

Emission-aware Energy Storage Scheduling for a Greener Grid

Abstract—Reducing our reliance on carbon-intensive energy sources is vital for reducing the carbon footprint of the electric grid. Although the grid is seeing increasing deployments of clean, renewable sources of energy, a significant portion of the grid demand is still met using traditional carbon-intensive energy sources. In this paper, we study the problem of using energy storage deployed in the grid to reduce the grid’s carbon emissions. While energy storage has previously been used for grid optimizations such as peak shaving and smoothing intermittent sources, our insight is to use distributed storage to enable utilities to reduce their reliance on their less efficient and most carbon-intensive power plants and thereby reduce their overall emission footprint. We formulate the problem of emission-aware scheduling of distributed energy storage as an optimization problem, and use a robust optimization approach that is well-suited for handling the uncertainty in load predictions and intermittent renewables. We evaluate our approach using a state of the art neural network load forecasting technique and real load traces from a distribution grid. Our results show a reduction of >0.5 million kg in annual carbon emissions — equivalent to a drop of 23.3% in our electric grid emissions.

I. INTRODUCTION

A key sustainability goal of the UN is to attain a zero carbon economy in order to tackle climate change, while maintaining our current standards of living. Doing so involves immense challenges, since it requires changing our energy consumption behavior, while also transitioning the electric grid to carbon-neutral or zero-carbon energy sources. Over the last decade, there has been an increasing deployment of clean, renewable energy sources such as solar and wind that are already contributing positively to reducing the grid’s overall carbon footprint. The leveled cost of these renewable technologies is now on par or below traditional carbon-intensive generation sources, and their carbon footprint is near zero.

However, due to their intermittent nature, the increasing penetration of these energy sources has increased the stochasticity and uncertainty in grid energy supply. Consequently, energy storage has emerged as a related grid technology to counter this stochasticity [1]. Energy storage batteries can act as “energy buffers” that smooth out the intermittent supply from renewable sources. The cost of energy storage has continued to fall, much like that of renewables, and their deployments have begun to increase. For instance, Green Mountain Power, a small utility in Vermont, USA, now leases Powerwall batteries to residential customers for just \$15/month, while allowing utility control of the battery during peak periods [2]. Such a distributed deployment of energy storage with utility control forms a type of Virtual Power Plant (VPP) that the utility can control for various grid optimizations.

Much of the recent work on energy storage-driven grid optimization has focused on *demand-side* optimizations such as cost arbitrage [3], peak load shaving, demand response [4], and ancillary services [5]. Peak load shaving is a grid optimization of particular interest to utilities and involves operating batteries during peak demand periods in order to reduce the grid stress and the reliance on peaking power plants that are operated solely to meet peak load. Peak shaving brings economic and cost benefits, since peaking power plants tend to be less efficient and hence the cost of supplying electricity during peak periods is much higher than at other times.

Although energy storage-based peak shaving has been studied from a cost reduction perspective, it also brings implicit greening benefits—peaking power plants are not only less efficient and costly to operate, they come with a high pollution and carbon cost. Despite the implicit greening benefit from reducing the use of peaking power plants, the problem of peak load reduction using energy storage *does not* directly translate to the problem of reducing the grid’s carbon footprint. This is because not all peak demand is met using “dirty” peaking power plants. In some cases, for instance, peak demand can be met using pumped hydro storage, which is a clean energy source, and operating energy storage batteries during such periods will not yield any emission reductions.

Thus, reducing the grid’s carbon footprint can not be achieved by naïvely using prior methods on energy storage-based peak load reduction. This problem of emission reductions, a *supply-side* optimization, is not only different, but also more challenging than peak load reduction. Since grid demand is directly observable, energy storage can be activated when peak demand occurs. Unlike grid demand, the grid’s carbon emissions are not directly observable and must be inferred through other means, which is a pre-requisite for scheduling energy storage whenever the grid’s emission footprint peaks.

The use of energy storage for explicitly optimizing the emission footprint of the grid has not been considered by prior work, with the sole exception of [6], where it was considered as part of a broader multi-objective optimization to reduce cost, emissions, etc., and by only considering small residential scale storage. Our work is more general since it addresses grid-scale storage at various levels of the grid network, and also more specific since it focuses on reducing emissions as a primary objective. Our work is motivated by the observation that a utility typically uses a mix of generation sources to fulfill its daily demands. Different generation sources have different cost and emission footprints—for example, while coal, oil and natural gas have high emission footprints, sources such

as nuclear, hydro and solar have zero emissions. The cost of generation also varies across these generation sources.

Utilities typically create a dispatch schedule that determines the order in which different generation sources are utilized to meet rising demand—more efficient energy sources are used more often or as base sources, while less efficient ones are often used only during high demand or peak demand periods. Our insight is that these dispatch schedules and marginal analysis of energy prices can be used to infer the carbon cost to produce the next unit of electricity at a particular demand level. Since demand is observable, we can combine this information with time-varying demand to compute the overall emission footprint at different points in the demand curve and then intelligently activate energy storage whenever the emissions footprint is high; as noted, the emissions footprint depends on the energy fuel sources used and not on the demand, yielding a different schedule for operating energy storage than that from peak shaving. Such emission-aware scheduling of energy storage can provide significant benefits in greening the grid.

However, there are many challenges in designing algorithms for emission-aware scheduling of energy storage. First, the daily electricity demand is stochastic and time-varying and also depends on weather conditions. Second, as the penetration of “clean” renewables increases in the grid, it results in lower emission footprint for the generation mix but also increases uncertainty in supply due to intermittency of generation. Third, the emission footprint of various sources, represented by the marginal cost of generation, is itself time-varying due to changes in prices of inputs and other factors. Finally the energy storage deployment will be distributed and heterogeneous with batteries of various sizes and technologies deployed in different parts of the grid.

