peak days of the year was studied for Ontario, Canada using a probabilistic algorithm [14]. Our peak day selection problem is analogous, but is performed at monthly granularity and is based on machine learning. We quantitatively compare our algorithm to that in [14] in Section 6. A different problem of predicting the peak hours of each day was studied in [24]—the problem is different from peak day selection since [24]

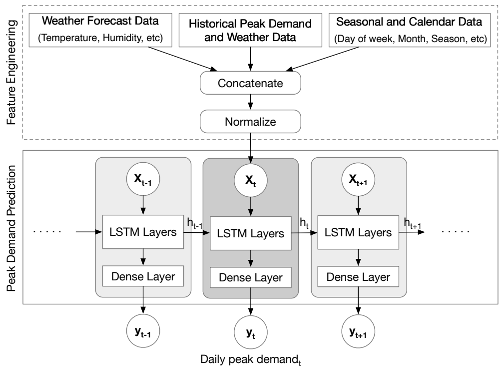

Figure 3: LSTM based Peak Demand Prediction Model.

focuses on predicting hourly peaks during a single day, rather than n days from a month or year with the greatest demand.

Mitigating prediction errors. The peak day selection problem is based on future predictions of peak demand and requires a selection algorithm to determines if the next day's peak is sufficiently high to qualify as a top-n day. Since future demand exhibits stochastic fluctuations as a result of human behavior, any predictions of whether the next day will be a peak day will inherently suffer from prediction errors, yielding both false positives and false negatives. For example, if multiple days of a month see extreme temperatures that are similar to each other, there will be several days with similar demand making it challenging to distinguish an actual peak day from a non-peak day with similar, but slightly lower demand. Small stochastic variations in the predicted peak will also result in prediction errors. In our case, false negatives are more harmful—missing a peak day entirely by not selecting it implies the grid sees the "full" impact of the peak demand, and can have a monetary impact (e.g. higher demand charges), and stress on grid resources.

To mitigate the impact of such prediction errors, our approach selects up to $(n + \Delta)$ peak days in each month, rather than exactly n days, where Δ is a small integer. The premise of our approach is that the top-$(n + \Delta)$ peak days in each month are *more likely* to contain the actual top-n days than if one were to select exactly n days, thereby reducing prediction errors. For instance, if the utility wishes to shave the peak demand on the top-3 days, it select the top-5 days for shaving (here $\Delta = 2$) which makes it more likely that chosen five days contain the top 3 days. In contrast, if exactly three days were selected, prediction errors may cause one or more of the top 3 days to not get selected. Prior work (e.g., [14]) has not considered such explicit methods to mitigate prediction errors.

As shown in Figure 1, our peak day selection algorithm has two components: an LSTM-based machine learning model to predict the next day peak, and a selection algorithm that uses this predicted peak to classify the following day as a peak or non-peak day.

4.2 Machine learning-based Peak day Selection

4.2.1 LSTM Model. VPeak uses a Long Short Term Memory (LSTM) model, which is a Recurrent Neural Network, to forecast the daily peak demand of the next day. A key advantage of using LSTM is its ability to encode both long and short term dependencies. VPeak utilizes the stacked LSTM model architecture. Stacking multiple LSTM layers allows the model to learn more complex tasks which is suitable for peak demand prediction that has many features. The LSTM model, which we depict in Figure 3, takes as input (i) weather forecast for the next day, (ii) observed historical demand over the past week, (iii) observed weather over the past week, and (iv) seasonal information such as month, date and day of the week. It uses this information to predict the peak demand that will be seen during the following day. We note that the LSTM model can be trained either for an individual transformer (to predict peak demand at overloaded transformers for selective peak shaving) or for the distribution grid as a whole (to predict the grid-wide peak demand).

Feature selection: The LSTM model uses three categories of features for its peak demand prediction. The first category is weather forecast data for the next day, which includes temperature, humidity, pressure, wind speed, cloud cover, dew point, and precipitation probability. Of these, temperature and humidity are the most important features, with the minimum, maximum and average forecast values for the next day. The second category is historical demand and weather features. In this case, the model uses the observed peak demand for the past seven days as well as the observed weather parameters for the past week. Historical weather data and weather forecasts are obtained from the open DarkSky API . The final category of features is seasonal and calendar data, including month, day of the week, week number, season, and whether the next day is a weekend or holiday. In all, the LSTM model uses 151 features.

Model training and tuning Our LSTM model comprises two layers with 32 and 16 neurons respectively. We use four years of training data to train the model and use one year for validation. We use k-fold cross validation over this five year dataset to train the model. A grid search is used to find the best hyperparameters (dropout rate, learning rate, bias, regularization penalty, and activation function). As a result, Adam with a learning rate of 0.005 and decay rate of 0.001 is used as an optimizer. Dropout rate of 11.5% is used during the training.

