

UC Berkeley

UC Berkeley Previously Published Works

Title

Distributed battery dispatch for uncertainty mitigation in renewable microgrids

Permalink

<https://escholarship.org/uc/item/47w7g7z5>

Authors

Sharma, Sunash B

Lee, Jonathan T

Callaway, Duncan S

Publication Date

2024-10-01

DOI

10.1016/j.epsr.2024.110671

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Distributed Battery Dispatch for Uncertainty Mitigation in Renewable Microgrids

Sunash B. Sharma^{*}, Jonathan T. Lee^{†,‡}, Duncan S. Callaway^{*,‡}

^{*}Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, United States

[†]New Sun Road, P.B.C., Richmond, United States

[‡]Energy and Resources Group, University of California, Berkeley, United States

{sbsharma, jtlee, dcal}@berkeley.edu

Abstract—In this work, we propose and analyze battery-level approaches to mitigate system welfare losses due to total energy forecast errors in energy-constrained microgrids. We present a receding horizon approach, develop a data-based method and heuristic approaches that leverage the receding horizon environment, and then analyze performance on a system based on real data. We find that the data-based method offers modest mitigation of the effects of total energy error, and a heuristic approach based on l_2 regularization offers similar mitigation. We also find that the choice of load utility model affects mitigation strategies’ profitability. Finally, we show how these methods can be used in a distributed framework.

Index Terms—distributed energy resources, distributed optimization, microgrids, power system economics, uncertainty management.

I. INTRODUCTION

Microgrids can improve the reliability of electricity for customers with unreliable grid connections, for example by continuing to operate in an islanded (i.e. stand-alone) mode during central grid outages [1], [2]. They can also serve as an alternative to grid extension in remote areas. Microgrids supplied exclusively by solar plus storage systems have the added benefit of zero point-source emissions and fuel transport costs; they are also becoming cost-competitive on a total cost basis with fuel-based microgrids as solar and storage costs decline. This is because decreasing costs of solar and storage allow a solar plus storage microgrid to be sized with a larger capacity for the same cost. As such, the cost of a solar plus storage microgrid approaches that of a fuel-based microgrid with the same capacity as these component costs decrease. Furthermore, several works have discussed the ongoing reduction of the costs of solar and storage in the context of renewable electrification of rural areas, indicating we can expect solar plus storage microgrid costs to continue to decrease [3], [4], [5], [6]. Many microgrid systems, includ-

ing fully renewable microgrids without fuel-based generation, have already been deployed [7], [8].

However, without a dispatchable generation source powered by a stored fuel supply (such as a gasoline or diesel generator), solar plus storage microgrids are typically “energy-constrained” over daily time horizons. In this case, forecasting energy generation and managing energy supply become essential not only for instantaneous power balance, but also for efficiently allocating total energy to loads. Errors in total energy forecasts could lead to a suboptimal dispatch in which, for example, low-value loads are served in an early time step, while high-value loads are not served in later time steps due to a subsequent unexpected shortfall in total energy. This could lead to consumer dissatisfaction, perceptions of energy mismanagement, and loss of critical loads, which are negative outcomes we model as a loss of welfare.

This paper proposes and analyzes approaches to managing microgrid energy uncertainty, with a focus on battery state of charge management. To enable microgrid modularity, and to preserve the ability of individual asset owners to make independent decisions and maintain privacy, we focus on strategies that can be implemented in a distributed coordination setting. There is an existing body of research in this space. For example, Crespo-Vazquez *et al* [9] consider community-based local energy trading with independent agents scheduling supply and demand under uncertainty in generation and storage forecasts. Their model assumes that uncertainty distributions on solar and storage are known and available to all participants in the grid. Alizadeh *et al* [10] consider participant-level uncertainty mitigation through two-stage stochastic optimization, focusing on managing local uncertainty through local action. Our paper is also relevant to methods that enable batteries to respond to price signals at the grid scale [11], [12], [13]. However, these papers assume batteries can operate as price-takers, which is not generally true in islanded microgrids. To the authors’ best knowledge, the literature has yet to explore battery-level methods to mitigate system-wide uncertainty due to forecast errors on a distributed, energy-constrained microgrid, specifically in the setting in which the batteries’ actions can influence price, but uncertainty distributions are unknown to the batteries. This gap in the literature is what this paper seeks to address.

In this paper, we aim to develop and evaluate methods in

Submitted to the 23rd Power Systems Computation Conference (PSCC 2024). This research used the Savio computational cluster resource provided by the Berkeley Research Computing program at the University of California, Berkeley (supported by the UC Berkeley Chancellor, Vice Chancellor for Research, and Chief Information Officer). This work was supported by Enphase Energy, Inc.

which battery agents do not have access to centralized uncertainty distributions. Rather, we examine a setting in which batteries receive only price forecasts and realizations and take local action to mitigate system-wide (as opposed to local) uncertainty. We consider distributed coordination through a simplified version of the Alternating Direction Method of Multipliers (ADMM), a distributed optimization algorithm that has previously been used to coordinate distributed energy resources [14]. We make the following specific contributions in this work:

- 1) We develop a data-driven approach that can work in a distributed setting to address system welfare losses at the level of an individual battery. In this approach, a battery operator uses historical errors in price forecasts to better predict true prices in real-time while simultaneously accounting for the effects of changes in the battery's own dispatch on price. Depending on the model of load utility used, this approach can be profit maximizing for the individual batteries.
- 2) We also develop heuristic reserve approaches – much like rules of thumb for power reserves that are commonly used on unit commitment problems with non-zero marginal cost generators – in which some fraction of the battery's capacity is either held at a “reserve price cap” for use during unexpected shortfalls in energy supply (which correspond to energy price spikes), or in which the dispatch of this fraction of the battery is l_2 regularized to hedge against uncertainty.
- 3) We compare these methods to a deterministic receding-horizon approach in a simulation environment based on historical residential demand and solar data [15], under different assumptions about load utility and battery energy capacity.

We find that the data-driven approach improves system welfare relative to the deterministic approach, and improves battery profit depending on the load utility model used. Additionally, we find that while the price-cap-based heuristic strategy does not mitigate welfare losses in the system, the regularization-based heuristic mitigates welfare losses in the system to a similar extent as the data-driven approach. This indicates that regularization-based heuristic reserves may be sufficient to address energy uncertainty, with marginal improvements provided by data-based methods.

II. THEORY

In this section, we develop a model representing an energy-constrained microgrid in which batteries can mitigate welfare losses caused by energy forecast errors. First, we define a model for optimal energy allocation among energy consumers over a time horizon and then discuss the introduction of the total energy forecast error into the system. Finally, we consider losses in welfare due to total energy forecast error in a simple case to motivate mitigation in practical cases.

A. Models

First, we formulate a welfare optimization problem for the system and discuss models for the utility of the load, the forecast error, and the inelastic load.

