Preserving fairness in EV charging under time-varying congestion levels

Martin Zeballos
Universidad ORT Uruguay
mzeballos@uni.ort.edu.uy

Andres Ferragut
Universidad ORT Uruguay
ferragut@ort.edu.uy

Fernando Paganini
Universidad ORT Uruguay
paganini@ort.edu.uy

***Abstract*—This paper concerns facilities for charging Electrical Vehicles at parking lots. We work under the assumption that power capacity may be insufficient to simultaneously charge all stations, and thus some scheduling must be performed. In previous work we analyzed the case of a stationary customer demand, developing a fluid model which characterizes the performance of different scheduling policies in overload. A new policy termed Least Laxity Ratio was proposed to improve fairness in service.**

In this paper we wish to incorporate the fact that congestion levels are not stationary in practice, rather they obey daily use cycles. We study the behavior of the different policies with load obtained from a real set of parking lot data. Empirical results show that the conclusions of the stationary analysis remain valid. In particular, during intervals of congestion, LLR achieves the best results in terms of proportional fairness.

I. INTRODUCTION

The increasing penetration of Electrical Vehicles (EVs) must be accompanied by an adequate expansion in the charging infrastructure [1], [2]. While home charging will cover some of the demand, another attractive option is to provide charging stations at parking lots; for instance, a large corporation may provide this service to employee-owned EVs.

When dimensioning the power capacity of such a facility, it may not be practical to provision for the peak load (all chargers operating at full power). Instead, some scheduling/curtailment policy can be put in place to manage the available capacity, exploiting users' statistical multiplexing and their deferability of service. In this context, overload situations will occur, where users only obtain partial charge; the fairness in such resource allocation is the topic of this paper.

In previous work [3], we analyzed this problem mathematically under the assumption of a stationary load on the parking lot. A fluid model was developed that allows the analysis of reneged service for different scheduling policies; these results are reviewed in Section II. In particular, with fairness in mind we introduced the Least Laxity Ratio (LLR) policy, which achieves proportionality in reneged work at the fluid scale.

In this paper we wish to validate the analysis under more realistic conditions, in particular the non-stationarity which would be typical of the daily use of a real parking lot. For this purpose, we have obtained a data set from the parking lots of a major tech company, which describe arrivals, departures and charge amounts. Assuming now a facility with restricted capacity, we apply the different scheduling algorithms to these input traces; results are reported in Section III. We verify the same qualitative behavior of the stationary analysis, in particular LLR achieves the best fairness results. We also develop a quantitative index for comparison purposes. Conclusions are given in Section IV.

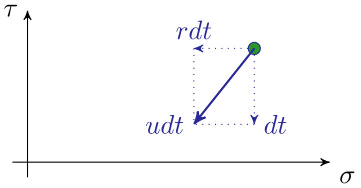

Fig. 1. Dynamics for each job.

A. Related work

The study of charging policies for fleets of EVs is a very active topic. We mention the work of [4], where future demand estimation is used to maximize the state of charge of vehicles upon departure. References [5], [6], [7] use dynamic programming to minimize reneged service, establishing properties of the optimal algorithm. Our work is closest to [8], where policies based on receding horizon optimization are proposed and tested with real data.

II. STATIONARY ANALYSIS OF EV CHARGING POLICIES

In this section we summarize the results of [3], where a fluid model for the car population in an EV parking lot was introduced, and used to characterize efficiency and fairness of different charge scheduling policies.

The starting point is a discrete queueing model in which each car in service is characterized by a point in the space (σ, τ) of service time - sojourn time, as depicted in Figure 1.

The residual sojourn time τ (time left before departure) is consumed at unit rate while the car is in the charging station. The residual service time σ will vary according to the charge received: we denote by $r(t)$ the power received *relative* to the maximum charging rate ($0 \leq r \leq 1$); this implies σ is consumed at rate r, as depicted in Figure 1.

As cars arrive at the parking lot, new points will appear in the (σ, τ) plane, at locations depending on their initial service requests and sojourn times. Similarly, when cars complete service or reach their deadline, they abandon the system[1].