4.2.2 Peak day Classification. Given the estimated peak demand for the next day, VPeak's peak day classifier then determines if the next day will be one of the n days with the highest peak. Recall that VPeak attempts to reduce prediction errors by selecting the top n' days, where $n' = n + \Delta$, rather than choosing exactly n days. Typically, Δ is a small integer such as 1, 2 or 3.

Two factors govern whether the next day will be classified as a top-n' peak day. First, the classifier compares the next day's peak to historical peak demand for this month that was seen in past years to determine if it is likely to be a top-n' day. Second, the classifier compares the next day's peak to the peaks seen thus far for the month. Both factors together yield an overall classification of the day as peak or non-peak day.

To compare the next day peak with historical values, VPeak constructs a distribution of past peak demand for each month of the year. A high percentile of this distribution is then chosen as a cutoff threshold—if the next day demand is above this threshold, it is classified as a peak day. Intuitively, this is analogous to sorting the historical peak demand seen for each of the 30 (or 31) days the month in decreasing order and choosing the n'-th value as the cutoff threshold; any value above this threshold is likely to end up as top-n' day. In line with this intuition, VPeak chooses $(1 - \frac{n'}{D}) * 100$ as the percentile that is the cutoff threshold for the historical probability distribution of peak demand for the month; D denotes the number of days in the month. For example, if $n' = 6$, we choose $(1 - \frac{6}{30}) * 100$ or the 80-th percentile as the cutoff threshold T.

In addition to comparing the next day peak to this threshold T, VPeak also compares this peak to the time series of daily peak demand seen thus far for the month. In particular, the threshold T is adjusted dynamically based on the number of peak days selected thus far compared with the expected number of peak days that should have been picked based on previous years data. Every five days, the number of peak days already picked is compared with the expected number. If the actual number is higher, the threshold T is increased to make the classification more conservative. If it is lower, the threshold T is decreased to make the classification more aggressive at selecting peak days. Then the adjustment is made every five days. Finally, the time series of daily peak demand is also used to eliminate clear false positives, such as when the next day peak is lower than peaks for n' other days seen thus far—in other words, the next day peak should be at least in the top-n' days seen thus far for the month, in addition to being above the threshold T.

4.3 Peak hour Prediction

If the next day is classified as a top-n' peak day, VPeak needs to determine which hours of the day will see peak demand and estimate the hourly demand during those peak hours. We refer to this problem as peak hour demand prediction. To do so, we use a modified version of the technique from [24], which uses an LSTM model to predict hourly demand during peak hours and tune the hyperparameters to make it perform well with our dataset.

Feature selection: The features used for peak hour demand prediction can be grouped into three categories: weather forecast data, seasonal data, and historical demand. The model uses 13 features per hour: temperature, humidity, apparent temperature, cloud cover, precipitation intensity, dew point, precipitation probability, day of the week, weekday, season, holiday, hour of the day, and historical demand. We use only two historical days that have the most impact on our predictions, which are the previous day and the same day from last week. For the predicted day, there are 12 features since its hourly demand is what we want to predict. These result in a total of 912 features as an input for the peak hour prediction model. All features are normalized before used.

Model training and tuning: Since the number of features for peak hour prediction is larger than peak day prediction, we expand our model to three layers: 128, 96, and 72 hidden nodes in each layer. We then add a dense layer to reduce it to 24 nodes to match with 24 hours output. Similar to the peak prediction model, we use grid search for hyperparameter tuning and Adam with adaptive learning rate with dropout during the training. The learning rate, decay rate, and dropout rate are set to 0.0006, 0.0005, and 0.4 respectively.

5 RESOURCE SELECTION FOR PEAK REDUCTION

This section describes VPeak's resource selection and optimization algorithm.

5.1 Selective and Grid-wide Peak Shaving

Once the next day is chosen to be a top-n' peak day, the goal of resource selection is two-fold: (i) choose a *subset* of volunteer resources to use for the next day, (ii) determine an *ordering* in which these resources should be considered until the desired level of demand reduction is achieved. Since each resource can be used only k times per month, it follows that on any selected peak day, no more than $\frac{k}{n+\Delta}$ fraction of the total volunteer resources can be used. Further, if a resource has already been used k times in the current month, it can no longer be considered.

For selective peak shaving, the resource selection algorithm uses a configurable utilization level U to determine all transformers that will see a peak utilization that exceeds U. It then only choses those transformers for peak shaving, and does so by only selecting volunteer resources connected to those transformers for the optimization algorithm. Specifically, it chooses $\frac{k}{n+\Delta}$ fraction of the resources at those transformers. For grid-wide peak shaving, all distribution transformers are chosen and a $\frac{k}{n+\Delta}$ fraction of resources at all transformers is chosen as the subset to be used. Unless specified, all experiments for selective and grid-wide peak shaving use a target reduction level of 50%

Next, an ordering is determined on the chosen subset using one of the following policies.