1) *Centralized System Model:* Our dispatch problem can be stated as a welfare-maximizing problem. Ignoring the rest of the system, we define the individual net benefit W_n for an agent n as a function of the energy quantity, which is represented as a vector $\mathbf{q} = \{q_t\}$ of quantities over time $t \in [1, \dots, T]$, where positive implies net consumption and negative implies net supply. Each individual can have a solar array, battery storage, and/or a load. The power consumptions of these resources are denoted \mathbf{p}^s , \mathbf{p}^b , and \mathbf{p}^l , which are all vectors of length T . The stored energy of the battery is s , a vector of length $T + 1$.

Given an allocated net quantity \mathbf{q} , an agent's net benefit is the maximum private benefit that can be achieved by setting decision variables \mathbf{p}^s , \mathbf{p}^b , \mathbf{p}^l , and s so that the power sums to \mathbf{q} and physical constraints are satisfied:

$$W_n(\mathbf{q}) := \max_{\mathbf{p}^l, \mathbf{p}^s, \mathbf{p}^b, s} \sum_{t=1}^T U_{n,t}(p_t^l) \quad (1a)$$

$$\text{s.t. } \nu : \mathbf{p}^l + \mathbf{p}^s + \mathbf{p}^b - \mathbf{q} = \mathbf{0} \quad (1b)$$

$$\lambda^l : \mathbf{0} \leq \mathbf{p}^l \leq \mathbf{P}_n^l \quad (1c)$$

$$\lambda^s : -\mathbf{P}_n^s \leq \mathbf{p}^s \leq \mathbf{0} \quad (1d)$$

$$\lambda^b : -\mathbf{P}_n^{b,-} \leq \mathbf{p}^b \leq \mathbf{P}_n^{b,+} \quad (1e)$$

$$\lambda^c : \mathbf{0} \leq s \leq \mathbf{S}_n \quad (1f)$$

$$\mu_t : s_t - (s_{t-1} + p_t^b \Delta T) = 0 \quad \forall t \in [1, \dots, T] \quad (1g)$$

$$\mu_0 : s_0 - s_{n,0} = 0 \quad (1h)$$

The vectors \mathbf{P}_n^l , \mathbf{P}_n^s , $\mathbf{P}_n^{b,+}$, $\mathbf{P}_n^{b,-}$, \mathbf{S}_n are the maximum load consumption, solar power generation, battery discharge rate, battery charge rate, and battery stored energy for agent n at time t . Note that any of these can be set to $\mathbf{0}$ if there is no load, solar array, or battery, respectively. The domain of W_n is the set of \mathbf{q} for which W_n is feasible given these constraints, denoted $\text{dom}(W_n)$. In the model above, the vectors ν , $\mu = \{\mu_t\}$ and $\lambda = [\lambda^l; \lambda^s; \lambda^{b,+}; \lambda^{b,-}]$ are the dual variables associated with their respective constraints. The parameter ΔT is the time step, the stored energy of the battery is in units of ΔT times the units of power, and the initial stored energy is $s_{n,0}$, a parameter of the problem, rather than a variable. Note that by our sign convention, solar power will be negative, and negative battery charge power is equivalent to the battery discharging into the system.

Given the individual benefit functions, the general optimal dispatch problem of allocating quantity to maximize total welfare is W :

$$W = \max_{\{\mathbf{q}_n\}} \sum_n W_n(\mathbf{q}_n) \quad (2a)$$

$$\text{s.t. } \pi : \sum_n \mathbf{q}_n = \mathbf{0} \quad (2b)$$

$$\mathbf{q}_n \in \text{dom}(W_n) \quad \forall n \quad (2c)$$

Equation (2b) is the power balance constraint for the system. The optimal value of its associated dual variable, π , can be interpreted as the optimal price for power.

2) *Load Utility Model*: We use horizontally shifted constant-elasticity marginal value functions to model the utility of the energy consumed by the load. Agents that have only battery storage and solar generation thus have no intrinsic utility for energy, at least at this stage in the formulation. For each load in each timestep, we construct a load utility function from an assumed price and elasticity, and an observed level of consumption taken from the data described in III-A. The utility function $U_{n,t}$ and the corresponding marginal utility function $g_{n,t}$ are given by

$$U_{n,t}(p_t^l) = \frac{\alpha_n \hat{\pi}_{n,t} \left((p_t^l + q_{\text{shift},t})^{\frac{1}{\alpha_n} + 1} - q_{\text{shift},t}^{\frac{1}{\alpha_n} + 1} \right)}{(\alpha_n + 1) (\hat{p}_t^l + q_{\text{shift},t})^{\frac{1}{\alpha_n}}} \quad (3a)$$

$$U'_{n,t}(p_t^l) := g_{n,t}(p_t^l) = \hat{\pi}_{n,t} \left(\frac{p_t^l + q_{\text{shift},t}}{\hat{p}_t^l + q_{\text{shift},t}} \right)^{\frac{1}{\alpha_n}} \quad (3b)$$

$$q_{\text{shift},t} = \hat{p}_t^l \left(\left(\frac{\hat{\pi}_{n,t}}{\pi_{\text{max}}} \right)^{\alpha_n} - 1 \right)^{-1}, \quad (3c)$$

where α_n is the elasticity of the load utility, $\hat{\pi}_{n,t}$ is the “observed” price of energy at time t , \hat{p}_t^l is the observed load at time t , and π_{max} is the maximum price of energy consumers would pay to serve elastic load. Figure 1 shows $U_{n,t}(p_t^l)$ for different α_n . This utility is based on the utility of consumption described in [16], with a few key differences. First, the shift $q_{\text{shift},t}$ is varying in time rather than constant, allowing $g_{n,t}(0)$ to equal π_{max} . Second, we do not compensate the elasticity α_n for the shift as [16] does. Although this is not the only way to model load utility, the fact that this utility function is strictly concave introduces the possibility of a welfare gap due to energy forecast error, as discussed in Section II-C.

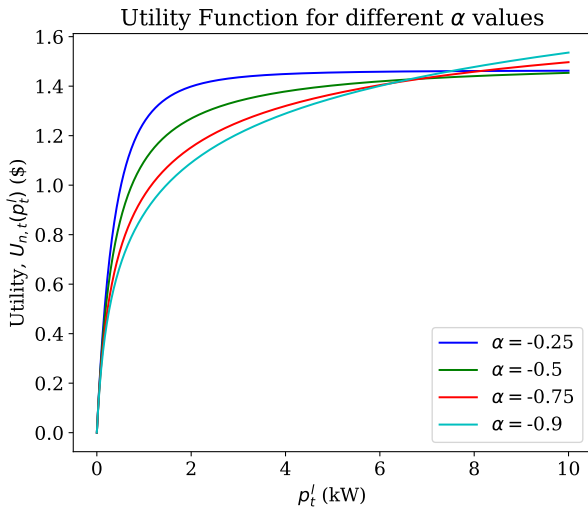


Fig. 1. Utility functions for different elasticity values α . The other parameters are $\hat{\pi}_t = 0.30$, $\pi_{\text{max}} = 4$, $\hat{p}_t^l = 1.0$.

