An Intelligent Energy Management System for Large-scale Charging of Electric Vehicles

Zhe Yu, Shiyao Chen, and Lang Tong, *Fellow, IEEE,*

***Abstract*—The problem of large-scale charging of electric vehicles (EVs) with consumer-imposed charging deadlines is considered. An architecture for the intelligent energy management system (iEMS) is introduced. The iEMS consists of an admission control and pricing module, a scheduling module that determines the charging sequence, and a power dispatch module that draws power from a mix of storage, local renewable energy sources, and purchased power from the grid. A threshold admission and greedy scheduling (TAGS) policy is proposed to maximize operation profit. The performance of TAGS is analyzed and evaluated based on average and worst-case performance measures and the optimality of TAGS is established for some instances. Numerical simulations demonstrate that TAGS achieves noticeable performance gains over benchmark techniques.**

***Index Terms*—Charging of electric vehicles; competitive ratio analysis; deadline scheduling; energy management systems.**

I. Introduction

LARGE scale adoption of electric vehicles (EVs) depends critically on the availability of convenient and economic charging facilities in both public and private settings. We consider the problem of EV charging at parking facilities where the charging of a large number of EVs can be managed centrally and efficiently. Typical settings include EV charging at parking garages, parking lots, and possibly street parking spaces.

For a charging service provider who operates charging of hundreds and up to thousands of EVs, an intelligent energy management system (iEMS) is essential to serve consumers in the most effective way. To this end, the design of an iEMS faces several challenges.

First, EV customers have diverse charging needs in terms of the amount of charging and the time by which charging needs to be completed. Second, the cost of charging may also be time varying and stochastic, especially if the charging facility has locally renewable energy sources or the cost of purchased electricity fluctuates with time. An efficient iEMS must have the ability to optimize its charging profile by taking advantage of the charging cost dynamics while satisfying consumer requirements. Third, since there are consumer-imposed charging deadlines and capacity constraints from the charging facility, it is necessary to have an optimized pricing and admission control strategy that yields economic return for the service provider.

Manuscript received August 13, 2015; revised January 4, 2015; accepted February 2, 2016. Date of publication March 30, 2016; date of current version January 10, 2016. This work is supported in part by the National Science Foundation under Grant CNS-1248079 and CNS-1135844.

Z. Yu, S. Chen, and L. Tong (corresponding author) are with the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853, USA. (Email: zy73@cornell.edu; sc933@cornell.edu; lt35@cornell.edu).

Finally, for large-scale charging, the underlying scheduling algorithm must be scalable with respect to the number of EVs, which rules out the algorithms that are based on direct applications of dynamic optimization principles. It is, therefore, necessary to consider structured algorithms that, although suboptimal in general, are optimal in some nontrivial special cases.

A. Summary of Results

We present a centrally managed iEMS architecture based on the concept of network switched charging. A network switch centrally controls chargers to charge selected EVs using the most economic mix of available local renewable sources and presumably more expensive power purchased from the grid. See a detailed description in Section II.

The algorithmic contribution is the development and analysis of threshold admission and greedy scheduling (TAGS) policies, which are online algorithms aimed at maximizing the overall operation profit. By online algorithms we mean that the admission and scheduling decisions at time t are made based on information received up to time t. To this end, we consider two types of performance measures: one is based on the average profit in a stochastic dynamic optimization framework; the other is based on the worst-case profit in a deterministic robust optimization setting.

In analyzing the performance, we show that TAGS policies are optimal when, at any time interval, only one EV is actively charged by the local renewables. Although the condition under which TAGS policies are optimal is somewhat limited, the fact that optimality can be achieved lends analytical support for TAGS in general cases when multiple EVs can be charged simultaneously. Indeed, it is well known that deadline scheduling problems involving multiple processors (EV chargers) are very challenging; the optimal scheduling algorithm is unknown even in absence of requiring admission control.

B. Related Work

Different from the centralized scheduling of EV charging framework considered here, there is a significant body of literature on decentralized EV charging problems. The work in [1] and [2] aims to minimize the load variance. It is shown that minimizing the load variance leads to the valley filling property that shapes the load to a more uniform profile. In fact, it is established in [1], [2] that the objective valley filling can be achieved in a decentralized fashion by iteratively adjusting the price of charging and the charging profile of the consumers. This approach pioneered in [1], [2] and followed through in

[3], [4] is particularly suitable for the day ahead planning of charging services since it requires that the aggregated charging demand and information of consumers within the planning horizon are known by the service provider ahead of time. In contrast, the real-time charging techniques considered in this work focus on real-time charging services and they are designed to exploit real-time demand and cost information.

If we ignore the admission policy and focus on the scheduling of charging for admitted EVs, the underlying scheduling problem is closely related to the classical deadline scheduling problem first considered by Liu and Layland [5]. In its most generic setting, jobs arrive with different sizes and deadlines, and they are processed by M identical processors. Each processor can process one job at a time, and its work can be preempted without cost. For the single processor case ($M = 1$), the results are quite complete. Simple online algorithms (with linear complexity) such as the earliest deadline first (EDF) [5], [6] and the least laxity first (LLF) [7] achieve the same performance as the optimal offline algorithm in the deterministic setting. In the stochastic setting, authors of [8] and [9] showed that EDF minimizes the unfinished work. A diffusion model and performance approximation of EDF are developed in [10], [11], and [12] under the assumption of heavy traffic. Lehoczky first introduced the approximation approach in [10], which included customer timing requirement into queueing models under the M/M/1 assumption. In the most recent work, Kruk and Lehoczky showed that EDF minimizes the fraction of the lost work and customers and gave an approximation of the fraction in a general open queueing setting [12].