[1]Charged cars that remain at the spot are no longer counted, we assume the number of parking spots is not the bottleneck.

The garage operator's decision is expressed in the choice of $r_k(t)$ for each car $k = 1, \ldots n(t)$ present and still requiring service, subject to the capacity constraint

$$\sum_{k=1}^{n(t)} r_k(t) \leqslant C. \tag{1}$$

C is the maximum number of chargers that could be simultaneously turned on at full rate; it is also possible to activate more than C chargers at a reduced rate.

We consider scheduling policies which discriminate cars only by their location (σ_k, τ_k) in the service-sojourn space. In particular, we analyzed in [3] the following policies:

- Processor Sharing (PS): all cars present in the system are charged at rate $r_k(t) = \min(1, \frac{C}{n(t)})$.
- Earliest Deadline First (EDF): the C cars with smallest τ_k are charged at full rate $r_k = 1$.
- Least Laxity First (LLF): the C cars with smallest *laxity* (spare time) $\ell_k := \tau_k - \sigma_k$ are charged at full rate.
- Least Laxity Ratio (LLR): the C cars with smallest *laxity ratio* $\theta_k := \frac{\tau_k}{\sigma_k} = 1 + \frac{\ell_k}{\sigma_k}$ are charged at full rate.

A discrete stochastic model of the system can be obtained by characterizing the arrival process and choosing one of the above service disciplines. In [3] we worked under the *stationary* assumption that cars arrive at the system as a Poisson process of intensity λ, with service and sojourn times (S, T) being random variables of joint density $f(\sigma, \tau)$. The system load was defined as

$$\rho = \lambda E[S] = \lambda \int_0^\infty \int_0^\infty \sigma f(\sigma, \tau) d\sigma d\tau.$$

For analytical tractability it is, however, more convenient to work instead with a fluid version of the dynamics.

A. Fluid model

The mathematical analysis of [3] applies to a scale where the number of points is large and characterized by a time-dependent density function $g(t, \sigma, \tau)$. New mass arrives into the system at rate $\lambda f(\sigma, \tau)$. Mass is transported along the vector field

$$u = -\begin{bmatrix} r(\sigma, \tau, g) \\ 1 \end{bmatrix}$$

of Figure 1, with r defined by the scheduling policy. The resulting dynamics of the population density g is the *advection equation*:

$$\frac{\partial g}{\partial t} + \nabla \cdot (gu) = \lambda f, \tag{2}$$

where $\nabla \cdot (\cdot)$ is the divergence operator on $\mathbb{R}_+^2$, i.e. on the variables σ, τ. This tool was used in [3] to study the equilibrium (steady state) configurations. In particular, we proved the following results:

- In underload $\rho < C$, all the above policies (indeed, any policy that is efficient in the sense of not wasting charging opportunities) has the same equilibrium, with $r = 1$ for all vehicles present in the system. In particular, all vehicles receive full service.
- In overload $\rho > C$, all efficient policies result in the same overall reneged work $W = \rho - C$.
- The *distribution* of reneged work among users depends on the policy. We elaborate on this next.

TABLE I
SUMMARY OF PERFORMANCE METRICS FOR THE DIFFERENT POLICIES.

Policy	Threshold equation	Reneged work (S_r)	Attained service ($S_a = S - S_r$)
EDF	$\lambda E[\min\{S, \tau^*\}] = C$	$(S - \tau^*)^+$	$\min\{S, \tau^*\}$
LLF	$\lambda E[(S - \sigma^*)^+] = C$	$\min\{S, \sigma^*\}$	$(S - \sigma^*)^+$
PS	$\lambda E[\min\{S, r^*T\}] = C$	$(S - r^*T)^+$	$\min\{S, r^*T\}$
LLR	$\lambda\theta^* E[S] = C$	$(1 - \theta^*)S$	$\theta^* S$

