

Online Station Assignment for Electric Vehicle Battery Swapping

Pengcheng You¹, Member, IEEE, John Z. F. Pang², Member, IEEE, and Steven H. Low², Fellow, IEEE

Abstract—This paper investigates the online station assignment for (commercial) electric vehicles (EVs) that request battery swapping from a central operator, i.e., in the absence of future information a battery swapping service station has to be assigned instantly to each EV upon its request. Based on EVs' locations, the availability of fully-charged batteries at service stations in the system, as well as traffic conditions, the assignment aims to minimize cost to EVs and congestion at service stations. Inspired by a polynomial-time offline solution via a bipartite matching approach, we develop an efficient and implementable online station assignment algorithm that provably achieves the tight (optimal) competitive ratio under mild conditions. Monte Carlo experiments on a real transportation network by Baidu Maps show that our algorithm performs reasonably well on realistic inputs, even with a certain amount of estimation error in parameters.

Index Terms—Electric vehicle, battery swapping, assignment, bipartite matching, online algorithm.

I. INTRODUCTION

A. Background, Motivation and Contributions

TRANSPORTATION consumes a tremendous amount of energy and meanwhile emits massive greenhouse gases, e.g., 28% of all energy in the US with over a quarter of greenhouse gas emission [1]. One way to drastically reduce its carbon footprint is electrifying transportation and increasing renewable generation of electricity. The rising global sales of EVs is promising; in the first quarter of 2017, sales jumped by 40% over the same period of 2016. China leads in both EV manufacturing and sales, with the first-quarter sale of 2017 grown by 41% over that of 2016 [2]. Notably, more than 150,000 commercial EVs were delivered in China in 2016, and most of them are fully electric taxis and buses.

Commercial EVs have been a driving force in EV adoption. They include taxis, buses, logistics vehicles, Zipcars (car-sharing service), etc. Unlike private cars, these commercial EVs are on the road most of the time. They consume a lot more energy, and hence have a greater potential in energy saving and

emission reduction if electrified, but also are far less flexible in refueling. The time-consuming process of battery charging is therefore a particularly severe pain point for commercial EVs.

Battery swapping, where an EV swaps a depleted battery at a service station for a fully-charged one, is more suitable for commercial EVs. This model has been well demonstrated in several pilot programs in China [3]. The advantages of battery swapping over charging are fourfold. *First*, it takes only minutes to swap a battery but often hours to recharge it. *Second*, the aggregation of charging loads reduces demand uncertainty compared with individual EV charging, simplifying power system operation. *Third*, the aggregation of charging loads allows the service stations greater flexibility in scheduling battery charging and providing ancillary services [4]. *Fourth*, batteries, as the most costly core of an EV, can be leased rather than purchased, tremendously lowering the expenditure for EV owners.

Battery swapping, however, also has its unique challenges. *First*, it requires standardization of vehicles, batteries, and swapping infrastructure, which has proven to be a difficult process [5]. *Second*, a new operation model is needed to address ownership, maintenance, and payment of shared batteries. These issues are less severe for commercial EVs which are typically under the management of a single entity. The State Grid of China, for instance, operates several city-size battery swapping based electric taxi programs, which exemplifies a vertically integrated system consisting of EVs, batteries and service stations.

A challenge in these systems is that EVs tend to swap their batteries at certain service stations that easily run short of fully-charged batteries to serve the demand. EVs occasionally have to wait in long queues to be served, which significantly degrades the efficiency of battery swapping. In order to alleviate the congestion at service stations, we propose to schedule EV battery swapping to fully utilize the (current/future) available fully-charged batteries in the system both temporally and spatially. Even in the centralized setting the scheduling of EV battery swapping is challenging in three aspects:

- It involves combinatorial optimization of station assignment that is computationally difficult.
- Modeling station congestion introduces nonlinear couplings.
- Station assignment has to be determined online, meaning without future information, we need to make an irrevocable service station assignment in real time to each requesting EV.

Manuscript received November 13, 2018; revised March 11, 2020 and July 21, 2020; accepted October 14, 2020. This work was supported by NSF under Grant CCF 1637598, Grant CPS ECCS 1739355, and Grant CPS ECCS 1932611. The Associate Editor for this article was B. De Schutter. (Corresponding author: Pengcheng You.)

Pengcheng You is with the Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, MD 21218 USA (e-mail: pcyou@jhu.edu).

John Z. F. Pang and Steven H. Low are with the Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: jzpang@caltech.edu; slow@caltech.edu).

Digital Object Identifier 10.1109/TITS.2020.3033731

To address these challenges, in this paper we consider a centralized system where EVs send their battery swapping requests to a central operator when their battery energy is running low. The operator assigns service stations in response to these requests based on locations of the requesting EVs and the availability of fully-charged batteries at service stations in the system. We design an efficient and implementable online algorithm for station assignment that aims to minimize cost to EVs and congestion at service stations. The *contributions* of this paper are highlighted as follows:

- We establish a practical system model and formulate an online station assignment problem.
- We first look at the offline version of the station assignment problem and reduce it to a standard bipartite matching problem, solvable in polynomial time.
- We then leverage online matching to derive an algorithm for the online station assignment problem that builds on the efficient offline solution. The online algorithm is proven to achieve a tight (optimal) competitive ratio of $2A - 1$ under mild conditions, where A is the number of battery swapping requests.
- Our online algorithm also performs well in average-case Monte Carlo tests that are simulated on a real transportation network by Baidu Maps, even with inexact parameter estimates.

B. Literature

The literature on EV battery swapping is relatively small, but has been growing rapidly in the past years. Zhang *et al.* [6] and Wang *et al.* [7] quantify certain aggregate characteristics of battery swapping stations. In particular, Zhang *et al.* [6] proposes to evaluate service capacities of a station by means of a Monte Carlo simulation approach while Wang and Wang [7] assesses the economic benefit for a station to provide frequency regulation service, both of which may provide guidelines for policymakers or investors. A queuing network model with an open queue of EVs and a closed queue of batteries is proposed in [8] to serve as a basic analytical framework for battery swapping service stations. Based on this model, an optimal charging strategy with quality of service guarantee is devised in [4]. Liu *et al.* [9] similarly models the charging and logistics of batteries in a battery swapping station using a closed-loop supply chain and optimally manages the system for maximum revenue. Gao *et al.* [10], Infante *et al.* [11] and Sarker *et al.* [12] handle the uncertainties in the operation of a battery swapping station, yet with different methodologies. Gao *et al.* [10] resorts to cutting-edge deep reinforcement learning algorithms to predict uncertain EV (electric buses in particular) visits and electricity prices. Both Infante *et al.* [11] and Sarker *et al.* [12] adopt a two-stage framework - Infante *et al.* [11] integrates in planning of battery resource allocation the stochastic EV station visits using clustered representative scenarios while Sarker *et al.* [12] focuses on day-ahead scheduling that robustly accounts for real-time uncertainty in battery swapping demand. In contrast, a deterministic operating model of a battery swapping station serving electric buses is

discussed in [13] and a highly efficient scheduling algorithm is designed for real-time implementation. This set of papers investigate battery swapping only on the station end.

Some other work takes the interaction between battery swapping stations and power grids into account. Both Sepetanc and Pandzic *et al.* [14] and Esmaeili *et al.* [15] establish a bi-level framework to capture this interaction. However, the focus is different - Sepetanc and Pandzic [14] models the decision on market participation of aggregated battery swapping stations as the upper level such that they behave strategically to maximize revenue; on the other hand, Esmaeili *et al.* [15] models the grid scheduling as the upper level that anticipates and accounts for the optimal decision making of battery swapping stations. Zheng *et al.* [16] focuses on the optimal layout design of battery swapping stations in distribution systems, and a criterion of life cycle cost is used to assess the placement and capacity of stations. Wang *et al.* [17] and Xie *et al.* [18] propose a strategy that utilizes controllable energy storage in battery swapping stations to supplement automatic generation control, such that a more reliable frequency regulation can be guaranteed. Kang *et al.* [19] devises a centralized charging strategy under a battery swapping scenario by considering optimal charging priorities and locations, with the aim of minimizing charging cost as well as power loss and voltage deviation of power networks. All the work above neglects *station assignment for battery swapping* from the EV perspective, which is the focus of this paper. A joint battery swapping and optimal power flow problem is solved in [20], [21] that first integrates the scheduling of battery swapping into the operation of a power network. However, the problem considered is local at a relatively large time scale that only involves one-slot optimization, leaving the scheduling issue in real-time operation unaddressed. Of high relevance to our work, Zhang *et al.* [22] considers a swapping service framework that allows reservations to alleviate service congestion. An empirical protocol is developed but it is generally difficult to derive analytical performance guarantee.

There is a large literature on online algorithm applications in the EV scheduling context. However, most, if not all, of these studies have been restricted to EV charging. For instance, a primal-dual online EV charging algorithm is adopted in [23] to deal with the volatility of real-time prices, such that a near-optimal charging cost can be guaranteed with little prior knowledge of future prices. An online auction protocol to coordinate EV charging is designed in [24] that allows EV owners to bid for charging power and time and provides high allocative efficiency. Guo *et al.* [25] investigates an online linear programming based algorithm for an adaptive EV charging network [26] (with a linear cost) and proposes conditions under which the algorithm achieves offline optimum or a bounded competitive ratio. Similarly, Zhao *et al.* [27] studies online EV charging with uncertain future EV arrivals and develops simple online algorithms that attain the optimal competitive ratios in terms of minimizing peak power consumption. The emerging technology of battery swapping, despite an alluring prospect, has not drawn much attention from the community and this paper complements

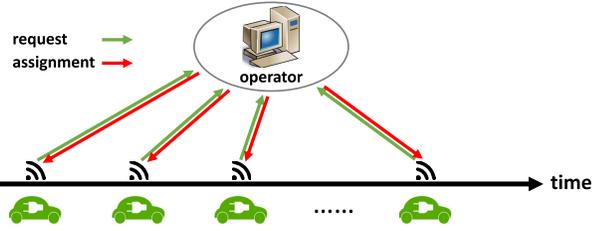


Fig. 1. Online battery swapping requests and assignments.

the existing literature by investigating the scheduling of EV battery swapping in a practical online setting.