For the multi-processor case ($M > 1$), optimal online algorithm is lacking. Derouzos and Mok [13] showed that even for the dual processor case, no on-line algorithm can guarantee 100% success. In [14], we developed a scalable index scheduling policy based on the Whittle's index and the performance was shown to be close to the upper bound.

One of the first works related to admission control is in the context of video on demand applications in [15] and is further developed in [16], [17], [18], and [19]. With admission control, online deadline scheduling becomes more challenging and existing results are limited.

C. Organization

This paper is organized as follows. The architecture of the iEMS is presented in Section II. In Section III, the charging problem is formulated as an admission and scheduling problem with hard deadlines under both average case and worst case. In Section IV, the threshold admission and greedy scheduling policies are proposed and the optimality of TAGS is shown in Section V. Simulations are presented in Section VI. Section VII concludes the paper.

II. The iEMS Architecture

The iEMS architecture is illustrated in Fig. 1. The hardware system of the proposed iEMS includes a dispatcher that delivers power from a mix of energy sources—local energy (e.g., renewable energy and local storage) and purchased electricity

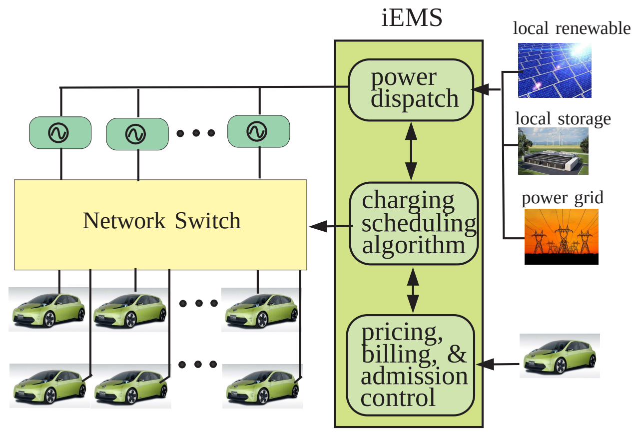

Fig. 1. Architecture for network switched charging and iEMS.

from the grid—to tens or hundreds of chargers. Through the network switch, the scheduler activates and deactivates the chargers connected to EVs admitted to the facility to serve more urgent or more profitable requests.

The iEMS is run by the software system that makes engineering and economic decisions. At the core of the software system for the iEMS is the charging scheduling algorithm, which is the focus of this paper. The scheduler 1) controls the power dispatcher to procure energy from available sources, 2) sets the connections of the switch so that a subset of EVs are charged by the available chargers, and 3) determines the admission of new EVs based on its charging demand and the system operating condition. The software system also has to handle billing, other ancillary services and possibly the forecast of available renewables in the future, which are not discussed in this paper.

III. Problem Formulation

A. EV Charging Job Description

The iEMS treats each EV as a job. Without loss of generality, the EVs are assumed to arrive at the system one by one. The charging request from the i^{th} EV J_i is specified by the tuple $J_i = (r_i, d_i, j_i, \bar{v}_i)$, as shown in Fig. 2, where $r_i \in \mathbb{R}^+$ is the arrival time, $d_i \in \mathbb{R}^+$ the deadline, $j_i \in \mathbb{R}^+$ the charging demand, and $\bar{v}_i \in \mathbb{R}^+$ the value of the EV (charging revenue collected from the EV). Assuming all chargers have the same fixed charging rate, the charging demand j_i is measured by the charging time. The deadline of EV J_i satisfies $d_i \geq r_i + j_i$ at its arrival. The leading time $T_i(t)$ and laxity $l_i(t)$ at time t are defined in (1), where $j_i(t)$ is the remaining charging demand to be completed at time t, and $a^+ = \max\{0, a\}$.

$$T_i(t) = (d_i - t)^+, l_i(t) = T_i(t) - j_i(t). \tag{1}$$

The input EV instance $I = \{J_1, \ldots, J_N\}$ is defined as an arrival sequence of N EVs. The number of EV N can vary across different instances.

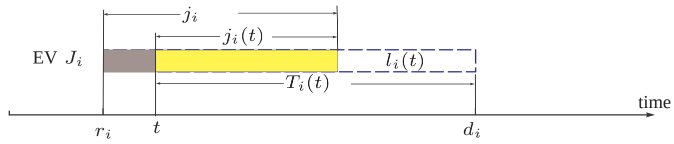

Fig. 2. EV attributes: arrival r_i, deadline d_i, charging demand $j_i(t)$, lead time $T_i(t)$, laxity $l_i(t)$.