3) *Forecast Error Model*: In the model described in (1), we did not specify how agent constraints and utilities are known for timesteps beyond the current one. We assume that the load and battery agents have perfect knowledge of their own utilities and constraints in the future, but we assume that solar agents use imperfect forecasts calculated once a day to specify $\mathbf{P}_{n,2:T}^s$, their peak generation for all timesteps beyond the current one.

We model solar forecast uncertainty by introducing multiplicative noise to measured solar generation data. For each day, we generate a scaling factor by sampling from a normal distribution with mean 1 and standard deviation σ and apply this to solar generation throughout the day to generate the forecast. Mathematically, the forecasted available solar power at the start of a given day is

$$\mathbf{P}_{n,1}^s = \hat{\mathbf{P}}_{n,1}^s \quad (4)$$

$$\mathbf{P}_{n,2:T}^s = \epsilon \hat{\mathbf{P}}_{n,2:T}^s \quad (5)$$

$$\epsilon \sim \mathcal{N}(1, \sigma^2), \quad \epsilon \geq 0. \quad (6)$$

The solar forecast has no error at the current timestep, and all subsequent solar forecast values are non-negative.

4) *Including Inelastic Load Implicitly*: Many systems have critical inelastic load that must be served before any elastic load, decreasing system flexibility. Because it is inelastic, it has no marginal utility. In a deterministic model, the inclusion of such a load would shift $U_{n,t}(p_{n,t}^l)$ to the right by the inelastic load $P_{t,n,i}^l$ and introduce an extra constraint to (1): $\mathbf{P}_{n,i}^1 \leq \mathbf{p}^1$. However, in a system with forecast error in total energy generation, this constraint may become infeasible due to the misallocation of energy in previous timesteps. To avoid technical modifications to the model to handle infeasibility, rather than formulating inelastic load as a constraint, we instead modify the system utility and forecast error to ensure that inelastic loads are always served first and have zero utility. Additionally, we assume that the system is sized to have sufficient capacity to serve the inelastic loads.

To achieve this, we introduce an inelastic load parameter $\gamma \in [0, 1]$, defined as the fraction of the consumed load that is inelastic. The inelastic load has zero utility when served, so we scale the utility function of the elastic load in (3a) horizontally by $1 - \gamma$:

$$U_{\gamma,n,t}(p_t^l) := U_{n,t}((1 - \gamma)p_t^l) \quad (7)$$

Although this introduces dispatch dependence of inelastic load, which is not technically correct in practice, this formulation approximates the characteristics of inelastic and flexible load.

We can also capture the notion that this load is served with priority in a system sized to always serve it. To capture this, we assume that there is always enough realized solar generation to serve inelastic load, and this solar is automatically and implicitly allocated to inelastic load in the dispatch. Thus, a fraction γ of solar generation is guaranteed to serve the inelastic load and is not forecasted. Therefore, the forecast error in solar generation only affects the utility and consumption of the elastic portion of the load, which is modeled by

applying all the solar forecast error to the remaining $1 - \gamma$ fraction of solar generation after the inelastic load is served. This allows us to build prioritization and sizing for inelastic load directly into the model without adding any additional potentially unsatisfiable constraints.

5) *Including Inelastic Load Explicitly*: Another method to include inelastic load in the consumption model without risking infeasibility follows. In Section II-A2, we used observed consumption \hat{p}_t^l to build a strictly concave load utility function at each timestep t . We may also consider this observed consumption as an estimate of total inelastic load at time t by including the constraint $\hat{p}_t^l \leq p_t^l$. However, to avoid potential infeasibility in the receding horizon dispatch from including this constraint on p_t^l (discussed in II-A4), we may instead penalize load being served below \hat{p}_t^l at some constant rate per kWh. We may conceptualize this rate as a value of lost load (VoLL), cost of fuel for dispatchable generation, or cost of energy from an expensive grid connection.

Specifically, we modify our problem formulation by adding a new variable to the load agent's local problem: \mathbf{p}^v , the vector of amounts of lost load over the dispatch horizon. Then, we modify the load agent's local maximization to include the VoLL:

$$\max_{\mathbf{p}^l, \mathbf{p}^s, \mathbf{p}^b, \mathbf{s}, \mathbf{p}^v} \sum_{t=1}^T U_{n,t}(p_t^l) - \pi_v \Delta T \mathbf{1}^\top \mathbf{p}^v \quad (8)$$

where π_v is our chosen VoLL in \$/kWh. We modify equations 1b and 1c respectively to

$$\nu : \mathbf{p}^l + \mathbf{p}^s + \mathbf{p}^b + \mathbf{p}^v - \mathbf{q} = \mathbf{0} \quad (9)$$

$$\lambda^l : \hat{\mathbf{p}}^l \leq \mathbf{p}^l \leq \mathbf{P}_n^l \quad (10)$$

and add the constraint $\lambda^v : \mathbf{0} \leq \mathbf{p}^v$. Finally, we modify the load utility to be

$$U_{\hat{p}_t^l, n, t}(p_t^l) := U_{n,t}(p_t^l - \hat{p}_t^l). \quad (11)$$

π_v is chosen to be $\geq U'_{\hat{p}_t^l, n, t}(\hat{p}_t^l)$. Now, we argue why including the variable \mathbf{p}^v as above is equivalent to penalizing a failure to serve \mathbf{p}^v units of the inelastic load $\hat{\mathbf{p}}^l$. If $\hat{p}_t^l \leq p_t^l$ and $p_t^v = 0$, then the true load consumed is p_t^l . However, if $p_t^v > 0$, the true load consumed is $p_t^l - p_t^v$, with a penalty of $\pi_v \Delta T p_t^v$ for incurring the VoLL. As the marginal utility of consumed energy is always less than π_v for consumed energy $\leq \hat{p}_t^l$, we can see that $p_t^v > 0 \implies p_t^l = \hat{p}_t^l$; that is, no elastic load is consumed until all inelastic load is consumed. Finally, note that since we shift the load utility right by \hat{p}_t^l , no utility is gained from serving inelastic load. Instead, the incentive to serve inelastic load comes from avoiding incurring a VoLL penalty, which is calculated from \mathbf{p}^v . So, we can include VoLL in the full objective without including it directly in each U_t by introducing the variable \mathbf{p}^v .