B. Fairness in reneged service

It is shown in [3] that in overload, each of the four policies listed above is characterized in equilibrium by a threshold condition in the distribution of (S, T) (arriving service and sojourn times). The position of an arriving EV relative to this threshold determines the attained (and reneged) service upon departure, as summarized in Table I. In particular:

- In EDF, a maximum service τ^* is provided to all EVs. Those with smaller service requests receive full service, the rest are truncated to this level.
- In PS, the thresholding condition is in S/T; the policy still favors smaller jobs, but now attained service is proportional to sojourn time.
- In LLF, a service amount σ^* is reneged to all EVs. Only those with larger service requests receive (reduced) service.
- In LLR, attained service is distributed *proportionally* to the service request. That is, all EVs leave the system with the same fraction θ^* of their required service.

We believe the proportional fairness of LLR is an attractive feature: when in overload, the system serves all users in the proportion allowed by the installed capacity.

C. Simulation results

In addition to the mathematical results, in [3], we simulated the discrete system under the different policies, to validate the predictions of our fluid model. For this purpose we developed a discrete-event simulator over Julia [9].

We reproduce some of the results here, focusing on EDF, LLF and LLR. The arrival rate is $\lambda = 120$ and sets the scale of the system. The job size S_k is exponentially distributed with rate $\mu = 1$. Each job arrives with an initial laxity L_k also exponentially distributed with rate $\gamma = 0.5$, independent of S_k, and we set $T_k = S_k + L_k$. The load of the system is thus $\rho = \lambda/\mu = 120 > C$ ($C = 80$). In Figure 2 we plot a snapshot of the system in steady-state for each policy, and the thresholds computed using the fluid model. It can be seen that the model predicts correctly the transition for all three policies.

Figure 3 shows the initial demand S and final reneged work S_r achieved by the jobs under the three policies, as

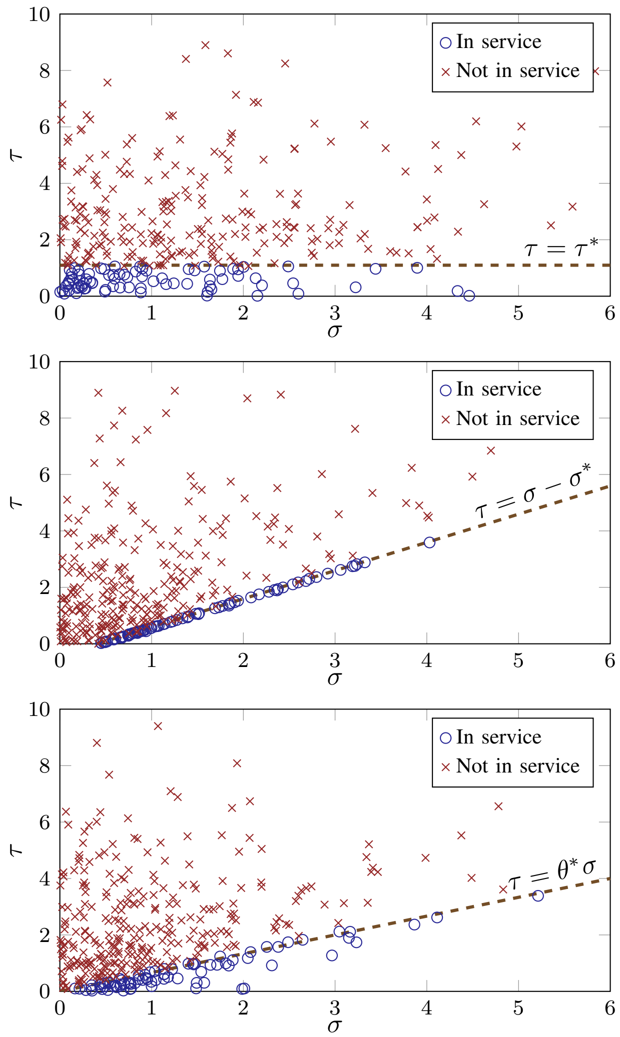

Fig. 2. Steady-state snapshot of a system under EDF (above), LLF (center) and LLR (below), with in service and not in service loads. The dotted line indicates the corresponding thresholds predicted by the fluid model.