II. SYSTEM MODEL

Consider an EV battery swapping program, exemplified by a city-size commercial taxi system. A central operator that also manages service stations assigns one station to each EV as it makes a battery swapping request. Suppose EVs are obligated to follow assignments and will instantly head towards the assigned stations to have their batteries swapped. Service stations accommodate battery swapping demand by unloading depleted batteries from EVs and providing them with fully-charged ones. Depleted batteries are recharged at the stations for future requests. We focus on real-time station assignment over a continuous time horizon $\mathbb{T} := [0, T]$, say an hour, where the supply of fully-charged batteries at each station is taken as a deterministic process, assuming that battery charging is exogenous and given.

Suppose battery swapping requests from EVs are received successively by the operator, as Fig. 1 shows. An EV a that requests battery swapping at time r_a will also reveal its current position σ_a and state of charge s_a to the operator. Without knowing future battery swapping demand, the operator has to determine an irrevocable station assignment for EV a immediately based on the EV state, station state as well as exogenous real-time traffic flows. Our goal is to devise a computationally efficient online station assignment algorithm with performance guarantee.

Let $\mathbb{A} := \{1, 2, \dots, A\}$ denote a given set of EVs that sequentially request battery swapping within \mathbb{T} . The whole set \mathbb{A} is available only in hindsight, while in practice EVs $a \in \mathbb{A}$, along with their states including request times r_a , positions σ_a and states of charge s_a , are observable sequentially only after their individual requests.¹ Let $\mathbb{I} := \{1, 2, \dots, I\}$ be the set of service stations managed by the operator. Each station i at location σ_i maintains a certain number of available (fully-charged) batteries to serve forthcoming battery swapping requests. Define $n_i(t) > 0$ as the number of available batteries at station i at time t . Let $n := (n_i(t), i \in \mathbb{I}, t \in \mathbb{T})$. Note that in the case of heavy demand when $n_i(t) < 0$, i.e., there is a deficit in available batteries to serve EVs, $-n_i(t)$ is also interpreted as the number of EVs waiting for batteries at station i at time t , which we shall term as *station congestion*.

An important practical observation is that EVs are not immediately at stations after receiving their station

assignments.² Given a city transportation network, we consider stable traffic conditions with edge-dependent travel time and distance. Let an external function $\tau(\sigma, \sigma', t)$, where σ and σ' are arbitrary positions, estimate the arrival time if a vehicle sets off from σ towards σ' at time t . This function is readily realizable by resorting to an off-the-shelf navigation service exemplified by Google Maps and Baidu Maps. Therefore, we assume this information is privy to the operator. Every time an EV a requests battery swapping, the operator gathers $\tau_{ai} := \tau(\sigma_a, \sigma_i, r_a)$, $i \in \mathbb{I}$. Let $\theta_{ai} := \tau_{ai} - r_a$ be the corresponding travel time. Here we use r_a to approximate the time when EV a sets off, which further requires that the station assignment be quickly determined. Meanwhile, τ_{ai} corresponds to a planned route with the travel distance denoted by d_{ai} . Suppose station i is assigned to EV a . Then it will arrive at time τ_{ai} and reduce $n_i(\tau_{ai})$ by one.³

The key decision variable is the station assignment to be determined upon each battery swapping request. Let $M := (M_a, a \in \mathbb{A})$ represent the station assignment variables for all EVs, where $M_a := (M_{ai}, i \in \mathbb{I})$ is the individual station assignment for EV a with

$$M_{ai} = \begin{cases} 1 & \text{if station } i \text{ is assigned to EV } a \\ 0 & \text{otherwise} \end{cases}$$

We require that only one station be assigned to each EV, i.e.,

$$M_{ai} \in \{0, 1\}, \quad a \in \mathbb{A}, i \in \mathbb{I} \quad (1a)$$

$$\sum_{i \in \mathbb{I}} M_{ai} = 1, \quad a \in \mathbb{A} \quad (1b)$$

At station i at time t , $n_i(t)$ increases by 1 when a charging battery becomes fully-charged (available), or decreases by 1 when an available battery is removed by an EV, as captured by

$$\Delta n_i(t) = c_i(t) - \sum_{a \in \mathbb{A}} M_{ai} \cdot 1(t = \tau_{ai}), \quad i \in \mathbb{I}, t \in \mathbb{T} \quad (2)$$

where $c_i(t)$ is the number of batteries that become newly available at station i at time t (which is assumed known *a priori*), and $1(z)$ is an indicator function for the predicate z . The second term on the right-hand side of (2) summarizes the impact of station assignments on the numbers of available batteries at stations. Therefore, $n_i(t)$ evolves according to

$$n_i(t)^+ = n_i(t)^- + \Delta n_i(t), \quad i \in \mathbb{I}, t \in \mathbb{T} \quad (3)$$

where $n_i(t)^+ := \lim_{y \rightarrow t^+} n_i(y)$ and $n_i(t)^- := \lim_{y \rightarrow t^-} n_i(y)$. Intuitively, the value of $n_i(t)$ steps up and down instantly in the continuous time horizon \mathbb{T} .

We assign stations to EVs for battery swapping with the goal of minimizing cost to EVs and congestion at service stations. Therefore, the system cost over the time horizon \mathbb{T} has two components. First, certain EV cost α_{ai} will be incurred for EV a to swap its battery at the assigned station i . For instance,

²This is one of the key points that distinguish our station assignment problem from canonical assignment problems since the impact of assignments here is hysteretic and implicit.

³Here we implicitly assume that EVs can have batteries swapped immediately when there are available batteries at a station.

¹Suppose ties are broken arbitrarily.

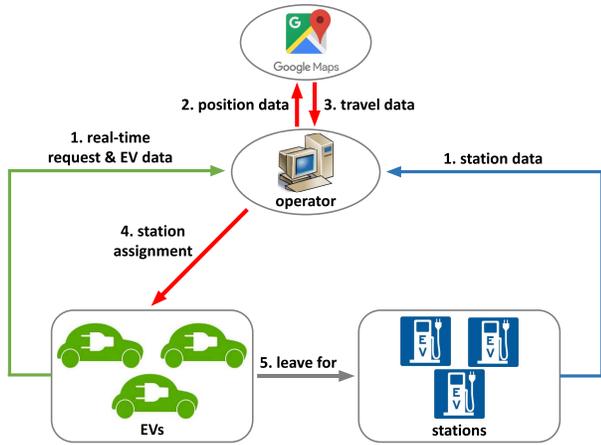


Fig. 2. Sequential information exchanges. A central operator (a) collects EV data along with a real-time battery swapping request as well as station data, (b) translates position data into travel data (traffic parameters) by means of a navigation service, and (c) determines a station assignment and instantly returns it to the requesting EV, which then leaves for the assigned station.

the cost can account for the corresponding travel distance, time, or a weighted sum of both, as well as the constant battery swapping service time. However α_{ai} is defined, we require that the planned route of EV a to station i minimizes it.⁴ The cost of EV a is then $\sum_{i \in \mathbb{I}} \alpha_{ai} M_{ai}$. Second, as explained above, the congestion of station i at time t can be represented by $\langle -n_i(t) \rangle^+$, where $\langle y \rangle^+ := \max\{y, 0\}$. Its integral congestion over \mathbb{T} is $\int_{t \in \mathbb{T}} \langle -n_i(t) \rangle^+ dt$. The goal of station assignment is to minimize both costs over all EVs and stations. The extra challenge is that they have to be determined in an online fashion:

Online station assignment problem

online decision variables: $M_a, a \in \mathbb{A}$

constraints: (1)

station state equations: (2)(3)

cost: $C(M) := \sum_{a \in \mathbb{A}} \sum_{i \in \mathbb{I}} \alpha_{ai} M_{ai} + \sum_{i \in \mathbb{I}} \int_{t \in \mathbb{T}} \langle -n_i(t) \rangle^+ dt$

Any weight to make the two cost components comparable can be incorporated in α_{ai} . The information exchanges involved in the system is illustrated in Fig. 2. Note that we ignore driving ranges of EVs for simplicity to focus on online assignment. In practice, it is possible to set a minimal driving range for requesting battery swapping so that feasibility will not be an issue.

In terms of the system cost, a set of station assignments M determined online can only be evaluated at the end of \mathbb{T} when the final system state n is observed. We first show in the next section an efficient solution to obtain offline optimal station assignments. Inspired by that, we then propose our online station assignment algorithm and analyze its performance.

III. OFFLINE SOLUTION

Assume global information, i.e., all the battery swapping requests in \mathbb{T} , is known *a priori*. The corresponding

⁴If it is an electric taxi with passengers, or a destination, α_{ai} could be interpreted as the extra EV cost to take a detour.

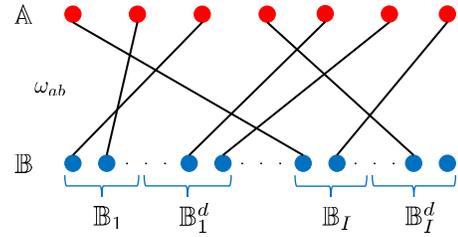


Fig. 3. A bipartite matching.

offline optimal station assignment problem is

$$\begin{aligned} \min_{M,n} \quad & \sum_{a \in \mathbb{A}} \sum_{i \in \mathbb{I}} \alpha_{ai} M_{ai} + \sum_{i \in \mathbb{I}} \int_{t \in \mathbb{T}} \langle -n_i(t) \rangle^+ dt \\ \text{s.t.} \quad & (1)(2)(3) \end{aligned} \quad (4)$$

Note that even solving the offline problem (4) is challenging. It is seemingly a mixed-integer program with nonlinearity and couplings across EVs, stations and time. We show however it is solvable in polynomial time by viewing the problem from a bipartite matching perspective.

