Energy Storage in Time Saves Nine: A Case for a Greener Smart Grid

Rishikesh Jha Srinivasan Iyengar Prashant Shenoy
College of Information and Computer Sciences
University of Massachusetts
Amherst, MA 01003, USA
{rishikeshjha,srini,shenoy}@cs.umass.edu

Abstract

Our greenhouse gas emissions vary with the mix of fuel types used in the electricity generation infrastructure at any given time. Shifting load from high polluting time to low polluting time can lead us to a greener grid. A possible avenue to adjust our demand is to leverage distributed energy storage that is being deployed to maintain grid stability. In this paper, we propose an algorithm to enable this demand-side management approach. First, we propose neural network based day-ahead time series regression models for the energy demand seen at the transformer-level. Second, we employ a linear programming-based approach that minimizes the carbon intensity of our energy usage by scheduling the charging/discharging of the distributed energy storage. This approach considers various grid and battery constraints and reduces the total greenhouse gas emission in the grid while fostering grid stability. We evaluate our approach on a real-world dataset of 776 transformers servicing over 9495 electric meters in a city in the New England region of the US. Specifically, our analysis showed a reduction of >10 million pounds in annual carbon emissions — equivalent to a drop of 14.48% in our electric grid emissions.

Introduction

Power grids are the backbone of any modern society as they provide the energy required for various human activities. A typical grid consists of power plants using diverse sources of energy — each contributing to greenhouse emissions to varying degrees. A combination of these power plants is in operation based on the energy demand at any give time. Consequently, our greenhouse gas emissions vary throughout the year with the mix of fuel types used in the electricity generation. Reducing our carbon emissions necessitates considering the fuel mix used in electricity production.

One approach to mitigate greenhouse emissions is to increase the mix of renewable sources in our overall energy generation. However, renewable sources, such as wind and solar, are highly intermittent as they are dependent on weather conditions. Such energy sources cannot be expected to respond to changing human energy demands. Further, hydropower needs a favorable geographic location with the suitable height difference between inflow and outflow of water. Therefore, renewable sources have significant limitations in completely offsetting our carbon emissions.

Another solution to reducing our emissions is to shift our energy consumption from a high polluting energy mix (coal, natural gas, etc.) to low polluting energy mix (nuclear, hydro, etc.) using energy storage. Fortunately, grid operators around the world are investing in energy storage. With the advent of the *smart grid*, there is a focus on installing distributed energy storage to provide grid stability by mitigating transformer overloads (Nourai 2007) and avoiding supply-demand mismatch (Nunna and Doolla 2013). The popularity of distributed storage is expected to increase in the coming years (Bloomberg NEF 2017). In this work, we aim to leverage this energy storage infrastructure, to mitigate greenhouse emissions by reducing our reliance on high polluting sources. Distributed storage can be used to prevent reliance on peaky power plants which are often the high polluting ones therby reducing emissions. Further, we need to ensure that shifting our grid energy usage does not violate the original mandate of installing energy storage — i.e. maintaining grid stability.

In prior work, energy storage have also been used for price arbitrage (Mishra et al. 2012; Walawalkar, Apt, and Mancini 2007). With real-time prices becoming more prevalent, cost arbitrage would involve future price prediction to infer battery schedules. However, electricity markets, akin to stock markets, are adaptive systems that respond to predictions. Modeling their behavior is hard as the predictions are available to market participants to respond to. Whereas, in our work, we need to predict energy usage to infer battery operation schedules. Human energy usage (like weather) is not an adaptive system as it does not respond to predictions. Thus, we are fundamentally solving a tractable problem.

In this paper, we present an energy storage scheduling scheme to reduce carbon emissions while maintaining grid stability. Our key contributions are as follows:

Transformer-level load forecasting. We propose an autoregressive neural network architecture that utilizes the historical transformer-level energy demand data exogenous variables to perform load forecasting. Specifically, we show that our approach improves over the state-of-the-art technique for load forecasting (Kong et al. 2017) on average 2.7%.

Distributed energy storage scheduling. We present a linear programming based schedule for the distributed energy storage. The generated schedule uses the load forecasts at the transformer-level to minimize the greenhouse emissions from electricity generation. Further, we incorporate several constraints that maintain the grid stability.

Grid-scale evaluation. We evaluate our approach on a grid-scale energy usage data from 776 transformers spread across a city in the New England region of the United States. We report a carbon emissions savings of $>$10 million lbs over a period of a year. This reduction is equivalent to 14.48% of overall emissions from the electric grid.