In this paper, we leverage robust optimization [7], [8] (RO) to tackle the uncertainty of the daily electricity demand. Classic stochastic optimization approaches require stochastic modeling of uncertain parameters, and deviations from the models, may degrade the performance of the proposed solutions substantially. In contrast, RO does not rely on an underlying probability distribution of the uncertain input, and only requires a limited information of the uncertain data, including mean, and interval predictions, i.e., upper and lower bounds of the uncertain data. When compared to probability distributions, mean and interval prediction values are much simpler to estimate. In addition, RO always calculates a solution that is *guaranteed* to be feasible within all possible realizations of the uncertainty sets. Note that RO and competitive algorithm design [9] are two approaches in the literature that do not rely on any stochastic modeling of the uncertain data. While a competitive approach is too conservative since it guarantees worst-case performance, the additional interval prediction data can result in better performance in RO. In designing and evaluating our emission-aware storage scheduling approach, our paper makes the following contributions.

Problem Formulation. We present a detailed formulation of the problem of emission-aware storage scheduling as an optimization problem. Our formulation is sufficiently general

to incorporate a range of possibilities, including heterogeneous storage deployments, distributed renewable generation, and time-varying marginal costs.

Emission-aware distributed energy storage scheduling using robust optimization. We use robust optimization, a state of the art stochastic optimization approach, to solve the energy-aware storage scheduling problem. As noted above, the use of robust optimization allows us to find a solution that is guaranteed to be feasible within all the possible realizations of the input in a predetermined uncertainty set.

Grid-scale evaluation. We present a grid-scale evaluation of our approach using real traces from a distribution grid comprising 100 transformers. Since our optimization approach requires load predictions, we also use a state of the art autoregressive neural network algorithm for transformer load forecasting, which is then utilized by our optimization approach. Our results show carbon emissions savings of >0.5 million kg over a period of a year. This reduction is equivalent to 23.3% of overall emissions from the electric grid. We also show that even at 50% storage penetration level we can achieve up to 13.87% reduction in carbon emissions.

II. BACKGROUND

In this section, we provide background on electric grids, generation sources, and energy storage.

A. Electric grids

As is well known, today’s electric grids comprise three components: generation, transmission, and distribution (see Figure 1). A key goal of the grid operator is to ensure that demand and supply are matched at all times for proper functioning of the grid. Since electricity demand changes continuously over the course of a day (see Figure 2(a)), the generation must be matched to changing demand via a *dispatch schedule* [10]. The dispatch schedule determines the order and schedule for activating and deactivating various generators that are at the disposal of the grid operator and specifies how the supply should be ramped up or down to match time-varying demand.

Typically, power plants and generators used for the dispatch schedule fall into three categories: (i) *Base load generators*: These are generators at power plants that operate at all times to support the base demand; generators at large-scale power plants such as nuclear, natural gas, coal, and biomass fall into these categories; (ii) *Load following generators*: These are generators that are activated during the high demand periods within each day (such as morning and evening) to support demand beyond the base load; (iii) *Peaking generators*: These are standby generators that are activated when the overall demand hits seasonal peaks. They may operate for only a few days of the year when the hot or cold weather causes the demand to peak for the season. In general, peaking generators tend to be older, less-efficient generators within the overall mix that are kept on standby for infrequent use; old coal and oil generators that are nearing the end of their lifetime are examples of peaking generators. Note that high or peak




Fig. 1. Electric grid architecture comprising generation, transmission and distribution. Our work assumes a heterogenous energy storage within the distribution grid.

demands can also be met through other means, such as pumped hydro storage, and hence, *the emission footprint may not always rise with demand*.

In contrast to traditional sources of electricity generation, renewable sources such as solar, wind and hydro are non-polluting in nature and have zero carbon emissions. Distributed renewable energy sources such as solar tend to be part of the distribution network and often net-meter their power output directly into the distribution grid. Further, renewable sources such as solar and wind are assumed to be uncontrollable due to their intermittent nature and thus not dispatchable.

B. Emission from generation sources

As discussed below, the carbon intensity, and the resulting emission footprint, of the grid varies continuously over the course of the day. If the emission footprint from generation were directly observable, we could simply schedule energy storage whenever emissions peak during each day. Since the emission footprint is not directly observable, we need to infer it through other means for our emission-aware scheduling approach. Two factors need to be considered for doing so: the *average carbon intensity* and *marginal carbon intensity*.

The *average carbon intensity* of an electric grid is defined as the weighted average of *emission factors* of the available fuel types, in which *emission factor* of each fuel type is defined as its carbon emission by generating one unit of electricity. Table I lists the values of emission factor for the available generation types in ISO New England [11]. The average carbon intensity is the weighted average of emission factors for the energy mix used by the grid. For example, if an electric grid produces electricity from coal, natural gas, nuclear, and hydro in equal proportions, then the average carbon intensity would be 339.49 kg/MWh ($962.97 \times 0.25 + 395.53 \times 0.25 + 0 \times 0.25 + 0 \times 0.25$), in the above example.

In general, however, the reduction or increase in generation, and consequently, carbon emission, due to changes in electric demand (dictated by the dispatch schedule) is not the same across all power plants. Most of the changes occur in the load following power plants, and occasionally in the peaker power plants, which we collectively refer to as *marginal power plants*. These are generators that can be ramped up or down at short notice to respond to changes in demand as determined by the dispatch schedule. Consequently, when

TABLE I
CARBON EMISSION BY DIFFERENT GENERATION TYPES, DATA PUBLICLY AVAILABLE FROM [11]

Generation Type	Emission Factor (CO ₂ kg/MWh)
Coal	962.97
Natural Gas	395.53
Oil	933.94
Nuclear	0
Hydro	0
Solar and Wind	0

attempting to reduce the emissions from the generation mix, we must consider the *marginal carbon intensity*, which is the emissions from generating the next unit of electricity; in our case, it is related to the operation marginal power plants. In the above example, if the marginal power plant uses natural gas as fuel, the marginal carbon intensity is 395.5 kg/MWh, which is higher than the average carbon intensity in our example. Since there are notable differences in the emission factors associated with the different fuel types, there is significant potential for reducing the carbon footprint of the overall electricity generation by optimizing the marginal carbon intensity through the use of energy storage.

