

The effect of deadline uncertainty on EV charge scheduling

Lucas Narbondo
Universidad ORT Uruguay
Montevideo, Uruguay
narbondo@fi365.ort.edu.uy

Andres Ferragut
Universidad ORT Uruguay
Montevideo, Uruguay
ferragut@fi365.ort.edu.uy

Fernando Paganini
Universidad ORT Uruguay
Montevideo, Uruguay
paganini@fi365.ort.edu.uy

Abstract—We consider an EV charging infrastructure where scheduling algorithms determine the service priority, based on the departure deadline. The focus is on the Earliest Deadline First algorithm, and the impact on its performance stemming from uncertainty in the deadlines. Through numerical simulations we evaluate the resource allocation achieved by the algorithm in two versions, depending on whether or not vehicles are curtailed after their declared deadline expires. Results are compared with a theoretical analysis of the idealized EDF algorithm. Experiments with real-world traces are also included.

Index Terms—EV charging, scheduling, deadlines.

I. INTRODUCTION AND BACKGROUND

This paper concerns charging infrastructures for Electrical Vehicles (EVs) which are increasingly being deployed to support and promote this alternative mode of transportation. The scenario is a parking lot at a certain corporate or school site, where EVs stay for a considerable time and may thus be conveniently recharged. While individual chargers may be provided across the facility, power provisioning need not cover the peak load, due to its expense and the possible disruption to the grid. Rather, it appears sensible to dimension the installation for a fraction of peak load, and use the time flexibility to *schedule* the sharing of this limited capacity.

The performance of such installations for a stream of arriving EVs, each with service requirement and departure deadline, was analyzed in [1], for various scheduling policies. In particular the Earliest-Deadline-First (EDF) policy was considered, which fills the capacity with EVs closest to their departure. For brevity we focus on this particular policy for the present paper, but several of the considerations discussed here can be extended to other scheduling policies.

The analytical model in [1] represents individual vehicles as particles in the space of residual service-sojourn times, and uses a fluid approximation to understand the resource allocation attained by the scheduling policy. This fluid approximation is suitable to model the behavior of a large system. In particular it was found that the attained service for EDF is approximated by $S_a = \min\{S, \tau^*\}$, where S is the requested service (in units of time to fully charge the vehicle) and τ^* is a threshold, essentially the remaining time to deadline expiration when a vehicle acquires priority. The value of τ^* emerges

from a fixed point equation that involves the system load. In particular, if the load is below the system capacity, the threshold $\tau^* \rightarrow \infty$ and every vehicle gets fully served. In the more interesting case where the system is in overload, the EDF policy truncates the largest jobs. We expand on this explanation in Section II.

The EDF policy is based on the assumption that the system operator knows the vehicle's departure time in advance; in practice, this requires that users declare their length of stay upon arrival. However, users may not be aware of their exact departure times, or can mislead the system with their reports. This brings up the question: what is the effect on *uncertainty* in the deadline information on the resulting allocation of resources? In this paper we explore this issue through modeling and numerical experiments.

In Section II we consider the effect of purely random deviations, attributed to users' own uncertainty about their departure time. For an EDF policy that schedules based on the reported deadline taken at face value, we find under homogeneous arrivals that the mean behavior of the assigned service still follows the theoretical result even under uncertainty; however there is an increased variance in the assigned service. A closer look reveals a correlation between assignment and uncertainty: vehicles under-reporting their deadline tend to be favored because they are prioritized by the EDF algorithm.

There is thus an incentive to deliberately misrepresent one's urgency to obtain priority service. Given this fact, we consider in Section III a curtailing variant to the EDF policy: when the reported deadline expires, so does the vehicle's priority. Experiments with this version show an aggregate behavior similar to the previous case, but individual EV performance is affected. Indeed, now the optimal strategy for an arriving EV is to report an unbiased estimate of the departure time. We also outline an analytical argument that explains these observations from the point of view of the model in [1].

In Section IV we compare the results obtained with homogeneous synthetic traffic with those obtained from real-world data on EV charging demand, obtained from the Caltech Adaptive Charging Network (ACN) [2]. To represent interesting conditions we consider a charging installation of reduced capacity. We find that the qualitative conclusions of the preceding study remain valid for this real world situation. Conclusions are given in Section V.