B. Receding Horizon Control

In Section II-A3 we specified that in this model the energy forecast error originates from solar generation forecasts. In this section, we explain how hourly dispatch is updated with actual

solar generation under receding-horizon control. Figure 2 and Algorithm 1 describe the strategy. At each step τ of




Fig. 2. Receding horizon control scheme.

Algorithm 1 Receding Horizon Control

Require: $T_{tot} \geq T$

$\forall n, \mathbf{q}_n^* \pi^* \leftarrow \text{empty}$

$\tau \leftarrow 1$

while $\tau \leq T_{tot} - T$ **do**

$P_{n,\tau}^s \leftarrow \hat{P}_{n,\tau}^s$

$\forall n, \mathbf{q}_{n,\tau}^*, \pi_\tau^* \leftarrow \text{run optimization } T \text{ ahead}$

$\tau \leftarrow \tau + 1$

end while

return $\mathbf{q}_n^* \forall n$ and π^*

the $(T_{tot} - T)$ -timestep receding horizon simulation, solar generation forecasts at the current timestep are converted to realizations. Then, all agents update their dispatches to preserve optimality under true solar generation at τ by solving (1) for timesteps $t \in [\tau, \dots, \tau + T]$. The optimal price and quantities realized for the current timestep are stored in \mathbf{q}^* and π^* , which are vectors of length $T_{tot} - T$. In this sense, the receding horizon dispatch can be considered a system-wide uncertainty mitigation strategy.

In receding-horizon control, the price vector π , the dual variable of the energy balance constraint, may also have inaccurate forecasts and thus will also update in the current time step to reflect the solar generation realized, producing a price forecast error that batteries may use to perform mitigation, described in Section II-D1.

C. Welfare Gap in a Simple System

We now analyze the effects of total energy forecast errors on total system welfare in a simple system with load, solar generation, and a battery. The concave load utility is $U(p_t^l)$ for T timesteps and 0 beyond that, while solar generates p^s for the first N timesteps and 0 beyond that, with $T > N$ the lookahead horizon. The battery's capacity is larger than the total energy in the system. Without uncertainty or inefficiencies of the battery, the optimal total welfare is $W = TU(Np^s/T)$, representing equal consumption and price at all times (enabled by the battery). Now, say that for all timesteps in the lookahead window ahead of the current one, solar forecasts are $p^s(1 + \epsilon)$, with $\epsilon \sim N(0, \sigma^2)$. We can exactly derive the consumption of the load at each timestep by noting that at any timestep, the

battery will plan to equally distribute the remaining inaccurately forecasted energy in the system across all timesteps.

Under the assumptions of this simple system, we can show the contribution to total welfare at each timestep τ is given by:

$$W_\tau = U\left(\frac{Np^s}{T} + \epsilon p^s \kappa_\tau\right) \quad \tau \in [1, \dots, N]$$

$$W_\tau = U\left(\frac{Np^s}{T} - \epsilon p^s \mu\right) \quad \tau \in [N+1, \dots, T]$$

μ and $\kappa := \{\kappa_\tau\}$ are constants that depend on simulation parameters N and T , and obey the relationship

$$\mu - \sum_{\tau=1}^N \kappa_\tau = 0 \quad (12)$$

We define the welfare gap ΔW of this system as the difference in welfare between the system operating with perfect forecasts and the system operating with imperfect forecasts:

$$\Delta W = TU\left(\frac{Np^s}{T}\right) - \sum_{\tau=1}^N U\left(\frac{Np^s}{T} + \epsilon p^s \kappa_\tau\right) - MU\left(\frac{Np^s}{T} - \epsilon p^s \mu\right), \quad (13)$$

$$(14)$$

where $M = T - N$. Note that if U is linear, $\Delta W = 0$. Positive welfare gaps arise only with concave load utilities, which represent the idea of decreasing marginal value of energy with increasing consumption. Also, the welfare gap increases with variance σ^2 of the forecast error when U is concave.

1) *Perfect Battery Uncertainty Mitigation:* With perfect knowledge of true solar generation p^s and N , the battery operator can guide the system to optimal dispatch in some cases by modifying its own utility. Specifically, we can show that when $\epsilon > 0$, if the battery gives itself a utility of energy consumption at all timesteps $U'(Np^s/T)p^b = \pi_c p^b$, where p^b is the charging power of the battery, the system will converge to the welfare-optimizing solution and the battery will increase its profit. Intuitively, this is because now the battery will charge instead of the load overconsuming when a surplus of energy is forecasted. When true forecasts are realized to be lower, the battery will serve the load with its extra stored energy, thus equalizing the price to its own marginal utility. This is a simple example of perfect uncertainty mitigation in a system with a perfect battery.

The lack of foreknowledge of true system price or error distributions prevents this strategy from being viable in practice. However, this example motivates methods to more accurately predict receding horizon price. This will be discussed in the next section.

D. Uncertainty Mitigation Methods

1) *Agent-based Data Driven:* In both centralized and distributed settings in which price is the coordination signal, battery operators have access to both forecasted and realized system prices. Batteries may use this price data to predict

shortfalls or overabundance of energy in the system by learning both historical trends in price errors and price error dependence on the battery's own dispatch. In this method, a battery operator trains an autoregressive model for the price error time series based on past price error data. Precisely, the price error vector is defined as the difference between the 1-timestep ahead forecasted price and the realized price:

$$\hat{e}_\tau = \hat{\pi}_\tau - \pi_{\tau-1,\tau} \quad (15)$$

where $\pi_{\tau-1,\tau}$ denotes the price prediction for time τ made at time $\tau-1$. Because the battery operator does not have access to the price in the current receding horizon control timestep, the operator instead uses the price error of the previous timestep to predict the error of the next timestep. A first-order price error autoregressive model can be expressed as:

$$e_{\tau+1} = \beta_h \hat{e}_{\tau-1} \quad (16)$$

where β_h is the trained coefficient, which varies by the hour h . We then supplement the total welfare of the system, Eq. (1a), with the battery's utility function, defined as:

$$U_{n,\tau}(\mathbf{p}^b) = \sum_{t=\tau}^{\tau+T} -e_t p_t^b \quad (17)$$

where $e_0 = 0$ (we assume the current timestep has no price error) and e_t is defined as in (16), using a dataset of historical errors for only the first unknown error (and reapplying the autoregressive model for the rest in the receding horizon window). This utility modification adds, for example, a penalty for charging in timesteps when the battery predicts a higher price (i.e. positive price error) than what the receding horizon forecasts.

Because the battery is operating on a microgrid, a battery operator can benefit from accounting for the influence, if any, of the battery action on the price error. Before taking any mitigation strategies, battery operators collect a data set of past price errors as well as past alterations to the battery's own dispatch in response to those price errors. In such a setting, one may expect a positive system price error (price higher than forecasted) to cause a battery to discharge more or charge less. As such, if a battery operator learns a linear model relating observed price error to a change in battery action, one would expect the correlation coefficient to be negative.

