Carbon-Aware Charging Strategies for Electric Vehicle Battery Swapping Stations

Lingdong Zhou
Computer Science and Engineering
University of California, Riverside
Riverside, USA
lzhou010@ucr.edu

Zuzhao Ye and Nanpeng Yu
Electrical and Computer Engineering
University of California, Riverside
Riverside, USA
zye066@ucr.edu and ericyu@ucr.edu

***Abstract*—Battery swapping stations (BSS) provide operational flexibility that can help reduce the environmental impact of electric vehicle (EV) charging. This paper presents a carbon-aware strategy that minimizes BSS emissions by scheduling battery charging based on predicted grid carbon intensity (CI) and EV swap demand. We employ online update with transformer models for CI and demand forecasting and a rolling horizon mixed integer linear programming (MILP) model for scheduling optimization. Simulations using real-world data demonstrate that the average carbon emission reductions compared to a baseline immediate-charging approach ranges from 1.8% (for the 5-slot cases) to 7.5% (for the 20-slot cases). These findings highlight the potential of predictive, carbon-aware control strategies to improve the sustainability of BSS.**

***Index Terms*—Electric vehicle, Battery swapping station, Carbon emission, Carbon intensity, Online algorithm.**

I. Introduction

Electric vehicles (EVs) have emerged as a crucial component in global strategies for reducing greenhouse gas emissions and combating climate change [1]. Driven by international agreements and ambitious national or regional environmental policies aiming for decarbonization, many jurisdictions worldwide are mandating or incentivizing the transition to zero-emission transportation and actively expanding the necessary charging infrastructure to support this shift.

While EVs produce no tailpipe emissions, their overall carbon footprint is significantly influenced by the electricity used for charging, which depends on the power grid's real-time carbon intensity (CI) [2], [3]. EV charging emissions vary widely based on the generation mix; charging with renewable energy results in substantially lower emissions compared to charging during periods dominated by fossil fuels. Even in regions with increasingly renewable energy, optimizing charging times to align with low-carbon periods remains essential for maximizing the environmental benefits of electrification [4].

An innovative approach to charging EVs is using Battery Swapping Stations (BSS) [5]. These stations allow for rapid replacement of depleted batteries with fully charged ones, typically in under five minutes, thereby addressing common concerns about range anxiety and long charging times associated with conventional plug-in charging [6]. Unlike traditional charging, which must often meet immediate demand when a vehicle is present, battery swapping stations offer unique operational flexibility because the batteries can be charged independently of vehicle use.

This operational flexibility of BSS offers an important opportunity to reduce its environmental impact. By strategically timing the charging of the station's battery inventory to coincide with periods when the grid has lower carbon emissions, BSS can effectively store cleaner energy for later use through battery swaps, even when grid emissions are high during the actual swap event. Harnessing this potential through carbon-aware scheduling is therefore a key pathway to further lowering the life-cycle carbon footprint of EVs [7].

This paper proposes and evaluates such a carbon-aware charging strategy specifically designed for EV BSS, aiming to minimize emissions by leveraging fluctuations in grid CI. We begin by reviewing existing studies in Section II. Section III formally defines the optimization model as a Mixed Integer Linear Programming (MILP) problem. The core methodology is detailed in Section IV, encompassing the online update framework used for predicting EV swap demand and grid CI, and the rolling-horizon optimization technique. Section V presents the experimental setup and evaluates the effectiveness of the proposed strategy. Finally, Section VI concludes the paper, summarizing key findings, and suggesting directions for future research.

II. Related Works

Optimizing EV charging schedules has attracted significant research interest, primarily focusing on minimizing charging costs for users and mitigating adverse impacts on the power grid. Common approaches include decentralized control methods for scalability [8], Model Predictive Control (MPC) for incorporating forecasts [9], and various techniques addressing demand or price uncertainty [10]. Data-driven methods, particularly Reinforcement Learning (RL), have also been employed for real-time scheduling [11] adaptable to grid conditions and user preferences [12]. While effective for cost and grid management, these studies often do not directly prioritize the minimization of carbon emissions associated with charging.

Recognizing that the environmental benefit of EVs is tied to the source of electricity, a growing body of work investigates strategies to reduce charging-related emissions. Some studies evaluate smart charging using real-world data

to align charging with periods which have higher renewable generation (and potentially lower emissions) [13]. More targeted approaches specifically consider the time-varying CI of the grid. For instance, researchers have proposed emissions-responsive charging strategies based on life-cycle assessments [14], incorporated marginal carbon emission factors alongside pricing signals [15], and developed scheduling algorithms that are explicitly aware of low CI periods [16], [17].

BSS offer a distinct operational model with enhanced flexibility, as battery charging can be decoupled from the presence of an EV [18]. Research specific to BSS has explored optimal scheduling of charging and swapping operations [19], network planning considering user satisfaction [20], and managing stochastic demand [21]. Advanced scheduling techniques using machine learning, such as LSTM [22] and RL [23], have been proposed to handle the dynamic nature of BSS operations. However, much of the BSS optimization literature focuses on economic viability or operational efficiency under uncertainty, rather than explicitly using the station's charging flexibility to reduce its carbon footprint.