B. Admission and EV Charging Policies

Assume that the charging facility has hundreds of chargers and the low cost energy from the local storage or renewable sources can support M of them. While M may be time varying in general, we assume that it is fixed during the scheduling horizon for simplicity. We also assume that all EVs admitted to the facility need to be finished by the deadlines. When local energy may not be sufficient to meet all the demand, the scheduler will need to purchase power from the grid with possibly high cost. The action of using expensive sources is referred to as outsourcing.

A policy π for the iEMS has two components: 1) admission policy, 2) charging policy. Upon the arrival of a new EV, the admission policy decides whether to accept it. Once admitted, the newly arrived EV enters a queue. The charging policy decides which EV(s) in the queue to be charged at any time interval.

Formally, policy π can be defined as follows. Let $\mathcal{R}(t)$ be the set of EVs arrived up to time t. Let $\mathcal{H}_t$ be the collection of past decisions at time t. An admissible online policy π is a mapping,

$$\pi : \mathcal{R}(t) \times \mathcal{H}_t \to \{0,1\}^{|\mathcal{R}(t)|} \times \{-1,0,1\}^{|\mathcal{R}(t)|} \quad (2)$$

that outputs two decision vectors $\mathcal{A}(t)$ and $\mathcal{C}(t)$. In particular, the admission decision $\mathcal{A}(t)$ is a binary vector where EV J_i is admitted if $\mathcal{A}_i(t) = 1$ or declined if $\mathcal{A}_i(t) = 0$. The charging decision is specified as: if $\mathcal{C}_i(t) = 0$, EV J_i is not charged and remains in the queue; if $\mathcal{C}_i(t) = 1$, J_i is charged by the local renewable energy; and if $\mathcal{C}_i(t) = -1$, J_i is outsourced, i.e., charged by expensive energy purchased from the grid.

For a policy π to be admissible, we impose the following constraints:

1) For an admission policy, once the admission decision is made at the arrival time r_i, it cannot be changed, i.e., $\mathcal{A}_i(t) = \mathcal{A}_i(r_i)$ for all $t \geq r_i$.
2) For a charging policy, the number of EVs charged by local energy is bounded by M at any time.

C. Performance Measure

Without loss of generality, the value of an EV (charging revenue collected from EV consumers) is assumed proportional to the charging time, i.e., $\bar{v}_i = p \times j_i$. The marginal cost of a local energy source is denoted by c_{l} and the marginal cost of the outsourcing energy is c_{o}. These values are normalized such that $c_{\text{l}} = 0, p = 1, c_{\text{o}} = c$.

Given an admissible policy π, the profit of the service provider by completing an admitted EV is the charging revenue collected from the EV minus the cost of the expensive energy paid to grid. The value function of the policy for a given instance I is then

$$V_\pi(I) = \sum_{\{i: i \in I, A_i(\infty)=1\}} (j_i - c \times o_i) \quad (3)$$

where o_i is the amount of time that EV $J_i \in I$ is charged using the expensive outsourced energy.

We treat two types of performance measures, the expected total profit for the stochastic setting and the competitive ratio for the deterministic setting.

1) Expected Total Profit: Under the stochastic setting, we are interested in the total expected profit. Denote $\mathcal{I}_N$ the set of all possible instances consists of N EVs. The optimal value function is the supreme of the total expected profit:

$$V^* = \sup_\pi \mathbb{E}_{I \in \mathcal{I}_N} V_\pi(I). \quad (4)$$

If the optimal value can be achieved, the policy π^* that achieves the optimality is considered as optimal:

$$\pi^* = \arg\max_\pi \mathbb{E}_{I \in \mathcal{I}_N} V_\pi(I). \quad (5)$$

2) Competitive Ratio: Under the worst setting, the performance of an online policy π is measured by the competitive ratio against the optimal offline (clairvoyant) scheduler that has information of future arrivals. Denote the profit obtained by the online (optimal offline) scheduler by $V_\pi(I)$ ($V^*_{\text{offline}}(I)$) for some instance I. The competitive ratio is defined as follows.

Definition 1: An online policy π is α-competitive if $\inf_I \frac{V_\pi(I)}{V^*_{\text{offline}}(I)} \geq \alpha$ for all instance $I \in \mathcal{I}$ where $\mathcal{I}$ is the collection of all instances with finite number EVs.
For any instance, an α-competitive online algorithm is guaranteed to achieve at least α fraction of the profit of the optimal offline algorithm.

IV. TAGS: Threshold Admission and Greedy Scheduling

After the problem formulation, we propose threshold admission and greedy scheduling policies for both stochastic and deterministic settings. The TAGS algorithm has two parts. For the charging policy, a greedy scheduling and outsourcing policy is proposed following the EDF principle. Based on the charging policy, threshold admission policies are presented under stochastic and deterministic settings.

A. Greedy Scheduling Policy: EDF-LMO

As shown in Fig. 1, each accepted EV is attached to a charger and the online scheduler determines when to activate the charger and whether to require outsourcing energy. The greedy scheduling and outsourcing policy is simply two conditions to activate chargers. Whether EV J_i 1) is among the M EVs with the earliest deadlines, or 2) whether it has a non-positive laxity as defined in (1). If either of these two conditions is satisfied, the scheduler activates the charger attached to EV J_i and purchases energy from grid if necessary.