However, for uncertainty mitigation strategies, battery operators are interested in how their actions affect the price error rather than how the price error affects their actions. One may reason that, in a system with price error, a battery discharging more/charging less than forecasted would tend to decrease positive price error.

We have found that empirically, for small microgrids in which the battery can significantly affect the price through a unilateral change in its own actions, negating the learned correlation coefficient between price error and change in battery action yields better uncertainty mitigation than using only the autoregressive model. The augmented model is:

$$e_{\tau+1} = \beta_h \hat{e}_{\tau-1} - \eta(p_{\tau,\tau+1}^b - p_{\tau-1,\tau+1}^b) \quad (18)$$

with η the learned correlation coefficient between price error and change in battery action. This predicted error is input to the battery utility function (17).

2) *Heuristic Reserve*: Akin to power reserves in the central grid, a battery can hold some portion of its capacity as reserve at a certain price cap. This is especially helpful if the battery expects certain price dynamics to be indicative of energy scarcity, as in the simple setting of Section II-C1. A given battery is split into a main and reserve battery, and the total system's welfare, Eq. (1a), is supplemented with the reserve battery's utility function, which is set equal to

$$U_{n,\tau}(\mathbf{p}^b) = \sum_{t=\tau}^{\tau+T} \pi_c p_t^b, \quad (19)$$

where π_c is the price cap for reserve. Note that in contrast to the data-based method, the heuristic price cap and reserve capacity do not change during the simulation. They are parameters that can be learned or swept in simulation.

3) *Regularization*: Another heuristic we may use is inspired by the fact that the η learned by the aforementioned data-driven strategy in simulations (discussed in Section III) is negative. Ignoring terms in the error in (18) that are constant in \mathbf{p}^b , we arrive at the battery utility

$$U_{n,\tau}(\mathbf{p}^b) = \sum_{t=\tau}^{\tau+T} -e_t p_t^b = \sum_{t=\tau}^{\tau+T} \eta (p_t^b)^2 \quad (20)$$

This is akin to l_2 regularization of the battery's dispatch with regularization weight η . Applying this utility to the aforementioned reserve portion of the battery will encourage that portion to behave more conservatively in both charging and discharging for all timesteps, with larger magnitude η encouraging more conservative behavior in future timesteps. The coefficient η is a heuristic; in this work, we set it equal to the relative RMSE error in solar forecasts, 0.25.

Because we set the error predicted by the data-driven strategy in the first timestep of the receding horizon to zero, we set the regularization term for the first timestep to zero as well, reflecting the idea that there is no uncertainty, and thus no need for the battery to act conservatively, in the current timestep.

E. Distribution of Mitigation Methods

Distributed microgrid coordination offers better resilience, computational scaling, and privacy preservation [17]. For these reasons, we also apply the uncertainty mitigation methods to a distributed setting using a simplified version of the Alternating Direction Method of Multipliers (ADMM) known as proximal exchange [18]. Coordination between agents is done via a price signal corresponding to the dual variable of the energy balance constraint. A schematic of the ADMM process is shown in Figure 3. We can show that ADMM will yield identical prices and welfare to the centralized setting. The data-based battery uncertainty mitigation method uses only the price signal, which still is broadcast to the battery in the ADMM process. Furthermore, the battery can still hold reserves at a price cap or

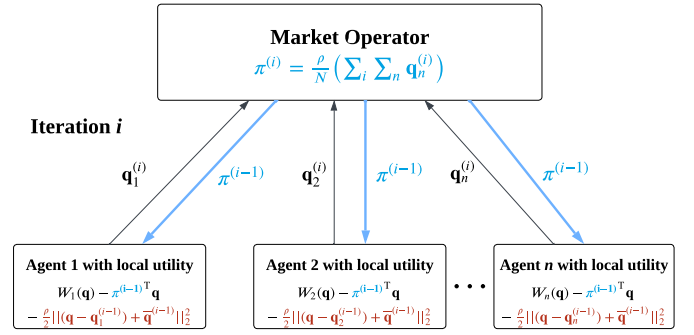


Fig. 3. Schematic of an ADMM iteration. Each agent submits an offer that optimizes a tradeoff between individual utility, energy price (blue), and step size (red). The market operator then forms a new price from the energy imbalance of all the offers over all the iterations.

add a regularization term to its utility in the distributed setting. Therefore, all of the aforementioned uncertainty mitigation strategies apply in the distributed setting.

1) *ADMM Formulation*: This subsection formulates the welfare optimization problem defined in (2) as a distributed optimization problem using the proximal exchange algorithm, a simplification of ADMM. Formulated in this manner, we can use existing results in distributed optimization to show that the iterations converge to optimal quantities and prices [18]. In [18], it is shown that the proximal exchange algorithm applies to the problem W defined in (2). This algorithm decomposes the problem into an independent step-penalized update for each agent, followed by a centralized step to compute the average surplus/shortage, which is used to generate the next point for the proximal operator. Applying the proximal operator in [18] to the local agent update, the agent-level decision then simplifies to

$$\mathbf{q}_n^{(i+1)} = \arg \min_{\mathbf{q}} -W_n(\mathbf{q}) + \frac{\rho}{2} \|\mathbf{q} - \mathbf{q}_n^{(i)} + \bar{\mathbf{q}}^{(i)}\|_2^2 + \pi^{(i)\top} \mathbf{q}, \quad (21a)$$

where $\bar{\mathbf{q}}^{(i)}$ is the energy imbalance averaged over the N agents; the price $\pi^{(i)}$ is centrally updated and is related to a system-wide integral error of total energy imbalance:

$$\pi^{(i+1)} = \pi^{(i)} + \frac{\rho}{N} \sum_n \mathbf{q}_n^{(i+1)} \quad (22a)$$

Applying this decision iteratively means, in words, updating the quantity that maximizes net benefits subject to a per unit price π , with an additional cost incurred for changing the quantity. We can show this converges to an equilibrium optimal price and quantity as $i \rightarrow \infty$. To be at equilibrium, it must be true that $\pi^{(i+1)} - \pi^{(i)} = \frac{\rho}{N} \sum_n \mathbf{q}_n^{(i+1)} = 0$, meaning that the solution is globally feasible, and the price is equal to an integral error scaled by ρ .