It should also be noted that the marginal carbon intensity will vary over time due to several factors. For example, if the generation from hydro plants has to be decreased during periods of little rain, generation from other (less green) sources will have to make up the shortfall, potentially increasing the marginal carbon intensity. Fuel prices of raw materials such as natural gas and oil may fluctuate over time, and dispatch schedules may be optimized to use cheaper sources. The dispatch schedule itself varies over the course of a season based on seasonal demand. All of these factors cause the marginal carbon intensity to vary, and any approach that seeks to optimize marginal emissions must account for such temporal variations. Figure 2(b) illustrates the variations of the marginal carbon intensities of different fuel types.

C. Energy Storage

Grid energy storage technologies in the form of batteries have been gaining traction in recent years. Companies such as Tesla have deployed both small- and large-scale energy storage within the grid in many different countries and demonstrated




Fig. 2. (a) Daily transformer load (b) Marginal factor of different fuels for a day.

the feasibility and benefits of using such storage for grid optimizations. This work assumes a distributed deployment of energy storage batteries within the distribution grid. The deployment is assumed to be heterogeneous — the sizes of batteries and the level within the grid where they are deployed are assumed to be different for different batteries. Some batteries may be small batteries, akin to the Tesla Powerwall, deployed adjacent to small neighborhood distribution transformers. Other batteries may be larger in size and deployed near larger transformers (potentially at the feeder or substation levels) that supply electricity to a larger number of homes. The penetration of energy storage within the distribution grid can differ from one scenario to another, and our work is designed to handle different penetration levels.

The distributed network of batteries is assumed to be under the control of the grid operator. However, rather than using them for grid optimizations, our work seeks to operate this distributed set of batteries to minimize the aggregate carbon emissions of the grid, given the time-varying marginal carbon intensity values—by operating them to reduce reliance on the dispatchable sources with high marginal carbon intensities. We formulate this problem of emission-aware scheduling of energy storage more formally in the next section.

III. PROBLEM FORMULATION

Consider a distribution grid comprising a network of substations, feeders and neighborhood transformers. Assume that energy storage is deployed at various points within this grid, i.e., at a subset of the transformers. In a typical grid, there is significant variation in the capacity of transformers and the number of customers it serves. Thus, the energy storage must be sized according to the transformer capacity to enable grid optimizations at that location. Further, our work assumes that this heterogeneous collection of batteries is under the control of the grid operator and the operator can control the charging and discharging of a distributed network of batteries.

As discussed in the previous section, there are several fuel types in a typical electric grid, each with a different level of emission factor and different time-varying marginal carbon intensity. The emission footprint of the grid at any instant depends on the mix of generation fuels used to satisfy the current demand. Grid operators must match supply and demand at all times by constructing a *dispatch schedule* in

TABLE II
SUMMARY OF NOTATIONS

Inputs	
T	The number of time slots, $T \geq 1$
\mathcal{T}	Set $\mathcal{T} = \{1, 2, \dots, T\}$
n	The number of transformers
m	The number of fuel types
C_i	The capacity of transformer i
B_i	The capacity of storage system at transformer i
ρ_i	Charge and discharge rate limit of storage system i
w_f	Emission factor of fuel type f
$\lambda_f(t)$	marginal factor of fuel type f at t
$l_i(t)$	Mean value of day-ahead forecast load at transformer i at t
Optimization variables	
$x_i(t)$	The charge/discharge amount of storage i at t
$x(t)$	The aggregate charge/discharge at t , i.e., $x(t) = \sum_{i=1}^n x_i(t)$
$b_i(t)$	The storage level (state of charge) of storage i at the end of t

advance. The dispatch schedule is typically computed a day in advance by first estimating the demand curve for the following day and then determining an order in which different generation sources are activated (or deactivated) to meet the predicted rise and fall in demand. This problem becomes a cost minimization problem since the next unit of rise in demand should be satisfied using the generation source with the cheapest marginal price [10]. This problem is also referred to as the unit commitment problem in power system literature [12]; in solving this dispatch schedule (i.e., unit commitment), time-varying marginal prices are computed at each step to select the least cost source for each unit change in demand. Since this problem is solved one day in advance (usually through day-ahead energy markets), we assume that time-varying marginal factors can be obtained when the dispatch schedule is finalized at the start of each day.

The key insight behind our approach is to take the demand seen within the distribution grid (which is directly observable) and these computed marginal prices at different times of the day to infer the emission footprint of the grid over time; the scheduling problem is then to intelligently schedule the batteries during peak emissions periods, subject to various demand constraints within the distribution grid. This leads us to the following problem: *Given the dynamics in the marginal factor of available fuel types and in the distribution-level demand, what is the optimal scheduling of energy storage that minimizes grid-wide carbon emissions and respects the operational constraints of the grid and energy storage systems?*

A. System Model

In this section, we formulate the *offline* version of the emission-aware storage scheduling problem (EASS) assuming that the entire load data is available in advance, and in the next section, we present the *online* formulation that takes into account the uncertainty of load.

We assume that the time horizon is divided into T real-time settlement intervals, indexed by t , each with fixed length. Time

slots are set according to the real-time settlement intervals in the U.S.-based electricity markets, e.g., 5 minutes in CAISO and NYISO, and 15 minutes in ERCOT [13]. The main notations are summarized in Table II. In what follows, we explain the details of the system model.

Assume there are n transformers in the system, each indexed by i . Further assume that there is an energy storage battery at each transformer. In practice, the operator may only deploy energy storage at a subset of transformers, which can be easily modeled by setting the sizes of batteries at all other transformers to zero. The scheduling decisions are assumed to be made at the transformer level. Let C_i be the capacity of transformer i , and B_i be the storage capacity deployed at transformer i . Let ρ_i be the maximum charging and discharging rate of storage i . Our formulation could be extended to the case with different charge and discharge rate constraints.

We assume there are m different fuel types in the grid, each indexed by f . Let w_f be the emission factor in kg/MWh of fuel type f . Further, let $\lambda_f(t)$ be the marginal factor, as the contribution of fuel type f at time t to the marginal increase or decrease in energy demand. The values of w_f are fixed and given in advance. The values of $\lambda_f(t)$ change over time based on solving grid dispatch and unit commitment problems.