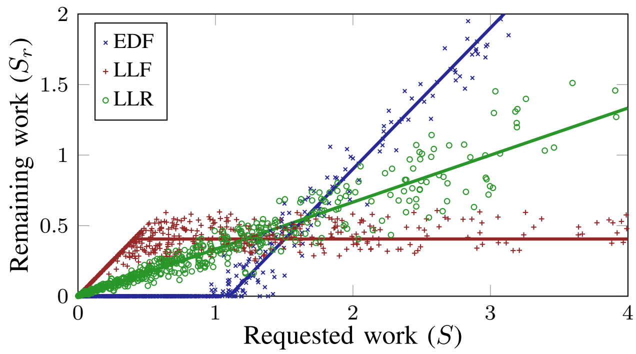

Fig. 3. Requested and reneged work for each job under the different policies and fluid prediction.

well as the fluid model predictions. As we can see, the simulated observations follow the predicted behavior closely, and the EDF and LLF policies discriminate against large and small jobs respectively. In contrast, our proposed LLR policy achieves the desired linear relationship, imposing proportional fairness across jobs.

III. Real time-varying congestion levels

The main limitation of the preceding results is the stationarity assumption in the parking lot load, which appears unrealistic. In practice, arrival patterns and charging requests reflect daily use cycles. And, since sojourn times are long with respect to these variations, we cannot assume that each EV sees an approximately steady situation.

Ideally, in a typical day with varying levels of congestion, the scheduling policy should only play a major rule during the intervals of overload, and impose during those times the appropriate fairness in the curtailed service.

Establishing this property mathematically is, however, very challenging, even with the fluid approximation (2). For this reason we turn to an experimental evaluation of our policies under time-varying conditions, using real EV charging data.

A. *Data set and its characteristics*

A dataset of about 100 EV parking lots of a major tech company was obtained, where each parking lot has a capacity between 10 and 70 vehicles. The data contains the values of: arrival times, sojourn times, consumed energy and the maximum charging power rate for each car, during a period of about 50 days. We carried out the following adjustments. The consumed energy was normalized, dividing by the charge power rate so as to be measured in time, providing the values of the requested service time S for this study. We filtered out the few cars (52 out of more than 25000) which presented unfeasible data: i.e., S larger than the sojourn time T.

In order to recreate a large scale parking lot, for example an office building, a shopping mall or a supermarket, data from 20 of these parking lots was merged creating a dataset of about 500 cars arriving through the day and reaching a maximum of 170 cars parked at the same time. Figure 4 shows the number of cars in the parking lot within 48 hours and shows the variability of congestion levels through this period.

As expected there are more cars in the system in working hours, while there is almost no car during the night. This is in line with what Figure 5 shows where the histogram of arrival time for all the entire dataset is presented. There is a peak at 8 AM and a relatively homogeneous distribution during the rest of the working hours.

The data clearly presents a time-varying arrival rate and thus, a time-varying requested work and power levels. We will assume that this load is applied to a parking lot with individual charging stations, but where the maximum number of chargers that can be simultaneously turned on at full rate is $C = 30$.

To simulate the system with this real data, our simulator was adapted to include the exogenous event data.

B. *Results and comparison*

In Figure 6 we plot a snapshot of the system in a condition of high congestion (with more than 100 cars in the parking

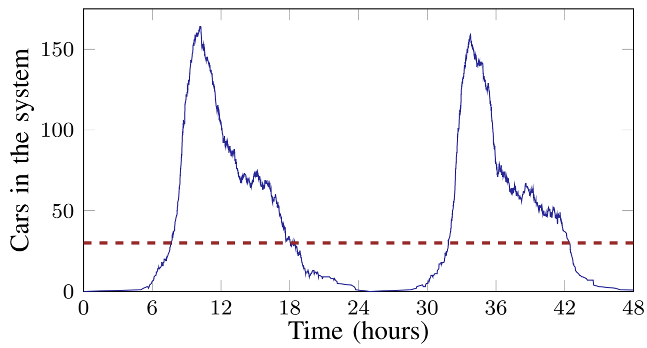

Fig. 4. Number of cars in the system within two typical days. The dotted line indicates the maximum capacity.