Background

In this section, we elucidate the relevant details surrounding electric grids and the problem addressed in the paper.

Energy sources - Characteristics and constraints

An electric grid has three main components: i) Generation ii) Transmission and iii) Distribution. At any time, demand and supply must be matched for the proper functioning of the grid. Since demand changes over the course of a day, the generation must be matched via a *dispatch schedule*. Different generation units have different properties with regards to response to change in load, and they can be categorized into three main types: i) *Base load plants*: Plants such as nuclear, large-scale coal and biomass belong to this category. They operate at maximum capacity and reduce power only to perform maintenance. ii) *Load following power plants*: Plants such as Hydro-power units operate during high demand periods within a day (morning and early evening). iii) *Peaker plants*: These generation units run only when there is a peak demand, which occurs during the period of extreme (high or low) temperatures to provide human heating and cooling requirements. They are the most inefficient power plants and tend to be the most polluting ones. Power plants such the ones based on Natural Gas and Oil belong to this category.

Average and Marginal Carbon Intensity

Estimating carbon intensity of the energy mix used for power generation is pivotal in developing energy management strategies geared towards reducing our greenhouse emissions. Specifically, we will explain two terms — i.e., *average carbon intensity* and *marginal carbon intensity* — that are important to understanding changes in greenhouse emissions with changing electricity demand in the grid.

Average carbon intensity of an electric grid is defined as the total carbon emissions per unit of electricity generated. Table 3 shows the emission factors for the different generation types, according to ISO New England (Hoedl 2016). Calculating *average carbon intensity* entails a weighted average of carbon emissions by energy mix used. For example, if an electric grid produces electricity from coal, natural gas, nuclear and hydro in equal proportion, then the *average carbon intensity* would be 748.75 lbs/MWh $(2123 \times 0.25 + 872 \times 0.25 + 0 \times 0.25 + 0 \times 0.25)$.

However, the reduction or increase in generation due to changes in electric demand is not averaged across all power plants. The change in the generation will occur in the *marginal power plant* — the power plant that can respond quickly to changes in electricity demand. Therefore, while designing an energy storage charging/discharging schedule

Generation Type	Emission Factor (CO_2 lbs/MWh)
Coal	2123
Natural Gas	872
Oil	2059
Nuclear	0
Hydro	0
Solar and Wind	0

Table 1: Carbon emission by different generation types.

one must consider *marginal carbon intensity*, which is the emission associated with the marginal power plant. Thus, in the above example, if the marginal power plant is using natural gas as fuel, the *marginal carbon intensity* is 872 lbs/MWh, which is higher than the *average carbon intensity* calculated earlier. As we have notable differences in the emission factors associated with the different fuel types, there is tremendous potential to reduce our carbon footprint by modulating our energy demand using energy storage.

Setup and Problem Statement

In this paper, we assume that the energy storage is deployed adjacent to each *distribution transformer*. In a typical grid, there is significant variation in the transformers capacity and the number of electric meters it supports. Thus, the energy storage is sized according to the transformer capacity to protect it from overloads. Further, we assume that a central authority, such as the grid operators, can control the charging-discharging of these distributed energy storage. We assume Coal, Natural Gas and Oil based power plants to be the marginal power plants i.e any fluctuation in grid load due to our charging or discharging action will be compensated by a combination of these power plants. Based on this setup, our problem statement is as follows - *Using the anticipated future marginal carbon intensity, to reduce grid-level emissions by modulating our consumption through scheduling energy storage available at the distribution transformer.*

Algorithm

In this section, we present our algorithm to reduce grid-level carbon emissions using the setup described earlier. Figure 1 shows the complete pipeline of our algorithm. It is imperative to have an accurate estimate of the future load at the transformer level to make optimal decisions for scheduling distributive storage. Thus, the first step is **load forecasting**, where we use historical transformer-level energy data along with exogenous variables to generate day-ahead forecasts. Follwoing this, the forecasts are given as input to **emissions-aware scheduling**, which considers anticipated marginal carbon intensity values and various grid-level parameters to generate day-ahead energy storage schedules. As part of **online implementation**, we make adjustments to the day-ahead schedule to account for real-time grid constraints and produce the final operational schedule. We cover these steps in detail in the section below.