Finally, let $l_i(t)$ be the day-ahead load forecast of transformer i at time t . In addition, let $\hat{l}_i(t)$ be the actual values of load in real-time. Note that $l_i(t)$ and $\hat{l}_i(t)$ might be different since there is always some error between the forecast day-ahead and actual values. The problem formulation in this section is an *offline* version that takes into account the day-ahead load values. In Section IV, we extend the formulation to include the uncertainty of actual load in real time.

B. Problem Formulation

a) *The Optimization Variables:* The charging/discharging amount of storage i at t is represented as $x_i(t)$. Positive values, i.e., $x_i(t) > 0$ indicate the charging of the storage, whereas negative values, i.e., $x_i(t) < 0$ indicate discharging. The aggregate change in load due to charge/discharge of different storages observed at the grid level at time t is represented by $x(t)$, i.e., $x(t) = \sum_{i=1}^n x_i(t)$. Finally, let $s_i(t)$ be the state of charge of storage i at time t , and we will obtain its evolution over time in the following by formalizing the constraints.

b) *Constraints:* The change in load observed at the grid level is the sum of all the charging/discharging decisions made at each transformer, i.e.,

$$x(t) = \sum_{i=1}^n x_i(t), \quad \forall t, \quad (1)$$

The evolution of storage is represented as the following constraint

$$s_i(t+1) = s_i(t) + x_i(t), \quad \forall t \text{ and } \forall i, \quad (2)$$

The other constraints regarding the physical limit of storage are presented as follows. 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. For simplicity, we assume the maximum charging and discharging rate to be equal. We represent the constraints regarding the state of the charge of storage as follows

$$s_i(t) \leq B_i, \quad \forall t \text{ and } \forall i, \quad (3)$$

$$s_i(t) \geq 0, \quad \forall t \text{ and } \forall i, \quad (4)$$

$$-\rho_i \leq x_i(t) \leq \rho_i, \quad \forall t \text{ and } \forall i. \quad (5)$$

In order to maintain the 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

$$-x_i(t) \leq l_i(t), \quad \forall t \text{ and } \forall i. \quad (6)$$

Note that in reality, it is possible to have peak load beyond the capacity of transformers. 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. Hence, by defining parameter η as the threshold violation level of transformer capacity, we express the this constraint as

$$C_i - x(t) - l_i(t) \geq \eta, \quad \forall i \text{ and } \forall t. \quad (7)$$

The value of η is set by the operator of the grid, and in experiments we set it to 1% of the transformer capacity.

c) *Objective Function:* The eventual goal is to minimize the carbon emission of the grid by managing the charging and discharging of the energy storage. More specifically, we aim to minimize the following objective function:

$$\sum_{t=1}^T \sum_{f=1}^m w_f \lambda_f(t) x(t), \quad (8)$$

where $x(t)$ represents the aggregate change in load by scheduling the storage observed at grid level as described in (1). Recall that $\lambda_f(t)$ represents the marginal factor of fuel type f at time t , and w_f is the emission factor of fuel type f .

d) *Optimization Problem Formulation:* Putting together, we formulate the emission-aware energy storage scheduling (EASS) problem as

$$\begin{aligned} \text{EASS :} \quad & \min \quad \sum_{t=1}^T \sum_{f=1}^m w_f \lambda_f(t) x(t) \\ & \text{subject to:} \quad \text{Equations (1) – (7),} \\ & \text{variable:} \quad x_i(t) \in \mathbb{R}, i \in \{1, \dots, n\}, t \in \{1, \dots, T\}. \end{aligned}$$

The EASS problem is linear in nature that could be solved optimally if the entire input to the problem, i.e., load values and emission parameters, are given in advance. In practice, however, these values are uncertain, and as we will show in Section VI-A, future predictions of load are never 100% accurate. Consequently, in the next section, we introduce the robust optimization formulation to tackle the uncertainty that arises when solving this problem *online* in real-world settings.

IV. ROBUST OPTIMIZATION APPROACH

In this section, we present the robust optimization version of EASS (called EASS-RO) by taking into account the uncertainty due to the imbalance between the forecast and actual real-time load values. Robust Optimization (RO) [8] is a well-established framework for general scenarios of decision making under uncertainty. In this paper, we leverage the RO framework for emission-aware storage scheduling under the uncertainty of electricity load. As compared to the traditional stochastic optimization approaches, problems formulated in an RO framework are typically computationally tractable and do not require the knowledge of a probability distribution over the uncertain input.

The first challenge in formulating an RO counterpart of EASS is to define an *uncertainty set* which bounds the upper and lower bounds that the uncertain input, i.e., load, can take. The classic approach in RO is to optimize for the worst case value in the uncertainty set, but that might be too conservative leading to a suboptimal solution. In this paper, we follow another variant of RO framework, called the price of robustness, proposed in [8]. More specifically, Bertsimas, et. al. [8] develop a generic uncertainty set that can be used to formulate a robust linear counterpart of an uncertain linear program. In this approach, the level of *robustness* can be controlled by parameter Γ known as the *budget of uncertainty*. Then they proved that Γ can be chosen based on the level of robustness desired by the operators such that the probability that the constraint is satisfied is $1 - \epsilon$. In what follows, we present the robust counterpart of the EASS problem using the robust framework in [8],

a) Robust Counterparts of Uncertain Linear Constraints:

In this section, we present the robust counterparts of the constraints that include load values.

First, we state the robust counterpart of constraint (6). Recall that constraint (6) is enforced to ensure that discharge from the battery should be less than the load observed at the transformer. The detailed steps toward stating constraint (6) in a robust framework is the following. First, we construct the uncertainty set associated with the transformer level loads. The uncertainty set for actual load of transformer i at time t can be represented as

$$\mathcal{L}_i(t) = [l_i(t) - \sigma_i(t), l_i(t) + \sigma_i(t)], \quad (9)$$

where $\sigma_i(t)$ is the deviation from the expected value. The values $\sigma_i(t)$ and accordingly $\mathcal{L}_i(t)$ should be obtained by using a forecast model of the transformer level load. In Section V-A, we present our forecasting approach based on state-of-the-art neural networks for predicting the load. The values of $\sigma_i(t)$ will be used to construct the robust constraint.