To analyze the priority of EVs and amount of outsourcing energy, we can intuitively to view the system as there are in total M local chargers (powered by local renewable energy) and infinitely many outsourcing chargers (powered by expensive outsourcing energy). The scheduler assigns M EVs to local chargers following the EDF principle and other EVs to outsourcing chargers when their laxity is non-positive. The outsourcing policy is referred to as last minute outsourcing (LMO).

1) EDF Scheduling: At time t, the pending set $\mathcal{P}_t$ is defined as the collection of EVs that satisfies 1) it has arrived and been admitted ($r_i \leq t$, $\mathcal{A}_i(r_i) = 1$), 2) the deadline has not yet passed ($d_i > t$), and 3) there is positive remaining charging demand ($j_i(t) > 0$). Upon the admission of an arriving EV, the newly arrived EV joins the pending set $\mathcal{P}_t$.

When there is only one local charger ($M = 1$), at any time t, according to the EDF rule, the EV with the earliest deadline in the pending set $\mathcal{P}_t$ gets assigned to it. When $M > 1$, the scheduler first sorts the admitted EVs in the ascending order of deadlines. Then the scheduler finds the local charger which will firstly finish charging the EVs scheduled to it and assigns the EV in the head of the queue to that charger. The amount of charging time planned for the newly scheduled EV is defined in the following subsection of LMO. The scheduler keeps doing so till all admitted EV is assigned to some local charger. In this way, the pending set $\mathcal{P}_t$ is divided into disjoint sets: $\mathcal{P}_t = \mathcal{P}_t^1 \cup \ldots \cup \mathcal{P}_t^M$, where $\mathcal{P}_t^k$ is the set of EVs assigned to local charger $k = 1, \ldots, M$.

2) Last Minute Outsourcing: To compute the amount of outsourcing energy planned for each EV, we first define the notion of real-laxity as follows.

Definition 2: At time t, assume 1) EVs $J_1, \ldots, J_n \in \mathcal{P}_t^k$, 2) $d_1 \leq \cdots \leq d_n$, and 3) the scheduler plans charging time $\hat{j}_1(t), \ldots, \hat{j}_n(t)$ by the local charger k for each EV. The real-laxity L_i of EV J_i at time t is defined by $L_i = d_i - t - \sum_{m=1}^{i} \hat{j}_m(t)$.

Assume at time t, charger k has pending EVs $J_1, \ldots, J_n$ with remaining planned local charging time $\hat{j}_1, \ldots, \hat{j}_n$, and the real-laxity vector $(L_1, \ldots, L_n) \geq 0$. Suppose that a newly arrived EV J with charging demand j is accepted and assigned to this charger, which leads to a new real-laxity vector $(L'_1, \ldots, L'_{n+1})$ with some negative components. Last Minute Outsourcing will plan for the newly arrived EV J (and only for it) an outsourcing charging time

$$o = -\min_{1 \leq i \leq n+1} L'_i,$$

and a local charging time $j - o$, i.e., reducing the local charging time of EV J by the amount o.

Take Table I as a single local charger example. Upon the arrival of EV 3, the online scheduler assigns $2, 2, 1$ local charging time to EV $0, 1, 2$ respectively. After accepting EV 3, the real-laxity is $L'_3 = -1$. The scheduler will assign an outsourcing charging time as $o_3 = 1$ and a local charging time as $j_3 - o_3 = 1$.

Under LMO policy, when outsourcing energy is needed, it is always assigned to the newly arriving EV, and the amount of the local charging time of other EVs remains unchanged.

Proposition 1 shows that EDF-LMO scheduling policy maintains the feasibility of the local energy sources charging plan.

TABLE I
EXAMPLE OF AN INPUT EV SEQUENCE

EV Index	r_i	j_i	d_i
0	0	2	3
1	1	3	4
2	2	1	6
3	3	2	6
4	4	1	5
5	7	1	10

Proposition 1: Under EDF-LMO,

1) local chargers always finish planned local charging workload for all accepted EVs;
2) the amount of outsourcing energy is the minimum for all charger activation policies with EDF rule to guarantee finishing all accepted EVs.

Proof: See [20]. ■

B. Threshold Admission Policy: Average Case

We now present the threshold admission policy for the average case, which is denoted by TAGS-A. Under the stochastic setting, the initial charging demand, laxity, and the inter-arrival time of EVs are independent and identically distributed (I.I.D.). Upon the arrival of EV J_i with charging demand j_i, suppose it is accepted and assigned to local charger k (both hypothetically) according to EDF-LMO. Denote the outsourcing charging time assigned to J_i by o_i.

The maximum profit the scheduler can collect from EV J_i is stated as $v_i \triangleq j_i - c \times o_i$. The admission policy is proposed as a threshold structure on the maximum profit v_i: if and only if $v_i \geq \nu_i$, J_i is accepted, where ν_i is the threshold we need to optimize according to the randomness in EVs and the electricity prices.

C. Threshold Admission Policy: Worst Case

In this subsection, we present the threshold admission algorithm for the worst case, which is named as TAGS-W. We divide the pending EVs into two types in Definition 3 according to the tightness of the deadlines.