Load Forecasting

Load forecasting is a widely studied problem in the smart grid domain. A detailed survey of several approaches can be

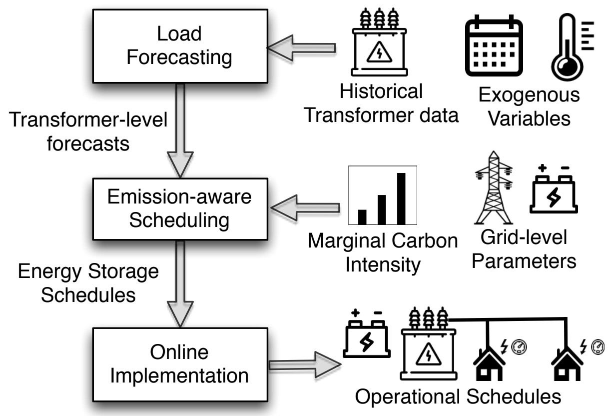

Figure 1: Algorithm for minimizing grid carbon emissions.

found in (Kyriakides and Polycarpou 2007). More recently, LSTM based techniques have achieved state-of-the-art accuracy (Kong et al. 2017). These techniques were used to get future predictions of energy usage at *grid-scale* using historical data. However, in our work, we want to perform load forecasting to predict *transformer-level* energy usage, where along with historical data we use exogenous variables. We know that human energy needs vary with change in temperature. Also, energy usage is dependent on human work schedules (Iyengar et al. 2016). Thus, we use both temperature and day of the week as exogenous variables.

Formally, in load forecasting, we need to learn a function $f^{[d]}$ $\forall d \in D$ (i.e. the set if all the transformers in consideration), which uses the load seen during the previous t timestamps, given by $l_1^d, l_2^d, ..., l_t^d$ to forecast the load for the next k timestamp over the 24-hour window, given by $l_{t+1}^d, l_{t+2}^d, ..., l_{t+k}^d$. The function $f^{[d]}$ can additionally leverage temperature estimates for the future k timestamps, given by $\tau_{t+1}, \tau_{t+2}, ..., \tau_{t+k}$ and δ, which is the day of the week. In summary, the regression function uses historical energy data along with exogenous variables and is represented as -

$$l_{t+1}^d, ..., l_{t+k}^d = f^{[d]}(l_1^d, ..., l_t^d, \tau_{t+1}, ..., \tau_{t+k}, \delta) \quad \forall d \in D$$

We evaluated several existing load forecasting approaches (see evaluation section). Apart from these, we propose two regression techniques. We describe them as follows -

Autoregressive Neural Network (Diaconescu 2008): Transformer level load data is noisy in nature. However, we observe that historical load data contains daily and weekly seasonality. It is important to extract the seasonality in the historical data for accurate forecasting. In order to forecast load for time $t+1$ we use the past p_1 timesteps as the input along with loads on past p_2 days at time $t+1$ as well as loads at past p_3 weeks at time $t+1$. Along with the historical load, we include one-hot encoded day of the week exogenous variable as part of feature vector of the neural network. We also use the forecasted temperature at time $t+1$ as an external regressor. Our forecast at time $t+1$ is -

$$l_{t+1}^d = \beta \cdot \tau_{t+1} + NN(l_t^d, l_{t-1}^d, .., l_{t-p_1}^d, l_{t+1-h \cdot p_2}^d, l_{t+1-h \cdot p_3}^d, \delta)$$

To forecast k steps in the future, we use rolling forecasting using past predictions.

Seasonal LSTM: As mentioned earlier, historical load data has multiple seasonalities. In order to capture them, we use different LSTMs for each seasonal component as they can have varying number of lagged values as input. We concatenate the encodings from each LSTM unit to get an encoded embedding. The encoded embedding is then passed through a feed forward layer to get the forecast one step ahead in the future. Similar to Autoregressive Neural Network, we use daily and weekly seasonality along with past p_1 timesteps. Below, we describe our LSTM model.

$$e_1 = LSTM(l_t^d, l_{t-1}^d, .., l_{t-p_1}^d)$$

$$e_2 = LSTM(l_{t-h}^d, l_{t-2h}^d, .., l_{t-p_2.h}^d)$$

$$e_3 = LSTM(l_{t-h}^d, l_{t-2h}^d, .., l_{t-p_3.h}^d)$$

$$l_{t+1}^d = NN([e_1, e_2, e_3])$$

Emissions-aware Storage Scheduling

Next, we use the load forecasts and other grid-level parameters to generate a 24 hour ahead energy storage schedule. We formulate the problem of reducing emissions using storage as a Linear Programming problem.