A. Related work

The analysis and design of charging policies for EV fleets is an active research topic. In the scheduling domain, [3] proposes an algorithm that estimates future demands to maximize the state of charge of vehicles upon departure. In [4] the authors obtain competitive-ratio bounds for several scheduling policies with deadline constraints, in particular EDF. Another relevant reference is [5], where the scheduling of a large aggregate of deferrable loads is studied, and many of the relevant policies such as earliest-deadline or least-laxity-first are brought into the smart grid context.

Our line of work is based on a queuing approach and its fluid limits: for instance [6], [7] use such methods for a simple equal sharing policy. In comparison, [1], [8], [9] cover more general scheduling policies when subject to deadlines, such as EDF and the ones introduced in [5]. We consider the fluid scale as an idealized reference model from which deviations can be studied, in particular the uncertainty on reported deadlines. We refer to [10], [11] for a comprehensive analysis of fluid limits for earliest deadline policies and many server queues. Empirical studies of EV charging with real-world data are presented in [2], [12].

II. EDF UNDER UNCERTAINTY

We focus on the analysis of an EV charging lot with limited capacity, under the EDF policy and uncertain deadlines. We begin by briefly recalling how EDF works and the attained service obtained by vehicles, following [1].

Assume vehicles arrive into the parking lot at rate λ (vehicles/hour) as a Poisson process. Each vehicle k has two random characteristics: a required *service time* S_k (i.e. the energy requested divided by the charging power), and a *sojourn time* T_k , which is the time until the car leaves the parking lot. We use S, T for a generic vehicle; these random quantities follow general distributions, and we assume that $T \geq S$ with probability 1, i.e. the demand of each EV is a priori feasible at the charging station.

The garage operator must assign to each EV a *charging rate* $r_k(t)$, interpreted as a fraction of nominal power. Since capacity is finite, the charging rate should verify:

$$0 \leq r_k(t) \leq 1 \text{ for every } k, t, \text{ and } \sum_{k=1}^{n(t)} r_k(t) \leq C,$$

where $n(t)$ is the number of EVs present in the garage which still require service. C can be interpreted as the maximum number of chargers that could be simultaneously turned on at full power.

The EDF policy assumes that the system operator has full knowledge of T_k for every vehicle present, and chooses to serve at full power the C vehicles closer to their deadlines: the rate is given by the indicator function

$$r_k(t) = \mathbf{1}_{\{T_k - t \leq \tau^*(t)\}}, \quad (1)$$

with $\tau^*(t)$ is such that the number of vehicles simultaneously served is $\leq C$ to satisfy the system constraint. A simple example is given in Fig. 1.

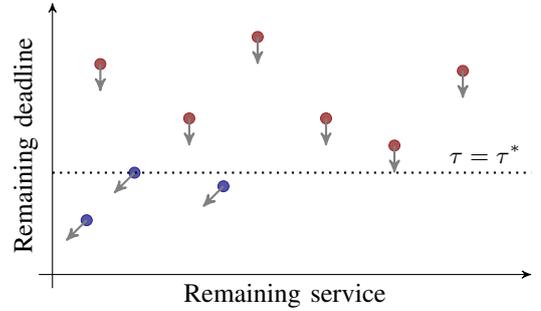


Fig. 1. EDF policy behavior for $n(t) = 9$ and $C = 3$. Vehicles in service reduce their remaining service and deadline. Those that are not served are only consuming their remaining deadline.

The *system load* is defined as:

$$\rho := \lambda E[S], \quad (2)$$

i.e. the average service time of the typical vehicle multiplied the arrival rate. The interpretation of the load is the number of chargers active on average when there is infinite capacity. If the real system capacity is $C < \rho$ the system is in *overload*. In that case, we have the following:

Proposition 1 ([1], Prop. 3): Under the EDF policy in overload, the attained service S_a per user is given by $\min\{S, \tau^*\}$, where the threshold τ^* satisfies

$$\lambda E[\min\{S, \tau^*\}] = C. \quad (3)$$

The above proposition is better understood by following the trajectory of a typical vehicle, as depicted in Fig. 2. A typical vehicle does not get any service until its remaining deadline is equal to τ^* , the threshold value. Once it reaches this value, it gets priority and begins service, but only receives charge for a period of τ^* units, and therefore $S_a = \tau^*$, and the vehicle leaves the system with unfinished demand. If instead the requested service is less than τ^* , then the vehicle can be fully served and $S_a = S$.