While carbon-aware strategies exist for plug-in EVs [14]–[16], their application to optimize the charging schedules of BSS remains underexplored. This paper addresses this critical gap, presenting a novel approach specifically designed to minimize the carbon footprint of BSS by leveraging real-time CI predictions, thereby highlighting an important pathway towards more sustainable EV infrastructure.

III. Problem Formulation

This work aims to determine an optimal charging schedule for swappable batteries in an EV BSS over a finite time horizon T. Our goal is to leverage this flexibility to minimize the total carbon emissions associated with the BSS operation, which includes emissions from charging batteries within the swapping slots, as well as the emissions resulting from plug-in charging for unmet swap demands. To reflect drivers' preference for fast, convenient service, we prioritize assigning incoming EVs to battery swapping whenever possible, and any EV swap demand arriving at time t that cannot be immediately satisfied by a fully charged battery from the BSS slots will be redirected to utilize on-site regular plug-in charging. We further assume that a sufficient number of plug-in chargers are available to accommodate any unmet swap demand. The charging process for these plug-in chargers is assumed to follow a greedy policy (i.e. immediate charging at maximum power upon EV connection). The optimization for the swappable batteries explicitly considers the time-varying CI of the electricity grid (CI_t) and the anticipated EV swap demands (D_t) to make scheduling decisions. Fig. 1 illustrates the operational setup in this study.

Other assumptions include that the target state for all charging considered in the emissions calculation – both for charging of swappable batteries within BSS and for plug-in charging – is 100% State-of-Charge (SOC). Furthermore, it is assumed that all batteries involved in the system, both within the BSS slots and on the arriving EVs, are homogeneous, possessing the same maximum capacity denoted as e^{max}.

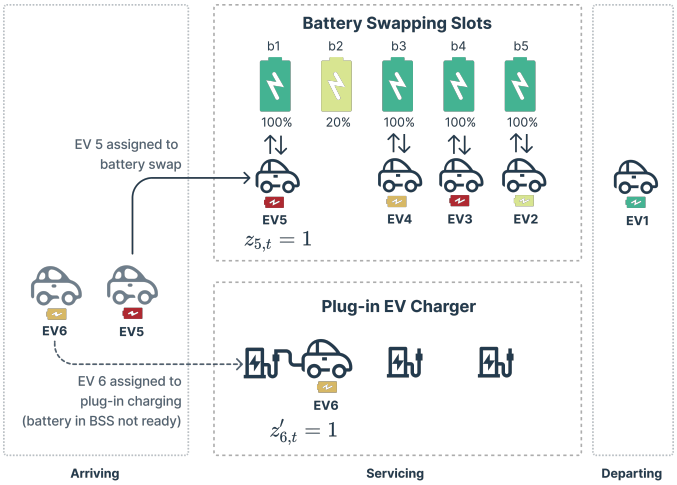

Fig. 1. Overview of the Battery Swapping Station operation.

A. Key Variables and Parameters

Our model features the following index sets. $\mathcal{B} = \{1, 2, ..., B\}$: the set of charging slots available in the BSS, indexed by b, where B is the total number of slots. $\mathcal{T} = \{1, 2, ..., T\}$: the set of discrete time steps in the optimization horizon, indexed by t, where T is the horizon length. $\mathcal{I}_t = \{1, 2, ..., D_t\}$: the set representing the individual EV swap demands anticipated to arrive during time step t, indexed by i, where D_t is the anticipated demand count for step t.

We model the state and decisions for each charging slot $b \in \mathcal{B}$ over the time steps $t \in \mathcal{T}$. The key decision variables are introduced below. $y_{b,t} \in \{0, 1\}$: binary variable indicating if the battery in charging slot b is charging during time step t. $z_{b,t} \in \{0, 1\}$: binary variable indicating if the (fully charged) battery in charging slot b is selected for a swap at the beginning of time step t. $SOC_{b,t} \in [0, 1]$: continuous variable representing the SOC of the battery residing in charging slot b at the *end* of time step t. $p^{btr}_{b,t} \in [0, p^{btr,max}]$: continuous variable representing the charging power to the battery in charging slot b during time step t. $\delta_{b,i,t} \in \{0, 1\}$: binary variable indicating if the battery from charging slot b is assigned to serve the i-th EV demand ($i \in \mathcal{I}_t$) arriving during time step t. $z'_{i,t} \in \{0, 1\}$: binary variable indicating if the i-th EV demand ($i \in \mathcal{I}_t$) arriving during time step t receives plug-in charging (i.e., is not served by the BSS).

Let ΔT denote the duration of a single time step (e.g., $\Delta T = 1$ hour). Key system parameters governing charging rates are $p^{btr,max}$, the maximum charging power for a BSS slot, and $p^{plug,max}$, the assumed constant charging power for the regular on-site plug-in chargers. Let SOC^{init}_b denote the known initial SOC of the battery in slot b before the first time step ($t = 1$). For notational convenience in the constraints, we define $SOC_{b,0} \equiv SOC^{init}_b$. Let $SOC^{ev,arr}_{i,t}$ be the assumed arrival SOC of the battery from the i-th EV arriving at time t. Note that this is treated as an input parameter.