III. EXPERIMENTS

A. Experimental Setting

We take the load and solar data from the Pecan Street dataset [15], split those houses with solar into monthly data,

and use the same 88 house-month sample for all experiments. We run a centralized receding horizon simulation with a 24-hour window and a 1-hour time step. Unless otherwise specified, we use the same parameters for each experiment: a load utility function with $\alpha_n = -0.5$, $\pi_{\max} = 4\$/\text{kWh}$, $\gamma = 0$. \hat{p}_t^l comes from the dataset, and $\hat{\pi}_{n,t}$ is assumed to be the same TOU tariff used in [16], and we assume an explicit inelastic load of $0.75 \cdot \hat{p}_t^l$ at each time step, and we enforce a VoLL of $4\$/\text{kWh}$. For solar, we use a noise level of $\sigma = 0.25$ and get the true \hat{p}_t^s from the dataset. We use two perfectly efficient batteries, each with a capacity of 3.36 kWh, and a max charge and discharge rate of 3 kW. In each experiment with 2 batteries, we have one perform uncertainty mitigation strategies and the other dispatch as usual. For the heuristic methods, we use 15% of the battery capacity as reserve. We set the price cap for the first reserve method to 1.00 $\$/\text{kWh}$. For the regularized heuristic method, we use a coefficient $\eta = -0.25$ for all but the current timestep (for which $\eta = 0$).

B. Characterizing the Welfare Gap

The welfare gap is the difference in the welfare of the system with and without forecast error, determined by running a receding horizon simulation for both cases. In this section, we make this comparison while varying different experiment parameters. We only consider one battery for these experiments, whose size and charge rates are doubled from the base battery mentioned in section III-A. Generally, total welfares are on the order of $-\$100$ to $\$200$ a month, increasing with lower elasticity, higher battery sizes, and higher VoLL.

1) *Varying Battery Size:* We examine the welfare gap for different battery sizes. We scale the battery from 0.25 to 8 times base size (6.72 kWh, 6kW charge/discharge rate), and plot the welfare gaps in Figure 4. We find that for small and large batteries, the welfare gap is low, whereas at around twice the battery base size, the median gap reaches a peak, around $\$8$. This indicates that batteries that are too small are so often constrained by their physical characteristics that there are no large consequences for the mismanagement of energy, while batteries that are too large can store so much energy that the only limiting factor becomes the presence of solar. However, at a certain, moderate size, batteries are large enough to shift energy significantly and improve the welfare of the system, but not so large that they are never constrained by their physical characteristics. At this size, the welfare gap is largest.

2) *Varying Solar Size:* We scale solar generation from 0.1 to 4 times base amount \hat{P}_t^s , and plot the welfare gap in Figure 4. We find that, similar to what we saw when varying battery scale, for small and large solar arrays, the welfare gap is low, whereas at around the solar base size, the gap reaches a peak, with a median around $\$6$. This indicates that at low solar, there is not much energy to mismanage, while with a large amount of solar, there is ample energy such that mismanagement does not lead to significant suboptimality. It is worth noting that by varying the amount of solar generation in the system, we are in effect varying the energy scarcity of the system. So, we would expect varying solar size while




Fig. 4. Welfare gap for different values of simulation parameters. From left to right, top to bottom: battery scale, solar scale, load elasticity, and value of lost load.

keeping inelastic load constant and varying inelastic load (in the opposite direction) while keeping solar constant to exhibit similar trends in welfare gap.

3) *Varying Load Elasticity:* We first consider the welfare gap for different levels of load elasticity. We vary the load elasticity from -0.25 to -0.75 , and plot the welfare gaps in Figure 4. We see that the gap grows with decreasing load elasticity; namely, there are far more large outliers as elasticity decreases. This is because the less elastic a load is, the higher the value received from serving the load. Additionally, the general levels of welfare in the system are higher with more inelastic load.

4) *Varying Value of Lost Load:* We vary the value of lost load (VoLL) from $\$1$ to $\$9$, and plot the welfare gaps in Figure 4. We observe that a higher VoLL yields a higher welfare gap because a higher VoLL means that energy in general is valued more by consumers, so energy mismanagement is more costly than in a low-VoLL setting.

C. Uncertainty Mitigation via Batteries

In this subsection, we test 4 uncertainty mitigation strategies discussed in section II-D: a simple autoregressive model, described in (16), a data-based linear model for price errors, described in (18), heuristic price-cap-based reserve, described in (19), and regularization, described in (20). The runs of each method over different varied parameter values are plotted

as a boxplot, with the y-axis the improvement to welfare (over doing nothing to mitigate uncertainty) offered by the strategies. In Figures 5 and 6, we show the mitigation of the strategies for varying load elasticities and battery scales, respectively. We choose to explore load elasticity and battery scale based on the results in Figure 4. We choose to vary these parameters due to the interesting effects these parameters have on the welfare gap: increasing battery size first increases, then decreases the gap, and decreasing elasticity increases the gap nonlinearly. Furthermore, the reasons for these effects are not as easily deduced as for the effects of varying solar scale and VoLL, but it is worth noting that solar scale and VoLL are also parameters that could be varied in evaluating uncertainty mitigation strategies. Notably, we find that scaling VoLL proportionally scales the mitigation amount of each strategy. Finally, it is also worth noting that mitigation strategies are unnecessary in systems with no welfare gap, such as a system with very low solar noise or battery size.

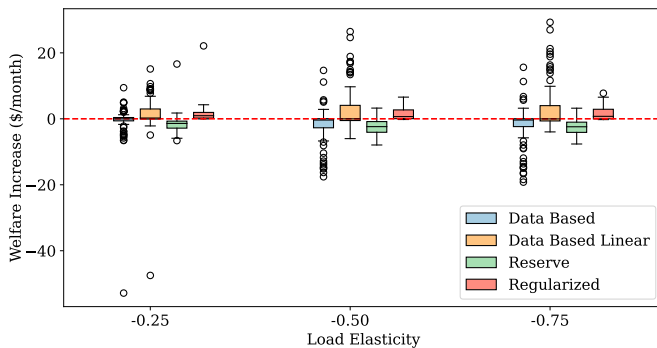


Fig. 5. Welfare loss mitigation at different load elasticities for the mitigation strategies.

When we vary the load elasticity, we can observe in Figure 5 that the data-based linear strategy performs well: while the median improvement is zero, the distribution is skewed positively, with positive mitigation ranging up to about \$25/month, and negative effects only down to -\$5/month aside from one outlier. The simple autoregressive data-based strategy may occasionally improve the welfare, but benefits tremendously from the addition of dependency of error on battery’s change in dispatch in the utility model. The price-cap-based reserve tends to hurt the performance of the system across all elasticity levels, but the regularized reserve tends to help performance consistently, although it misses some of the large opportunities for gap improvement that the data-based linear strategy exploits. All these trends seem to hold across different elasticities. When we vary battery scaling, we observe that for a larger battery, only the regularized reserve mitigates uncertainty consistently, although in some cases the linear model still offers greater mitigation. On the other hand, with a small battery, the mitigation strategies offer at most very small improvements to the system, even in a relative sense. This agrees with the intuition developed in III-B1, that when the battery is already very constrained by




Fig. 6. Welfare loss mitigation at different battery scales for the mitigation strategies.

its physical characteristics, giving additional consideration to energy management has low benefit.