By defining Γ as the *budget of uncertainty* [8], we re-express constraint (6) as

$$-x_i(t) - l_i(t) + \beta_i(l_i(t), \Gamma) \leq 0 \quad \forall i \text{ and } \forall t, \quad (10)$$

where $\beta_i(l_i(t), \Gamma)$ represents deviation of $l_i(t)$ from its expected value, given Γ as the budget of uncertainty. Note that

Γ could be readily extended to be defined for each transformer separately. In the following, we explain how to calculate the value of $\beta_i(l_i(t), \Gamma)$.

In the original robust optimization framework under the paradigm of price of robustness [8], the budget of uncertainty is defined for each uncertain constraint separately, and its goal is to provide a trade-off between the robustness against the performance of the solution. More specifically, the value of Γ determines that for each constraint how many elements should be robust against violation; the higher the value of Γ , the higher the robustness, the lower the performance. In other words, the solutions with the higher values of Γ might be suboptimal since it is too conservative for the sake of ensuring robustness.

Since constraint (7) is independent for each transformer load at each time slot, following the original approach in [8] requires us to have $n \times T$ separate values robust constraints each for one instance of (6). This approach does not provide any flexibility to determine the level of robustness and limits us to the case with the maximum robustness in solution. To provide flexibility for trade-off between robustness and performance, we slightly change the original framework by considering a common budget of uncertainty for the entire time horizon of each transformer, i.e., grouping all the constraints of each transformer over time.

More specifically, let $\mathbf{z}^* = [z_i^*(t)]_{t \in \{1, 2, \dots, T\}}$ be the optimal value of the following optimization problem [8, Section 3, Proposition 1]:

$$\mathbf{z}^* = \arg \max_{z_i(t) \in [0, 1]} \sum_{t=1}^T \sigma_i(t) z_i(t), \quad \text{s.t.} \quad \sum_{t=1}^T z_i(t) \leq \Gamma. \quad (11)$$

Note that the above problem should be solved for each transformer i separately. Then, we calculate the values of each $\beta_i(l_i(t), \Gamma)$ as follows:

$$\beta_i(l_i(t), \Gamma) = \sigma_i(t) z_i^*(t).$$

Using this formulation one can see that $\beta_i(l_i(t), \Gamma)$ ranges between $[0, \sigma_i(t)]$, thus $l_i(t)$ ranges between $[0, l_i(t) + \sigma_i(t)]$. We can safely ignore the set $[l_i(t) - \sigma_i(t), l_i(t)]$ as we want to make our solution robust to the worst case scenario. Using the budget of uncertainty parameter Γ we can control the level of robustness across time in terms of the following optimization problem.

Intuitively, we can see how the value of Γ controls the deviation from $l_i(t)$. In case we set $\Gamma = 0$ the EASS-RO problem reduces to EASS as all the deviations from $l_i(t)$ will be 0. On the other hand, increasing Γ increases the deviation of $l_i(t)$ thereby we have more robustness in the solution. When $\Gamma = |T|$ each $z_i(t)$ will be set to 1, thus EASS-RO is the most conservative formulation. Last, in EASS, constraint (7) is involved with the uncertainty of the load. Hence, the same procedure as for constraint (6) should be done to have its robust counterpart.

V. EVALUATION SETUP

In this section, we discuss our experimental setup and methodology.

A. Load forecasting

Our approach requires load forecast as input and internally deals with its uncertainty. While there are have been several research on forecasting demand [14], most approaches focus on predicting the aggregate grid demand, which is often smooth and predictable. However, transformer load sees higher variations depending on the number of homes the transformers feed electricity [15]. As such, it is more challenging to provide accurate forecasts and has higher uncertainty in prediction. Formally, forecasting transformer loads requires learning a function F_i that predicts future loads based on input parameters stated as follows

$$l_i(t+1, t+2, \dots, t+k) = F_i(l_i(1, 2, \dots, t), \tau) \quad \forall i.$$

where F_i predicts future load for the next k time steps, and τ is a vector that represents exogenous feature inputs such as temperature, day of the week.

Load at the transformer level show both diurnal and weekly patterns [15]. For example, load during mid-day will differ from load seen at night. Similarly, weekday load differs from weekend load patterns. We use this insights to model our load. Specifically, our approach is based on an *Autoregressive Neural Network* [16]. To forecast load for time $t+1$, in addition to temperature and day of the week, we use past m loads seen at previous time slot, previous day load at time $t+1$, and previous week load at time $t+1$ as input features. We also use one-hot encoding to capture the day of the week in the input feature vector.

B. Experimental Datasets

a) Load Dataset: For evaluating the efficacy of our load forecasting techniques along with the distributed storage schedule, we use a grid-scale dataset obtained from an utility company in the Northeastern US containing energy data from 1,341 smart meters connected to 100 transformer. This data is available at a 5-minute granularity over a period of 2 years. On an average, each transformer is connected to 13.4 smart meters (ranging from 5 to 85). Likewise, the transformer capacity varies between 25 and 750 kVA.

b) Marginal Carbon Intensity: Additionally, our scheduling scheme requires marginal carbon intensity as the input to the problem. This data is not directly available for the New England region. Hence, we use the method specified in [17] that estimates the marginal power plants in operation using the hourly locational marginal price (LMP) of electricity generation and the monthly fuel prices available through [18] and [19]. The approach uses a symmetric Gaussian membership function in (12) that maps the LMP values to fuel types.

$$M_f(t) = e^{-\frac{(p(t) - \mu_f)^2}{2\nu_f^2}} \quad (12)$$

TABLE III
PARAMETER SETTINGS OF OUR APPROACH.