The key observation is that the offline problem (4) is reducible to a problem of minimum weight maximum bipartite matching. The matching is not defined between EVs and stations. Instead, we let EVs and individual batteries be vertices on each bipartition and endow each possible matchup with a weight based on its contribution to the system cost.

Define a bipartite graph $\mathbb{G} = (\mathbb{A} \cup \mathbb{B}, \mathbb{E})$, where \mathbb{A} and \mathbb{B} are the bipartition of the vertex set, representing respectively the EV set and the battery set, and $\mathbb{E} = \mathbb{A} \times \mathbb{B}$ is the set of all possible edges between \mathbb{A} and \mathbb{B} with endowed weights $\omega := (\omega_{ab}, (a, b) \in \mathbb{E})$, as shown in Fig. 3. The goal is to match each vertex in \mathbb{A} with one unmatched vertex in \mathbb{B} (as we will show later, $|\mathbb{A}| \leq |\mathbb{B}|$) such that the collective weight is minimized. Therefore, the minimum weight maximum matching problem defined on \mathbb{G} is

$$\min_x \quad \sum_{(a,b) \in \mathbb{E}} \omega_{ab} x_{ab} \quad (5a)$$

$$\text{s.t.} \quad \sum_b x_{ab} = 1, \quad a \in \mathbb{A} \quad (5b)$$

$$\sum_a x_{ab} \leq 1, \quad b \in \mathbb{B} \quad (5c)$$

$$x_{ab} \in \{0, 1\}, \quad (a, b) \in \mathbb{E} \quad (5d)$$

where $x := (x_{ab}, (a, b) \in \mathbb{E})$ is the matching to be decided.⁵ $x_{ab} = 1$ if vertex a is matched with vertex b , and 0 otherwise. An instance of the matching problem (5) is specified by the graph \mathbb{G} and weights ω .

Given an instance of the offline problem (4), we now construct an instance of the matching problem (5) such that an optimal matching of (5) yields an optimal station assignment of (4).

In \mathbb{G} , \mathbb{A} is the same set of EVs in (4), while the battery set \mathbb{B} is inferred from the parameters $(n_i(0), i \in \mathbb{I})$ and $(c_i(t), i \in \mathbb{I}, t \in \mathbb{T})$. Let $\mathbb{B} := \bigcup_{i \in \mathbb{I}} (\mathbb{B}_i \cup \mathbb{B}_i^d)$. The composition

⁵ x can also be deemed as the set of edges that represent matchups in x .

of \mathbb{B} reveals one characterization for each battery b , i.e., the station i_b which it belongs to. A battery b in \mathbb{B}_i corresponds to a real battery at station i that turns available anytime within \mathbb{T} . A battery b in \mathbb{B}_i^d is a dummy battery that arises from the possibility that an EV has not had its battery swapped by the end of \mathbb{T} , i.e., $n_i(T) < 0$ for some station i . In this case, we add dummy batteries as matching options to guarantee feasibility and consistency with the offline problem (4). In addition, a battery b is also characterized by the time when it turns available, denoted by ρ_b , e.g., $\rho_b = 0$ for batteries that are already available at the beginning. Let $\rho_b = T$ for dummy batteries, which contributes to the uniform definition of ω_{ab} , as we will show later. Each battery $b \in \mathbb{B}$ is then labeled with a tuple (i_b, ρ_b) .

Remark 1: \mathbb{B}_i^d is used to make up the shortfall when there is a larger battery swapping demand than the supply of real batteries at station i in \mathbb{T} . $|\mathbb{B}_i^d| := \max\{|\mathbb{A}| - |\mathbb{B}_i|, 0\}$ is defined in order to maintain the consistent solution space when mapping x back to M .

Given \mathbb{G} , the weight ω_{ab} of each matchup (a, b) is interpreted as the incremental system cost it contributes to if EV a is matched with battery b , i.e., station i_b is assigned to EV a . To determine ω_{ab} , the main idea is to translate station congestion into the waiting time each EV suffers. Note that station congestion results from the fact that some EVs arrive but have to wait for batteries to turn available. Regardless of whether the battery that an EV is matched with is real or dummy, its waiting time constitutes a period of congestion. In other words, the temporal integral congestion of a station is equivalent to the aggregate waiting time of all EVs that the station is assigned to. This intuition guides us to set weights ω as follows⁶: define $\omega_{ab} := \alpha_{ab} + \max\{\rho_b - \tau_{ab}, 0\}$, $(a, b) \in \mathbb{E}$. Here, α_{ab} is the EV cost and $\max\{\rho_b - \tau_{ab}, 0\}$ is the duration for which EV a waits until battery b turns available. Particularly, since \mathbb{T} is the whole time horizon, EVs matched with dummy batteries will wait until the end after their arrivals and thus the weight reduces to $\omega_{ab} = \alpha_{ab} + (T - \tau_{ab})$ if battery b is dummy. The definition of ω_{ab} implies that the goal of station assignment is to optimally trade off EVs' cost on the road and waiting time at stations for battery swapping.

From above, given an instance of the offline problem (4), a well defined matching problem (5) requires $O(N^2)$ time to establish \mathbb{G} and ω , where $N := |\mathbb{B}|$. On the other hand, suppose we have an optimal matching x^* of (5), an optimal station assignment M^* of (4) is straightforward - assign to each EV the station which its matched battery belongs to:

$$M_{ai}^* = \sum_{b \in \mathbb{B}_i \cup \mathbb{B}_i^d} x_{ab}^*, \quad a \in \mathbb{A}, i \in \mathbb{I} \quad (6)$$

which is obtainable in $O(N)$ time.

So far we have reduced the offline optimal station assignment problem (4) to the standard minimum weight maximum matching problem (5), which is well known to be polynomial-time solvable, e.g., by the Hungarian algorithm [28], [29]. Therefore, the offline problem (4) can be efficiently handled via the above bipartite matching approach.

⁶Without ambiguity, $\alpha_{ab} := \alpha_{ai_b}$, $\tau_{ab} := \tau_{ai_b}$ and $d_{ab} := d_{ai_b}$ are applied.

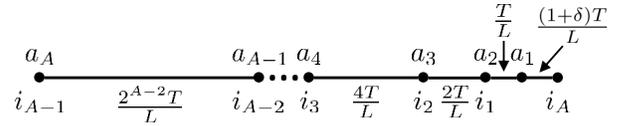


Fig. 4. An example where greedy approaches perform poorly.

This inspires us to come up with an online station assignment algorithm by building on the online version of the minimum weight maximum matching problem (5), as we explain in the next section.

IV. ONLINE STATION ASSIGNMENT

An offline solution to (4) is not implementable in general, but can serve as an optimal baseline for any online station assignment algorithm. The performance of an online algorithm is commonly measured by its *competitive ratio*, the worst possible ratio between its performance and the offline optimum over all input instances. For the online station assignment problem, an algorithm that outputs M^{on} is *competitive* if its competitive ratio $\gamma_M := \sup \frac{C(M^{\text{on}})}{C(M^*)}$ is bounded, where the supremum is taken over all input instances.

One possible online algorithm is a greedy approach that discretizes the time horizon \mathbb{T} into slots and optimizes the station assignments for all in-slot battery swapping requests simultaneously at the end of each slot. The time granularity of a slot becomes significant as it leads to a tradeoff between the approximation to offline optimum and implementability. The larger the time granularity is, the better the local optimal station assignment problem at each slot approximates (4), which is however harder to implement. In the case of a continuous time horizon as we consider here, a greedy approach will simply make a best station assignment for each individual battery swapping request. However, the following example illustrates that the competitive ratio of a greedy approach is lower bounded by $2^A - 1$.

Consider a network of a line topology where A stations, each with only one fully-charged battery, are located at nodes as Fig. 4 shows. The EV cost incurred to travel through each edge is also marked therein. Suppose in total A pending battery swapping requests are indexed in chronological order. Put EV a_1 at the node in between stations i_1 and i_A , which is slightly closer to station i_1 by $\frac{\delta T}{L}$ with δ and L being both constants. A greedy approach assigns station i_1 to EV a_1 myopically. An adversarial parameterization then puts EV a_k exactly where station i_{k-1} is, $k = 2, 3, \dots, A$. Note that as long as $L \gg 2^{A-1} + \delta$ and $r_{a_A} \rightarrow 0$, which basically mean assigning one station to two EVs will incur huge cost, a greedy approach always looks for an unassigned station with the least EV cost incurred, i.e., station i_k is assigned to EV a_k , $k = 1, 2, \dots, A$, with the total system cost of $\frac{(2^A - 1 + \delta)T}{L}$, while the offline optimum, which achieves the minimal system cost of $\frac{(1 + \delta)T}{L}$, assigns station i_A to EV a_1 and tells other EVs to stay still. In this specific case, $\frac{C(M^{\text{on}})}{C(M^{\text{off}})} = \frac{2^A - 1 + \delta}{1 + \delta} \rightarrow 2^A - 1$ as $\delta \rightarrow 0$, which suggests $\gamma_M \geq 2^A - 1$, i.e., the competitive ratio of a greedy approach for the online station assignment problem is lower bounded by $2^A - 1$.

From above, a greedy approach is usually unsatisfactory and yields a poor competitive ratio, since a local optimum can incur a high cost globally. Instead, we aim to exploit the specific problem structure and design a *competitive* online station assignment algorithm.