Battery-first: The battery-first policy is designed to *minimize user inconvenience.* From the subset of chosen resources, it first selects all available energy storage batteries across homes as the highest priority group. Note that when batteries are programmed to discharge during peak hours, they will absorb the demand and users do not see any deferred loads. In the event that battery capacity is not adequate to shave the peak, the policy then chooses EV loads as its next priority group, following by HVAC loads in the last priority group).

User Priority-based: The priority-based policy is designed to *respect user-specified preferences*. It assumes that each user assigns an order to their volunteered resources that indicates the preference for which resources should be used first before the next one is used. In this case, the selection policy selects the highest priority resource from all chosen homes and groups them in the first priority group. It then chooses the next resource listed by each user, where available, and groups them in the next priority group and so on. The policy allows each user to specify a different order based on personal preferences.

5.2 Optimization Problem

The goal of our iterative optimization algorithm is to iteratively choose each group of volunteer resources, in the specified priority order, until the desired target level of demand reduction is achieved (with the objective of using the least amount of volunteer resources to perform the desired amount of shaving).

Let $B^\tau = \{b_1, b_2, b_3 \cdots b_m\}$ denote the set of batteries connected to transformer τ, each indexed by i. Let $A^\tau = \{a_1, a_2, a_3 \cdots a_n\}$ denote the set of AC units connected to transformer τ, each indexed by j. Further, let $E^\tau = \{e_1, e_2, e_3 \cdots e_p\}$ denote the set of EVs connected to transformer τ, each indexed by k. In the total time period during which a transformer is experiencing peak usage, we assume a time-slotted model where in each time slot t and transformer τ, we have the transformer's peak power usage denoted by $p^{\tau}_{peak}(t)$, and the target power usage denoted by $p^{\tau}_{target}(t)$.

As the optimization variables, let $x^{\tau}_{i}(t)$ denote the discharge rate of battery i, $y^{\tau}_{j}(t)$ denote the per-degree power rating of AC unit j, and $z^{\tau}_{k}(t)$ denote the charging rate of EV k, all connected to transformer τ at time t. Also, let F^{τ}_{j} denote the maximum change in setpoint temperature for HVAC unit j connected to transformer τ.

Although *all* volunteer resources are specified in our optimization algorithm, since only a subset is chosen for the next day and a priority order is specified, we use binary variables for each resource to indicate whether it should be used in the current round of optimization. A resource is only considered in the current round for peak shaving if its binary variable is set to 1 (this also allows iterative optimization, when progressively larger groups of resources are considered in the specified order). Let $\alpha^{\tau}_{i}(t)$ denote a binary variable indicating whether battery i can be discharged during time t, $\beta^{\tau}_{j}(t)$ denote a binary variable indicating whether AC unit j can have temperature control during time t, and $\gamma^{\tau}_{i}(t)$ denote a binary variable indicating whether the charging of EV k can be postponed during time t, all connected to transformer τ. Then, the demand constraint indicates that the sum of battery power discharged, AC power reduced and EV charging foregone should be greater than or equal to the difference between the peak and target power, and we have

$$\sum_{i=0}^{m} x_i\alpha_i(t)+\sum_{j=0}^{n}(y_j\times F_j)\beta_j(t)+\sum_{k=0}^{p} z_k\gamma_k(t) \geq p_{peak}(t)-p_{target}(t) \tag{1}$$

Let $l^{\tau}_{i}(t)$ denote the level of charge in battery i at the beginning of time t. The total power discharged from the battery during the duration of time slot t should be less than or equal to the initial level of charge in the battery, and we have

$$x_i \times t \leq l_i(t) \quad \forall i, \forall t \tag{2}$$

Additionally, the discharge rate of any battery cannot exceed its maximum discharge rate d_i, and we have

$$x_i(t) \leq d_i \quad \forall i, \forall t \tag{3}$$

In the succeeding time slot, the flow conservation constraint indicates that the level of charge in battery i at the beginning of the next time slot is the difference between the level of charge in battery i at the beginning of the current time slot minus the energy discharged during that time slot, and we have

$$l_i(t+1) = l_i(t) - (x_i \times t) \quad \forall i, \forall t \tag{4}$$

Note that since the utility does not own the resources in the VPP, it only has a limited amount of time during which it can use them for peak shaving. Let $u^{\tau}_{i}(t)$ denote the controllable time for battery i connected to transformer τ at time t. Similarly, let $v^{\tau}_{j}(t)$ denote the controllable time for AC unit j at time t, and $w^{\tau}_{k}(t)$ denote the controllable time for EV k at time t, both connected to transformer τ. The flow conservation constraint indicates that the remaining controllable time at the beginning of the next time slot must be equal to the controllable time at the beginning of the previous time slot minus the length of the time slot, and for all resources, we have