Parameters	Value
Charge/Discharge Rate Limit	60 mins
Marginal Fuel Sources	Coal, Oil, Gas
Emission Factor (kg/MWh)	Refer to Table 1
η in Equation (7)	1% of C_i
Γ in EASS-RO	[10,20]

where μ_f and ν_f^2 is the average cost and variance of fuel type f , and $p(t)$ is LMP of the market at time t . Subsequently, we compute the marginal factor $\lambda_f(t)$ of fuel type f as follows.

$$\lambda_f(t) = \frac{M_f(t)}{\sum_f M_f(t)} \quad (13)$$

C. Experimental Settings and Baseline

a) Parameter Settings: We set the time horizon to one day, and the length of each slot is 5 minutes, hence $T = 12 \times 24 = 288$. We initialize the storage level 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, i.e.,

$$s_i(t) = \frac{B_i}{2}; \quad \text{if } t = 1 \text{ and } t = T = 288, \forall i.$$

While evaluating our robust approach, we use load values and fuel type parameters directly read from the dataset described above. The additional parameters are described in Table III.

1) Baseline Algorithms: We compare the performance of EASS-RO with the following approaches.

- 1) *Optimal Offline Solution:* The optimal offline approach assumes complete knowledge of future load and provides the best achievable schedule to minimize carbon emissions. Although not practical, it serves as a best theoretical upper bound to compare with.
- 2) *Online Linear Programming:* In this approach, we use the day-ahead forecast load as input to solve the linear program EASS and determine the schedule. However, the day-ahead charge/discharge schedule may violate real-time grid constraints as the actual load at time t may differ from predicted load at time t . To ensure that all grid constraints are satisfied, the day-ahead schedule is adjusted as follows. Let $\hat{l}_i(t)$ be the actual load observed at time t , and $\hat{x}_i(t)$ be the modified storage charge/discharge value at t to ensure feasibility.

Transformer Constraints

$$\hat{x}_i(t) = 0, \quad \text{if } C_i - x_i(t) - \hat{l}_i(t) \leq \eta \quad (14)$$

$$\hat{x}_i(t) = C_i - \hat{l}_i(t), \quad \text{if } x_i(t) \geq C_i - \hat{l}_i(t) \quad (15)$$

Storage Constraints

$$\hat{x}_i(t) = B_i - s_i(t), \quad \text{if } x_i(t) + s_i(t) \geq B_i \quad (16)$$

$$\hat{x}_i(t) = -s_i(t), \quad \text{if } x_i(t) + s_i(t) \leq 0 \quad (17)$$




Fig. 3. Efficacy of load forecasting methods. The figure shows there can be significant uncertainty in forecasts.

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

$$\hat{x}_i(t) = -\hat{l}_i(t), \text{ if } \hat{l}_i(t) \leq -x_i(t). \quad (18)$$

- 3) *The PreDay Algorithm*: This approach uses the previous day's load and emissions factor as input to the linear program to determine the emission-aware schedule. We use a similar approach and modify the schedule as above to ensure that constraints are not violated.

VI. EXPERIMENTAL RESULTS

In this section, we evaluate our approach and compare it to the optimal approach and other heuristic approaches.

A. Load Forecasting Uncertainty

First, we evaluate the efficacy and accuracy of our proposed load forecasting method in Section V-A. We compare our forecasting method based on state-of-the-art neural network approach with two popular statistical time series techniques — ARIMA [20] and TBATS [21]. Figure 3 compares the performance of the proposed regression technique with the two baseline approaches. The results show the distribution of mean absolute percentage error (MAPE) values for load forecast at all transformers evaluated over a period of one year. Based on our analysis, TBATS has the highest average MAPE of 34.17%, while the MAPE of ARIMA was 21.5%. The performance of the autoregressive neural network with exogenous variables outperformed all other techniques and has the lowest average MAPE value of 20.14%. In our experiments, we observed that including exogenous variables improves the accuracy of our forecast significantly. Despite its higher accuracy, we observe that the forecast still contains error — indicating uncertainty in prediction. The presence of such error is a motivation for leveraging robust optimization for emission-aware storage scheduling.

B. Emission-aware Storage Schedule

Figure 4 depicts our emission-aware storage scheduling approach in action. The figure shows the impact of the 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, e.g., between 6 am to 9 am, which represents the high polluting hours of the day. Conversely, charging occurs when




Fig. 4. Battery charge and discharge based on our emission-aware energy schedule. Our emission-aware algorithm discharges battery when marginal carbon intensity (MCI) is high.

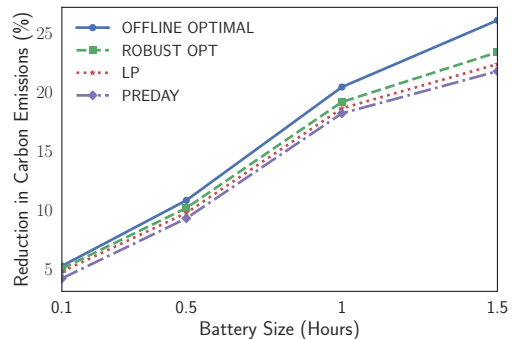


Fig. 5. Carbon emissions reduction for different battery sizes. The battery size is computed as the number of hours it can sustain the maximum load of the transformer.

the marginal emissions are low, usually between 1 to 4 pm. 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 17.5 kg of carbon emissions. Thus, emissions-aware distributed energy storage has significant potential to reduce carbon emissions at the grid-scale.

C. Benefits of Emission-aware Scheduling

We analyze the change in carbon savings¹ by varying the size of the energy storage. We size the battery as the number of hours it can support the annual maximum load of the transformer. Thus, one hour of battery capacity indicates that it can support the maximum load of the transformer for an hour. Figure 5 shows the reduction in carbon emissions with increasing battery size for different algorithms. We observe that the carbon emissions reduces with increasing battery size. This is because a larger battery has more flexibility in shifting transformer load, where batteries can charge during lower emissions and discharge during high carbon emissions period. Even with a battery size of 0.5 hours, we observe that our emissions-aware algorithm achieves 10.16% reduction in carbon emissions. Further, a battery size of 1.5 hours

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

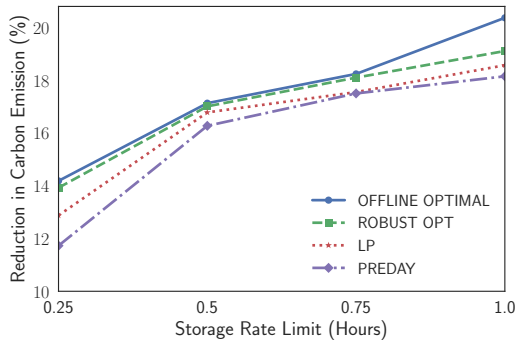


Fig. 6. Carbon emissions reduction for different charge/discharge rate and a battery size of one hour.

can annually save >0.5 million kg of carbon emissions — equivalent to a reduction of 23.3% in electric grid emissions.