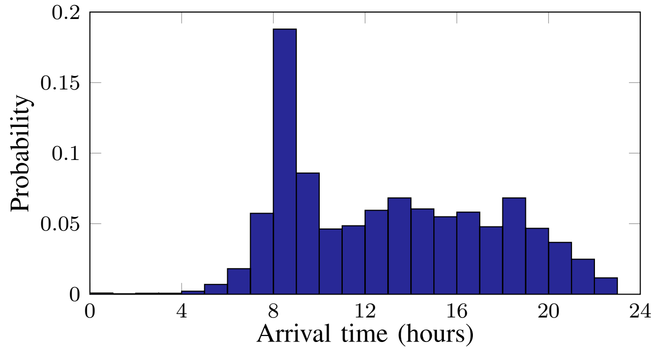

Fig. 5. Histogram of arrival time of the cars to the system.

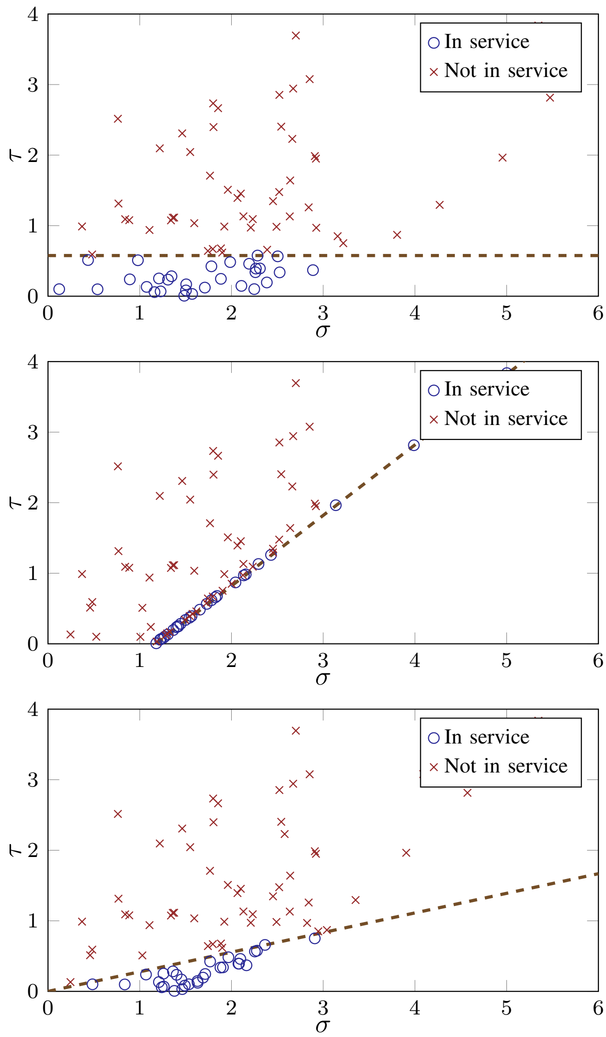

Fig. 6. Snapshot of the system under EDF (above), LLF (center) and LLR (below), with in service and not in service loads. The dotted line indicates the corresponding empirical threshold between the two classes.

lot), for each of the three policies, EDF, LLF and LLR. We classify cars according to whether they are in service, and for illustration we mark the empirical threshold between the two classes. As we can see, the transition has the same characteristics of that showed in the stationary simulated system (Figure 2).

In Figure 7 we plot the initial demand and final reneged work achieved by the jobs under the three policies. It is seen that the empirical time-varying observations present the same trend as the simulated stationary system. EDF and LLF policies discriminate against large and small jobs respectively, but with variable thresholds due to the time-varying congestions levels. On the other hand, while the slope of the threshold varies, the LLR policy achieves a linear relationship, appearing to impose the same proportional fairness across jobs as in the fluid limit and stationary simulations.