1) Pressing vs non-pressing EVs: Upon the release of EV J at time r, suppose EV J is accepted and conducted following EDF-LMO scheduling policy (both hypothetically). The pressing and non-pressing EVs are defined as follows.

Definition 3: At arrival of EV J, assume 1) $J \in \mathcal{P}_t = \{J_1, \cdots, J_n\}$ with $d_1 \leq \cdots \leq d_n$, and 2) the real-laxity vector satisfies $L_i = 0$ and $L_{i+1}, \cdots, L_n > 0$. Then EVs $J_1, \ldots, J_i$ are classified as *pressing EVs*, and EVs $J_{i+1}, \ldots, J_n$ are classified as *non-pressing EVs*.

The distinction between the two types of EVs lies in the tightness of the deadline: the pressing EVs have relatively tight deadlines and little slack time (since the i^{th} EV J_i has laxity vector component 0, there is no slack time before its deadline d_i, and any EV in $J_1, \ldots, J_i$ cannot afford any delay in charging), while the non-pressing EVs have slack time before deadline. Since TAGS-W follows the EDF rule, at any instant, if the set of pressing EVs is not empty, the local chargers will not work on any non-pressing EVs.

As shown in Table I, at the arrival of EV 3, after hypothetical admission of EV 3, the first three EVs (EV 0, 1, and 2) are pressing, and EV 3 is non-pressing.

Remark 1: Once an EV is classified as pressing, it always remains so afterwards under TAGS-W, even if the EV is classified as pressing due to a newly released EV J that is immediately declined and disappears from the queue (recall EV J is only hypothetically admitted).

2) Pressing vs non-pressing busy intervals: After the classification of pressing and non-pressing EVs, we define below *pressing and non-pressing busy intervals*.

Definition 4: A pressing busy interval (denoted by $\mathbf{B}^p$) is a continuous time period in which some local charger is planed to charge pressing EVs. A non-pressing busy interval ($\mathbf{B}^n$) is a continuous time period in which no local charger is planned for pressing EVs and some local chargers are planned for some non-pressing EVs.

3) Threshold admission policy: After hypothetically admitting the newly arrived EV J and computing the real-laxity vector, two scenarios may occur:

1) J is classified as non-pressing, then it is accepted;
2) J is classified as pressing, then TAGS-W computes two tentative profit values for the "accepting" and "declining" options and then makes the admission decision by comparing the profit ratio (profit associated with accepting divided by profit associated with declining) against a threshold $1+\beta$: if the profit ratio is no greater than threshold $1+\beta$, the EV is declined. β is the parameter that we will optimize according to EV instance and prices.

At the arrival of a pressing EV J, the tentative profit of "accepting" and "declining" options is calculated in a pressing busy interval $\mathbf{B}^p$:

1) The tentative profit for accepting J is computed as the total value of the accepted pressing EVs (including J) in $\mathbf{B}^p$ after hypothetically admitting EV J, less the outsourcing energy cost accumulated and planned so far for these EVs (including J);
2) The tentative profit for declining J is computed as the total value of the accepted pressing EVs (excluding J) in $\mathbf{B}^p$ after hypothetically admitting EV J, less the outsourcing energy cost accumulated and planned so far for these EVs (excluding J);

The difference of the tentative profit for accepting and declining options is the value of the newly arrived EV J less the planned outsourcing energy cost for it. As in Table I, upon the arrival of EV 3, all four EVs (EV 0, 1, 2, and 3) are all pressing. The outsourcing charging time for EV 0, 1, 2, and 3 is 0, 1, 0, and 1, respectively.

1) The tentative profit associated with accepting EV 3 is $(1+3+1+2)-c\times(0+1+0+1)$.
2) The tentative profit associated with declining EV 3 is $(1+3+1)-c\times(0+1+0)$.

Then the ratio of the two tentative profits is compared against the threshold $1+\beta$ to finally render the admission decision for EV 3.

V. Optimality of TAGS: $M=1$

In this section, EDF-LMO is shown as an optimal scheduling policy under any admission policy. Then TAGS is proved optimal in average and worst case respectively.

A. Optimality of EDF-LMO

In [12], the authors studied the performance of EDF policy without admission control. The fraction of work that can not be finished by the only processor on time is denoted as reneging load, which in our setting is finished by outsourcing energy. Theorem 5.1 of [12] shows that EDF minimizes the reneging load for single local charger queue which gives the optimality of EDF-LMO.

Theorem 1: (Theorem 5.1 [12]) Let π be a scheduling and outsourcing policy and $o_\pi(t)$ be the amount of outsourcing energy up to time t. Let $o_{\text{EDF-LMO}}(t)$ be the amount of outsourcing energy of policy EDF-LMO . Then for any $t\geq 0$, $o_{\text{EDF-LMO}}(t)\leq o_\pi(t)$.

Theorem 1 holds for any instance I. So EDF-LMO is optimal under any admission policy in both average case and worst case when $M=1$.

B. Average Case

1) Low outsourcing cost ($c\leq 1$): When the outsourcing charging cost is less than the charging price, any EV is profitable even if it is fully charged by the outsourcing energy. The optimal admission policy is all-accept policy. The threshold ν_i can be set to $-\infty$ and TAGS-A policy is optimal.