We now turn our attention to a more realistic scenario where the system operator *does not know* T_k in advance for every vehicle. Instead, the owner of an EV will estimate its own departure time and report a perturbed deadline T'_k . If the system operator takes this at face value, and applies EDF scheduling, we have an *Uncertain EDF* scenario.

To analyze the performance of such a system, we create a Poisson stream of EVs with a defined arrival, departure and workload, as specified by equations (4). Charging times are assumed to be exponentially distributed. For this stream of arrivals we perform two simulations: the first one considers the real deadlines and will be our benchmark. The second one adds uncertainty in the deadlines by adding an error in the departure time, modeled as a Gaussian distribution with standard deviation σ . All simulations were performed in Julia 1.6 using the library `EVQueues` [13], that we developed.

For our simulations we chose the the following parameters:

$$\lambda = 60; E[S] = 1; E[T] = 2; C = 40 \quad (4)$$

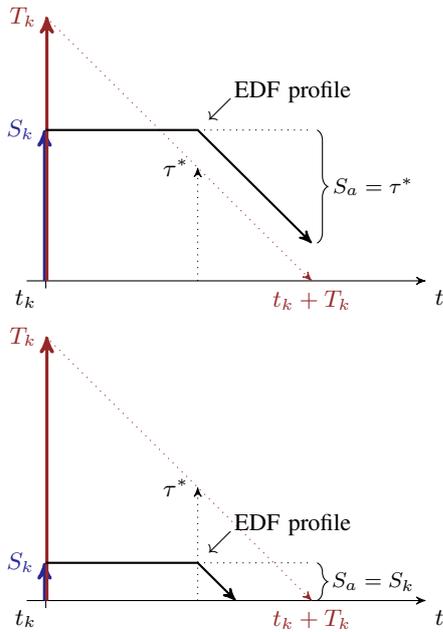


Fig. 2. Typical vehicle trajectories under EDF: large demand (above) and small demand (below).

The system is in overload, since $\rho = \lambda E[S] = 60 > C$.

First, we analyze the overall performance of the algorithm. In Fig. 3 we show the ratio between average charge delivered and requested energy as a function of σ . As a reference, we show the benchmark performance for EDF when considering the real deadlines, which is very close¹ to the expected $C/\rho = 66.6\%$. As the reporting error increases, the overall energy delivered worsens, but the loss is small.

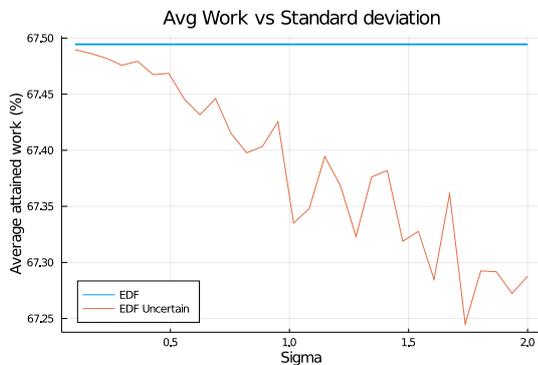


Fig. 3. Average Work vs Standard deviation

We now analyze individual vehicle performance: Fig. 4 shows the energy received as a function of the energy requested for an uncertainty parameter $\sigma = 0.2$. In this case (3) can be solved and the threshold satisfies $\tau^* = 1.1$. Adding uncertainty does not affect the threshold value, meaning that on average the vehicles receive the same service even under

¹Differences are attributed to transient effects.

uncertain conditions. The only effect is an increased variability on individual performance around the threshold.

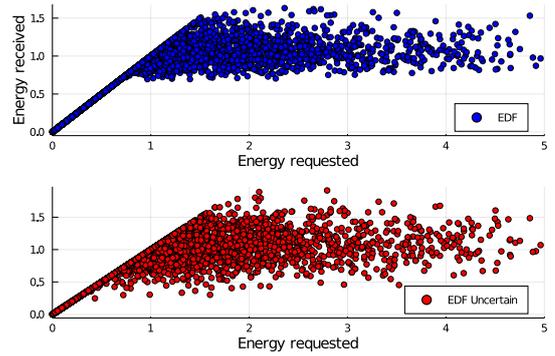


Fig. 4. Energy Requested and received for EDF, $\sigma = 0.2$.

