Scheduling EV charging with uncertain departure times

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

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

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

ABSTRACT

In an EV charging facility, with multiple vehicles requesting charge simultaneously, scheduling becomes crucial to provide adequate service under vehicle sojourn time constraints. However, these departure times may not be known accurately, and typical policies such as Earliest-Deadline-First or Least-Laxity-First are affected by this uncertainty in information. In this paper, we analyze the performance of these policies under uncertain deadlines, using a mean-field approach. We characterize the deviation in individual attained service as a function of the uncertainty. Since incentives appear to under-report deadlines in order to be prioritized, we analyze a simple modification of the policies to enforce incentive compatibility. Simulation experiments are carried out with a practical data set.

Keywords

EV charging, Scheduling, Deadlines.

1. INTRODUCTION

The penetration of Electrical Vehicles (EVs) currently under way is demanding the deployment of an adequate charging infrastructure [4, 7]. In this regard, an attractive option is to have centralized parking lots with support for EV charging, for instance at large workplaces. Since the power demand of EVs is significant, when this mode of transportation becomes ubiquitous, it seems reasonable to rely on statistical multiplexing in order to provide adequate service without the need to design for peak power consumption. In fact, since EVs may tolerate some deferral of service, *scheduling* becomes crucial [8].

In such a facility, a scheduling policy takes into account users' sojourn times and charge requirements, and makes decisions on who receives service at any given time. There is a rich literature on the scheduling of deadline constrained tasks, particularly in processor task scheduling [9]; more recently the problem has received renewed attention in a smart-grid context [6, 11] given the presence of deferrable loads. The main difference with classical scheduling is that, while deadlines are strict, partial service (charge) has value.

*This work was partially supported by ANII Uruguay under grant FSE_1_2018_1_153050. The authors wish to thank Steven Low and the NetLab team at Caltech for facilitating the real EV dataset.

IFIP Performance '21 November 8–12, 2021, Milan, Italy

In recent work, a queueing approach has been proposed for this kind of system, akin to a many server queue with deadlines. In [1, 2], the authors analyze the rate-limited processor sharing policy, and how to extend it to the network case. In our prior research [12, 13], a general fluid model for deadline-based policies was proposed, and analytical results given for their behavior in overload, including in particular the classical Earliest Deadline First (EDF) and Least Laxity First (LLF), among others.

The main limitation of such deadline-based policies is that they require advance knowledge of the exact sojourn time of customers. In a practical scenario, users have uncertainty on their departure time, or may even misreport it to affect their scheduling priority. In this paper, we build on the mean-field approach of [13] and extend it to include deadline uncertainty. We provide detailed analysis for the EDF and LLF policies, in particular of how each EV's attained service depends on the reporting error. It is shown that strategic users may game the system in order to benefit from a larger service share. We thus also provide a simple variant of the classical policies that removes those incentives.

The paper is organized as follows: in Section 2 we discuss our model and summarize the main results for EDF and LLF from [13]. In Section 3 we discuss the inaccuracy in reported departure times, deriving explicit results for attained service under a parametric model. Section 4 analyzes incentives. Simulations with a practical data set are presented in Section 5, and conclusions provided in Section 6.

2. SYSTEM MODEL

Consider a parking lot providing EV homogeneous charging stations at every spot, with a nominal power rating (maximum charging rate). We assume that the size of the parking lot is large and never fills, but electrical power is limited so that at most C chargers can be on simultaneously. The scheduling policy must allocate these limited resources among the EV clients currently present, considering their energy needs and planned departure times. In this sense, the system behaves as a many-server queue with deadlines.

Assume that vehicles arrive as a Poisson process of intensity λ, each one having two random characteristics: a required *service or charging time* S_k, and a *sojourn time* T_k, which is the time until the car leaves the parking lot. The pairs (S_k, T_k) follow general distributions and we impose $T_k \geqslant S_k$ with probability 1, which ensures the demand of each EV is a priori feasible. We denote by S, T the random variables representing the characteristics of a general vehicle. The system load is given by $\rho := \lambda E[S]$.