C. *Proportional fairness measurement*

While empirical results show graphically a qualitative match with our theoretical predictions on fairness, we would like to have a more quantitative validation and comparison between the different policies.

For this purpose, we resort to the classical Jain's fairness index [10], developed originally in the context of bandwidth sharing in congested telecom networks. Given n users, each assigned a quantity $x_i, i = 1, \dots n$ of a shared resource, Jain's index is defined by:

$$J(x_1, x_2, ..., x_n) = \frac{(\sum_{i=1}^{n} x_i)^2}{n \sum_{i=1}^{n} x_i^2}. \qquad (3)$$

This quantity $J \in [0, 1]$ reaches unity under an egalitarian distribution. In our case, to measure proportional fairness, we will apply the index to the ratio $x_i = \frac{S_{a_i}}{S_i}$ between the attained and requested service. Since by definition this quantity can only be evaluated at the end of service, to obtain a time-varying measurement of fairness we will compute Jain's index at time t between the set of cars which have finished service in a window of time $[t - t_0, t]$.

In Figure 8 we plot the Jain fairness index computed across a typical day considering a half-hour window, i.e. with t_0 equal to half an hour. As expected the index is 1 when there is no congestion in the system, where $x_i = 1$ for all EVs.

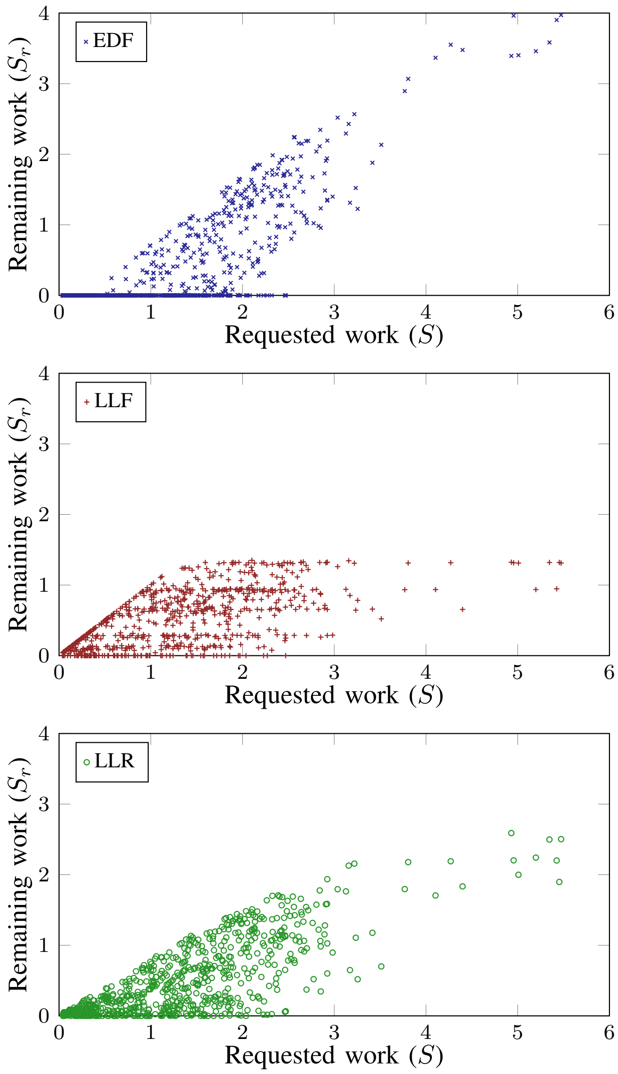

Fig. 7. Requested and reneged work for each job under EDF (above), LLF (center) and LLR (below) under time-varying congestion levels.