$$u_i(t+1) = u_i(t) - t \quad \forall i, \forall t \tag{5}$$

$$v_j(t+1) = v_j(t) - t \quad \forall j, \forall t \tag{6}$$

$$w_k(t+1) = w_k(t) - t \quad \forall k, \forall t \tag{7}$$

At the end of each time slot, our optimization must determine whether a resource will be available for peak shaving in the next time slot. A resource is available for peak shaving in a particular time slot if the remaining controllable time is greater than zero (we use 0.001 in our optimization constraints). If this condition is true, the binary variable controlling the resource's availability for peak shaving is set to true for the next time slot. If the condition is false, the binary variable for the resource is set to false, and it will not be used for peak shaving in that time slot. We therefore use the $big - M$ constraint method to ensure that there exists remaining controllable in the next time slot. Let $M = 8760$ denote the upper bound for controllable time for any resource. We use 8760 because it represents the true upper bound for all hours in a year, and it is therefore impossible to have a value greater than this for controllable time. Therefore, for battery i connected to transformer τ at time t, we have

$$u_i(t+1) \geq 0.001 - M(1-\alpha_i)(t+1) \quad \forall i, \forall t \tag{8}$$

$$u_i(t+1) \leq 0.001 - M\alpha_i(t+1) \quad \forall i, \forall t \tag{9}$$

$$\alpha_i(t) \in \{0,1\} \quad \forall i, \forall t \tag{10}$$

Similarly, for AC unit j and EV k, both connected to transformer τ at time t, we have

$$v_j(t+1) \geq 0.001 - M(1-\beta_j)(t+1) \quad \forall j, \forall t \tag{11}$$

$$v_j(t+1) \leq 0.001 - M\beta_j(t+1) \quad \forall j, \forall t \tag{12}$$

$$\beta_j(t) \in \{0,1\} \quad \forall j, \forall t \tag{13}$$

$$w_k(t+1) \geq 0.001 - M(1-\gamma_k)(t+1) \quad \forall k, \forall t \tag{14}$$

$$w_k(t+1) \leq 0.001 - M\gamma_k(t+1) \quad \forall k, \forall t \tag{15}$$

$$\gamma_k(t) \in \{0,1\} \quad \forall k, \forall t \tag{16}$$

Given the above constraints, our goal is to minimize the sum of power drawn from discharging batteries, reducing AC loads, and foregoing EV charging. Hence, the optimization problem that determines the minimum power drawn from all resources in the VPP across all transformers is formally formulated as

$$Miminize \sum_{i=0}^{m} x_i(t) + \sum_{j=0}^{n}(y_j \times F_j)(t) + \sum_{k=0}^{p} z_k(t) \quad \forall t \tag{17}$$

At each time slot, we use the predicted peak and target power for each transformer to compute the minimum set of resources that provide the desired reduction from the peak to the target power for all transformers. As noted above, the optimization is invoked iteratively using a progressively larger number of the selected resources, in priority order, until we obtain the desired reduction.

6 EXPERIMENTAL EVALUATION

We have implemented a prototype of our VPeak system and in this section, we evaluate its efficacy using real-world datasets.

6.1 Datasets and Workloads

Distribution Grid Dataset. Our grid dataset, which has been made available through our utility partners, consists of power usage data of 15,242 smart meters distributed across an entire city. The distribution grid consists of 1,297 transformers that serve these smart meters. The power usage data is recorded at 5 minute granularity and is available for 6 years (i.e. 2014 to 2019). To compute the load on each distribution transformer, we sum up the power drawn by each smart meter connected to it to obtain a time series of transformer load 5-minute granularity. Table 3 summarizes the characteristics of this dataset.

Electric Vehicle Dataset. Since our peak shaving algorithm leverages EV loads to shave grid peaks, we use the openly-available Dataport EV data[1] to augment our electric load dataset with real world EV loads. The EV dataset consists of real-world EV power consumption data gathered from 91 EVs. The data is recorded at 5-minute granularity, and is available for the year 2016. The data includes the power drawn, time of charge, and duration the car was connected to the power outlet. The data is gathered from a heterogeneous mix of three popular EV car models i.e. 12 Tesla Model S cars, 18 Nissan Leafs, and 61 Chevy Volt cars. To integrate EVs into the grid dataset, we randomly assign one of the 91 EVs to randomly chosen homes up to a particular penetration level and then overlay a year long trace for that EV onto the home's energy usage.