We follow the same pipeline in Sec. III using the bipartite matching approach. First, focus on the online version of the minimum weight maximum weight matching problem (5): suppose the full graph \mathbb{G} in Fig. 3 is not known a priori, yet \mathbb{B} is available to the operator. \mathbb{A} reveals one EV vertex as well as its incident edges with weights when a battery swapping request emerges, upon which the operator selects an unmatched battery vertex as the matchup with an objective of minimizing the aggregate weight.

Online matching problem

online decision variables: $x_a, a \in \mathbb{A}$

constraints: (5b)-(5d)

aggregate weight: $W(x) := \sum_{(a,b) \in \mathbb{E}} \omega_{ab} x_{ab}$

Similarly, let $\gamma_x := \sup \frac{W(x^{\text{on}})}{W(x^*)}$ be the competitive ratio of an algorithm for the online matching problem, where x^{on} is its output. Unfortunately, γ_x can not be bounded for the general online matching problem [30]. However, as we will show below, our problem possesses specific structural properties that contribute to a theoretical performance bound and meanwhile fit well in a practical setting.

Recall that the EV cost previously defined is quite general. Now we pin it down explicitly as a weighted sum of travel time and distance, i.e., $\alpha_{ab} := \alpha_1 \theta_{ab} + \alpha_2 d_{ab}$. Therefore, according to the analysis in Sec. III, the weights on the incident edges are

$$\omega_{ab} = \alpha_1 \theta_{ab} + \alpha_2 d_{ab} + \max\{\rho_b - \tau_{ab}, 0\}, \quad (a, b) \in \mathbb{E}$$

We propose our online station assignment algorithm. Let $a_k, k = 1, 2, \dots, A$, be the k^{th} EV in chronological order that requests battery swapping (also the k^{th} EV vertex that emerges in the graph \mathbb{G}). Denote the *interim* offline minimum weight maximum matching right after a_k 's emergence as x_k^{off} , i.e., x_k^{off} consists of offline optimal matchups for (a_1, a_2, \dots, a_k) assuming there is no future arrival, and the corresponding online matching by our algorithm as x_k^{on} . Let M_k^{off} and M_k^{on} be the corresponding station assignments with respect to x_k^{off} and x_k^{on} , which can be computed by (6). Denote the final decision outputs as $x^{\text{off}} := x_A^{\text{off}}, x^{\text{on}} := x_A^{\text{on}}, M^{\text{off}} := M_A^{\text{off}}$ and $M^{\text{on}} := M_A^{\text{on}}$. Hence x^{off} and M^{off} are the true offline optima to benchmark against when all A EVs have emerged. The online station assignment algorithm is explicitly demonstrated in Algorithm 1.

Here are a few notes about Algorithm 1. In step 5, x_k^{off} is efficiently computed by solving (5) with the first k EVs. \oplus in step 6 is the operator of symmetric difference. Basically, the symmetric difference of two sets is the set of elements (edges here) that are in either of the sets and not in their intersection. The remedy of x_k^{off} in step 7 is trivial, as explained in Property 1 below. Algorithm 1 tracks the difference between two consecutive offline matchings instead of being greedy. It has the following two properties that help guarantee a bounded competitive ratio under mild conditions:

Algorithm 1: Online Station Assignment Algorithm

- 1 **Input:** $(a_k, k = 1, 2, \dots, A)$, and $(\omega_{ab}, (a, b) \in \mathbb{E})$ after a_k emerges;
- 2 **Output:** M^{on} ;
- 3 **Initialization:** $k \leftarrow 1, x_0^{\text{off}} \leftarrow \emptyset, x_0^{\text{on}} \leftarrow \emptyset, M_0^{\text{on}} \leftarrow \emptyset$;
- 4 **while** a_k emerges **do**
- 5 Compute x_k^{off} ;
- 6 $P_k \leftarrow x_k^{\text{off}} \oplus x_{k-1}^{\text{off}}$;
- 7 Remedy x_k^{off} such that P_k consists of only an odd-length path from a_k to a battery vertex b_k ;
- 8 Match a_k with b_k to form x_k^{on} from x_{k-1}^{on} ;
- 9 Assign station i_{b_k} to EV a_k to obtain M_k^{on} from M_{k-1}^{on} ;
- 10 $k \leftarrow k + 1$;
- 11 **end while**
- 12 **return** $M^{\text{on}} \leftarrow M_A^{\text{on}}$

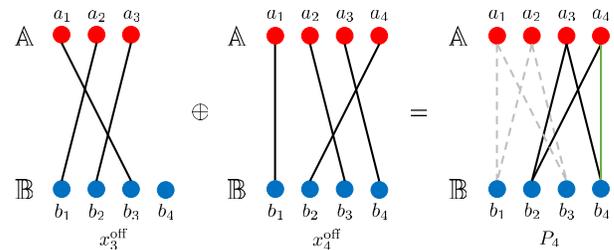


Fig. 5. An illustrative example of P_k : P_4 consists of a cycle (grey dash lines), which can be removed through the remedy of x_4^{off} , and a path from a_4 to b_4 (black solid lines).

Property 1: Without loss of generality, $P_k = x_k^{\text{off}} \oplus x_{k-1}^{\text{off}}$ can be taken as an odd-length path from a_k to b_k .

Proof: The potential rematching in x_k^{off} that differs from x_{k-1}^{off} arises from the emergence of a_k . Conceivably, their symmetric difference P_k covers vertices that have either one (matched only in one matching) or two (matched in both matchings with different matchups) incident edges that link the opposite bipartition. Then P_k must be the disjoint union of paths or cycles, or both. See Fig. 5 for an illustrative example.

Since a_k is only matched in x_k^{off} , it has only one incident edge and cannot be contained in a cycle (as vertices in cycles all have two incident edges), meaning that a_k must be one end of a path. If this is an even-length path, the other end will be another EV vertex $a_\kappa, \kappa \in \{1, 2, \dots, k-1\}$, which is impossible as $a_\kappa, \kappa = 1, 2, \dots, k-1$, has been matched in both x_k^{off} and x_{k-1}^{off} – they are either excluded in P_k or linked by two incident edges. As a result, P_k contains a unique odd-length path that starts from a_k and ends with a battery vertex b_k .

In addition, cycles (of even length), if any, will have the identical aggregate weight on their odd edges and even edges, which respectively belong to x_k^{off} and x_{k-1}^{off} ; otherwise, one of the offline matchings can be improved by exchanging the edges and is not optimal. Therefore, we can exchange the edges in x_k^{off} for those in x_{k-1}^{off} without losing optimality and get rid of cycles. Similarly, even-length paths, if any, which can only start and end at the battery bipartition, will also have

the identical aggregate weight on their odd edges and even edges. Again, we can remedy x_k^{off} by exchanging the edges to remove even-length paths. Therefore, P_k is an odd-length path from a_k to b_k without loss of generality. \square

Property 2: b_k is not matched in x_{k-1}^{on} .

Proof: All the vertices except a_k and b_k along the odd-length path have two incident edges that represent their matchups in x_k^{off} and x_{k-1}^{off} , respectively. Therefore, the path alternates the edges in x_k^{off} with those in x_{k-1}^{off} (as it starts with a_k which is only matched in x_k^{off}). Due to its odd length, the only incident edge of b_k belongs to x_k^{off} , i.e., b_k is not matched in x_{k-1}^{off} .

By induction it can be verified that the battery vertices matched in x_{k-1}^{on} are the same as those matched in x_{k-1}^{off} ; suppose it holds for x_κ^{off} and x_κ^{on} , $\kappa = 0, 1, \dots, k-2$; when $x_{\kappa+1}^{\text{off}}$ incorporates an extra $b_{\kappa+1}$ on the basis of x_κ^{off} , $x_{\kappa+1}^{\text{on}}$ also includes $b_{\kappa+1}$ according to Algorithm 1. Therefore, b_k is not matched in x_{k-1}^{on} either. \square

Property 2 suggests the battery vertices matched in x^{on} by Algorithm 1 is consistent with those matched optimally in hindsight. This is the key principle that enables Algorithm 1 to hedge against worst-case parameterization and avoid unbounded performance loss.

Theorem 1: Algorithm 1 achieves a competitive ratio of $\gamma_x = 2A - 1$ for the online matching problem, i.e., $\frac{W(x^{\text{on}})}{W(x^{\text{off}})} \leq 2A - 1$, if either of the following conditions holds:

- 1) EV cost consists of only travel time, i.e., $\alpha_2 = 0$, and moreover, $\alpha_1 \geq 1$;
- 2) There is a positive correlation between travel time and distance, i.e., for $\forall a, a' \in \mathbb{A}, b \in \mathbb{B}$, $\theta_{ab} \leq \theta_{a'b} \iff d_{ab} \leq d_{a'b}$, and moreover, $d_{a'b} - d_{ab} \geq \frac{1-\alpha_1}{\alpha_2}(\theta_{a'b} - \theta_{ab})$.

Proof: We first make some key observations.

Observation 1: $W(x_1^{\text{off}}) \leq W(x_2^{\text{off}}) \leq \dots \leq W(x_A^{\text{off}})$.

Otherwise, there exist matchings that can be improved and are not offline optimal.

Observation 2: EV cost satisfies $\alpha_{ab} \leq \alpha_{a'b} + \alpha_{ab'} + \alpha_{a'b'}$, $\forall a, a' \in \mathbb{A}, b, b' \in \mathbb{B}$.

By definition, α_{ab} is the minimal EV cost to travel from the position of EV a to station i_b in terms of routing. If Observation 2 doesn't hold, α_{ab} is not minimal because EV a can take a detour $a \rightarrow b' \rightarrow a' \rightarrow b$ that bypasses b' and a' with strict less cost.