Heuristic price-cap-based reserve performs even worse as battery size increases, possibly indicating that the current formulation of this reserve mismanages energy, so the more capacity the battery has, the more energy will be misallocated through the price cap. On the other hand, regularization performs better as battery size increases, indicating such a reserve formulation may be making appropriate use of the capacity of the battery.

It is worth noting that no data-based strategy ever uniformly improves welfare, as being data-based, the strategies naturally fail to capture all sources of error, especially with simple linear or autoregressive models. The regularized heuristic rarely fails to improve welfare, which makes it the best approach when consistency is required.

1) *Battery Profit via Mitigation:* While the strategies outlined in this paper are primarily intended to improve the welfare of the entire system through local battery action, we next explore whether these methods are profit-maximizing to individual battery operators. In Figure 7, we see that a battery operator is not likely to profit from performing uncertainty mitigation (except the data-based strategy, which has little mitigating effect, as seen in Figures 5 and 6). Meanwhile, the non-mitigating battery operator tends to profit when another battery operator is doing mitigation. In such a setting, a distributed mitigation approach would require side payments – akin to those paid to power reserve providers in ancillary services markets – to overcome losses incurred from performing mitigation. However, in Section III-D we will see that details of the load utility model can change this conclusion. Furthermore, if many batteries started performing mitigation in different ways, we could enter a game-theoretic setting, which is not discussed in this paper but is an item for future work.

D. Effect of Load Utility Model on Battery Profit

While Section III-C1 indicated that uncertainty mitigation may not be profitable for a battery operator unless a separate market for mitigation were involved, in this section, we explore a model of load utility for which uncertainty mitigation is




Fig. 7. Increases in profits for the batteries at different battery scales and with one battery exercising various uncertainty mitigation strategies. The left plot is the profit increase for the battery attempting uncertainty mitigation, the right is the profit increase for the battery doing nothing to mitigate uncertainty.

indeed profitable. In particular, this load utility model does not have explicit inelastic load, circumventing VoLL penalties and including inelastic load implicitly as in section II-A4. We take $\gamma = 0.9$ and $\pi_c = 4$ and analyze the performance of mitigation strategies in this system.

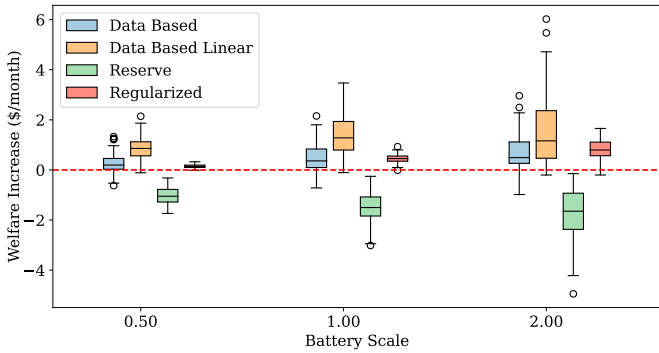


Fig. 8. Welfare loss mitigation at different battery scales for the mitigation strategies, with implicit inelastic load and no VoLL.

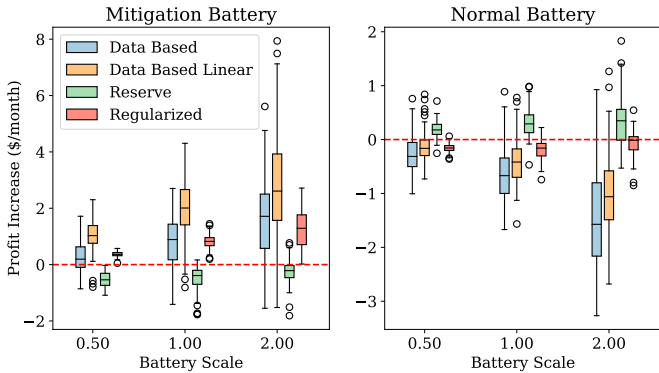


Fig. 9. Increases in profits for the batteries at different battery scales and with one battery exercising various uncertainty mitigation strategies. The left plot is the profit increase for the battery attempting uncertainty mitigation, the right is the profit increase for the battery doing nothing to mitigate uncertainty. No VoLL, implicit inelastic load.

Figures 8 and 9 show that mitigation strategies perform better with the new load model than with the model with VoLL, and uncertainty mitigation strategies are now profitable for a battery operator to employ. Figure 9 also shows that the battery operator that did not adopt any mitigation strategies generally loses money once the other battery operator employs an effective mitigation strategy. Thus, with this load utility, a battery operator performing mitigation can both increase their own profits and avoid personal profit losses due to mitigation by other batteries.

As modeling choices for utility of consumption of energy affect conclusions about mitigation strategy profitability, aligning load utility models with real consumer preferences will affect decisions about market structure. For example, load utility with VoLL might require side payments to battery operators performing mitigation. On the other hand, utilities without VoLL may naturally encourage battery operators to perform mitigation as part of their profit-maximizing objective. Alternatively, if details of load utility models are decided by system designers rather than consumers, these results may motivate choosing load utilities that do not lead to side payments for profitable mitigation. These questions motivate further study of load utility models.

E. Convergence and Performance Verification of ADMM

We run ADMM as the internal optimization in the receding horizon simulation for 6 houses with the parameters specified in III-A and an energy imbalance tolerance of 10^{-5} . We find that ADMM converges. Each house is a month of hourly simulation with a 24-hour look ahead, yielding about 720 ADMM solves per house. We find that each ADMM solve takes an average of 195.3 iterations (standard deviation 152.4) at 0.03 seconds per iteration, with a maximum of 1007 iterations encountered. We find that although ADMM prices deviate an average of 5% from centralized prices, all calculated system welfares over the month are within 0.4% of the centralized simulation, indicating that we can expect the welfare gap results of the centralized simulations to apply in the distributed setting as well.

IV. CONCLUSIONS

In energy-constrained microgrids, welfare losses due to energy generation forecast error can occur depending on the load utility functions, solar and storage scales, forecast error level, and inelastic load in the system, even in a receding horizon setting. In this paper, we propose and analyze the performance of methods that individual batteries may use to mitigate these welfare losses.

We find that data-based uncertainty mitigation generally reduces the welfare gap due to total energy uncertainty in the system, but the profitability of these strategies for the battery operator depends on the load utility models in the system. Welfare improvements are generally robust to changes in model parameters that preserve the existence of a welfare gap, with the exception of large battery scale reducing improvements from data-based mitigation. In part due to the simplicity of

the total energy forecasting method, the data-based method needs only about two weeks of data to perform well. Future work may study mitigation methods in the context of system planning, such as how utilization of computationally cheap uncertainty mitigation strategies may allow a smaller system capacity while preserving welfare.