In [1], the authors develop a mean-field model for such a system using the Processor Sharing policy in a large system setting: namely $\lambda \to \infty$ while keeping ρ/C constant. This mean field model was extended in [12, 13] to encompass all policies that depend on the *residual service time* σ_k and *residual sojourn time* τ_k of each vehicle present. In this paper we will focus on analyzing the mean field model equilibria as given in [12, 13]. For a relationship between the mean field model and the stochastic system we refer the reader to [2].

Consider as an example the Earliest-Deadline-First (EDF) scheduling policy: here the EVs are ranked in increasing order of their residual sojourn time τ, and the C most urgent vehicles are served at any given time (with preemption)[1], In [12, 13], it is shown that, if the system is in overload, i.e. $\rho > C$, the equilibrium solution of the mean field model exhibits a threshold behavior: any EV arriving into the system acquires priority only when its residual sojourn time reaches a critical threshold τ_0^*. After that point, it is served at full rate up to full charge (as represented on the top of Fig. 1) or deadline expiration. As a consequence, the *attained service* for a typical vehicle, S_a^0, is given by:

$$S_a^0 = \min\{S, \tau_0^*\}, \tag{1}$$

where the threshold τ_0^* follows from imposing the capacity condition:

$$\lambda E[S_a^0] = \lambda E[\min\{S, \tau_0^*\}] = C, \tag{2}$$

which has a unique solution when $\rho > C$ since $E[\min\{S, \tau_0^*\}] \uparrow E[S]$ as $\tau_0^* \uparrow \infty$.

A similar analysis applies for LLF. In this policy, a vehicle arrives with a certain *laxity* $L = T - S$. As time elapses, the remaining laxity is given by $\ell = \tau - \sigma$, and vehicles are ranked in increasing order of ℓ_k. In this case, ℓ may become negative, meaning that the EV is already late for full service.

In [12, 13], it was established that the equilibrium solution for the mean field limit in overload displays again a threshold behavior. A fixed threshold ℓ_0^* determines the laxity at which vehicles start service. In overload, $\ell_0^* < 0$ meaning that all vehicles are behind schedule. A typical trajectory is depicted in the second diagram of Fig. 1. An EV arriving at time t consumes its laxity and begins service at time $t + L - \ell_0^*$, receiving a total service time $T - L + \ell_0^* = S + \ell_0^*$. If a vehicle has $S + \ell_0^* < 0$, its laxity never becomes urgent enough and it receives no service at all. Defining $\sigma_0^* = -\ell_0^*$, the attained service is given by:

$$S_a^0 = (S - \sigma_0^*)^+, \tag{3}$$

Again the threshold is determined by the capacity condition:

$$\lambda E[S_a^0] = \lambda E\left[(S - \sigma_0^*)^+\right] = C, \tag{4}$$

which has a unique solution since the left hand side decreases monotonically from ρ to 0 as $\sigma_0^* \uparrow \infty$.

It turns out that EDF and LLF operate as dual policies in the mean field regime: while the former prioritizes small jobs, which are served to completion, and caps large jobs to a threshold τ_0^*, the latter provides service to larger jobs and equalizes their *reneged* work to σ_0^*. We analyze next how these behaviors are perturbed when the system does not have access to the true deadlines.

[1] All the policies considered here are preemptive, otherwise the model of [1, 12, 13] is not applicable.

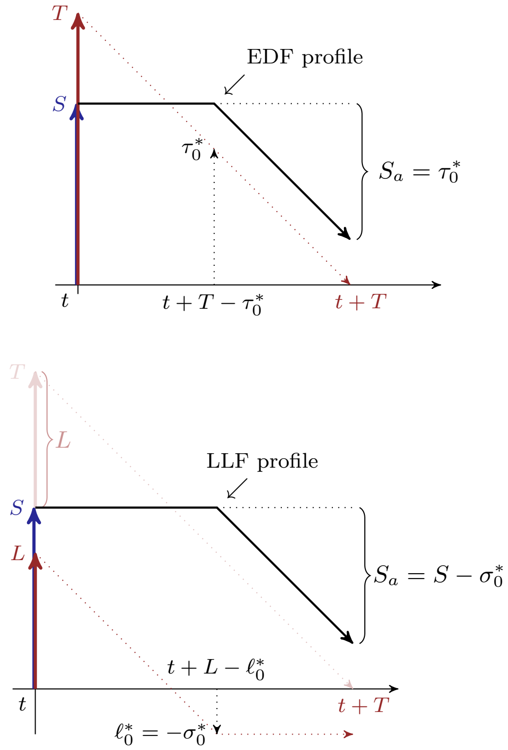

Figure 1: Charging profiles under EDF (above) and LLF (below) in the mean field limit.