Based on the above observations, we show if a emerges later than a' (EV a requests battery swapping later than EV a'), i.e., $r_a > r_{a'}$, then the weights on their incident edges satisfy $\omega_{ab} \leq \omega_{a'b} + \omega_{ab'} + \omega_{a'b'}$, $\forall b, b' \in \mathbb{B}$, when either of the above conditions hold.

Given either condition, it is trivial to have $\theta_{ab} \leq \theta_{a'b} + \theta_{ab'} + \theta_{a'b'}$, which leads to two possible cases:

- θ_{ab} is strictly larger than $\theta_{a'b}$. Conspicuously, $\tau_{ab} = r_a + \theta_{ab} > r_{a'} + \theta_{a'b} = \tau_{a'b}$, i.e., EV a arrives at station i_b later than EV a' . Therefore, EV a will suffer no longer waiting time than EV a' to snatch battery b , i.e., $\max\{\rho_b - \tau_{ab}, 0\} \leq \max\{\rho_b - \tau_{a'b}, 0\}$. Then naturally comes

$$\begin{aligned} \omega_{ab} &= \alpha_{ab} + \max\{\rho_b - \tau_{ab}, 0\} \\ &\leq \alpha_{a'b} + \max\{\rho_b - \tau_{a'b}, 0\} + \alpha_{ab'} + \alpha_{a'b'} \\ &\leq \omega_{a'b} + \omega_{ab'} + \omega_{a'b'} \end{aligned} \quad (7)$$

- θ_{ab} is no larger than $\theta_{a'b}$. If EV a arrives at station i_b no earlier than EV a' , this is basically the same as the first case. Thus we care about situations where EV a arrives earlier than EV a' , i.e., $\tau_{ab} < \tau_{a'b}$. As expected, EV a' will suffer no longer waiting time than EV a , but note that $\max\{\rho_b - \tau_{ab}, 0\} - \max\{\rho_b - \tau_{a'b}, 0\} \leq \tau_{a'b} - \tau_{ab}$. Moreover,

$$\begin{aligned} r_a &= \tau_{ab} - \theta_{ab} > \tau_{a'b} - \theta_{a'b} = r_{a'} \\ \implies \theta_{a'b} - \theta_{ab} &> \tau_{a'b} - \tau_{ab} \\ &\geq \max\{\rho_b - \tau_{ab}, 0\} - \max\{\rho_b - \tau_{a'b}, 0\} \end{aligned} \quad (8)$$

If the first condition holds,

$$\begin{aligned} \alpha_{a'b} - \alpha_{ab} &= \alpha_1(\theta_{a'b} - \theta_{ab}) > \tau_{a'b} - \tau_{ab} \\ &\geq \max\{\rho_b - \tau_{ab}, 0\} - \max\{\rho_b - \tau_{a'b}, 0\} \\ \implies \omega_{ab} &< \omega_{a'b} \leq \omega_{a'b} + \omega_{ab'} + \omega_{a'b'} \end{aligned} \quad (9)$$

If the second condition holds,

$$\begin{aligned} d_{a'b} - d_{ab} &\geq \frac{1-\alpha_1}{\alpha_2}(\theta_{a'b} - \theta_{ab}) \\ \implies \alpha_1(\theta_{a'b} - \theta_{ab}) &+ \alpha_2(d_{a'b} - d_{ab}) \\ &\geq \theta_{a'b} - \theta_{ab} > \tau_{a'b} - \tau_{ab} \\ &\geq \max\{\rho_b - \tau_{ab}, 0\} - \max\{\rho_b - \tau_{a'b}, 0\} \\ \implies \omega_{ab} &< \omega_{a'b} \leq \omega_{a'b} + \omega_{ab'} + \omega_{a'b'} \end{aligned} \quad (10)$$

From above, $\omega_{ab} \leq \omega_{a'b} + \omega_{ab'} + \omega_{a'b'}$, $\forall b, b' \in \mathbb{B}$, is satisfied for $\forall a, a' \in \mathbb{A}$ with $r_a > r_{a'}$.

Corollary 1: If EV a is the latest one that requests battery swapping, for $\forall b \in \mathbb{B}$, ω_{ab} is no larger than the aggregate weight of a path (composed of the existing edges in \mathbb{G}) that connects vertices a and b .

Therefore, $\omega_{a_k b_k} \leq W(P_k) \leq W(x_k^{\text{off}}) + W(x_{k-1}^{\text{off}}) \leq 2W(x_A^{\text{off}})$, $k = 2, 3, \dots, A$. Since the first step is greedy, $\omega_{a_1 b_1} = W(x_1^{\text{off}}) \leq W(x_A^{\text{off}})$. Adding them all up leads to $W(x^{\text{on}}) = \sum_{k=1}^A \omega_{a_k b_k} \leq (2A - 1)W(x^{\text{off}})$, i.e., $\frac{W(x^{\text{on}})}{W(x^{\text{off}})} \leq 2A - 1$. \square

Remark 2: Under the conditions in Theorem 1, the competitive ratio scales linearly with the number of battery swapping requests (see Appendix A for an exponential competitive ratio in the general case without the conditions). These conditions are easy to enforce in practice. For instance, set $\alpha_1 \geq 1$ (weight on travel time) and then either let $\alpha_2 = 0$, e.g., when travel time and distance are irregularly correlated during heavy traffic congestion (long travel time but short travel distance), or assume the positive correlation between travel time and distance holds and properly set an α_2 , e.g., when traffic flows are smooth.

Although the competitive ratio attained by Algorithm 1 seems unduly pessimistic, it is tight as the following theorem states.

Theorem 2: No deterministic online algorithm is able to achieve a competitive ratio better than $2A - 1$ for the online matching problem.

Proof: Consider the following case: we know there are in total A pending battery swapping requests, indexed in chronological order. A star transportation network of one central node and A identical leaf nodes is shown in Fig. 6.

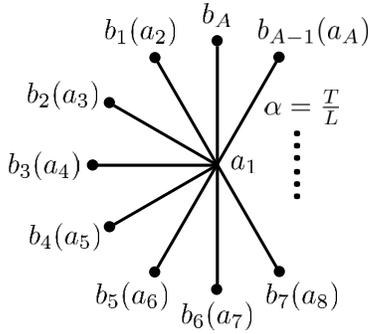


Fig. 6. The competitive ratio is tight in this case.

All edges represent unique paths that connect two corresponding nodes and are homogeneous in the sense that the same amount of EV cost $\alpha = \frac{T}{L}$ is incurred to travel through each edge, where L is a constant. Assume at each leaf node there is a station with only one fully-charged battery. Certainly $A - 1$ dummy batteries can be added at each station to accommodate all EVs. An EV has to wait until the end of the time horizon \mathbb{T} if matched with a dummy battery, while it can have its battery swapped immediately upon arrival if matched with a real battery.

First assume EVs will prioritize real batteries. Suppose EV a_1 is at the central node, a deterministic online algorithm matches it with an arbitrary real battery at a leaf node, say b_1 without loss of generality, thus $\omega_{a_1 b_1} = \frac{T}{L}$. Then the worst-case parameterization is to put EV a_k exactly where real battery b_{k-1} is, $k = 2, 3, \dots, A$, as Fig. 6 shows, thus forcing $\omega_{a_k b_k} = \frac{2T}{L}$, $k = 2, 3, \dots, A$. As a result, any deterministic online algorithm will achieve an aggregate weight of $\frac{(2A-1)T}{L}$ at minimum, while the offline optimum is $\frac{T}{L}$ by matching a_1 with b_A and a_k with b_{k-1} , $k = 2, 3, \dots, A$. Therefore, this case indicates $\gamma_x \geq 2A - 1$ by any deterministic online algorithm.

One would argue that EVs are not necessarily matched with real batteries. Then we show it will do no better by allowing matching EVs with dummy batteries. To avoid being trapped in the above case, the best countermeasure is to match a_1 with an arbitrary dummy battery and a_k with b_{k-1} , $k = 2, 3, \dots, A$. This matching achieves an aggregate weight of $\frac{T}{L} + (T - \tau_{a_1})$ at minimum when at least one dummy battery is matched, where τ_{a_1} is the arrival time of EV a_1 at any station (due to homogeneity). However, if $L \rightarrow \infty$ and $\tau_{a_1} \rightarrow 0$ (an extreme case where EV a_1 and all stations gather closely enough), $\gamma_x \rightarrow L + 1$, which means the competitive ratio can be unboundedly worse as $L \rightarrow \infty$.

Hence, the competitive ratio $\gamma_x = 2A - 1$ of Algorithm 1 is tight as no deterministic online algorithm is able to perform better. \square

It turns out the competitive analysis of Algorithm 1 in terms of the online matching problem is also extendable to the primal online station assignment problem:

Theorem 3: Algorithm 1 achieves a tight competitive ratio of $\gamma_M = 2A - 1$ for the online station assignment problem, given the same conditions in Theorem 1.

Proof: From Sec. III, it is obvious that $C(M^{\text{off}}) = W(x^{\text{off}})$. However, the online station assignment M^{on} , attained from x^{on} by (6), may be executed in a different way from the expected matching x^{on} due to the fact that EVs do not necessarily obtain the batteries matched with them.⁷

As characterized in the online station assignment problem, an EV that follows the station assignment will have its battery swapped immediately if any battery is available at the station upon its arrival. For instance, consider one single station, suppose EV 1 requests battery swapping earlier than EV 2, and is matched with battery 1, which turns available at time 3, while it arrives at time 5. EV 2, which arrives at time 3, is however matched with battery 2, which turns available at time 6, due to its late request. Apparently, $W(x^{\text{on}}) = \alpha_{11} + \alpha_{22} + 3$ while $C(M^{\text{on}}) = \alpha_{12} + \alpha_{21} + 1$, where $\alpha_{11} = \alpha_{12}$, $\alpha_{21} = \alpha_{22}$. The matching is revocable in practice, not as assumed when we analyze through the bipartite matching approach.