Parameters: In our setup, we have a distributed energy storage at every transformer. The scheduling decisions are made at the transformer level. The capacity tc_d of the transformer and the storage capacity b_d of the d^{th} transformer are known beforehand and are parameters to our Linear Programming formulation. Let r_d be the maximum charging and discharging rate of the d^{th} storage unit. Further, day-ahead load forecasts l_t^d at time t for each transformer d are available. The emission factor $e_f$lbs/MWh of the f^{th} fuel type is used to compute the marginal emissions. We define marginal factor m_t^f as the contribution of the f^{th} fuel type at time t to the marginal increase or decrease in energy demand.

Variables: The variables in our formulation are the state of charge s_t^d of each distributed storage unit at time t. The charging/discharging schedule of each distributed storage at time t for the d^{th} transformer is represented as c_t^d. $c_t^d > 0$ indicates the charging of the storage whereas $c_t^d < 0$ indicates discharging. The overall change in load observed at the grid level at time t is represented by the variable λ_t.

Constraints: Our constraints represent the operating level constraints of the grid, distributed storage units, and the transformer. The output of the linear program is the scheduling decisions of storage units for the next 24 hours. The change in load observed at the grid level will be the sum of all the charging/discharging decisions made at the transformer level. This constraint is represented by equation 1.

$$\lambda_t = \sum_{d=1}^{D} c_t^d, \qquad \forall t \in T \tag{1}$$

The state of charge of the storage after a charge/discharge action is represented as:

$$s_{t+1}^d = s_t^d + c_t^d, \quad \forall t \in T \text{ and } \forall d \in D \tag{2}$$

In order to maintain demand and supply relationship, the discharge from the storage unit should not be greater than the load observed at the transformer at any time t. Thus, we have

$$c_t^d \leq l_t^d, \quad \forall t \in T \text{ and } \forall d \in D \tag{3}$$

In case of the load at the transformer at time t is greater than the capacity of the transformer, we do not want to worsen the situation by charging the storage at that time leading to transformer overload. We represent this as -

$$c_t^d \leq 0; \quad \text{if } l_t^d \leq tc_d, \forall t \in T \text{ and } \forall d \in D \tag{4}$$

We initialize the storage levels at half its total capacity to allow both charging and discharging starting with time $t = 0$. We also constrain the storage capacity at the end of the day to half its capacity so as to have the same state of charge for the next day.

$$s_t^d = \frac{b^d}{2}; \quad \text{if } t = 0 \text{ and } t = T, \forall d \in D \tag{5}$$

The scheduling decisions should be taken within the operating constraints of the battery. For example, the maximum charge of the storage unit should not exceed the storage capacity while it is in operation. Also, we assume the maximum charging and discharging rate to be equal. We represent battery level constraints as follows -

$$s_t^d \leq b^d; s_t^d \geq 0; c_t^d \geq -r^d; c_t^d \leq r^d \quad \forall t \in T \text{ and } \forall d \in D \tag{6}$$

Objective Function Our objective function defined below represents the sum of *marginal carbon emissions* at every time t.

$$min \sum_{t=1}^{T} \sum_{f} ef^f . mf_t^f . \lambda_t \tag{7}$$

where λ_t represents the change in load observed at grid level as described in equation 1. mf_t^f represents the marginal factor as defined in parameters. ef^f is the emission factor of the f^{th} fuel type.

Online Implementation

The day-ahead schedule of charging-discharging may violate real-time grid constraints as the observed load at time t may deviate from the forecasted load. To avoid such a scenario, we generate the online schedule by adjusting the day-ahead schedule using the following constraints -

Transformer Constraints

$$\alpha_t^d = 0; \text{ if } c_t^d > 0 \text{ and } tc_d \leq l_d^t \tag{8}$$

$$\alpha_t^d = tc_d - l_t^d \text{ if } c_t^d \geq tc_d - l_t^d \tag{9}$$

Storage Constraints

$$\alpha_t^d = b_d - s_t^d \text{ if } c_t^d + s_t^d \geq b_d \tag{10}$$

$$\alpha_t^d = -s_t^d \text{ if } c_t^d + s_t^d \leq 0 \tag{11}$$

As indicated earlier, we would like to avoid excessive storage discharging during low energy demand periods. This constraints is represented as:

$$\alpha_t^d = -l_t^d \text{ if } l_t^d \leq -c_t^d \tag{12}$$

Parameters	Value
Min. Charge/Discharge Time	200 mins
Battery Fraction	[.25, .5, 0.75, 1, 1.5]
Marginal Fuel Sources	Coal, Oil, Gas
Emission Factor (lbs/MWh)	Refer Table 1