3. MODELING DEADLINE UNCERTAINTY

We now extend the previous model to include the possibility that sojourn times are not exactly known by the scheduler. This is the case, for example, when users are asked to declare their departure time upon arrival, an estimate subject to uncertainty or misreporting.

We thus introduce a *declared sojourn time* T' available to the scheduler, perturbation of the real sojourn time T. S is the requested service time as before. Also let τ' denote the remaining declared sojourn time of a given vehicle. When working in the mean field limit, again a threshold behavior emerges but *relative to the declared sojourn time.* We now analyze both policies under this assumption.

3.1 EDF under uncertainty

Consider first the EDF algorithm where users are served in increasing order of τ_k'. The typical trajectory in such a scenario is depicted in the first diagram of Fig. 2. We have the follwong:

PROPOSITION 1. *Consider an EDF charging system with priorities based on declared sojourn times T', servicing traffic with service-sojourn times (S, T). For an overload situation $\rho > C$ and in the mean field regime, the attained service of a given vehicle is given by*

$$S_a = \min\{S, (T - T' + \tau^*)^+\}, \tag{5}$$

where the threshold τ^* satisfies the balance equation:

$$\lambda E[\min\{S, (T - T' + \tau^*)^+\}] = C. \tag{6}$$

PROOF. In the mean field regime, a threshold τ^* emerges and a vehicle arriving at time t becomes prioritized whenever

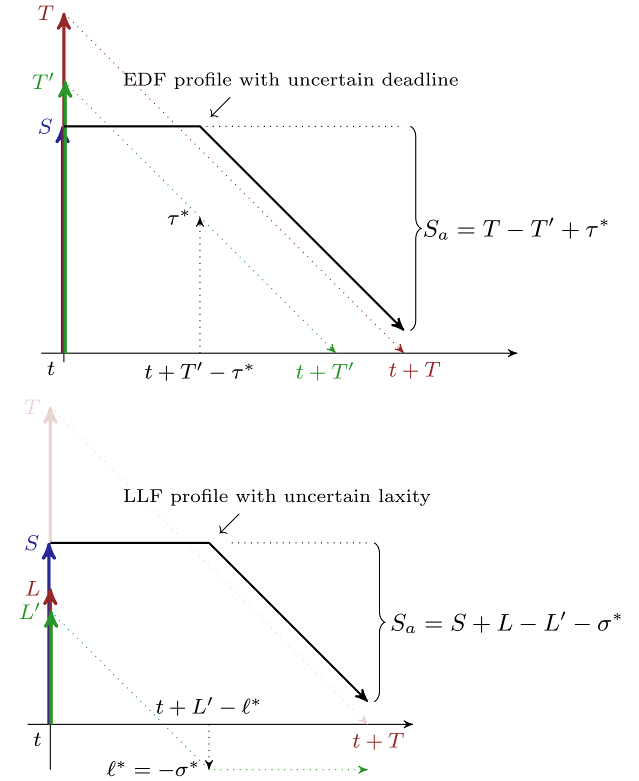

Figure 2: Charging profile of EDF (above) and LLF (below) under deadline uncertainty.

$\tau' < \tau^*$, i.e. at time $t + T' - \tau^*$ since arrival. Its service ends whenever it reaches full charge or departs at time $t+T$. Therefore in this case its service time is $\min\{S, T-T'+\tau^*\}$. If $T-T'+\tau^* < 0$ it departs before getting any service, thus proving (5). Eq. (6) follows from the capacity condition. $\square$

Note that eq. (6) always has a unique solution when $\rho > C$, since $E[\min\{S,(T-T'+\tau^*)^+\}]$ grows monotonically from 0 to $E[S]$ as $\tau^* \uparrow \infty$.