Since both simulations are fed with the same arrival process, it is interesting to compare how each individual EV fares in both cases. Fig. 5 plots the energy received for each individual EV, in both simulations. Black dots represent vehicles that were completely charged with both algorithms, red dots were only completely charged when using the real deadlines, yellow dots when using the reported deadlines and blue dots are EVs that were not fully charged independently of the deadline considered. Again, while each vehicle receives different service, the overall performance of the system is robust to uncertainty.

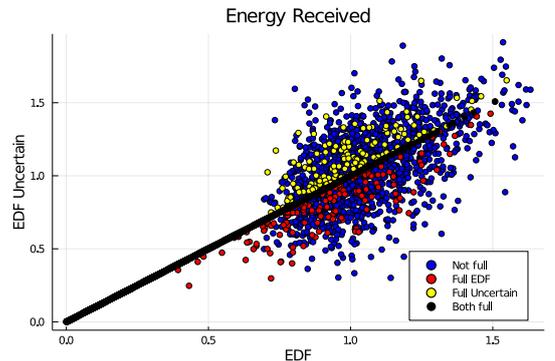


Fig. 5. Energy Received considering real and reported deadlines, $\sigma = 0.2$.

We now take a closer look at the relationship between individual EV performance and its own error in deadline declaration. Fig. 6 reveals an unsurprising correlation between assignment and uncertainty: vehicles under-reporting their deadline tend to be favored because they get priority in the EDF algorithm. Since the scheduling continues to charge vehicles after their reported deadline has expired, there is an incentive to deliberately misrepresent one's urgency to obtain priority service. We analyze how to correct this issue in the next Section.

III. ENFORCING INCENTIVES IN DEADLINE REPORTING

The main conclusion of the preceding analysis is that the overall performance of the EDF algorithm is robust under

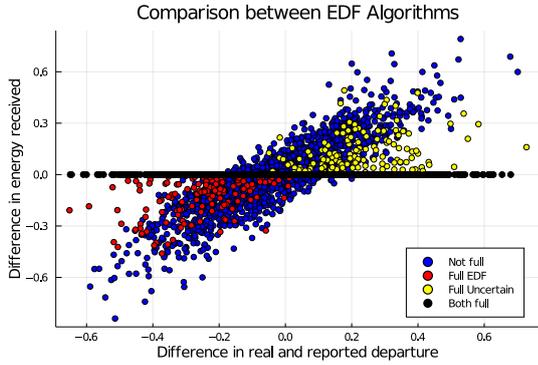


Fig. 6. Received energy vs. difference in reported deadline $T_k - T'_k$. A positive value means that the vehicle under-reported their departure time.

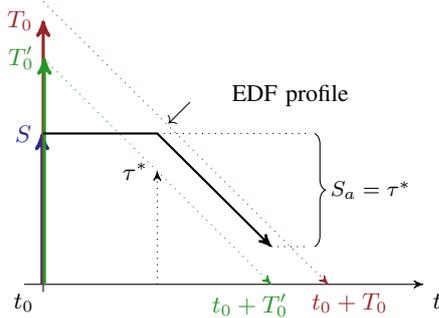


Fig. 7. Typical vehicle trajectory under the curtailed version of EDF. The green arrow represents an under-reported deadline.

deadline uncertainty, but an incentive for under-reporting develops. Therefore, it is also of interest to consider an algorithm that, while keeping the average performance of EDF, rewards users that report their true deadlines.

We thus propose to use the following *curtailed EDF* algorithm: the system operator at any point in time t computes the remaining estimated deadlines:

$$\tau'_k(t) = T'_k - t,$$

for all vehicles present in the system. The EVs where $\tau'_k(t)$ is *positive* are served in EDF fashion using τ'_k as an estimation of their remaining sojourn time. Those EVs that have *negative* τ'_k have exceeded their estimated time in the system, and thus are moved to the back of the priority queue and only charged when there are less than C vehicles with positive $\tau'_k(t)$. In practice, in an overload scenario, vehicles that under-report their deadline are effectively curtailed when they reach their estimated deadline.

In Fig. 7 we show the typical service trajectory of an EV that under-reports its deadline. The net effect of this under-reporting is that the vehicle gets priority sooner (their remaining estimated deadline reaches the threshold before). However, since it gets curtailed at the end of its reported sojourn time, the overall service received does not increase.

For vehicles that over-report their length of stay, the situation is more complicated, and indeed they can incur some loss

in service due to the fact that they lose priority. However, this is not an issue since the strategic behavior we are trying to curb is to avoid under-reporting. Over-reporting would never be strategically beneficial.