Table 2: Linear Programming Settings

Evaluation Setup

Dataset

For evaluating the efficacy of our load forecasting techniques along with the distributed storage schedule, we use a grid-scale dataset containing energy data from 9495 smart meters connected to 776 transformer (Iyengar et al. 2016). This data is available at a 5-minute granularity over a period of 2 years. On average, each transformer serves 12.22 meters (ranging from 5 to 161). Likewise, the transformer capacity varies with the number of associated meters between 25 and 750 kVA. Additionally, our scheduling scheme requires marginal carbon intensity. This data is not directly available for the New England region. Hence, we use the method specified in (Rogers et al. 2013) that estimates the marginal power plants in operation using the current fuel and electricity generation prices available through (EIA 2018) and (ISO-NE 2018). Moreover, we use temperature as an exogenous variable in our regression technique.

Regression Techniques

We compare the results of the two proposed regression techniques with the state-of-the-art approach for load forecasting that uses LSTMs with exogenous variables such as temperature and day of the week (Kong et al. 2017). Additionally, we also compare our forecasts with two popular statistical time series techniques — ARIMA (Brockwell, Davis, and Calder 2002) and TBATS (De Livera, Hyndman, and Snyder 2011).

Experimental Settings

Regression Settings: As discussed in the previous section, load forecasting is done every 24 hours. Thus, we perform sliding window rolling evaluation of our load forecasting techniques i.e we train the model on past one years data and evaluate its performance on the next 24 hour window.

Linear Programming Settings: While evaluating our LP formulation, we use several parameters that are directly read from the dataset described above. Some of the additional parameters are described in table 2. For example, the battery size (represented in kWh) is dependent on the transformer capacity (represented in KVA). As shown below, we experiment with 5 battery sizes relative to the transformer capacity.

Experimental Results

Now, we present experimental results detailing the the performance of our algorithm. In particular, we assess - (i) the accuracy of our load forecasting models, (ii) the impact of our storage schedules on carbon emissions, and (iii) the changes in carbon emissions across various transformers with varying battery size.

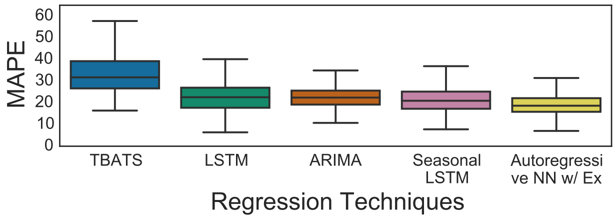

Figure 2: Comparison of the proposed regression techniques

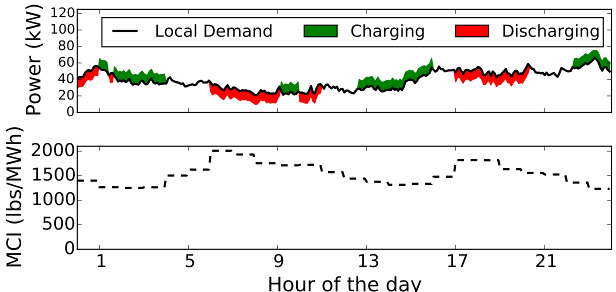

Figure 3: Battery action at a transformer on Feb 13^{th} 2015.

Load Forecasting results

Load forecasts accuracy directly impacts the efficacy of our schedules. Figure 2 compares the performance of our two proposed regression techniques with two popular statistical time series baselines and a state-of-the-art approach. This figure shows the distribution of MAPE values for the load at all our 776 transformers evaluated over a period of one year. Based on our analysis, TBATS performed the worst with an average MAPE of 34.17%. Whereas, the performance of the autoregressive neural network with exogenous variables outperformed all other techniques and had the lowest average MAPE value of 20.14%. The performance of all other methods was comparable with an average MAPE between 21.5% and 22.8%. Our autoregressive neural network method combined modeling of multiple seasonality along with the exogenous variables to generate the most accurate forecasts.