We wish to understand the effect of the uncertainty on system performance. Note that in an overload situation, the *mean* attained service across EVs is insensitive to uncertainty: it is fixed by the balance equation (6). Our focus is on the individual EV performance, as a function of the misreport in sojourn time, which we denote by $U = T' - T$, assumed independent of the service S. The following expression is an immediate consequence of (5):

$$E[S_a \mid U] = E[\min\{S, (\tau^* - U)^+\} \mid U], \tag{7}$$

the expected service attained by vehicles that misreport their deadline by U.

To make the discussion concrete we calculate the above expressions for a parametric model to obtain explicit expressions that are suitable to analysis. In particular, if we assume that the charging time $S \sim \exp(\mu)$ we have a useful Lemma.

Lemma 1. *If S is an exponential random variable with parameter μ and $x \geqslant 0$ then:*

$$E[\min\{S,x\}] = \int_0^x e^{-\mu s} ds = \frac{1}{\mu}(1 - e^{-\mu x}).$$

The following proposition follows directly from Lemma 1 and Proposition 1:

Proposition 2. *In an EDF system in overload, with $S \sim \exp(\mu)$ service times and independent uncertainty $U = T' - T$ in declared deadlines, the attained service for a given uncertainty is:*

$$E[S_a \mid U] = \frac{1 - e^{-\mu(\tau^* - U)^+}}{\mu}, \tag{8}$$

where τ^* satisfies the fixed point equation (6).

To proceed further we now assume that $U \sim \text{Uniform}[-\theta, \theta]$; here θ acts as an uncertainty parameter.[2] Let $X = (\tau^* - U) \sim \text{Uniform}[\tau^* - \theta, \tau^* + \theta]$. Assuming $\tau^* \geq \theta$ so $X \geq 0$ *a.s*[3], we have:

$$E[S_a] = E[E[\min\{S, X\} \mid X]] = \frac{1}{\mu}\left(1 - e^{-\mu\tau^*}\frac{\sinh(\mu\theta)}{\mu\theta}\right).$$

With the above formula, we can compute the threshold for a given λ, C by solving (6) to yield:

$$\tau^* = -\frac{1}{\mu}\log\left(\frac{\mu\theta}{\sinh(\mu\theta)}\left(1 - \frac{C}{\rho}\right)\right). \tag{9}$$

As $\theta \to 0$, the uncertainty disappears, $\frac{\mu\theta}{\sinh(\mu\theta)} \to 1$ and we have the expression $\tau_0^* = -\frac{1}{\mu}\log\left(1 - \frac{C}{\rho}\right)$ for the EDF threshold studied in [12].

We can further combine (8) with the threshold in (9) to quantify the uncertainty impact in that case. Specifically, we compute the *relative gain* (RG_{EDF}) in attained service for a given uncertainty level with respect to S_a^0, the attained service in the perfect information case:

$$RG_{EDF}(U) = \frac{E[S_a - S_a^0 \mid U]}{E[S]} = \frac{E[S_a \mid U] - E[S_a^0]}{E[S]}.$$

From eq. (2) we know that $E[S_a^0] = C/\lambda$ in the mean field limit, and $E[S_a \mid U]$ follows from eq. (8). For the uniform uncertainty case, we solve this explicitly to yield:

$$RG_{EDF}(U) = \left(1 - \frac{C}{\rho}\right)\left(1 - \frac{\mu\theta}{\sinh(\mu\theta)}e^{\mu U}\right). \tag{10}$$

As discussed above, the relative gain is decreasing in U, penalizing those cars that overreport their sojourn times.

3.2 LLF under uncertainty

As before, an analogous situation occurs when working with the LLF policy. We denote by $L' = T' - S$ the uncertain initial laxity and by $\ell' = \tau' - \sigma$ the remaining uncertain laxity. The analogue of Proposition 1 is:

Proposition 3. *Consider an LLF charging system working under deadline uncertainty. Assume that the system is in overload, i.e. $\rho > C$. Then in the mean field regime the attained service of a given vehicle is given by*

$$S_a = (S + L - L' - \sigma^*)^+, \tag{11}$$

where the threshold σ^* satisfies the balance equation:

$$\lambda E[(S + L - L' - \sigma^*)^+\}] = C. \tag{12}$$

[2] Similar computations can be performed for different distributions. We choose the uniform case for its simplicity.
[3] The alternate case $\tau^* < \theta$ can be handled similarly by a corrected calculation.