2) High outsourcing cost ($c>1$): When $c>1$, the outsourcing energy is costly. The scheduler needs to determine whether to accept a particular EV at its arrival. TAGS-A is proved to be optimal for the identical charging demand scenario, i.e., $j_i=j$.

Theorem 2: The threshold admission policy is optimal for the identical charging demand case ($j_i=j$).

Proof: The proof consists of two steps. We first show that J_i should be accepted if $v_i\geq j$, and declined if $v_i<0$. Then we show the profit difference of acceptance and decline is a monotone and continuous function of v_i in $[0,j]$ so there is a zero point ν_i when acceptance and decline are equivalent. The detailed proof can be found in [20]. ∎

C. Worst Case

Denote the optimal competitive ratio by $C^*(M,c)$ where M is the number of local chargers and c the marginal outsourcing energy cost. In Theorem 3 we establish an explicit characterization of the optimal competitive ratio and show that TAGS-W achieves optimality.

Theorem 3: The optimal competitive ratio is given by

$$C^*(1,c)=\begin{cases}1 & \text{if } c\leq 1;\\ (\sqrt{c}-\sqrt{c-1})^2 & \text{if } c>1.\end{cases}\qquad(6)$$

Proof: When $c\leq 1$, all EVs should be accepted. In Section V-A, it is shown that the scheduling policy EDF-LMO minimizes the outsourcing energy cost among all online and offline scheduling policies under any admission policy. Thus TAGS-W achieves the maximum profit and $C^*(1, c) = 1$.

The case when $c > 1$ is more complicated. The proof is composed of two parts. First we show that $(\sqrt{c} - \sqrt{c-1})^2$ is an upper bound of the optimal competitive ratio using an adversary game argument. Then we analyze TAGS-W and show it achieves the upper bound. The detailed proof can be found in [20]. ■

VI. Simulation

In the simulation, we assume that the outsourcing energy is costly ($c > 1$) and the average performance measure is employed.

A. Single Local Charger: M=1

In this subsection, the EDF conservative (EDFC) algorithm is employed as a baseline. EDFC applies conservative admission policy: only when the EV can be finished purely using local renewable energy following EDF principle, it will be admitted. For TAGS-W, a set of training data is used to find the parameter β that gives the maximum profit. Then the simulation is carried out using the obtained threshold $1 + \beta$. For TAGS-A, the optimal threshold ν_i is different for each EV, which is determined by the distribution of the traffic and the order of EVs. In the simulation, a lower bound approximation of TAGS-A is carried out. Instead of different threshold ν_i for each EV, an optimal uniform threshold ν^* is obtained using training data such that any EV with $v_i \geq \nu^*$ will be accepted. Then the simulation is carried out with the obtained threshold ν^*.

The performance upper bound is obtained based on the result from [12]. In [12], the performance of EDF policy without admission control is approximated under the heavy traffic assumption. In particular, we have the approximation as follows.

Theorem 4: (Theorem 1.1 in [12]) The fraction of demand charged by outsourcing chargers can be approximated as $R^*_W \approx e^{-\delta \bar{D}}[(1-\rho)/(\rho - \rho e^{-\delta \bar{D}})]$, where the inter-arrival time distribution has mean $1/\lambda$, variance α^2; the charging demand has mean $1/\mu$, variance γ^2; the traffic intensity is defined as $\rho = \lambda/\mu$; $\delta = 2(1-\rho)/(\lambda(\alpha^2 + \gamma^2))$, and $\bar{D}$ is the mean of the initial leading time T_i.

According to Theorem 4, under the all-accept policy, the fraction of the demand charged by the local charger is approximately $1 - R^*_W$ and this will generate 1 profit per unit charging time; while the fraction of demand charged by outsourcing energy is R^*_W and this will generate $1 - c$ profit per unit charging time. The total unit time profit is approximated by $(1 - R^*_W + (1-c)R^*_W)\mathcal{J}/\mathcal{T} = (1 - cR^*_W)\lambda/\mu$ where $\mathcal{J}$ is the total charging demand of all EVs, $\mathcal{T}$ the total simulation time and $\mathcal{J}/\mathcal{T}$ indeed the traffic intensity λ/μ.

When the outsourcing energy cost is high ($c > 1$), the admission policy is active and not all the EVs are accepted. Theorem 4 shows that $(1 - R^*_W)\mathcal{J}$ is the demand charged by the local charger if all EVs are accepted. This is an upper bound of the total charging time of the local charger for any admission policy because of possible idleness caused by declining. Thus $(1 - R^*_W)\lambda/\mu$ serves as an upper bound of the unit time reward collected by the local charger. This is an upper bound of the unit time profit of the online scheduler since we neglect the cost of the outsourcing charger part.