The simple case indicates that the revocation is spontaneous to rematch EVs that have already arrived with currently available batteries in real implementation. From the system point of view, the revocation, if any, always improves performance, i.e., $C(M^{\text{on}}) \leq W(x^{\text{on}})$. Therefore, $\frac{C(M^{\text{on}})}{C(M^{\text{off}})}$ is upper bounded by $\frac{W(x^{\text{on}})}{W(x^{\text{off}})}$. Moreover, the tightness of $\gamma_M = 2A - 1$ still holds, as verified by the case in Fig. 6. \square

As indicated above, although the bipartite matching approach enables us to obtain an analytical competitive ratio, the actual performance may be better. The potential improvement arises from the fact that the online matching tends to idle available batteries at a station when EVs with later arrivals request earlier, which however can be spontaneously revoked by the online station assignment. Essentially, the online matching problem is more specific yet easier to tackle. In practice, EVs only care about the stations assigned rather than the batteries matched. This suggests that, when we analyze using the bipartite matching approach, a non-optimal online decision is acceptable and sometimes beyond expectation as it may still correspond to the optimal station assignment.

V. NUMERICAL RESULTS

Our online station assignment algorithm through bipartite matching is justified analytically to be *competitive* under realistic conditions. Theoretically, it is able to accommodate all possible input instances and bound the worst-case performance tightly. Therefore, future uncertainties are well addressed. Even though the worst-case performance of our algorithm is linear in A , its average-case performance seems to scale much more slowly, as our simulation results suggest.

We use the Monte Carlo method to randomly generate test cases and examine the distribution of performance ratios by our algorithm compared with offline optima. The basic setup includes a continuous time horizon of 100 minutes, 5 stations uniformly sited in an urban area as Fig. 7 shows, and 100 pending battery swapping requests. The uniform randomness is

⁷In this paper we use the bipartite matching approach for analysis only and therefore do not require the matching to be executed in practice. However, it is also executable with certain reservation mechanisms. In that case, the station assignment problem is specified to be a battery matching problem.

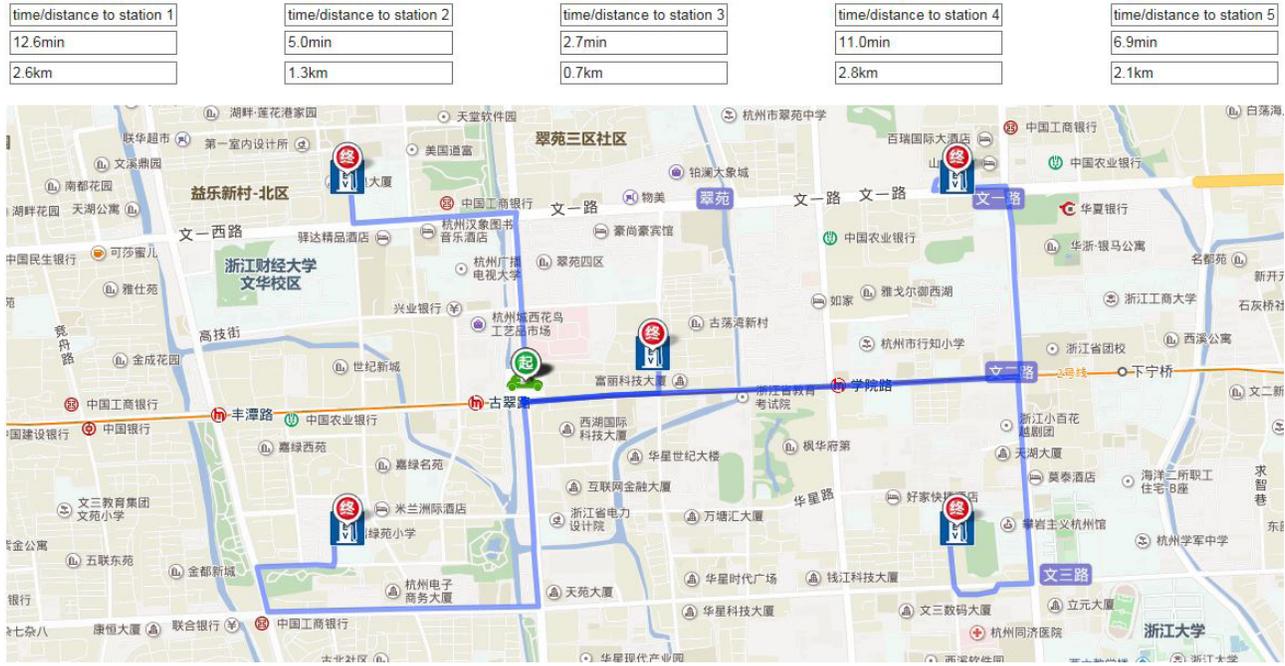


Fig. 7. An instance of parameter input through external routing after an EV requests battery swapping.

employed to generate the request times ($r_a \leq 85$ minutes) and positions of EVs in the square area, the number of real batteries at each station ($10 \leq |B_i| \leq 15$) and the times when these batteries turn available ($\rho_b \leq 100$ minutes). By means of the routing service provided by Baidu Maps [31], the parameters of travel time and distance to all the stations are readily attainable after an EV requests battery swapping, as shown in Fig. 7. For convenience, let $\alpha_1 = 1$ and $\alpha_2 = 0$, i.e., the first condition in Theorem 1 holds.

Our algorithm is tested on 1000 randomly generated cases, and the performance ratios $\frac{W(x^{on})}{W(x^{off})}$ of online matching w.r.t offline matching are recorded in Fig. 8 (top). Overall, the mean ratio is 1.255 with a variance of 0.003 and our algorithm maintains the performance ratios below 1.3 in 79.2% of the test cases. As Theorem 3 suggests, the actual performance ratios $\frac{C(M^{on})}{C(M^{off})}$ of online station assignment w.r.t offline station assignment may be better, which is confirmed by Fig. 8 (middle). Indeed, in 81.0% of the test cases, the performance ratios are maintained below 1.3, with a mean of 1.251 and a variance of 0.003. This improvement is captured in Fig. 8 (bottom). On average, a 0.30% relative performance improvement is observed over the performance ratios $\frac{W(x^{on})}{W(x^{off})}$.

Our algorithm depends on external traffic information and therefore its performance could be heavily affected by inaccurate estimation of the arrival time function $\tau(\cdot)$. In the above 1000 cases we add uniform white noise to simulate estimation error in EVs' arrival times at stations and test the robustness of our algorithm. Fig. 9 shows that the algorithm performance is not sensitive to the error in estimated arrival times (or travel times). Specifically, four scenarios are considered where the function $\tau(\cdot)$ suffers from $\pm 5\%$, $\pm 10\%$, $\pm 15\%$ and $\pm 20\%$ estimation error, respectively. In each scenario our algorithm is implemented based on the inaccurate arrival times, and the

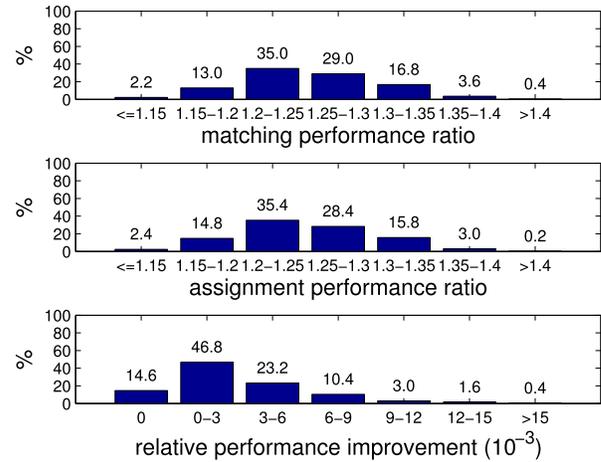


Fig. 8. Monte Carlo tests with perfect parameter estimation.

actual performance ratios are displayed in Fig. 9. Compared with the perfect-estimation scenario in Fig. 8 (middle), overall performance degrades as estimation error increases, as shown in Fig. 10 and Table I. However, even with $\pm 20\%$ estimation error, the mean performance ratio is only 1.267 (relative degradation 1.354%) with a variance of 0.003, and still in 74.4% of the test cases the ratios are maintained below 1.3, as Fig. 9 (bottom) shows. Fig. 11 displays the distribution of the relative performance degradation for all individual Monte Carlo test cases. Surprisingly, it indicates that in a high percentage of the test cases (43.4%, 40.6%, 38.2% and 36.2% for the four scenarios, respectively) the performance of our algorithm does not degrade at all or even improves using the inaccurate estimation. The intuition is that the arrival sequence of EVs at a station has a larger impact on performance than the precise

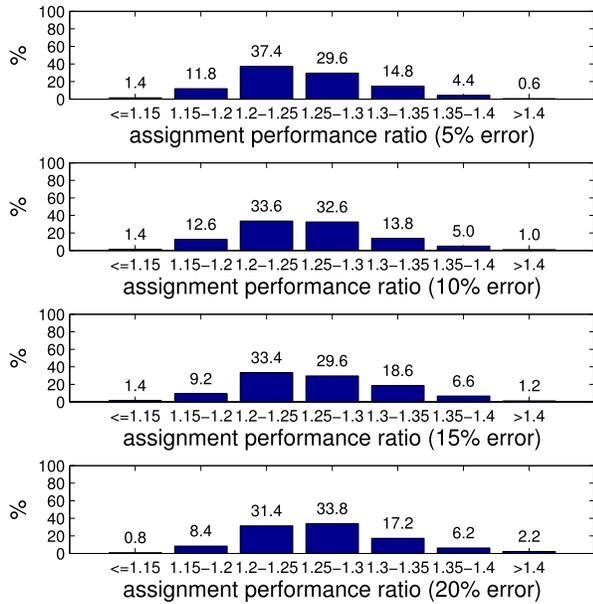


Fig. 9. Monte Carlo tests with estimation error in $\tau(\cdot)$.