AC Load Data. The AC load for each home are estimated using a well-known thermal model [2]. Since our electric grid dataset also includes house information such as size and year built, we use the process described in Cooling Load Calculation [2] to model the heating and cooling energy usage of each home.

Weather Data. We use historical as well as next day weather forecast data from hourly data available via the Dark Sky API.

6.2 Peak Day Prediction

We begin with an evaluation of VPeak's peak day selection algorithm described in Section 4.2. Our experiments assume that all prediction approaches have 5 years of training data and validation data (2014 to 2018), and use the 2019 data for making predictions. We first compare our approach to two baseline methods: (i) temperature-threshold, which uses a historical cutoff temperature to classify next day as a peak, and (ii) demand-threshold, which uses a historical cut-off demand for classification. We use *precision* and *recall* as our metrics. Figure 5a shows the efficacy of the three approaches for predicting top-n days from $n = 3$ to $n = 10$ (here, $\Delta = 0$). As shown, VPeak outperforms the other two approaches substantially, with recall values ranging from 0.44 to 0.68. Next, we compare the benefits of using Δ and the use of VPeak's dynamic threshold method. Figure 5b compares the recall values for static and dynamic thresholds for various values of Δ, with $n = 5$. The figure shows that the use of dynamic threshold always improves performance over static threshold. Further, as even use of small non-zero Δ of one or two days yields significant improvements; the recall increases from 0.63 to 0.87 when we use $\Delta = 2$. The corresponding precision values range from 0.44 to 0.85. Finally, Figure 6 compares VPeak to the state-of-the-art probabilistic method used in Ontario [14]. We modify the approach of [14] to predict monthly top-n peaks using a three day lookahead. Figure 6 shows that VPeak's machine learning approach outperforms the state of the art probabilistic approach for both precision and recall, even though VPeak is using a one day lookahead compared to the 3-day lookahead for the probabilistic method. As top-n' is varied from 5 to 7 (i.e., $n = 5$, $\Delta = 0$ to 2), VPeak's recall approaches 0.87 with a precision of 0.8, which is an improvement of 20 and 33% respectively over the state of the art.

[1] https://www.pecanstreet.org/dataport/

Number of smart meters	15,242
Number of transformers	1,297
Smart meter data granularity	5 minutes
Duration	2014 to 2019

Table 3: Grid dataset

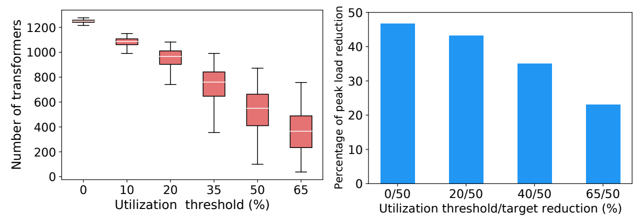

(a) (b)

Figure 4: Selective Peak Shaving: (4a) Number of transformers whose load is shaved for different utilization thresholds, and (4b) Percentage of load reduction for different utilization thresholds.

6.3 Peak Load Shaving

Next, we evaluate the efficacy of VPeak for selective and grid-wide peak shaving. Like before, we predict the top-n' days in each month for 2019 and use our optimization approach to perform peak shaving.

6.3.1 Selective Peak Shaving. For selective peak shaving, once the next day is classified as a top-n' peak day, we use a high load threshold U to select a subset of the transformers that will see peak utilization greater than U; only the peak load at these selective transformers is shaved. Figure 4a depicts the number of transformers chosen for selective shaving as U is varied from 0 to 65% utilization. As can be seen, at a high cutoff of 65% load, an average of 400 transformers are chosen on each peak day. The number of transformers increases as this cutoff is lowered, with $U = 0$ selecting all 1,297 transformers in the distribution grid. Figure 4b shows the amount of total load reduction for varying degrees of selective peak shaving. As shown, for $U = 65\%$, VPeak selectively reduces stress on the most overloaded transformers and yet reduces 23% of the overall grid load. Overall, as the threshold is reduced, more transformers participate in selective shaving, which causes the overall reduction in grid-wide peak load to rise to 46.8% as the threshold goes to zero. Overall, the results show the flexible nature of our approach and its ability to perform targeted shavings at localized hotspots.

6.3.2 Grid-wide Peak Shaving. Next, we evaluate VPeak's grid-wide peak shaving benefits, where all transformers are chosen for peak shaving on each peak day. Figure 7a provides a visual depiction of the original grid load for the entire year and the load after peak shaving is performed. We notice that the tallest peaks in each month

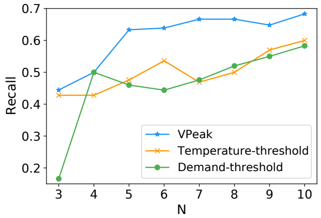

(a)

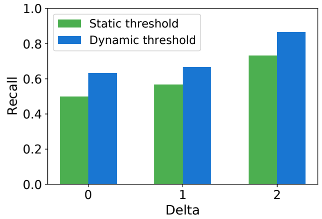

(b)

Figure 5: (5a) Comparing VPeak's peak day selection to baseline methods, and (5b) Impact of Δ and dynamic threshold for $n = 5$.