B. Objective Function

The BSS's operation objective is to minimize the total carbon emissions over the horizon $\mathcal{T}$, by choosing the optimal values for the decision variables:

$$\begin{aligned}\min_{\substack{y_{b,t},z_{b,t},p^{btr}_{b,t},\\ \delta_{b,i,t},z'_{i,t}}} \quad & \sum_{t\in\mathcal{T}}\sum_{b\in\mathcal{B}} CI_t \cdot p^{btr}_{b,t}\cdot\Delta T \\ & + \sum_{t\in\mathcal{T}}\sum_{i\in\mathcal{I}_t}\left(\sum_{h=t}^{T} CI_h\cdot\Delta e^{ev}_{h,i,t}\right)\cdot z'_{i,t} \\ & + M\sum_{t\in\mathcal{T}}\sum_{i\in\mathcal{I}_t} z'_{i,t}, \end{aligned} \tag{1}$$

where $y_{b,t}, z_{b,t}, \delta_{b,i,t}, z'_{i,t}$ are binary decisions, $SOC_{b,t}$ is the state of charge, and $p^{btr}_{b,t}$ is the charging power, across all slots b, demands i, and time steps t. The first term represents the emissions from charging batteries within the station's slots, calculated from the energy consumed ($p^{btr}_{b,t}\cdot\Delta T$) and carbon intensity CI_t. The second term accounts for emissions from EVs using plug-in charging. The parameter $\Delta e^{ev}_{h,i,t}$ denotes the amount of energy added to the i-th EV (which arrived at time t) during a subsequent time h (where $h \geq t$). To compute $\Delta e^{ev}_{h,i,t}$, we first determine the total energy needed to charge the ith EV from $SOC^{ev,arr}_{i,t}$ to 100%. This total energy is then allocated across time slots starting from t according to a greedy rule: at each time step h, the EV receives either $\Delta e^{ev}_{h,i,t} = p^{plug,max}\cdot\Delta t$ or the remaining amount needed to complete the charge, whichever is smaller. The third term applies a large penalty M for each EV assigned to plug-in charging ($z'_{i,t}=1$), aiming to prioritize fast turnaround by satisfying demand via battery swaps whenever feasible and discouraging plug-in charging.

C. Constraints

The model incorporates the following constraints to ensure proper operation:

$$y_{b,t} + z_{b,t} \leq 1, \quad \forall b\in\mathcal{B},\ \forall t\in\mathcal{T} \tag{2}$$

$$p^{btr}_{b,t} \leq p^{btr,max}\cdot y_{b,t}, \quad \forall b\in\mathcal{B},\ \forall t\in\mathcal{T} \tag{3}$$

$$p^{btr}_{b,t}\cdot\Delta T \leq (1-SOC_{b,t-1})\cdot e^{max}\cdot y_{b,t}, \quad \forall b\in\mathcal{B},\ \forall t\in\mathcal{T} \tag{4}$$

$$z_{b,t} \leq SOC_{b,t-1}, \quad \forall b\in\mathcal{B},\ \forall t\in\mathcal{T} \tag{5}$$

$$\sum_{b\in\mathcal{B}}\delta_{b,i,t} + z'_{i,t} = 1, \quad \forall t\in\mathcal{T},\ \forall i\in\mathcal{I}_t \tag{6}$$

$$\sum_{b\in\mathcal{B}}\delta_{b,i-1,t} \geq \sum_{b\in\mathcal{B}}\delta_{b,i,t}, \quad \forall t\in\mathcal{T},\ \forall i\in\mathcal{I}_t,\ i\geq 2 \tag{7}$$

$$\begin{aligned} SOC_{b,t}\cdot e^{max} = {} & (1-z_{b,t})\cdot(SOC_{b,t-1}\cdot e^{max} + p^{btr}_{b,t}\cdot\Delta T) \\ & + z_{b,t}\cdot\sum_{i\in\mathcal{I}_t}\left(SOC^{ev,arr}_{i,t}\cdot e^{max}\cdot\delta_{b,i,t}\right), \\ & \forall b\in\mathcal{B},\ \forall t\in\mathcal{T} \end{aligned} \tag{8}$$

$$\sum_{b\in\mathcal{B}} z_{b,t} + \sum_{i\in\mathcal{I}_t} z'_{i,t} = D_t, \quad \forall t\in\mathcal{T} \tag{9}$$