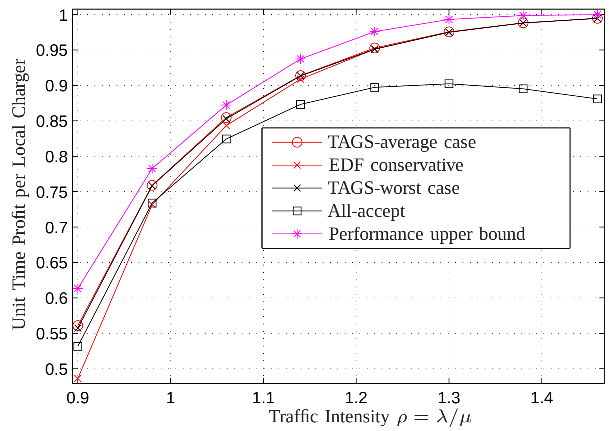

Fig. 3. $M = 1$. Parameter of inter-arrival: $1/\lambda = 2$; initial charging demand: tail index $k = 1/64$; initial laxity: tail index $k = 1/64$, scale $\sigma = 0.0015$, location $\theta = 0.0984$; prices: $p = 1$, $c = 1.25$.

In Fig. 3, the Pareto distributed charging demand, Pareto distributed laxity, and the exponential distributed inter-arrival time EV sequence is simulated and the expected unit time profit versus traffic intensity is shown. When the traffic is light, TAGS admits almost every EV and the performance is close to all-accept policy. When the system gets busy, the TAGS performance is similar to EDFC. One explanation is that TAGS balances the outsourcing cost of charging EVs with tight deadlines and the risk of local charger idleness. When the incoming EVs are dense, we can always reject urgent EVs and some less urgent ones will come in soon enough and keep the local charger busy. The performance upper bound is shown in Fig. 3. The gap between the upper bound and the TAGS is less than 5%. TAGS-W is developed for the worst case while the performance is reasonably well under the average measure.

B. Multiple Local Charger: $M > 1$

In this subsection, we extend TAGS to multiple local charger case, in which no optimal scheduling policy is known in either average or deterministic cases so far. We modify TAGS based on least laxity first (LLF) principle. EVs with least laxity are given high priorities and assigned to local chargers preferentially. Thus when computing the real-laxity in Definition 2 and looking for pressing and non-pressing EVs in Definition 3, the admitted EVs are sorted by the laxity l_i rather than the deadlines d_i. Last minute outsourcing and threshold admission policies remain unchanged.

In Fig. 4, the simulation of Pareto distributed charging demand, Pareto distributed laxity, and the exponential distributed inter-arrival time with multiple local chargers is presented. When the traffic is heavy ($\lambda/\mu = 1.2M$), all-accept policy is too aggressive and LLF conservative is too backwards. TAGS-A rejects EVs appropriately and maintains a high utility rate

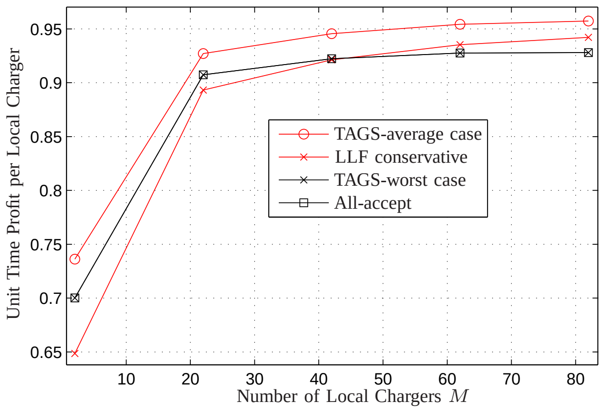

Fig. 4. $M > 1$. Parameter of inter-arrival: $1/\lambda = 1/(3\sqrt{M})$; charging demand: $1/\mu = 0.4\sqrt{M}$, tail index $k = 1/64$; initial laxity: tail index $k = 1/64$, scale $\sigma = 0.0015$, location $\theta = 0.0984$; prices: $p = 1$, $c = 1.25$

of local chargers. This simulation shows the ability of TAGS-A to be extended to multiple local charger cases. The average computation time of TAGS-A for each EV is 0.1 ms per local charger on a PC platform with Intel 2.67 GHz quad core and 8 GB memory.

VII. Conclusion

The admission and scheduling problem for EVs with deadlines and choice of energy is considered. When the local renewable energy can support only one EV ($M = 1$), in the average case, the problem is formulated as a stochastic dynamic programming and the optimality of TAGS is obtained for identical charging demand cases. In the worst case, the upper bound of the competitive ratio is developed by construction and TAGS is proposed to achieve the upper bound. When the local renewable energy can support charging of multiple EVs ($M > 1$), TAGS is extended and simulation suggests notable performance improvement over benchmark techniques.

References

[1] L. Gan, U. Topcu, and S. Low, "Optimal decentralized protocols for electric vehicle charging," *IEEE Transactions on Power Systems*, vol. 28, no. 2, pp. 940–951, 2013.

[2] Z. Ma, D. S. Callaway, and I. Hiskens, "Decentralized charging control of large populations of plug-in electric vehicles," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 1, pp. 67–78, 2013.

[3] E. L. Karfopoulos and N. D. Hatziargyriou, "A multi-agent system for controlled charging of a large population of electric vehicles," *IEEE Transactions on Power Systems*, vol. 28, no. 2, pp. 1196–1204, 2013.

[4] A. Sheikhi, S. Bahrami, A. Ranjbar, and H. Oraee, "Strategic charging method for plugged in hybrid electric vehicles in smart grids; a game theoretic approach," *International Journal of Electrical Power & Energy Systems*, vol. 53, no. 1, pp. 499–506, 2013.