We also compare our robust optimization with baseline approaches described in Section V-C1. The optimal approach provides the maximum carbon savings that can be achieved. However, the optimal needs the full information in advance that is not practical. We observe the gap between the optimal and our robust optimization approach is less than 1.2% having battery size less than or equal to one hour. Further, we note that robust optimization *consistently* performs better than the other baseline approaches.

D. Impact of Storage Parameters

So far, we have assumed that the energy storage can discharge at the maximum load of the transformer, i.e., no charge/discharge rate limit. Next, we study the effects of different rate limits on carbon emission reduction as shown in Figure 6. We fix the energy storage size such that it can sustain the maximum load at the transformer for one hour and vary the charge/discharge rate. The charge/discharge rate is set such that the fraction indicates the percentage of the maximum load at the transformer the energy storage can charge or discharge. Thus, a 0.25 hour charge/discharge rate can discharge at one-fourth the maximum load at the transformer. As seen in the figure, with a charge/discharge rate of 0.25 hour, our robust optimization approach achieves a carbon emission reduction of 13.9%. However, an increase in charge/discharge rate further reduces carbon emissions. This is because a higher discharge rate is able to reduce demands thereby minimizing the need to utilize generation sources with high emission footprints. In particular, we observe that the reduction in carbon emissions increases by 37.2% (from 13.93% to 19.12%) when the rate is increased from 0.25 to 1 hour.

E. Impact of Storage Penetration

We study the benefit that comes from installing energy storage at only a fraction of the transformers in the grid. Like before, in this experiment, we fix the energy storage size such that it can sustain the transformer load at its maximum for one hour and select transformers at random, where batteries are installed. Figure 7 shows the reduction in carbon emissions

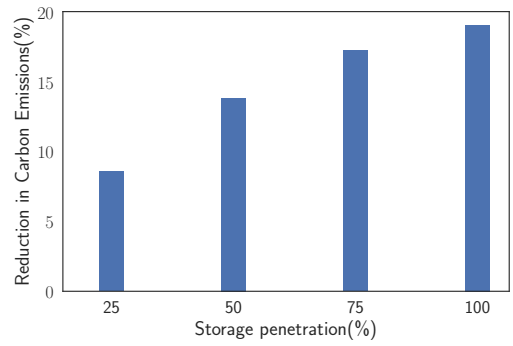


Fig. 7. Carbon emissions reduction for different levels of storage penetration across transformers. A 50% penetration indicate half the transformers have storage units installed.

for different storage penetration levels. An energy storage penetration of 25% can achieve 8.5% carbon emission reduction. However, if 50% of the transformers install energy storage, the carbon emission reductions improves to 13.9%, a 63.5% improvement in emissions reduction. This is because higher energy storage penetration can offset more loads that have high emissions footprint. Further, if all the transformers have energy storage, the reduction in carbon emission is 19.12%.

VII. RELATED WORK

Energy Storage Systems in the Electric Grid. There has been significant work on using energy storage in the electric grid [22], [23]. However, the majority of work has focused on improving grid stability or cost arbitrage. This paper focuses on using energy storage to reduce carbon emissions of the grid by shifting the demand from high polluting periods to low polluting periods.

Additionally, shifting the energy demand has been suggested in the literature by introducing flexibility in loads through a mechanism called *demand response* [24], [25]. Monetary incentives are set aside to compensate for 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.

Load Forecasting. In smart grids research, load forecasting is a widely studied problem. The regression techniques used to solve this problem range from traditional time series approaches such as ARIMA [26] to neural network [27]. 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 [27].

Robust Optimization for Scheduling in Smart Grid. Robust optimization has been extensively used to solve different problems in different application domains that deal with uncertainty, including smart grid. Some examples are generator placement [28], [29], EV charging scheduling [30], storage

sizing [31]. As compared to the other stochastic approaches it has several advantages: (1) it does not require stochastic modeling of uncertain parameters in terms of probability distribution functions; (2) by defining the notion of budget of uncertainty [8], it provides a design space to trade-off between the robustness and performance of the decision making. Note that the notion of uncertainty set has been used in other theoretical approaches such as competitive analysis [32], however, the algorithm approach used in [32] is problem specific and cannot be applied to the emission-aware scheduling scenario in this paper.

VIII. CONCLUSION

The benefits of distributed energy storage have been previously studied for grid optimizations such as peak shaving, price arbitrage, and demand-response. However, in this work, we focus on using distributed energy storage to reduce the emission footprint of electricity generation. Our main insight is that energy storage can help utility companies reduce the reliance on less efficient and most carbon-intensive power plants, shifting electric demand from high polluting periods to low polluting periods. We formulated the problem of emission-aware scheduling as an optimization problem with the objective of minimizing the carbon emission, subject to transformer and storage operational constraints. Given the dynamics in transformer-level load, we leveraged robust optimization to handle the uncertainty in load predictions. We evaluated our emission-aware energy storage scheduling approach on a dataset containing 100 transformers connected to over 1,340 electric meters in a city in the Northeastern part of the US. Our analysis showed that our approach can offset >0.5 million kg in annual carbon emissions, which is equivalent to a 23.3% reduction in the electric grid emissions.