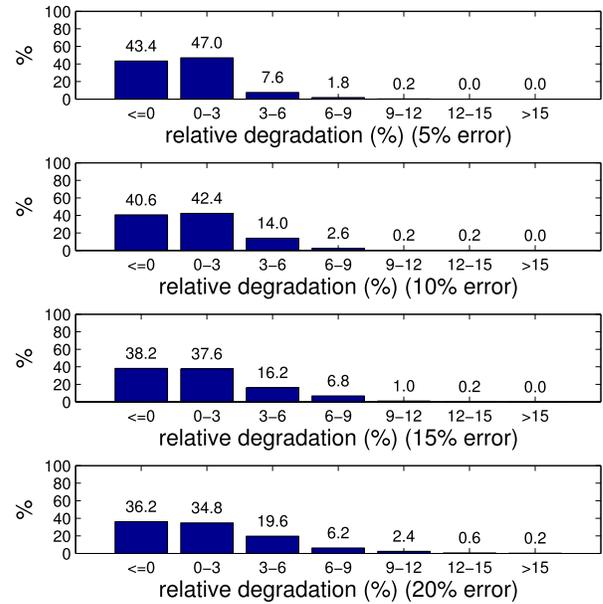


Fig. 11. Performance degradation with estimation error in $\tau(\cdot)$.

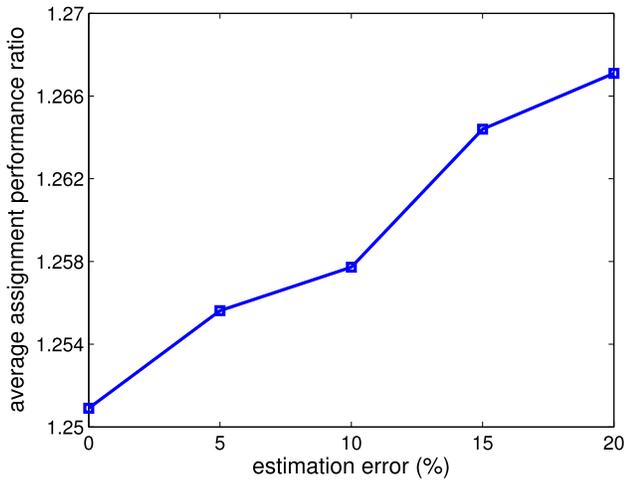


Fig. 10. Average performance with estimation error in $\tau(\cdot)$.

TABLE I
ALGORITHM PERFORMANCE DEGRADATION DUE TO ESTIMATION ERROR

Estimation error	Average performance ratio	Relative performance degradation
0	1.2509	0
$\pm 5\%$	1.2556	0.3992%
$\pm 10\%$	1.2577	0.5797%
$\pm 15\%$	1.2644	1.1220%
$\pm 20\%$	1.2671	1.3543%

arrival times, thus lowering the precision requirement on the function $\tau(\cdot)$.

Hence the performance gap between our online algorithm and offline optimum is acceptable in average cases, even with inaccurate estimation of arrival times. This suggests that our algorithm is simple, not conservative and robust.

VI. CONCLUSION

The online station assignment for (commercial) EVs is investigated that aims to minimize cost to EVs and congestion at service stations. We first develop an efficient offline solution via a bipartite matching approach to solve for the optimal station assignments by assuming global inputs are available. An online station assignment algorithm that builds on online matching is then proposed. We prove that the online algorithm achieves a tight (optimal) competitive ratio of $2A - 1$ under mild conditions, where A is the number of battery swapping requests. The algorithm is further proven to be simple, not conservative and robust, which maintains the performance ratios below 1.3 in over 80% of Monte Carlo test cases.

Limitations: First, the explicit station congestion model is simplified. In light of the service time of battery swapping, a queuing model with constrained service capability can be developed at each station, but a careful reevaluation of the analytical results is required under that circumstance. Second, taking the supply of fully-charged batteries at each station as an exogenous input does not fully capture the closed-loop behavior that decisions of online assignment have an impact on future battery supply. However, considering the multi-hour battery charging process, this cross-timescale dependence is trivial for real-time online decisions in practice. Therefore, we ignore this aspect to simplify analysis. Third, the tight competitive ratio of $2A - 1$ does not hold if additional driving range constraints are enforced for EVs. This condition seems to provide the worst-case adversary with excessive power/freedom in parameterization in the online optimization context. See Appendix B for an illustrative example.

APPENDIX

A. General Competitive Ratio of Algorithm 1

Theorem 4: Generally, for the online matching problem, Algorithm 1 attains a competitive ratio of $O\left(\left(\frac{\alpha_1+1}{\alpha_1}\right)^{A-1}\right)$.

Proof: Similar to part of the proof of Theorem 2, we start with Observation 2, i.e., $\alpha_{ab} \leq \alpha_{a'b} + \alpha_{ab'} + \alpha_{a'b'}$, $\forall a, a' \in \mathbb{A}, b, b' \in \mathbb{B}$, and look at cases where EV a requests battery swapping later than EV a' , i.e., $r_a > r_{a'}$. Generally,

$$\begin{aligned} & \max\{\rho_b - \tau_{ab}\} - \max\{\rho_b - \tau_{a'b}\} \\ & \leq \tau_{a'b} - \tau_{ab} \\ & = (r_{a'} + \theta_{a'b}) - (r_a + \theta_{ab}) \\ & \leq \theta_{a'b} - \theta_{ab} \\ & \leq \theta_{a'b} \end{aligned} \quad (11)$$

Therefore,

$$\begin{aligned} \omega_{ab} & = \alpha_{ab} + \max\{\rho_b - \tau_{ab}\} \\ & \leq \alpha_{a'b} + \alpha_{ab'} + \alpha_{a'b'} + \max\{\rho_b - \tau_{a'b}\} + \theta_{a'b} \\ & \leq \omega_{a'b} + \theta_{a'b} + \frac{\alpha_2 d_{a'b} + \max\{\rho_b - \tau_{a'b}\}}{\alpha_1} + \alpha_{ab'} + \alpha_{a'b'} \\ & \leq \frac{\alpha_1 + 1}{\alpha_1} \omega_{a'b} + \omega_{ab'} + \omega_{a'b'} \\ & \leq \frac{\alpha_1 + 1}{\alpha_1} (\omega_{a'b} + \omega_{ab'} + \omega_{a'b'}) \end{aligned} \quad (12)$$

which implies that the weight of a latest matchup is still bounded. We generalize (12) to an edge-connected path $P(l)$ of length l that starts from a to b as

$$\omega_{ab} \leq \left(\frac{\alpha_1 + 1}{\alpha_1}\right)^{\frac{l-1}{2}} W(P(l)) \quad (13)$$

As aforementioned, every time EV a_k requests battery swapping, Algorithm 1 matches a_k with b_k which is linked by a unique path P_k . P_k has an odd length of at most $2k - 1$, i.e., P_k passes all of a_1, a_2, \dots, a_{k-1} . Therefore, $\omega_{a_k b_k} \leq \left(\frac{\alpha_1 + 1}{\alpha_1}\right)^{k-1} W(P_k) \leq \left(\frac{\alpha_1 + 1}{\alpha_1}\right)^{k-1} (W(x_k^{\text{off}}) + W(x_{k-1}^{\text{off}})) \leq 2 \left(\frac{\alpha_1 + 1}{\alpha_1}\right)^{k-1} W(x_A^{\text{off}})$, $k = 2, 3, \dots, A$. In addition, $\omega_{a_1 b_1} = W(x_1^{\text{off}}) \leq W(x_A^{\text{off}})$. By adding them up, we have $W(x^{\text{on}}) = \sum_{k=1}^A \omega_{a_k b_k} \leq (1 + 2 \sum_{k=2}^A \left(\frac{\alpha_1 + 1}{\alpha_1}\right)^{k-1}) W(x^{\text{off}})$, or $\frac{W(x^{\text{on}})}{W(x^{\text{off}})} \leq 1 + 2 \sum_{k=2}^A \left(\frac{\alpha_1 + 1}{\alpha_1}\right)^{k-1}$, which suggests $\gamma_x = O\left(\left(\frac{\alpha_1 + 1}{\alpha_1}\right)^{A-1}\right)$. \square

B. A Toy Example With EV Driving Range Constraints

Consider two stations i_1 and i_2 , each with one already fully-charged battery. Suppose there are two pending battery swapping requests from EVs a_1 and a_2 , respectively, in a time horizon of length T . EV a_1 requests earlier and Algorithm 1 assigns station i_1 to it due to $\alpha_{a_1 i_1} < \alpha_{a_1 i_2}$. Say EV a_1 is quick enough to arrive at the station and take the battery. In terms of EV a_2 , assume it has a driving range constraint that restricts its option to only station i_1 . Then EV a_2 can only be assigned the station without any battery, which leads to the cost of $\alpha_{a_2 i_1} + T - \tau_{a_2 i_1}$, i.e., the total system cost of $\alpha_{a_1 i_1} + \alpha_{a_2 i_1} + T - \tau_{a_2 i_1}$. However, in case of $\alpha_{a_1 i_1} < \alpha_{a_1 i_2} \ll T - \tau_{a_2 i_1}$ and $\alpha_{a_2 i_1} \ll T - \tau_{a_2 i_1}$, the minimal total system cost of $\alpha_{a_1 i_2} + \alpha_{a_2 i_1}$ could be realized by assigning EV a_1 to station i_2 and EV a_2 to station i_1 . In this case,

the performance ratio of Algorithm 1 could be arbitrarily bad because of

$$\begin{aligned} & \frac{\alpha_{a_1 i_1} + \alpha_{a_2 i_1} + T - \tau_{a_2 i_1}}{\alpha_{a_1 i_2} + \alpha_{a_2 i_1}} \\ & = 1 + \frac{\alpha_{a_1 i_1} - \alpha_{a_1 i_2} + T - \tau_{a_2 i_1}}{\alpha_{a_1 i_2} + \alpha_{a_2 i_1}} \\ & \approx 1 + \frac{T - \tau_{a_2 i_1}}{\alpha_{a_1 i_2} + \alpha_{a_2 i_1}} \rightarrow \infty \end{aligned} \quad (14)$$