PROOF. The main argument is depicted in the lower diagram of Fig. 2. In the mean field limit a fixed threshold ℓ^* emerges and a vehicle arriving at time t starts service whenever $\ell' = \ell^*$, i.e. at time $t + L' - \ell^*$. Under this assumption, the vehicle would depart at $t + T$ so the total service time is $(t + T) - (t + L' - \ell^*) = S + L - L' + \ell^*$, provided the latter quantity is positive. Otherwise it will never be prioritized and receives no service upon departure. Identifying as before $\sigma^* = -\ell^*$ and combining the above remarks we have (11), and (12) follows from the capacity condition. □

Eq. (12) has a unique solution since $E[(S+L-L'-\sigma^*)^+]$ decreases monotonically as $\sigma^* \uparrow \infty$. Noting also that $L' - L = T' - T = U$, we arrive at the expression, analogous to (7) of the attained service conditioned on the uncertainty level:

$$E[S_a \mid U] = E[(S - (U + \sigma^*))^+ \mid U]. \tag{13}$$

Computations for the LLF case now proceed for the previously stated parametric models. For $S \sim \exp(\mu)$ service times, we have similar statements to Lemma 1 and Proposition 2:

LEMMA 2. *If S is an exponential random variable with parameter μ and $x \geqslant 0$ then:*

$$E[(S - x)^+] = \int_x^\infty e^{-\mu s} ds = \frac{e^{-\mu x}}{\mu}.$$

PROPOSITION 4. *In an LLF system in overload, with $S \sim \exp(\mu)$ service times and independent uncertainty $U = T' - T$ in declared deadlines, the attained service for a given uncertainty is:*

$$E[S_a \mid U] = \frac{e^{-\mu(U+\sigma^*)}}{\mu}, \tag{14}$$

where σ^* satisfies the fixed point equation (12).

For the uniform uncertainty case $U \sim \text{Uniform}[-\theta, \theta]$ we can compute:

$$E[S_a] = E[E[S_a \mid U]] = E\left[\frac{e^{-\mu(U+\sigma^*)}}{\mu}\right] = \frac{e^{-\mu\sigma^*}}{\mu} E\left[e^{-\mu U}\right]$$
$$= \frac{1}{\mu} e^{-\mu\sigma^*} \frac{\sinh(\mu\theta)}{\mu\theta}.$$

With this result we can explicitly solve for the threshold σ^* by substituting in (12):

$$\sigma^* = -\frac{1}{\mu} \log\left(\frac{\mu\theta}{\sinh(\mu\theta)} \frac{C}{\rho}\right). \tag{15}$$

As in the previous case, as $\theta \to 0$ we recover the threshold expression for σ_0^* from [12].

Finally, we can compute the relative gain in attained service to quantify the impact of uncertainty for our specific parametric model. Combining (14) with the threshold in (15) we obtain:

$$E[S_a \mid U] = \frac{e^{-\mu\sigma^*}}{\mu} e^{-\mu U} = \frac{\mu\theta}{\sinh(\mu\theta)} \frac{C}{\lambda} e^{-\mu U},$$

and from there we compute the relative gain, which again penalizes the users that overreport their sojourn times:

$$RG_{LLF}(U) = \frac{E[S_a \mid U] - C/\lambda}{1/\mu} = \frac{C}{\rho} \left(\frac{\mu\theta}{\sinh(\mu\theta)} e^{-\mu U} - 1\right).$$