From this discussion, we conclude that the average performance of the curtailed EDF policy should be near that of EDF in overload, but should curb the incentive to under-report. We now validate this through simulation. In Fig. 8 we plot the energy delivered using the same conditions than in Section II adding the curtailed version. The added curtailment indeed does not incur a large penalty in terms of overall efficiency.

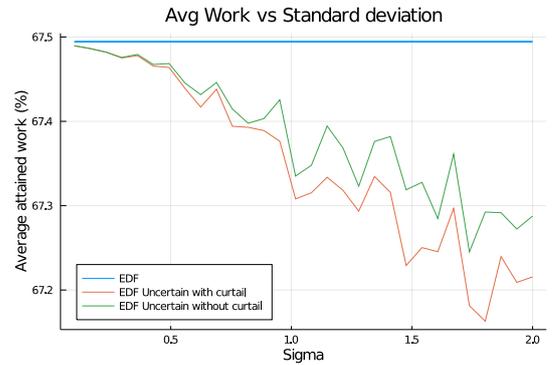


Fig. 8. Average Work vs Standard deviation - Curtail

As for the received energy, in Fig. 9 we compare the attained charge obtained by each EV under perfect information EDF and the curtailed version. We see that in this case, a similar threshold is attained and thus the average overall performance per vehicle presents no significant departures from the full information case.

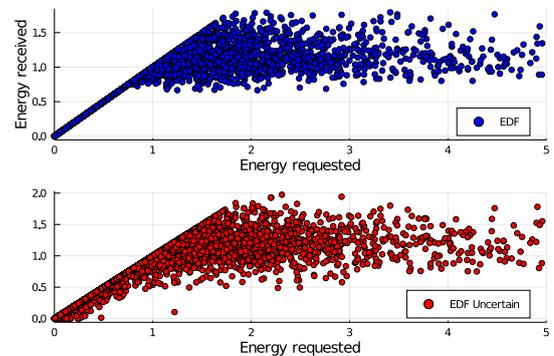


Fig. 9. Energy Requested and received - Curtail

Finally, we turn our attention to the incentives. In Fig. 10 we plot the difference in energy received as a function of the deadline under-reporting $T_k - T'_k$; large values of this difference means larger under-reporting. As compared to Fig. 6, we have reduced the correlation between under-reporting and received energy. In fact, the peak is attained when $T'_k \approx T_k$, meaning that the incentive for users is to declare their true deadlines upon arrival, since doing this maximizes the energy they receive.

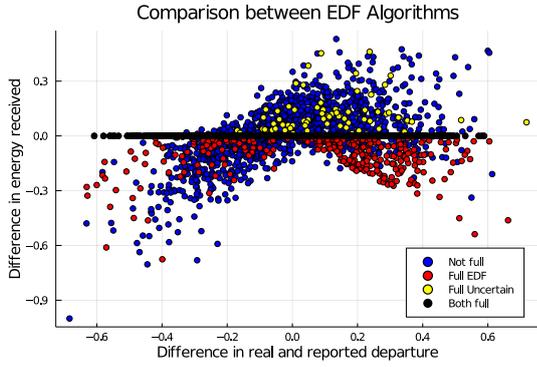


Fig. 10. Comparison between EDF with and without uncertainty - Curtail

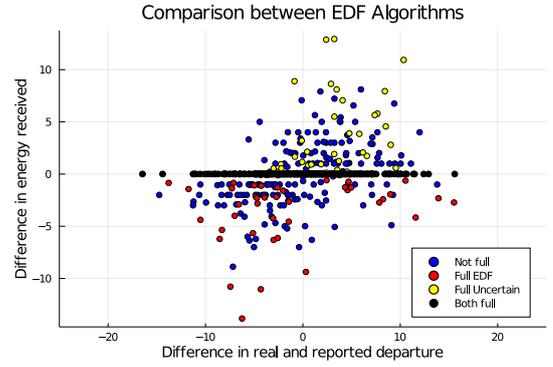


Fig. 12. Comparison between algorithms - Real world traces

IV. PERFORMANCE IN A REAL WORLD SCENARIO

In this Section we evaluate our algorithms using a real world scenario. We use data from Caltech’s ACN Portal [2], with information from the 5th of September to the 12th of September of 2018 in a Caltech parking lot. We focus on the performance of the curtailed version of the EDF algorithm.