REFERENCES

- [1] NAS. (2017). *How we use Energy: Transportation*. [Online]. Available: <http://needtoknow.nas.edu/energy/energy-use/transportation/>
- [2] (2017). *Global EV Sales for 2017 Q1*. [Online]. Available: <http://www.ev-volumes.com/country/total-world-plug-in-vehicle-volumes/>
- [3] Y. Li, C. Zhan, M. de Jong, and Z. Lukszo, "Business innovation and government regulation for the promotion of electric vehicle use: Lessons from Shenzhen, China," *J. Cleaner Prod.*, vol. 134, pp. 371–383, Oct. 2016.
- [4] B. Sun, X. Tan, and D. H. K. Tsang, "Optimal charging operation of battery swapping and charging stations with QoS guarantee," *IEEE Trans. Smart Grid*, vol. 9, no. 5, pp. 4689–4701, Sep. 2018.
- [5] M. Chafkin, "A broken place: The spectacular failure of the startup that was going to change the world," *Fast Company*, vol. 26, p. 2015, 2014. [Online]. Available: <https://www.fastcompany.com/3028159/a-broken-place-better-place>
- [6] T. Zhang, X. Chen, Z. Yu, X. Zhu, and D. Shi, "A Monte Carlo simulation approach to evaluate service capacities of EV charging and battery swapping stations," *IEEE Trans. Ind. Informat.*, vol. 14, no. 9, pp. 3914–3923, Sep. 2018.
- [7] X. Wang and J. Wang, "Economic assessment for battery swapping station based frequency regulation service," *IEEE Trans. Ind. Appl.*, vol. 56, no. 5, pp. 5880–5889, Sep. 2020.
- [8] X. Tan, B. Sun, and D. H. K. Tsang, "Queueing network models for electric vehicle charging station with battery swapping," in *Proc. IEEE Int. Conf. Smart Grid Commun. (SmartGridComm)*, Nov. 2014, pp. 1–6.
- [9] X. Liu, T. Zhao, S. Yao, C. B. Soh, and P. Wang, "Distributed operation management of battery swapping-charging systems," *IEEE Trans. Smart Grid*, vol. 10, no. 5, pp. 5320–5333, Sep. 2019.
- [10] Y. Gao, J. Yang, M. Yang, and Z. Li, "Deep reinforcement learning based optimal schedule for a battery swapping station considering uncertainties," *IEEE Trans. Ind. Appl.*, vol. 56, no. 5, pp. 5775–5784, Sep. 2020.
- [11] W. Infante, J. Ma, X. Han, and A. Liebman, "Optimal recourse strategy for battery swapping stations considering electric vehicle uncertainty," *IEEE Trans. Intell. Transp. Syst.*, vol. 21, no. 4, pp. 1369–1379, Apr. 2020.
- [12] M. R. Sarker, H. Pandzic, 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, Mar. 2015.
- [13] P. You, Z. Yang, Y. Zhang, S. H. Low, and Y. Sun, "Optimal charging schedule for a battery switching station serving electric buses," *IEEE Trans. Power Syst.*, vol. 31, no. 5, pp. 3473–3483, Sep. 2016.
- [14] K. Sepetanc and H. Pandzic, "A cluster-based operation model of aggregated battery swapping stations," *IEEE Trans. Power Syst.*, vol. 35, no. 1, pp. 249–260, Jan. 2020.
- [15] S. Esmaeili, A. Anvari-Moghaddam, and S. Jadid, "Optimal operation scheduling of a microgrid incorporating battery swapping stations," *IEEE Trans. Power Syst.*, vol. 34, no. 6, pp. 5063–5072, Nov. 2019.
- [16] Y. Zheng, Z. Y. Dong, Y. Xu, K. Meng, J. H. Zhao, and J. Qiu, "Electric vehicle battery charging/swap stations in distribution systems: Comparison study and optimal planning," *IEEE Trans. Power Syst.*, vol. 29, no. 1, pp. 221–229, Jan. 2014.
- [17] X. Wang, J. Wang, and J. Liu, "V2G frequency regulation capacity optimal scheduling for battery swapping station using deep Q-network," *IEEE Trans. Ind. Informat.*, early access, May 11, 2020, doi: 10.1109/TII.2020.2993858.
- [18] P. Xie, Y. Li, L. Zhu, D. Shi, and X. Duan, "Supplementary automatic generation control using controllable energy storage in electric vehicle battery swapping stations," *IET Gener., Transmiss. Distrib.*, vol. 10, no. 4, pp. 1107–1116, Mar. 2016.

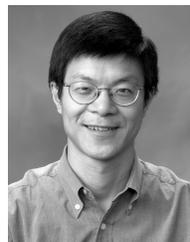
- [19] Q. Kang, J. Wang, M. Zhou, and A. C. Ammari, "Centralized charging strategy and scheduling algorithm for electric vehicles under a battery swapping scenario," *IEEE Trans. Intell. Transp. Syst.*, vol. 17, no. 3, pp. 659–669, Mar. 2016.
- [20] P. You *et al.*, "Scheduling of EV battery swapping—Part I: Centralized solution," *IEEE Trans. Control Netw. Syst.*, vol. 5, pp. 1887–1897, Dec. 2018.
- [21] P. You *et al.*, "Scheduling of EV battery swapping—Part II: Distributed solutions," *IEEE Trans. Control Netw. Syst.*, vol. 5, no. 4, pp. 1920–1930, Dec. 2018.
- [22] X. Zhang, Y. Cao, L. Peng, N. Ahmad, and L. Xu, "Towards efficient battery swapping service operation under battery heterogeneity," *IEEE Trans. Veh. Technol.*, vol. 69, no. 6, pp. 6107–6118, Jun. 2020.
- [23] R. Deng and H. Liang, "Whether to charge an electric vehicle or not? A near-optimal online approach," in *Proc. IEEE Power Energy Soc. Gen. Meeting (PESGM)*, Jul. 2016, pp. 1–5.
- [24] E. H. Gerding, V. Robu, S. Stein, D. C. Parkes, A. Rogers, and N. R. Jennings, "Online mechanism design for electric vehicle charging," in *Proc. 10th Int. Conf. Auton. Agents Multiagent Syst.*, vol. 2, 2011, pp. 811–818.
- [25] L. Guo, K. F. Eriksson, and S. H. Low, "Optimal online adaptive electric vehicle charging," in *Proc. IEEE Power Energy Soc. Gen. Meeting*, Jul. 2017, pp. 1–5.
- [26] Z. J. Lee *et al.*, "Large-scale adaptive electric vehicle charging," in *Proc. IEEE Int. Conf. Commun., Control, Comput. Technol. Smart Grids (SmartGridComm)*, Oct. 2018, pp. 1–7.
- [27] S. Zhao, X. Lin, and M. Chen, "Peak-minimizing online EV charging," in *Proc. 51st Annu. Allerton Conf. Commun., Control, Comput. (Allerton)*, Oct. 2013, pp. 46–53.
- [28] H. W. Kuhn, "The Hungarian method for the assignment problem," *Nav. Res. Logistics Quart.*, vol. 2, nos. 1–2, pp. 83–97, Mar. 1955.
- [29] P. You, J. Z. F. Pang, M. Chen, S. H. Low, and Y. Sun, "Battery swapping assignment for electric vehicles: A bipartite matching approach," in *Proc. IEEE 56th Annu. Conf. Decis. Control (CDC)*, Dec. 2017, pp. 1421–1426.
- [30] S. Khuller, S. G. Mitchell, and V. V. Vazirani, "On-line algorithms for weighted bipartite matching and stable marriages," *Theor. Comput. Sci.*, vol. 127, no. 2, pp. 255–267, May 1994.
- [31] BaiduMaps. (2018). *Hangzhou*. [Online]. Available: <http://map.baidu.com>



Pengcheng You (Member, IEEE) received the B.S. degree in electrical engineering and the Ph.D. degree in control from Zhejiang University, China, in 2013 and 2018, respectively. He is currently a Post-Doctoral Fellow with the Department of Electrical and Computer Engineering, Johns Hopkins University. His research interests include power systems and electricity markets.



John Z. F. Pang (Member, IEEE) received the B.S. degree in mathematics from Nanyang Technological University and the Ph.D. degree in computing and mathematical sciences from Caltech. He is currently a Research Scientist with the A*STAR Institute of High Performance Computing (IHPC). Before that, he was a Data Scientist with Schlumberger. His research interests include power systems, network economics, optimization, and approximation algorithms.



Steven H. Low (Fellow, IEEE) received the B.S. degree in EE from Cornell and the Ph.D. degree in EE from Berkeley. He is currently the Gilloon Professor with the Department of Computing and Mathematical Sciences and the Department of Electrical Engineering, Caltech, and an Honorary Professor with The University of Melbourne, Australia. Before that, he was with AT&T Bell Laboratories, Murray Hill, NJ, USA, and The University of Melbourne. His current research interest includes power systems.