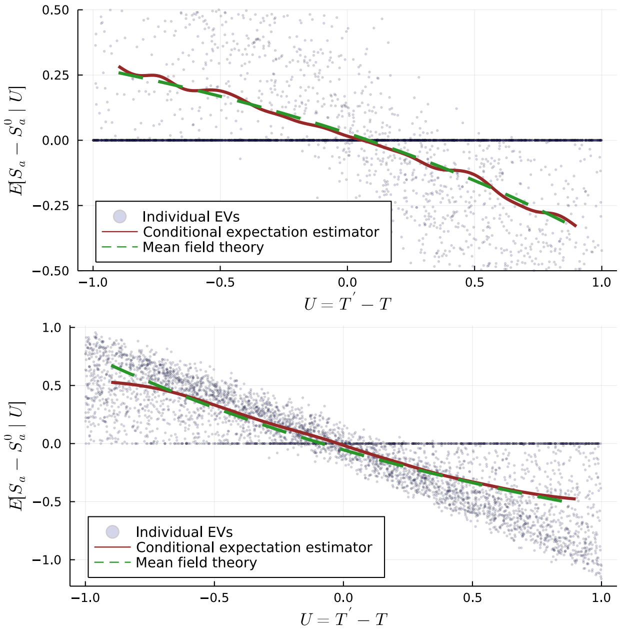

Figure 3: Attained energy difference between the uncertain and perfect information case. EDF algorithm (above) and LLF algorithm (below).

3.3 Simulation example

In Fig. 3, we show the results of a simulation experiment performed using the Julia library `EVQueues.jl` [3], based on a discrete event simulation of the parking lot. The parameters are $\lambda = 30$ EVs/h for the Poisson arrival rate, $C = 40$, $E[S] = 2$h and $\theta = 1$h. The total simulation time is 200h, with around 6000 vehicles being served. For the same arrival stream and demands, we compute the attained energy difference when using the uncertain deadline against the perfect information case, as a function of the uncertainty level $U = T' - T$. We do so for both EDF and LLF.

In solid lines, the average gain $E[S_a - S_a^0 \mid U]$ is estimated via the Nadaraya-Watson kernel regression estimator for the conditional expectation [10] applied to the observed set of points. This is compared in dashed lines with the corresponding theoretical expressions derived from the mean field limit (8) and (14), showing good fit.

A crucial remark from the previous analysis is that EVs *under-reporting deadlines* tend to receive a larger share of service. This is true regardless of the uncertainty model and the algorithm in use, because the expressions (7), (13) are *decreasing* in U.

4. CURBING INCENTIVES

Given the preceding observation, we wish to provide incentive compatibility to declare unbiased deadlines. A simple algorithm to enforce this is to serve EVs only up to departure or declared departure time, whichever happens first. We call this the *curtailed* version of the policy.

We begin by analyzing the curtailed EDF case, where we have the following:

PROPOSITION 5. *Consider an EDF charging system working under deadline uncertainty and curtailing users when their declared deadline expires. In the mean field limit with* *the system in overload ($\rho > C$) the attained service of a given vehicle satisfies:*

$$S_a = \min\{S, ((T - T')\mathbf{1}_{\{T<T'\}} + \tau^*)^+\}, \quad (16)$$

where the threshold comes from the fixed point equation:

$$\lambda E[\min\{S, ((T - T')\mathbf{1}_{\{T<T'\}} + \tau^*)^+\}] = C. \quad (17)$$

PROOF. The proof is analogous to the proof of Proposition 1, but noting that a vehicle arriving at time t will receive service up to $\min\{t+T, t+T'\}$, and its service time will be either S or:

$$\min\{t + T, t + T'\} - (t + T' - \tau^*) = (T - T')\mathbf{1}_{\{T<T'\}} + \tau^*,$$

provided that the above term is positive. Otherwise it will not receive service. This proves (16), and (17) follows from the capacity condition. Note that the indicator term in (16) only becomes active when an EV over-reports its deadline. $\square$

We now solve the above equations for the parametric case of exponential service times and uniform deadline uncertainty. The attained work for $\tau^* > \theta$ is given by:

$$S_a = \min\{S, \tau^* - U\mathbf{1}_{\{U>0\}}\}.$$

Invoking again Lemma 1 we arrive at:

$$E[S_a] = E[E[S_a \mid U]] = \int_{-\theta}^{\theta} \frac{1 - e^{-\mu(\tau^* - u\mathbf{1}_{\{u>0\}})}}{\mu} \frac{1}{2\theta} du$$
$$= \frac{1}{\mu}\left[1 - e^{-\mu\tau^*}\left(\frac{1}{2} + \frac{e^{\mu\theta} - 1}{2\mu\theta}\right)\right],$$

and we can solve for the threshold in the same way as before:

$$\tau^* = -\frac{1}{\mu}\log\left[\left(\frac{1}{2} + \frac{e^{\mu\theta} - 1}{2\mu\theta}\right)^{-1}\left(1 - \frac{C}{\rho}\right)\right]. \quad (18)$$

Analogous to Proposition 2, by invoking Lemma 1 we find an expression for the conditional expectation of attained service with respect to the uncertainty

$$E[S_a \mid U] = \frac{1 - e^{-\mu(\tau^* - U\mathbf{1}_{\{U>0\}})^+}}{\mu}, \quad (19)$$

where τ^* is given by (18).