Illustrative Energy Storage Schedule

Figure 3 depicts the impact of storage schedule on the load observed at a transformer for a sample day overlaid with the local demand. As shown, discharging action occurs when the marginal emissions are high, i.e., during the high polluting hours of the day. Conversely, charging occurs when the marginal emissions are low. Based on the overall energy usage and the mix of fuels used at different times of the day, the alternating charging and discharging actions at this transformer mitigates 38.58 lbs of carbon emissions. Thus, emissions-aware distributed energy storage has tremendous potential to reduce carbon emissions at grid-scale.

Carbon Emission Analysis

We analyze the change in carbon savings[1] by varying the size of the energy storage. Table 3 shows the relationship between battery size and the carbon savings based on the year-long evaluation of our algorithm. A larger battery would

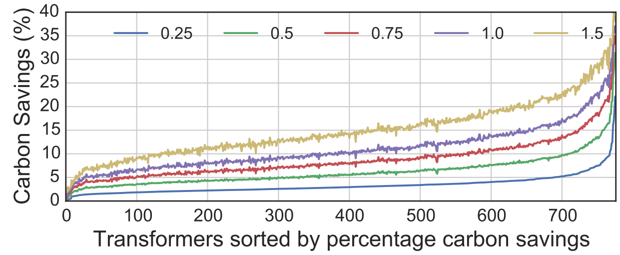

Figure 4: Transformer level carbon savings for 2015

Battery Fraction	Savings (%)	Savings (10^3 lbs)
0.25	3.153	2204.07
0.5	5.923	4140.671
0.75	8.397	5870.32
1	10.623	7426.354
1.5	14.483	10124.243

Table 3: Total Carbon Emission Saving for 2015.

have more flexibility in shifting transformer load. Although a larger battery translates to higher energy, we observe a diminishing return. As shown, a linear increase in battery size results in a sublinear reduction in emissions. However, a battery size that is $1.5\times$ the transformer capacity can annually save >10 millions lbs of carbon emissions — equivalent to a drop of 14.48% in our electric grid emissions.

Additionally, significant variations are observed across the different transformers. Figure 4 shows the percentage carbon savings across all 776 transformers. 146 transformers have less than 10% carbon savings even with the largest battery size. Whereas, 131 transformers have upwards of 20% carbon savings. Thus, grid operators can consider strategic placement of batteries at a select number of transformers.

Related Work

In smart grids research, load forecasting is a widely studied problem (Taieb et al. 2017). The regression techniques used to solve this problem range from traditional time series approaches such as ARIMA (Nguyen and Hansen 2017) to neural network based LSTMs (Kong et al. 2017). Traditionally, grid-level load forecasting was used to assess power systems security, schedule maintenance services, etc. Our regression model produces forecasts at the transformer-level and improves over the state-of-the-art technique (Kong et al. 2017). Further, we present a real-world case-study where our predictions are leveraged to schedule energy storage.

There has been significant work on using energy storage in electric grids (Nourai 2007; Nunna and Doolla 2013; Datta and Senjyu 2013; Hill et al. 2012; Jiang and Powell 2015). However, the majority of work has focussed on improving grid stability or performing cost arbitrage. In this paper, we focus on using energy storage to reduce greenhouse emissions from electric grids by shifting the electricity demand from high polluting periods to low polluting periods. To the best of our knowledge, we present the first work which explicitly aims at reducing greenhouse emissions while maintaining grid stability.

Additionally, shifting our energy demand has been sug-

[1]We use marginal carbon intensity and change in battery state compared to the previous day to calculate the daily carbon savings.

gested in the literature by introducing flexibility in loads through a mechanism called *demand response* (Gnauk, Dannecker, and Hahmann 2012; Scott et al. 2013). Monetary incentives are set aside to compensate the customers participating in demand response. However, demand response involves customer buy-in and often include installing specialized hardware on the electric loads, which may not always be feasible. On the contrary, grid operators around the world can readily employ our approach by utilizing carbon intensity values from the set of power plants they control.

Conclusion

Our greenhouse emissions vary with the mix of fuel types used in the electricity generation infrastructure at any given time. In this work, we presented an algorithm that leverage distributed energy storage to shifting electric demand from high polluting time to low polluting time. Initially, we proposed two neural network based day-ahead time series regression models and showed how they outperform state-of-the-art technique in load forecasting. We also employed a linear programming-based distributed energy storage scheduling scheme using the load forecasts while maintaining various grid constraints. We evaluated our approach on a large-scale dataset containing 776 transformers servicing over 9495 electric meters in a city in the New England region of the US. Our analysis showed a reduction of >10 million pounds in annual carbon emissions, which is equivalent to a 14.48% reduction in our electric grid emissions.