Eq. (2) requires that a slot b will either be charging ($y_{b,t}=1$) or be selected for a swap ($z_{b,t}=1$) within the same time step t. Eq. (3) limits the charging power not to exceed the maximum charging rate. Eq. (4) enforces energy capacity limits that the energy delivered during the time step cannot exceed the battery's remaining capacity from the previous state. Eq. (5) ensures a swap operation ($z_{b,t}=1$) can only be initiated if the battery is fully charged at the end of the previous time step. Eq. (6) assigns EV i arriving at time t to exactly one BSS slot b for swapping ($\sum_b \delta_{b,i,t}=1$) or plug-in charging ($z'_{i,t}=1$). Eq. (7) enforces that within the same time step t, an earlier-indexed EV will be given a higher priority for battery swap. Eq. (8) defines the state transition for battery energy in slot b: If no swap occurs ($z_{b,t}=0$), the battery energy at the end of t equals to the initial energy ($SOC_{b,t-1}\cdot e^{max}$) plus the energy charged during t ($p^{btr}_{b,t}\cdot\Delta T$). If a swap occurs ($z_{b,t}=1$), the battery energy in the slot is updated to the arrival energy of the assigned incoming EV ($SOC^{ev,arr}_{i,t}\cdot e^{max}\cdot\delta_{b,i,t}$). Eq. (9) ensures that the number of demands served by BSS swaps ($\sum_b z_{b,t}$) plus the number served by plug-in charging ($\sum_i z'_{i,t}$) must equal the total anticipated demand D_t for time t.

IV. Methodology

A. System Overview

Our proposed carbon-aware charging strategy integrates online forecasting and rolling horizon optimization. The system collects historical and real-time data (EV demand, grid CI, weather), utilizes online update models to predict future EV swap demand and grid CI, and employs a rolling horizon optimization framework based on these predictions to determine adaptive charging schedules for the batteries in BSS. Fig 2 depicts this overall architecture, showing the flow of information from data sources through the prediction models to the controller executing the optimized schedule.

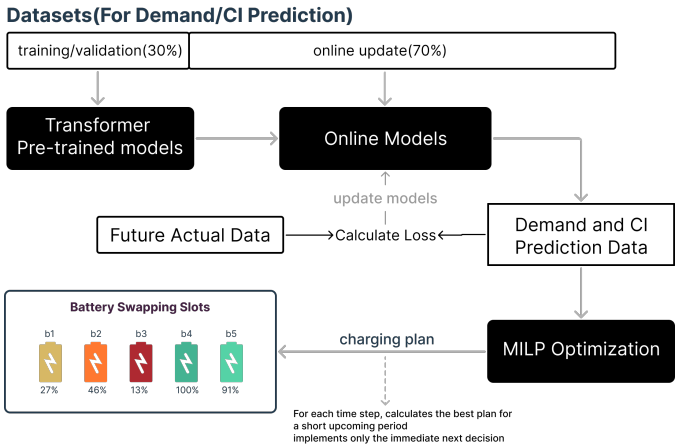

Fig. 2. System Overview: Data flows through online prediction models to inform the rolling horizon optimization controller.

B. Rolling Horizon Control and Online Prediction

Rolling Horizon Control Strategy: While Section III defines the optimization problem over a complete horizon T assuming perfect information, practical application must address the inherent uncertainty in forecasting future grid carbon intensity (CI_t) and EV demand (D_t). To manage this uncertainty and make use of the best available future information, we employ a rolling horizon optimization strategy, akin to Model Predictive Control (MPC). This approach repeatedly solves the optimization problem using the latest available forecasts for a near-term horizon, implements only the immediate next action, and then re-optimizes in the subsequent step with updated information and refreshed forecasts obtained from online prediction models. This allows the system to adapt to changing conditions dynamically.

The specific workflow of the rolling horizon control is described as follows. First, at each decision step (e.g., hourly), the system obtains the latest multi-step forecasts for CI_t and D_t (e.g., for the next 24 hours) from the online prediction models. Next, the MILP model (formulated in Section III) is solved using these forecasts over a defined optimization horizon (e.g., 6 hours). Crucially, only the decisions corresponding to the *immediate next* time step are actually implemented by the BSS. Finally, the system state is updated based on the implemented decisions and observed EV swaps/charges during the past time step, providing the initial conditions for the subsequent optimization cycle.

Demand and Carbon Intensity Prediction Models: The effectiveness of the rolling horizon strategy depends crucially on the quality of the forecasts. We generate the forecasts for both EV swap demand and grid CI using a unified methodology based on the Transformer architecture [24], chosen for its strength in modeling complex temporal dependencies. The models are initially trained offline and feature an online update component for continuous adaptation.

Input Features and Preprocessing: Models utilize historical target values (demand counts or CI values) along with relevant exogenous features known to influence the target variables. For EV charging demand prediction, inputs include historical demand counts derived from sources like the City of Palo Alto Open Data [25], various weather features obtained from the Global Forecast System (GFS) [26] (such as temperature, humidity, pressure, precipitation, snow, and irradiance), standard temporal indicators (e.g., time of day, day of week encoded cyclically), and holiday flags. For grid CI prediction, inputs include historical CI values, temporal indicators, and key grid operational forecasts sourced from platforms like the California Independent System Operator (CAISO) OASIS [27]. Specifically, we utilize Day-Ahead Market (DAM) data including system-wide load forecasts as well as wind and solar generation forecasts, which are primary drivers of the grid's real-time generation mix and thus its CI.

Standard preprocessing steps are applied to the input data for both models. Missing values are imputed using a forward-fill followed by a backward-fill strategy. Time-based features providing essential temporal context are engineered, including cyclical features generated using sine and cosine transformations for periodic inputs v with period P (e.g., hour of the day $v \in [0, 23], P = 24$; month of the year $v \in [1, 12], P = 12$). Subsequently, all input features undergo standardization (**Scaling**) to achieve zero mean and unit variance, which typically benefits model training.