We obtain 578 entries which we run into our algorithms, reducing the site’s capacity to two EVs in order to analyze the system in an overload condition. The average work requested for EVs is 8.7 kWh and the average sojourn time is 5.3 hours.

In Fig. 11 we plot the amount of energy received when running the Uncertain EDF algorithm using both the reported and real departure time. We observe the same trends as before: some EVs might receive considerable more energy depending on the reported departure.

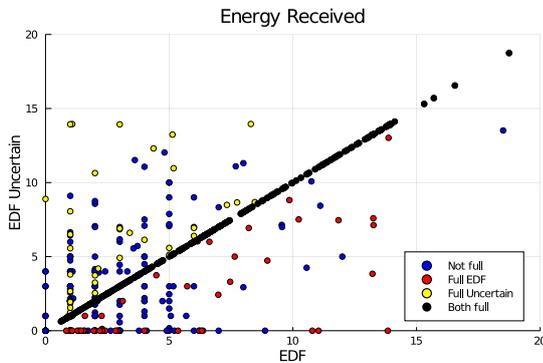


Fig. 11. Energy Received - Real world traces

Furthermore, we plot in Figure 12 which vehicles receive more energy as a function on the error in their reported deadlines. The pattern looks similar to the one in Fig. 10, where the peak received energy is attained by vehicles that report their true departure times. Since the parking lot is small, more variability is observed due to the fact that the fluid limit approximation is not as good in this case. However, we expect the algorithm to perform better in larger systems.

V. CONCLUSIONS

We studied EV scheduling algorithms for a centralized charging facility, focusing on the EDF policy, building on analytical models of attained service from [1]. We showed overall performance is robust against random errors in the deadline declaration, but may generate improper incentives. An alternative that removes these incentives was proposed and demonstrated to perform appropriately, in both synthetic data and real-world data.

REFERENCES

- [1] M. Zeballos, A. Ferragut, and F. Paganini, “Proportional fairness for ev charging in overload,” *IEEE Transactions on Smart Grid*, vol. 10, no. 6, pp. 6792–6801, 2019.
- [2] Z. Lee, T. Li, and S. H. Low, “ACN-Data: Analysis and applications of an open EV charging dataset,” in *Proceedings of the Tenth International Conference on Future Energy Systems*, ser. ACM e-Energy ’19, Jun. 2019.
- [3] 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.
- [4] S. Chen, L. Tong, and T. He, “Optimal deadline scheduling with commitment,” in *Communication, Control, and Computing (Allerton), 2011 49th Annual Allerton Conference on*. IEEE, 2011, pp. 111–118.
- [5] A. Subramanian, M. J. Garcia, D. S. Callaway, K. Poolla, and P. Varaiya, “Real-time scheduling of distributed resources,” *IEEE Transactions on Smart Grid*, vol. 4, no. 4, pp. 2122–2130, 2013.
- [6] A. Aveklouris, Y. Nakahira, M. Vlasίου, and B. Zwart, “Electric vehicle charging: a queueing approach,” *ACM SIGMETRICS Performance Evaluation Review*, vol. 45, no. 2, pp. 33–35, 2017.
- [7] A. Aveklouris, M. Vlasίου, and B. Zwart, “A stochastic resource-sharing network for electric vehicle charging,” *arXiv:1711.05561*, 2017.
- [8] 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.
- [9] —, “Preserving fairness in EV charging under time-varying congestion levels,” in *Proc. of IEEE 9th Power, Instrumentation and Measurement Meeting*, 2018.
- [10] L. Kruk, J. Lehoczky, K. Ramanan, and S. Shreve, “Double Skorokhod map and reneging real-time queues,” in *Markov Processes and Related Topics: A Festschrift for Thomas G. Kurtz*. Institute of Mathematical Statistics, 2008, pp. 169–193.
- [11] R. Atar, A. Biswas, and H. Kaspi, “Fluid limits of G/G/1+G queues under the nonpreemptive earliest-deadline-first discipline,” *Mathematics of Operations Research*, vol. 40, no. 3, pp. 683–702, 2014.
- [12] Y. Nakahira, N. Chen, L. Chen, and S. H. Low, “Smoothed least-laxity-first algorithm for EV charging,” in *Proc. of the 8th Intl. Conf. on Future Energy Systems*, 2017, pp. 242–251.
- [13] “EVQueues.jl, a Julia library for simulating EV charging policies.” <https://github.com/Grupo-MATE/EVQueues.jl>.