So far, price-cap-based heuristic reserves do not show mitigation of welfare losses in energy-constrained microgrids, meaning that the operation of these microgrids with strategies similar to conventional grid power reserves may be challenging. However, we acknowledge that there are many alternative ways to formulate these heuristics. These heuristics may also perform well under different system models, such as in a system with low-frequency extreme shortfalls that data-based methods cannot predict, or with high-value loads that must be served. The latter case may motivate heuristics at the system planning level, explicitly matching battery capacity to specific loads.

Regularization consistently mitigates uncertainty across system parameters. Additionally, regularization requires no training, and so may work on a larger variety of error distributions than data-based methods do. These characteristics make regularization particularly promising for practical implementation.

All of the mitigation strategies work in a distributed setting in which energy price is broadcast to agents as a coordinating signal. To demonstrate this, we verified that system welfares in a distributed setting do not deviate significantly from those in the centralized setting, indicating the observed performance of mitigation strategies in the centralized setting will transfer to a distributed setting.

Because the formulation of load utility has a large effect on the results of this work, more accurate representations of individual load utility and methods to deal with more general utility functions are desirable. For example, discrete loads may accurately model residential consumption and utility, and are solvable via mixed integer methods [19]. The study of mitigation strategies in such a framework is a direction for future work. A study of how the choice of load utility affects mitigation profitability is a natural first step in this exploration, given the observations in section III-D.

The data-based method may be improved via more complex models for price error, such as neural networks or higher-order autoregressive models, though such methods may take longer to train. The price error dynamics' relationship to battery dispatch lends itself to a reinforcement learning framework, which is another direction for future research.

We do not include any models of infrastructure in this work beyond fixed battery parameters. Future work could include network models and a power flow validation as part of the optimization, with energy uncertainty influencing power flow as well. Finally, we do not consider multiple uncertainty-mitigating batteries working non-cooperatively. Such a system may involve a game-theoretic approach in which one battery anticipates the mitigation strategies of others. This is a path for future work.

REFERENCES

- [1] P. M. Costa and M. A. Matos, "Economic analysis of microgrids including reliability aspects," in *2006 International Conference on Probabilistic Methods Applied to Power Systems*, 2006, pp. 1–8.
- [2] S. Wang, Z. Li, L. Wu, M. Shahidepour, and Z. Li, "New metrics for assessing the reliability and economics of microgrids in distribution system," *IEEE Transactions on Power Systems*, vol. 28, no. 3, pp. 2852–2861, 2013.
- [3] J. T. Lee and D. S. Callaway, "The cost of reliability in decentralized solar power systems in sub-Saharan Africa," *Nature Energy*, vol. 3, p. 960, Oct. 2018.
- [4] C. Breyer, D. Bogdanov, A. Gulagi, A. Aghahosseini, L. S. Barbosa, O. Koskinen, M. Barasa, U. Caldera, S. Afanasyeva, M. Child, J. Farfan, and P. Vainikka, "On the role of solar photovoltaics in global energy transition scenarios," *Progress in Photovoltaics: Research and Applications*, vol. 25, no. 8, pp. 727–745, 2017. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/pip.2885>
- [5] P. Bertheau, A. S. Oyewo, C. Cader, C. Breyer, and P. Blechinger, "Visualizing national electrification scenarios for sub-saharan african countries," *Energies*, vol. 10, no. 11, 2017. [Online]. Available: <https://www.mdpi.com/1996-1073/10/11/1899>
- [6] R. Wallsgrove, J. Woo, J.-H. Lee, and L. Akiba, "The emerging potential of microgrids in the transition to 100Energies," vol. 14, no. 6, 2021. [Online]. Available: <https://www.mdpi.com/1996-1073/14/6/1687>
- [7] "Redwood Coast Airport Microgrid - Redwood Coast Energy Authority — redwoodenergy.org," <https://redwoodenergy.org/rcam/1678915316342-18d5d04e-6542>.
- [8] T. S. Ustun, C. Ozansoy, and A. Zayegh, "Recent developments in microgrids and example cases around the world—a review," *Renewable and Sustainable Energy Reviews*, vol. 15, no. 8, pp. 4030–4041, 2011. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1364032111002735>
- [9] J. L. Crespo-Vazquez, T. AlSkaif, A. M. González-Rueda, and M. Gibescu, "A community-based energy market design using decentralized decision-making under uncertainty," *IEEE Transactions on Smart Grid*, vol. 12, no. 2, pp. 1782–1793, March 2021.
- [10] A. Alizadeh, I. Kamwa, A. Moeini, and S. M. Mohseni-Bonab, "Energy management in microgrids using transactive energy control concept under high penetration of renewables; a survey and case study," *Renewable and Sustainable Energy Reviews*, vol. 176, p. 113161, 2023.
- [11] Y. Bian, N. Zheng, Y. Zheng, B. Xu, and Y. Shi, "Predicting strategic energy storage behaviors," *IEEE Transactions on Smart Grid*, pp. 1–1, 2023.
- [12] J. Cao, D. Harrold, Z. Fan, T. Morstyn, D. Healey, and K. Li, "Deep reinforcement learning-based energy storage arbitrage with accurate lithium-ion battery degradation model," *IEEE Transactions on Smart Grid*, vol. 11, no. 5, pp. 4513–4521, 2020.
- [13] D. J. Harrold, J. Cao, and Z. Fan, "Data-driven battery operation for energy arbitrage using rainbow deep reinforcement learning," *Energy*, vol. 238, p. 121958, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0360544221022064>
- [14] E. Münsing, J. Mather, and S. Moura, "Blockchains for decentralized optimization of energy resources in microgrid networks," in *2017 IEEE Conference on Control Technology and Applications (CCTA)*, 2017, pp. 2164–2171.
- [15] I. Pecan Street. (2020) Pecan street dataport. [Online]. Available: <https://www.pecanstreet.org/dataport/>
- [16] J. T. Lee, R. Henriquez-Auba, B. K. Poolla, and D. S. Callaway, "Pricing and energy trading in peer-to-peer zero marginal-cost microgrids," *IEEE Transactions on Smart Grid*, vol. 13, no. 1, pp. 702–714, 2022.
- [17] E. Espina, J. Llanos, C. Burgos-Mellado, R. Cárdenas-Dobson, M. Martínez-Gómez, and D. Sáez, "Distributed control strategies for microgrids: An overview," *IEEE Access*, vol. 8, pp. 193 412–193 448, 2020.
- [18] N. Parikh and S. Boyd, "Proximal algorithms," *Foundations and Trends in optimization*, vol. 1, no. 3, pp. 127–239, 2014.
- [19] A. Ogunjuyigbe, T. Ayodele, and O. Oladimeji, "Management of loads in residential buildings installed with pv system under intermittent solar irradiation using mixed integer linear programming," *Energy and Buildings*, vol. 130, pp. 253–271, 2016. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0378778816307381>