Transformer Model Architecture: Both forecasting tasks utilize a model based on the Transformer architecture, adapted for multi-variate time-series prediction. The input sequence is first embedded, and then positional encoding (PE) is added to retain sequence order using standard sinusoidal functions.

A stack of Transformer encoder layers processes the sequence using multi-head self-attention mechanisms to capture temporal dependencies. A final output layer maps the processed representation to the required multi-step forecast horizon. Models are initially trained offline by minimizing the Mean Squared Error (MSE) loss function.

Online Update Process: To maintain prediction accuracy as data patterns drift over time, a unified online update mechanism is employed. As new demand and CI data become available, they are stored in historical buffers (e.g., covering the last 30 days). At regular intervals (e.g., daily), the models are fine-tuned using recent data from these buffers, typically minimizing the MSE loss on recent samples. This periodic updating allows the models to adapt to non-stationary patterns or distribution shifts, ensuring the forecasts provided to the rolling horizon controller remain relevant.

V. Numerical Studies

In this section, we present the numerical study results of our proposed model and baseline methods.The computational tasks in this study were executed on two distinct systems. Model training and online prediction were performed on a workstation running Ubuntu 22.04.5 LTS, equipped with an AMD Ryzen Threadripper 3970X 32-Core Processor and an NVIDIA GeForce RTX 2080 Ti GPU. All subsequent evaluation and optimization were conducted on an Apple Mac Mini featuring the M4 processor and 16 GB of unified memory.

A. Datasets and Evaluation Approach

We used hourly data derived from several sources. The weather forecasts are from GFS [26]. The EV swap demand counts (D_t) is from the Palo Alto EV charging records (July 2011-Dec. 2020) [25]. The electric load and total emissions from April 2018 to Dec. 2020 are obtained from CAISO [28].

The hourly grid CI (CI_t), a key input for our optimization, was calculated from the CAISO demand (P_t) and emission (E_t) data using the following formula with the unit of grams of CO_2 equivalent per kilowatt-hour (gCO_2eq/kWh):

$$CI_t[\text{gCO}_2\text{eq/kWh}] = \frac{E_t[\text{mTCO}_2\text{/h}]}{P_t[\text{MW}]} \times 1000$$

To simulate the state of depleted batteries arriving for a swap, the initial SOC for each demand instance ($\gamma_{i,t}$ in Eq. (8)) was sampled from a distribution designed to approximate real-world conditions. First, a raw value γ_{raw} was drawn from a

log-normal distribution, characterized by the mean (μ) and standard deviation (σ) of the variable's natural logarithm: $\gamma_{raw} \sim \text{LogNormal}(\mu = 3.5, \sigma^2 = 0.4^2)$. The final SOC value used in the simulation was then obtained by clamping γ_{raw} to a realistic range: $\gamma_{i,t} = \max(3\%, \min(\gamma_{raw}, 85\%))$. This approach provides typical EV battery levels observed upon arrival at charging or swapping facilities.

To avoid potential anomalies related to the COVID-19 pandemic, all prediction models and optimization evaluations used data only up to December 31, 2019. A multi-set evaluation methodology was employed. Starting from November 26, 2019, and working backward, we extracted 40 distinct 72-hour datasets. For each dataset, a 48-hour rolling horizon simulation was performed. All reported results represent the average performance across these 40 datasets.

B. Prediction Model Training Configuration

The Transformer models for both EV demand and grid CI prediction underwent an initial offline training phase followed by online adaptation. For the initial offline training phase, both Transformer models were optimized using the Adam optimizer. Key hyperparameters were set as follows: the input sequence length is 120 hours, the prediction horizon is 24 hours, the batch size is 128, and the initial learning rate is 0.0001. The Transformer architecture itself was configured with a model dimension (d_{model}) of 128, 8 attention heads (n_{head}), and 6 layers, utilizing a dropout rate of 0.2 for regularization. Training was conducted for a maximum of 200 epochs, with early stopping implemented based on validation performance with a patience of 20 epochs.

To maintain accuracy as data patterns evolve, the models were continuously adapted during the online update phase. This adaptation was performed via daily fine-tuning, utilizing recent historical data stored within a 30-day memory buffer. Distinct online update rates were used for this process: 0.001 for the EV demand model and 0.0003 for the CI model.

C. Prediction Model Performance

The online prediction models, evaluated on the pre-2020 test data, yielded the following performance: The Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE) of EV demand prediction are 1.66 counts/hour and 0.76 counts/hour, respectively. The RMSE and MAE of grid CI prediction are 22.31 gCO_2eq/kWh and 10.29 gCO_2eq/kWh, respectively. Note that this study does not aim to achieve the highest possible forecasting accuracy, but rather to use the prediction models as key components of the optimization framework. We skip the direct comparisons with studies that solely focused on forecasting, as the prediction accuracy achieved is sufficient to demonstrate the effectiveness of the proposed framework. Improving forecasting accuracy can be explored in future research.