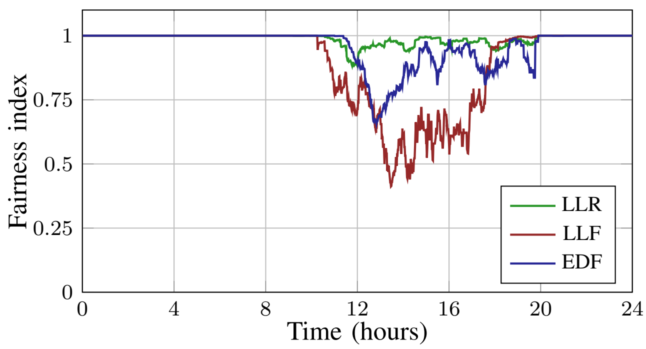

Fig. 8. Jain fairness index computed across a typical day considering a half-hour window.

During busy hours (see Figure 4) the fairness index decreases, with different behavior across policies. EDF shows a sudden drop in the fairness index but recovers to values near 1 quickly. In the case of LLF, the fairness index suffers the most, staying near 0.5 for several hours. On the other hand, for LLR the fairness index stays close to unity at all times, with far less variability, evidence of the desired proportional fairness behavior.

IV. Conclusion

In this work, we analyzed the performance of different scheduling policies for EV charging with load obtained from a real set of parking lot data that presents time-varying congestion levels. We show that the conclusions of the stationary analysis remain valid in the empirical case in terms of fairness in the distribution of service. We introduce a method for measuring proportional fairness in attained service, based on the Jain's fairness index; experiments show that the Least Laxity Ratio policy preserves fairness across jobs under time-varying congestion levels.

In future work, we plan a more extensive analysis of the fairness index for different scenarios. We are also interested analyzing the decision of when the garage operator should seek for more power in the electricity market.

Acknowledgment

The authors would like to thank Steven Low and his team of the California Institute of Technology who kindly provided us the EV charging dataset. This work was partially supported by ANII under project FSE_1_2016_1_131605.

References

[1] W. Su, H. Eichi, W. Zeng, and M.-Y. Chow, "A survey on the electrification of transportation in a smart grid environment," *IEEE Transactions on Industrial Informatics*, vol. 8, no. 1, pp. 1–10, 2012.

[2] D. B. Richardson, "Electric vehicles and the electric grid: A review of modeling approaches, impacts, and renewable energy integration," *Renewable and Sustainable Energy Reviews*, vol. 19, pp. 247–254, 2013.

[3] M. Zeballos, A. Ferragut, and F. Paganini, "Achieving fairness for EV charging in overload: a fluid approach," in *Proc. of the 36th IFIP Performance Conference*, 2018.

[4] W. Su and M.-Y. Chow, "Performance evaluation of an eda-based large-scale plug-in hybrid electric vehicle charging algorithm," *IEEE Transactions on Smart Grid*, vol. 3, no. 1, pp. 308–315, 2012.

[5] S. Chen, Y. Ji, and L. Tong, "Large scale charging of electric vehicles," in *Power and Energy Society General Meeting, 2012 IEEE*. IEEE, 2012, pp. 1–9.

[6] Z. Yu, Y. Xu, and L. Tong, "Large scale charging of electric vehicles: A multi-armed bandit approach," in *Communication, Control, and Computing (Allerton), 2015 53rd Annual Allerton Conference on*. IEEE, 2015, pp. 389–395.

[7] Y. Xu, F. Pan, and L. Tong, "Dynamic scheduling for charging electric vehicles: A priority rule," *IEEE Transactions on Automatic Control*, vol. 61, no. 12, pp. 4094–4099, 2016.

[8] Y. Nakahira, N. Chen, L. Chen, and S. H. Low, "Smoothed least-laxity-first algorithm for EV charging," in *Proceedings of the Eighth International Conference on Future Energy Systems*. ACM, 2017, pp. 242–251.

[9] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, "Julia: A fresh approach to numerical computing," *SIAM review*, vol. 59, no. 1, pp. 65–98, 2017.

[10] R. K. Jain, D.-M. W. Chiu, and W. R. Hawe, "A quantitative measure of fairness and discrimination," *Eastern Research Laboratory, Digital Equipment Corporation, Hudson, MA*, 1984.