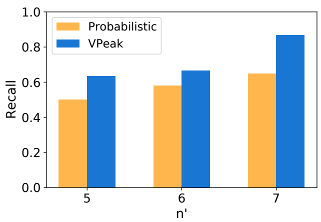

(a) Recall

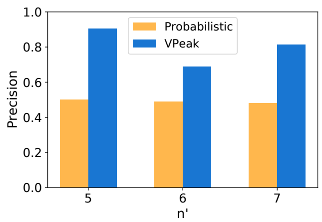

(b) Precision

Figure 6: Comparing VPeak's peak selection algorithm to the state-of-the-art probabilistic method.

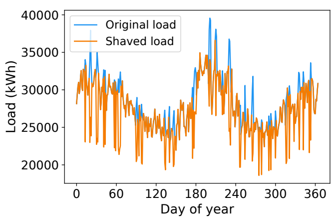

(a)

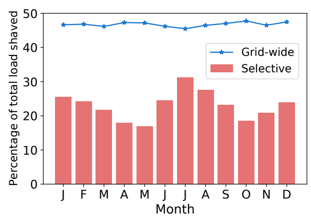

(b)

Figure 7: Grid-wide peak shaving: (7a) Original and shaved demand for the whole grid during the year, and (7b) Percentage of grid-wide load reduction by month of the year.

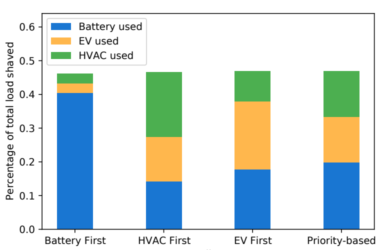

Figure 8: Impact of different resource selection policies on the contribution of each resource to the total peak shaved.

are reduced by VPeak. Figure 7b depicts the grid-wide peak reduction for each month of the year achieved by VPeak and compares it to selective shaving for $U = 65\%$. The figure shows that grid-wide shaving is able to reduce peak demand by around 46% for all months of the year. Interestingly, reductions from selective shaving vary by the season, with higher reductions in summer and winter when the peaks are higher. These months see a 26.2% peak reduction, on average, which is greater than the 19.8% peak reduction in Fall and Spring months. Also, the summer month of July sees the greatest peak reduction of 31.2%.

6.4 Impact of Resource Selection Policies

Next, we analyze the efficacy of different resource selection policies by analyzing the relative proportion of various resources (battery, HVAC, EV) that contribute to the peak reduction. Assuming 50% penetration of EV and batteries, we evaluate peak shaving using battery-first, HVAC-first, EV-first and random priority-based policies (our priority policy is used to construct HVAC-first and EV-first approaches by giving highest priority to those resources). Figure 8 depicts the results of this analysis over summer and winter months. As shown, all policies are able to achieve a similar degree (around 40%) of grid-wide peak shaving, but the relative contribution of resources that yield this reduction varies by policy. In the battery-first approach, most of the demand reduction is achieved by discharging batteries, with HVAC and EVs providing a small 2.9% reduction each—which minimizes user inconvenience by not deferring loads whenever possible. In contrast, HVAC-first and EV-first policies provide a higher contribution of demand reduction by deferring HVAC and EV loads first (yielding 19.2% and 19.8% peak reduction, respectively). The other two loads provide smaller, but meaningful contributions. In the random priority policies, we assume priorities are assigned randomly users, and all three loads yield similar reductions. Overall, our results show that VPeak's ability to order loads for peak shaving provides good flexibility in addressing higher-level goals such as addressing user convenience or user preferences.

6.5 Impact of Resource Penetration

Since VPeak relies on volunteered resources, the availability of these resources affects how much demand reduction can be achieved in practice. We evaluate the impact of varying levels of resource penetration, where penetration represents the percentage of homes that have volunteered a particular resource.

Figures 9a and 9b depict the impact of varying levels of battery and EV penetration on peak demand reduction. The figures shows that reduction in demand increases from 426,013 to 3,393,252 kWh (a 696% increase) as the battery penetration increases from 5% to 75%, and it increases from 171,018 to 1,665,222 kWh (a 873% increase) as EV penetration from 5% to 50%.

To evaluate the impact of HVACs on demand reduction, we vary the setpoint change that can be used by the utility—the higher the setpoint change, the higher the HVAC load reduction. As shown in Figure 9c, as the HVAC setpoint is changed by 3°F to 15°F, the demand reduction increases from 200,587 to 777,368 kWh (a 287% increase). Overall our results show that even modest amount of volunteer resources can yield non-trivial amounts of peak reductions and higher participation in such schemes have the potential to yield significant benefits.