[5] C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a hard-real-time environment," *Journal of ACM*, vol. 20, no. 1, pp. 46–61, 1973.

[6] M. Dertouzos, "Control robotics: the procedural control of physical processes," in *Proceedings of International Federation for information Processing Congress*, pp. 807–813, 1974.

[7] A. Mok, *Fundamental design problmes of distributed systems for the hard real-time environment*. PhD thesis, MIT, 1983.

[8] S. S. Panwar, D. Towsley, and J. K. Wolf, "Optimal scheduling policies for a class of queues with customer deadlines to the beginning of service," *Journal of the ACM (JACM)*, vol. 35, no. 4, pp. 832–844, 1988.

[9] D. Towsley and S. Panwar, "On the optimality of minimum laxity and earliest deadline scheduling for real-time multiprocessors," in *Proceedings of IEEE Euromicro 90' Workshop on Real-Time*, pp. 17–24, Jun. 1990.

[10] J. Lehoczky, "Real-time queueing theory," in *Proceedings of 17th IEEE Real-Time Systems Symposium*, pp. 186 –195, Dec. 1996.

[11] B. Doytchinov, J. Lehoczky, and S. Shreve, "Real-time queues in heavy traffic with earliest-deadline-first queue discipline," *Annals of Applied Probability*, vol. 11, no. 2, pp. 332–378, 2011.

[12] L. Kruk, J. Lehoczky, K. Ramanan, and S. Shreve, "Heavy traffic analysis for EDF queues with reneging," *Annals of Applied Probability*, vol. 21, no. 2, pp. 484–545, 2011.

[13] M. L. Dertouzos and A. K.-L. Mok, "Multiprocessor online scheduling of hard-real-time tasks," *IEEE Transactions on Software Engineering*, vol. 15, no. 12, pp. 1497–1506, 1989.

[14] Z. Yu, Y. Xu, and L. Tong, "Large scale charging of electric vehicles: a multi-armed bandit approach," in *Proceedings of the 52nd Annual Allerton Conferene on Communication, Control, and Computing*, Oct. 2015.

[15] A. Bar-Noy, J. A. Garay, and A. Herzberg, "Sharing video on demand," *Discrete Applied Mathematics*, vol. 129, no. 1, pp. 3–30, 2003.

[16] M. H. Goldwasser and B. Kerbikov, "Admission control with immediate notification," *Journal of Scheduling*, vol. 6, no. 3, pp. 269–285, 2003.

[17] J. Ding and G. Zhang, "Online scheduling with hard deadlines on parallel machines," in *Proceedings of 2nd Algorithmic Aspects in Information and Management International Conference*, pp. 32–42, Springer, 2006.

[18] J. Ding, T. Ebenlendr, J. Sgall, and G. Zhang, "Online scheduling of equal-length jobs on parallel machines," in *Proceedings of 15th Annual European Symposium*, pp. 427–438, 2007.

[19] T. Ebenlendr and J. Sgall, "A lower bound for scheduling of unit jobs with immediate decision on parallel machines," in *Proceedings of 6th Workshop on Approximation and Online Algorithms*, pp. 43–52, 2008.

[20] S. Chen, Z. Yu, and L. Tong, "Optimality of TAGS," July 2015. [Online]. Available: http://acsp.ece.cornell.edu/papers/ChenTong15TR.pdf

Zhe Yu received his B.E. degree from Department of Electrical Engineering, Tsinghua University, Beijing, China in 2009, and M.S. from Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, USA in 2010. He is currently working toward the Ph.D. degree in the School of Electrical and Computer Engineering, Cornell University. His current research interests are in smart grid, demand response, dynamic programming and optimization.

Shiyao Chen received his B.E. degree from Department of Automation, Tsinghua University, Beijing, China in 2009, and Ph.D. degree in the School of Electrical and Computer Engineering, Cornell University, Ithaca, NY, USA. His research interests are in smart grid, electricity market and demand response.

Lang Tong (S87,M91,SM01,F05) is the Irwin and Joan Jacobs Professor in Engineering of Cornell University and the site director of Power Systems Engineering Research Center (PSERC). He received the B.E. degree from Tsinghua University in 1985, and M.S. and Ph.D. degrees in electrical engineering in 1987 and 1991, respectively, from the University of Notre Dame. He was a Postdoctoral Research Affiliate at the Information Systems Laboratory, S-tanford University in 1991. He was the 2001 Cor Wit Visiting Professor at the Delft University of Technology and had held visiting positions at Stanford University and the University of California at Berkeley. Lang Tongs research is in the general area of statistical inference, communications, and complex networks. His current research focuses on inference, optimization, and economic problems in energy and power systems. He received the 1993 Outstanding Young Author Award from the IEEE Circuits and Systems Society, the 2004 best paper award from IEEE Signal Processing Society, and the 2004 Leonard G. Abraham Prize Paper Award from the IEEE Communications Society. He is also a coauthor of seven student paper awards. He received Young Investigator Award from the Office of Naval Research. He was a Distinguished Lecturer of the IEEE Signal Processing Society.