D. Optimization Configuration

The rolling horizon MILP simulations were conducted for each 48-hour test set using the following configuration. The optimization employed a 6-hour look-ahead horizon with a time step duration $\Delta T = 1$ hour. A homogeneous battery capacity of 75 kWh was assumed, and the maximum charging rate within the BSS slots ($p^{btr,max}$) was set equivalent to charging 50.0% SOC per hour (corresponding to 37.5 kW). Separate simulation runs were performed for BSS configurations with 5, 10, and 15 battery slots. Gurobi was used as the MILP solver for all runs, each simulating a 48-hour period, with a solver time limit of 3600 seconds and an aggressive presolve setting (Presolve=2).

E. Computation Time

To assess the real-world feasibility of the MILP-based approach, we analyzed its computation time. Table I summarizes the per-step optimization time, which are critical for real-time control.

TABLE I
PER-STEP OPTIMIZATION TIME (SECONDS)

Slots	Mean	Median	25th Pctl.	75th Pctl.
5	0.053	0.054	0.013	0.081
10	0.542	0.215	0.042	0.677
15	1.203	0.501	0.089	1.019
20	6.937	0.802	0.157	1.501

Pctl: Percentile

The per-step computation time is substantially shorter than the one-hour decision interval, demonstrating the practicality of our approach for real-world deployment. However, we observed that for larger configurations—such as the 20-slot case—certain steps required over 30 minutes to solve, highlighting the exponential growth in problem complexity. These observations suggest that while the current MILP-based approach is suitable for small-scale BSS, its scalability poses challenges for real-time operations in larger systems. This underscores the need for more computationally efficient methods, as further discussed in Section VI

F. Baseline Strategies and Evaluation Metrics

The performance of the proposed strategy, referred to as the **Rolling Horizon with Prediction**, was compared against three other strategies. The first is a baseline strategy, which represents a "greedy" operational approach. When a swap occurs, the returned depleted battery is placed in an available slot and immediately begins charging at the maximum rate until it reaches 100% SOC without considering grid carbon intensity. The second strategy is the **Price-Aware Strategy**, which uses the same optimization framework but aims to minimize the electricity cost based on predicted EV demand and time-of-use (TOU) prices [29], representing a common economic-driven approach. The third is the **Rolling Horizon with Actual Value**. This represents an idealized scenario with perfect foresight, utilizing the same optimization framework as the proposed strategy but operating with actual, known future demand and CI values rather than predictions.

Performance is evaluated based on the total carbon emissions (kg CO_2eq) averaged across the 40 test sets, and the carbon reduction (%) compared to the Baseline strategy.

G. Simulation Results

The performance of the Baseline strategy and the carbon-aware strategies (using either predicted data or actual data) across three BSS configurations is summarized in Table II.

TABLE II
EMISSIONS AND CARBON REDUCTION VS. BASELINE

Slots	Baseline (kg CO_2eq)	Price-Aware	Carbon-Aware (Prediction)	Carbon-Aware (Actual)
5	4897.93	4825.39 (1.6%)	4806.34 (1.8%)	4781.68 (2.5%)
10	4935.43	4768.22 (3.4%)	4727.35 (4.2%)	4688.42 (5.2%)
15	4894.19	4652.34 (4.9%)	4590.17 (6.2%)	4524.51 (7.8%)
20	4887.25	4606.6 (5.5%)	4508.92 (7.5%)	4390.98 (10.6%)

Values in each cell show avg. absolute emissions (kg CO_2eq) and the percentage reduction (%) vs. Baseline.

These results demonstrated the effectiveness of leveraging CI predictions to reduce the carbon footprint of BSS. The Rolling Horizon with Actual Value Strategy shows the potential emission reduction achievable with perfect information, while the gap between the Rolling Horizon with Actual Value and Prediction strategy reflects the impact of prediction inaccuracies. As detailed in Table II, while the economic-driven Price-Aware strategy also achieves carbon reductions because low-price periods often coincide with low-CI periods from high renewable generation, our proposed carbon-aware approach consistently outperforms it across all scenarios. Increasing the number of battery slots generally improved the average carbon reductions achieved by the proposed strategy compared to the Baseline Strategy, with reductions ranging from 1.8% for 5 slots to 7.5% for 20 slots. This highlights the increased flexibility for shifting charging load towards lower CI periods when more battery swapping slots are available.

To illustrate the operational dynamics of the proposed carbon-aware strategy, Fig. 3 presents a representative example from one 48-hour simulation period (using the 15-slot configuration). The figure compares the hourly charging power (in MW) of the Baseline strategy, the Rolling Horizon with Prediction, and the Rolling Horizon with Actual Value, against the actual grid CI profile for that period. It can be observed that both carbon-aware strategies successfully shift significant portions of their charging load away from high-CI periods towards hours with lower grid CI, compared to the Baseline strategy which charges immediately whenever a battery is available.