7 RELATED WORK

Prior work in this area falls into two categories: peak shaving and peak demand prediction. Peak load shaving has been studied extensively using centralized resources owned and controlled by the utility [19] as well as distributed storage resources installed at customers' homes [5, 6, 8]. However, past work has not focused much on targeted selective peak shaving or specifically on the use of volunteer resources to perform peak reductions.

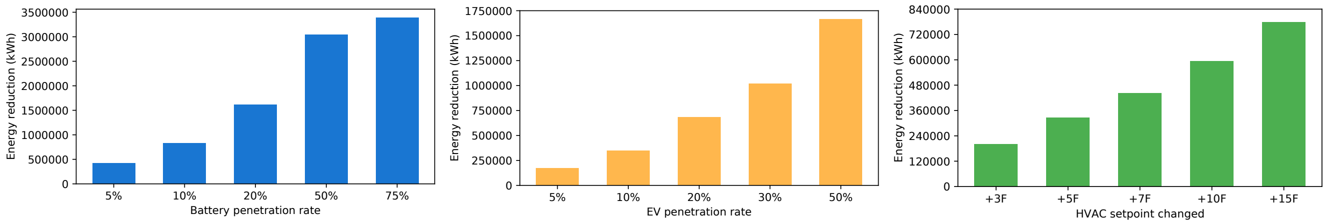

(a) (b) (c)

Figure 9: (9a) Impact of increased battery penetration on demand reduction, and (9b) Impact of increased EV penetration on demand reduction, and (9c) Impact of increased HVAC setpoint level on demand reduction.

Demand prediction in electric grids has also been extensively studied. Past approaches include the use of wavelet-based forecasting [16], the use of time series forecasting methods such as ARIMA [15, 21] as well as the use of deep learning models [10, 18, 26]. Hybrid methods the combine machine learning and ARIMA forecasting have also been studied [11]. Peak forecasting, which is complementary to demand prediction, has received much attention. The most relevant effort is a probabilistic method to forecast the top-k peak days of the year for Ontario, Canada [14]. VPeak differs from these efforts in its use of LSTM-based machine learning and its use of dynamic thresholds to mitigate the impact of forecasting errors.

8 CONCLUSIONS

In this paper, we presented VPeak, an approach that uses residential loads volunteered by homeowners for coordinated control by a utility for grid optimizations. Using machine learning, VPeak's peak day prediction algorithm was able to identify 86% of the monthly peak days over a year while outperforming other approaches. Using an iterative linear program optimization, VPeak implemented flexible peak shaving that can selectively target hotspots within the distribution grid while also supporting grid-wide peak shaving. VPeak also can intelligently select energy resources for shaving based on a given policy or preferences. Our results using a city scale electric showed that VPeak is able to reduce the overall grid demand by 23% for selective peak shaving and by 46.7% for grid-wide shaving.

9 ACKNOWLEDGEMENT

We thank our anonymous reviewers for their insightful comments. This research is supported by NSF grants 1645952, 2020888, 2021693, and 2136199.

REFERENCES

[1] [n.d.]. ConnectedSolutions | National Grid. https://www.nationalgridus.com/MA-Home/Energy-Saving-Programs/ConnectedSolutions/. (Accessed on 07/14/2021).

[2] [n.d.]. Cooling Load Calculation | The Engineering Mindset. https://theengineeringmindset.com/cooling-load-calculation-cold-room/. (Accessed on 02/14/2021).

[3] [n.d.]. Rush Hour Rewards 2013 Results White Paper for nest.com.pages. https://storage.googleapis.com/nest-public-downloads/press/documents/rush-hour-rewards-white-paper.pdf. (Accessed on 02/14/2021).

[4] David P. Anderson. 2020. BOINC: A Platform for Volunteer Computing. *Journal of Grid Computing* 18 (2020), 99–122.

[5] Sean Barker, Aditya Mishra, David Irwin, Prashant Shenoy, and Jeannie Albrecht. 2012. Smartcap: Flattening peak electricity demand in smart homes. In *2012 IEEE International Conference on Pervasive Computing and Communications*. IEEE, 67–75.

[6] Lukas Barth and Dorothea Wagner. 2019. Shaving peaks by augmenting the dependency graph. In *Proceedings of the Tenth ACM International Conference on Future Energy Systems*. 181–191.

[7] A Caramizaru and A Uihlein. 2020. Energy Communities: An Overview of Energy and Social Innovation. European Comission JRC Science for Policy Report.