REFERENCES

- [1] "Global Storage Market to Double Six Times by 2030," <https://about.bnef.com/blog/global-storage-market-double-six-times-2030/>, 2017, accessed: 2018-08-08.
- [2] "Green mountain power," <https://greenmountainpower.com/product/powerwall/>, 2019.
- [3] R. Walawalkar, J. Apt, and R. Mancini, "Economics of electric energy storage for energy arbitrage and regulation in new york," *Energy Policy*, vol. 35, no. 4, pp. 2558–2568, 2007.
- [4] H. K. Nunna and S. Doolla, "Energy management in microgrids using demand response and distributed storage - a multiagent approach," *IEEE Transactions on Power Delivery*, vol. 28, no. 2, pp. 939–947, 2013.
- [5] C. Ji, M. H. Hajiesmaili, D. F. Gayme, and E. Mallada, "Coordinating distribution system resources for co-optimized participation in energy and ancillary service transmission system markets," in *Proc. of IEEE American Control Conferences*, 2019.
- [6] Z. Olivieri, "Optimization of Residential Battery Energy Storage System Scheduling for Cost and Emissions Reductions," *RIT Scholar Works*, 2017.
- [7] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust optimization*. Princeton University Press, 2009, vol. 28.
- [8] D. Bertsimas and M. Sim, "The price of robustness," *Operations research*, vol. 52, no. 1, pp. 35–53, 2004.
- [9] A. Borodin and R. El-Yaniv, *Online computation and competitive analysis*. Cambridge University Press, 1998.
- [10] Y. Zhang, M. H. Hajiesmaili, S. Cai, M. Chen, and Q. Zhu, "Peak-aware online economic dispatching for microgrids," *IEEE Transactions on Smart Grid*, vol. 9, no. 1, pp. 323–335, 2018.
- [11] S. Hoedl, "Location-Based ISONE Pollutant Intensity Analysis," Harvard University's Office for Sustainability, Tech. Rep., 08 2016.
- [12] N. Padhy, "Unit commitment a bibliographical survey," *IEEE Trans. On Power Systems*, vol. 19, no. 2, pp. 1196–1205, 2004.
- [13] "Wholesale electricity market design initiatives in the united states: Survey and research needs," *EPRI, Technical Results*, available at <https://www.epri.com/pages/product/3002009273/>, 2016.
- [14] E. Kyriakides and M. Polycarpou, "Short term electric load forecasting: A tutorial," in *Trends in Neural Computation*. Springer, 2007, pp. 391–418.
- [15] S. Iyengar, S. Lee, D. Irwin, and P. Shenoy, "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*, 2016.
- [16] E. Diaconescu, "The use of narx neural networks to predict chaotic time series," *Wseas Transactions on computer research*, vol. 3, no. 3, pp. 182–191, 2008.
- [17] M. M. Rogers, Y. Wang, C. Wang, S. P. McElmurry, and C. J. Miller, "Evaluation of a rapid lmp-based approach for calculating marginal unit emissions," *Applied energy*, vol. 111, pp. 812–820, 2013.
- [18] "EIA Data Browser," <https://www.eia.gov/electricity/data/browser/>, 2018, accessed: 2018-08-08.
- [19] "ISO New England Web Services API v1.1," <https://webservices.iso-ne.com/docs/v1.1/>, 2018, accessed: 2018-08-08.
- [20] P. J. Brockwell, R. A. Davis, and M. V. Calder, *Introduction to time series and forecasting*. Springer, 2002, vol. 2.
- [21] A. M. De Livera, R. J. Hyndman, and R. D. Snyder, "Forecasting time series with complex seasonal patterns using exponential smoothing," *Journal of the American Statistical Association*, vol. 106, no. 496, pp. 1513–1527, 2011.
- [22] D. Irwin, S. Iyengar, S. Lee, A. Mishra, P. Shenoy, and Y. Xu, "Enabling distributed energy storage by incentivizing small load shifts," *ACM Transactions on Cyber-Physical Systems*, vol. 1, no. 2, p. 10, 2017.
- [23] C. A. Hill, M. C. Such, D. Chen, J. Gonzalez, and W. M. Grady, "Battery energy storage for enabling integration of distributed solar power generation," *IEEE Transactions on smart grid*, vol. 3, no. 2, pp. 850–857, 2012.
- [24] B. Gnauk, L. Dannecker, and M. Hahmann, "Leveraging gamification in demand dispatch systems," in *Proceedings of the 2012 Joint EDBT/ICDT workshops*. ACM, 2012, pp. 103–110.
- [25] P. Scott, S. Thiébaux, M. Van Den Briel, and P. Van Hentenryck, "Residential demand response under uncertainty," in *International Conference on Principles and Practice of Constraint Programming*. Springer, 2013, pp. 645–660.
- [26] H. Nguyen and C. K. Hansen, "Short-term electricity load forecasting with time series analysis," in *Prognostics and Health Management (ICPHM), 2017 IEEE International Conference on*. IEEE, 2017, pp. 214–221.
- [27] W. Kong, Z. Y. Dong, Y. Jia, D. J. Hill, Y. Xu, and Y. Zhang, "Short-term residential load forecasting based on lstm recurrent neural network," *IEEE Transactions on Smart Grid*, 2017.
- [28] Z. Wang, B. Chen, J. Wang, J. Kim, and M. M. Begovic, "Robust optimization based optimal dg placement in microgrids," *IEEE Transactions on Smart Grid*, vol. 5, no. 5, pp. 2173–2182, 2014.
- [29] Z. Wang, B. Chen, J. Wang, and M. M. Begovic, "Stochastic dg placement for conservation voltage reduction based on multiple replications procedure," *IEEE Transactions on Power Delivery*, vol. 30, no. 3, pp. 1039–1047, 2015.
- [30] N. Korolko and Z. Sahinoglu, "Robust optimization of ev charging schedules in unregulated electricity markets," *IEEE Transactions on Smart Grid*, vol. 8, no. 1, pp. 149–157, 2017.
- [31] F. Kazhamiaka, Y. Ghiassi-Farrokhfal, S. Keshav, and C. Rosenberg, "Robust and practical approaches for solar pv and storage sizing," in *Proceedings of the Ninth International Conference on Future Energy Systems*. ACM, 2018, pp. 146–156.
- [32] M. H. Hajiesmaili, C.-K. Chau, M. Chen, and L. Huang, "Online microgrid energy generation scheduling revisited: The benefits of randomization and interval prediction," in *Proceedings of the Seventh International Conference on Future Energy Systems*. ACM, 2016.