To evaluate the robustness of the proposed strategy under high uncertainty and potential stress conditions, we performed a sensitivity analysis on the prediction errors. We introduced structured noise into the forecasts, simulating scenarios with

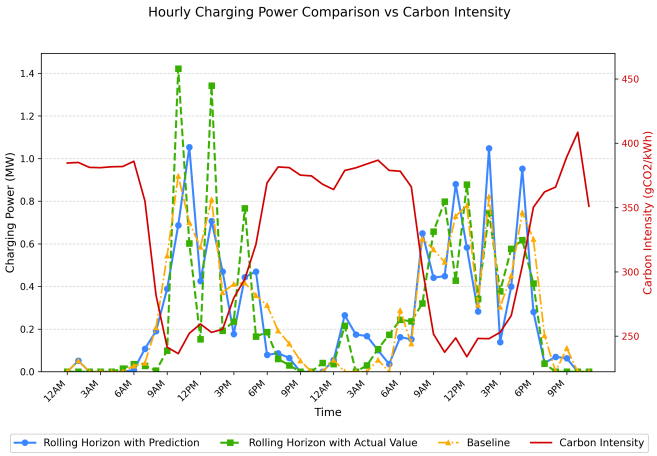

Fig. 3. Hourly charging power comparison versus actual carbon intensity for a representative 48-hour period (15-slot configuration). Illustrates how carbon-aware strategies shift charging to low-CI periods compared to the Baseline.

increasing levels of error. This noise was sampled from a normal distribution, with the standard deviation proportional to the historical pointwise absolute error of our model and scaled by a multiplier α. As summarized in Table III, the strategy maintains strong performance under moderate noise levels, exhibiting only a gradual reduction in carbon savings. For example, at $\alpha = 2.0$, the approach still achieves a 3.2% carbon reduction. However, under extreme prediction errors (e.g., $\alpha = 5.0$), performance can degrade significantly, even resulting in negative outcomes (i.e., "carbon leakage"). These findings highlight the importance of improving forecast accuracy to ensure the reliability of carbon-aware scheduling in volatile operational environments.

TABLE III
IMPACT OF PREDICTION ERROR ON CARBON EMISSION REDUCTION (20-SLOT CASE)

Scenario	Noise Multiplier (α)	Carbon Reduction (%)
Perfect Foresight	N/A	10.6%
Original Prediction	0.0	7.5%
Noise Level 1	1.0	6.6%
Noise Level 2	1.5	5.6%
Noise Level 3	2.0	3.2%
Noise Level 4	3.0	1.8%
Noise Level 5	5.0	-2.7%

VI. CONCLUSION

This paper developed a carbon-aware framework for scheduling charging at EV battery swapping stations (BSS). The framework uses online predictions for grid carbon intensity (CI) and EV demand, combined with a rolling horizon optimization to determine optimal charging schedules. The experiments, using real-world data, demonstrated that the proposed framework cuts emissions compared to the baseline greedy battery charging strategy. We found an average carbon reductions of 1.8% with 5 battery slots, and up to 7.5% with 20

battery slots, highlighting how increased scheduling flexibility with more battery slots can lead to greater emission reductions.

Our study also highlighted several limitations that open up avenues for future research. First, for larger BSS configurations, it is essential to develop more computationally efficient optimization methods (e.g., reinforcement learning) to overcome the limitation of commercial MILP solvers. Second, future studies should extend the optimization framework to be more comprehensive by examining cost-emission trade-offs through a multi-objective model. This involves creating a holistic cost model that includes not only time-varying electricity prices but also the economic impact of battery degradation and effects on service quality. A key outcome would be calculating the unit emission reduction cost for comparison against other technologies. Third, improving prediction accuracy for demand and CI is key, alongside validating the algorithm by integrating data from actual Battery Management Systems (BMS) in a hardware-in-the-loop setting to enhance real-world viability.

VII. Acknowledgment

This work was supported by the National Science Foundation award 2324940 and California Energy Commission award GFO-19-309.

References

[1] Z. Gao, H. Xie, X. Yang, L. Zhang, H. Yu, W. Wang, Y. Liu, Y. Xu, B. Ma, X. Liu *et al.*, "Electric vehicle lifecycle carbon emission reduction: A review," *Carbon Neutral.*, vol. 2, no. 5, pp. 528–550, 2023.

[2] G. Bieker, "A global comparison of the life-cycle greenhouse gas emissions of combustion engine and electric passenger cars," *Commun.*, vol. 49, no. 30, pp. 847 129–102, 2021.

[3] W. Wang, Y. Li, and N. Yu, "Predict locational marginal greenhouse gas emission factors of electricity with spatial-temporal graph convolutional networks," in *IEEE PES Innovative Smart Grid Technologies Europe (ISGT EUROPE)*, 2023, pp. 1–6.

[4] J. Li, G. Wang, X. Wang, and Y. Du, "Smart charging strategy for electric vehicles based on marginal carbon emission factors and time-of-use price," *Sustain. Cities Soc.*, vol. 96, p. 104708, 2023.

[5] F. Ahmad, M. Saad Alam, I. Saad Alsaidan, and S. M. Shariff, "Battery swapping station for electric vehicles: Opportunities and challenges," *IET Smart Grid*, vol. 3, no. 3, pp. 280–286, 2020.

[6] H. Wu, "A survey of battery swapping stations for electric vehicles: Operation modes and decision scenarios," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 8, pp. 10 163–10 185, 2021.