[8] Yasra Chandio, Aditya Mishra, and Anand Seetharam. 2019. Gridpeaks: Employing distributed energy storage for grid peak reduction. In *2019 Tenth International Green and Sustainable Computing Conference (IGSC)*. IEEE, 1–8.

[9] Weiwei Chen, Xing Wang, Jon Petersen, Rajesh Tyagi, and Jason Black. 2013. Optimal scheduling of demand response events for electric utilities. *IEEE Transactions on Smart Grid* 4, 4 (2013), 2309–2319.

[10] Alejandro J del Real, Fernando Dorado, and Jaime Durán. 2020. Energy demand forecasting using deep learning: applications for the French grid. *Energies* 13, 9 (2020), 2242.

[11] Ali Ghasemi, Hossien Shayeghi, Mohammad Moradzadeh, and Mohammad Nooshyar. 2016. A novel hybrid algorithm for electricity price and load forecasting in smart grids with demand-side management. *Applied energy* 177 (2016).

[12] Yashar Ghiassi-Farrokhfal, Catherine Rosenberg, Srinivasan Keshav, and Marie-Benedicte Adjaho. 2016. Joint optimal design and operation of hybrid energy storage systems. *IEEE Journal on Selected Areas in Communications* 34, 3 (2016), 639–650.

[13] Henrique Steinherz Hippert, Carlos Eduardo Pedreira, and Reinaldo Castro Souza. 2001. Neural networks for short-term load forecasting: A review and evaluation. *IEEE Transactions on power systems* 16, 1 (2001), 44–55.

[14] Yuheng Helen Jiang, Ryan Levman, Lukasz Golab, and Jatin Nathwani. 2014. Predicting peak-demand days in the Ontario peak reduction program for large consumers. In *Proceedings of the 5th international conference on Future energy systems*. 221–222.

[15] Yunsun Kim, Heung-gu Son, and Sahm Kim. 2019. Short term electricity load forecasting for institutional buildings. *Energy Reports* 5 (2019), 1270–1280.

[16] Abderrezak Laouafi, Mourad Mordjaoui, Farida Laouafi, and Taqiy Eddine Boukelia. 2016. Daily peak electricity demand forecasting based on an adaptive hybrid two-stage methodology. *International Journal of Electrical Power & Energy Systems* 77 (2016), 136–144.

[17] Shuhui Li and Dong Zhang. 2014. Developing smart and real-time demand response mechanism for residential energy consumers. In *2014 clemson university power systems conference*. IEEE, 1–5.

[18] Xueheng Qiu, Ye Ren, Ponnuthurai Nagaratnam Suganthan, and Gehan AJ Amaratunga. 2017. Empirical mode decomposition based ensemble deep learning for load demand time series forecasting. *Applied Soft Computing* 54 (2017), 246–255.

[19] Azar Rahimi, M Zarghami, M Vaziri, and S Vadhva. 2013. A simple and effective approach for peak load shaving using Battery Storage Systems. In *2013 North American Power Symposium (NAPS)*. IEEE, 1–5.

[20] Mauricio Restrepo, Jordan Morris, Mehrdad Kazerani, and Claudio A Canizares. 2016. Modeling and testing of a bidirectional smart charger for distribution system EV integration. *IEEE Transactions on Smart Grid* 9, 1 (2016), 152–162.

[21] Hossein Javedani Sadaei, Frederico Gadelha Guimarães, Cidiney José da Silva, Muhammad Hisyam Lee, and Tayyebeh Eslami. 2017. Short-term load forecasting method based on fuzzy time series, seasonality and long memory process. *International Journal of Approximate Reasoning* 83 (2017), 196–217.

[22] SETI [n.d.]. SETI@Home: The Search for Extra Terrestrial Intelligence. SETI Webite: https://setiathome.ssl.berkeley.edu.

[23] Yuanyuan Shi, Bolun Xu, Di Wang, and Baosen Zhang. 2017. Using battery storage for peak shaving and frequency regulation: Joint optimization for superlinear gains. *IEEE Transactions on Power Systems* 33, 3 (2017), 2882–2894.

[24] Akhil Soman, Amee Trivedi, David Irwin, Beka Kosanovic, Benjamin McDaniel, and Prashant Shenoy. 2020. Peak Forecasting for Battery-based Energy Optimizations in Campus Microgrids. In *Proceedings of the Eleventh ACM International Conference on Future Energy Systems*. 237–241.

[25] Peng Xu. 2006. Peak demand reduction from pre-cooling with zone temperature reset in an office building. (2006).

[26] Jian Zheng, Cencen Xu, Ziang Zhang, and Xiaohua Li. 2017. Electric load forecasting in smart grids using long-short-term-memory based recurrent neural network. In *2017 51st Annual Conference on Information Sciences and Systems (CISS)*. IEEE, 1–6.