References

2017. Global Storage Market to Double Six Times by 2030. https://about.bnef.com/blog/global-storage-market-double-six-times-2030/. Accessed: 2018-08-08.

Brockwell, P. J.; Davis, R. A.; and Calder, M. V. 2002. *Introduction to time series and forecasting*, volume 2. Springer.

Datta, M., and Senjyu, T. 2013. Fuzzy control of distributed pv inverters/energy storage systems/electric vehicles for frequency regulation in a large power system. *IEEE Transactions on Smart Grid* 4(1):479–488.

De Livera, A. M.; Hyndman, R. J.; and Snyder, R. D. 2011. Forecasting time series with complex seasonal patterns using exponential smoothing. *Journal of the American Statistical Association* 106(496):1513–1527.

Diaconescu, E. 2008. The use of narx neural networks to predict chaotic time series. *Wseas Transactions on computer research* 3(3):182–191.

2018. EIA Data Browser. https://www.eia.gov/electricity/data/browser/. Accessed: 2018-08-08.

Gnauk, B.; Dannecker, L.; and Hahmann, M. 2012. Leveraging gamification in demand dispatch systems. In *Proceedings of the 2012 Joint EDBT/ICDT workshops*, 103–110. ACM.

Hill, C. A.; Such, M. C.; Chen, D.; Gonzalez, J.; and Grady, W. M. 2012. Battery energy storage for enabling integration of distributed solar power generation. *IEEE Transactions on smart grid* 3(2):850–857.

Hoedl, S. 2016. Location-Based ISONE Pollutant Intensity Analysis. Technical report, Harvard University's Office for Sustainability.

2018. ISO New England Web Services API v1.1. https://webservices.iso-ne.com/docs/v1.1/. Accessed: 2018-08-08.

Iyengar, S.; Lee, S.; Irwin, D.; and Shenoy, P. 2016. Analyzing energy usage on a city-scale using utility smart meters. In *Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments*.

Jiang, D. R., and Powell, W. B. 2015. Optimal hour-ahead bidding in the real-time electricity market with battery storage using approximate dynamic programming. *INFORMS Journal on Computing* 27(3):525–543.

Kong, W.; Dong, Z. Y.; Jia, Y.; Hill, D. J.; Xu, Y.; and Zhang, Y. 2017. Short-term residential load forecasting based on lstm recurrent neural network. *IEEE Transactions on Smart Grid*.

Kyriakides, E., and Polycarpou, M. 2007. Short term electric load forecasting: A tutorial. In *Trends in Neural Computation*. Springer. 391–418.

Mishra, A.; Irwin, D.; Shenoy, P.; Kurose, J.; and Zhu, T. 2012. Smartcharge: Cutting the electricity bill in smart homes with energy storage. In *Proceedings of the 3rd International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet*, 29. ACM.

Nguyen, H., and Hansen, C. K. 2017. Short-term electricity load forecasting with time series analysis. In *Prognostics and Health Management (ICPHM), 2017 IEEE International Conference on*, 214–221. IEEE.

Nourai, A. 2007. Installation of the first distributed energy storage system (dess) at american electric power (aep). *Sandia Report SAND2007-3580*.

Nunna, H. K., and Doolla, S. 2013. Energy management in microgrids using demand response and distributed storage - a multiagent approach. *IEEE Transactions on Power Delivery* 28(2):939–947.

Rogers, M. M.; Wang, Y.; Wang, C.; McElmurry, S. P.; and Miller, C. J. 2013. Evaluation of a rapid lmp-based approach for calculating marginal unit emissions. *Applied energy* 111:812–820.

Scott, P.; Thiébaux, S.; Van Den Briel, M.; and Van Hentenryck, P. 2013. Residential demand response under uncertainty. In *International Conference on Principles and Practice of Constraint Programming*, 645–660. Springer.

Taieb, S. B.; Yu, J.; Barreto, M. N.; and Rajagopal, R. 2017. Regularization in hierarchical time series forecasting with application to electricity smart meter data. In *AAAI*, 4474–4480.

Walawalkar, R.; Apt, J.; and Mancini, R. 2007. Economics of electric energy storage for energy arbitrage and regulation in new york. *Energy Policy* 35(4):2558–2568.