[7] S. Wu, Q. Xu, Q. Li, X. Yuan, and B. Chen, "An optimal charging strategy for PV-based battery swapping stations in a DC distribution system," *Int. J. Photoenergy*, vol. 2017, no. 1, p. 1504857, 2017.

[8] L. Gan, U. Topcu, and S. H. Low, "Optimal decentralized protocol for electric vehicle charging," *IEEE Trans. Power Syst.*, vol. 28, no. 2, pp. 940–951, 2012.

[9] W. Tang and Y. J. Zhang, "A model predictive control approach for low-complexity electric vehicle charging scheduling: Optimality and scalability," *IEEE Trans. Power Syst.*, vol. 32, no. 2, pp. 1050–1063, 2016.

[10] Y. Zhou, D. K. Yau, P. You, and P. Cheng, "Optimal-cost scheduling of electrical vehicle charging under uncertainty," *IEEE Trans. Smart Grid*, vol. 9, no. 5, pp. 4547–4554, 2017.

[11] Z. Ye, Y. Gao, and N. Yu, "Learning to operate an electric vehicle charging station considering vehicle-grid integration," *IEEE Trans. Smart Grid*, vol. 13, no. 4, pp. 3038–3048, 2022.

[12] K. Park and I. Moon, "Multi-agent deep reinforcement learning approach for EV charging scheduling in a smart grid," *Appl. Energy*, vol. 328, p. 120111, 2022.

[13] S. I. Spencer, Z. Fu, E. Apostolaki-Iosifidou, and T. E. Lipman, "Evaluating smart charging strategies using real-world data from optimized plugin electric vehicles," *Transp. Res. Part D: Transp. Environ.*, vol. 100, p. 103023, 2021.

[14] Y. Tang, T. T. Cockerill, A. J. Pimm, and X. Yuan, "Reducing the life cycle environmental impact of electric vehicles through emissions-responsive charging," *iScience*, vol. 24, no. 12, 2021.

[15] J. Li, G. Wang, X. Wang, and Y. Du, "Smart charging strategy for electric vehicles based on marginal carbon emission factors and time-of-use price," *Sustain. Cities Soc.*, vol. 96, p. 104708, 2023.

[16] K.-W. Cheng, Y. Bian, Y. Shi, and Y. Chen, "Carbon-aware EV charging," in *Proc. IEEE SmartGridComm*, 2022, pp. 186–192.

[17] Z. Ye, N. Yu, R. Wei, and X. C. Liu, "Decarbonizing regional multi-model transportation system with shared electric charging hubs," *Transp. Res. Part C: Emerging Technologies*, vol. 144, p. 103881, 2022.

[18] S. Schmidt, "Use of battery swapping for improving environmental balance and price-performance ratio of electric vehicles," *ETransportation*, vol. 9, p. 100128, 2021.

[19] M. R. Sarker, H. Pandžić, and M. A. Ortega-Vazquez, "Optimal operation and services scheduling for an electric vehicle battery swapping station," *IEEE Trans. Power Syst.*, vol. 30, no. 2, pp. 901–910, 2015.

[20] J. Yang, F. Guo, and M. Zhang, "Optimal planning of swapping/charging station network with customer satisfaction," *Transp. Res. Part E: Logist. Transp. Rev.*, vol. 103, pp. 174–197, 2017.

[21] D. S. Nayak and S. Misra, "An operational scheduling framework for electric vehicle battery swapping station under demand uncertainty," *Energy*, vol. 290, p. 130219, 2024.

[22] A. A. Shalaby, M. F. Shaaban, M. Mokhtar, H. H. Zeineldin, and E. F. El-Saadany, "A dynamic optimal battery swapping mechanism for electric vehicles using an LSTM-based rolling horizon approach," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 9, pp. 15 218–15 232, 2022.

[23] J. Jin, S. Mao, and Y. Xu, "Optimal priority rule-enhanced deep reinforcement learning for charging scheduling in an electric vehicle battery swapping station," *IEEE Trans. Smart Grid*, vol. 14, no. 6, pp. 4581–4593, 2023.

[24] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *NeurIPS*, vol. 30, 2017.

[25] City of Palo Alto, "Electric vehicle charging station usage July 2011 - Dec 2020," City of Palo Alto Open Data Portal, 2020.

[26] National Centers for Environmental Prediction, National Weather Service, NOAA, U.S. Department of Commerce, "NCEP GFS 0.25 degree global forecast grids historical archive," Boulder CO, 2015. [Online]. Available: https://rda.ucar.edu/datasets/d084001/

[27] California Independent System Operator (CAISO), "OASIS - Open Access Same-time Information System," https://oasis.caiso.com/, Accessed 2024.

[28] ——, "Today's outlook - demand and emissions data," https://www.caiso.com/todays-outlook, Accessed 2024.

[29] City of Palo Alto, "Approval of electric rate schedules effective july 1, 2019," https://www.cityofpaloalto.org/files/assets/public/v/1/agendas-minutes-reports/reports/city-manager-reports-cmrs/year-archive/2019/id-10217-mini-packet-05.02.2019.pdf, 2019, city Manager Report ID 10217.