The expression in (19) is again non-increasing in U, but it is now *constant* whenever $U < 0$, i.e. $T' < T$. This curbs the incentive to under-report the deadline since no gain is obtained on average with respect to $T' = T$. For the parametric model of U under consideration, we obtain from (18) a formula for the relative gain in service:

$$RG_{EDFc}(U) = \left(1 - \frac{C}{\rho}\right)\left(1 - \frac{e^{\mu U\mathbf{1}_{\{U>0\}}}}{\frac{1}{2} + \frac{e^{\mu\theta}-1}{2\mu\theta}}\right). \quad (20)$$

A similar analysis can be performed for the curtailed LLF policy to yield:

PROPOSITION 6. *Consider an LLF charging system working under deadline uncertainty and curtailing users when their declared deadline expires. In the mean field limit with the system in overload ($\rho > C$) the attained service of a given vehicle satisfies:*

$$S_a = \left(S - (L - L')\mathbf{1}_{\{L<L'\}} - \sigma^*\right)^+, \quad (21)$$

where the threshold comes from the fixed point equation:

$$\lambda E[\left(S - (L - L')\mathbf{1}_{\{L<L'\}} - \sigma^*\right)^+] = C. \quad (22)$$

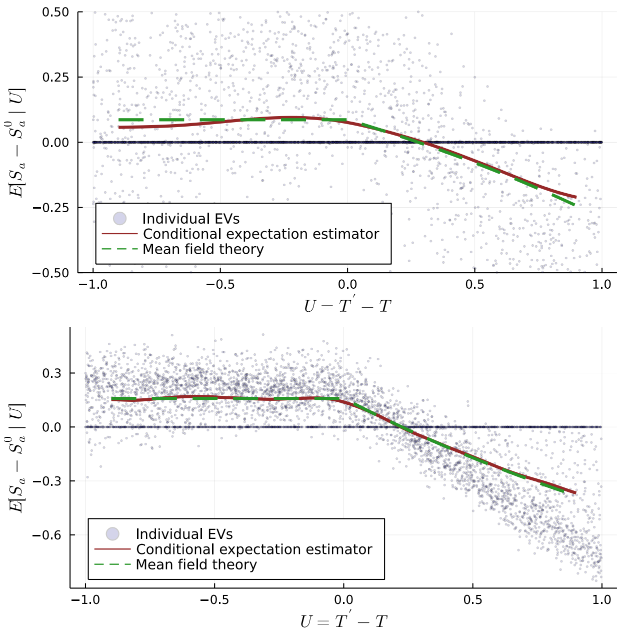

Figure 4: Attained energy difference between the uncertain and perfect information case when curtailing is applied. EDF (above) and LLF (below).

PROOF. A vehicle arriving at time t will become prioritized at time $t + L' - \ell^*$ and receive service up to $\min\{t + T, t + T'\}$ due to the curtailing. Therefore, its service is:

$$\min\{t + T, t + T'\} - (t + L' - \ell^*) =$$
$$S + \min\{L, L'\} - L' + \ell^* = S + (L - L')\mathbf{1}_{\{L<L'\}} + \ell^*,$$

provided the above term is positive. Otherwise, its laxity will never become low enough to get priority in the first place. Identifying again $\sigma^* = -\ell^*$ we get (21) and the proposition follows. $\square$

Working again in the parametric case, taking $L' - L = U$ then eq. (21) becomes:

$$S_a = \left(S - U\mathbf{1}_{\{U>0\}} - \sigma^*\right)^+.$$

Applying now Lemma 2 we can compute:

$$E[S_a] = E[E[S_a \mid U]] = \int_{-\theta}^{\theta} \frac{e^{-\mu(u\mathbf{1}_{\{u>0\}} + \sigma^*)}}{\mu} \frac{1}{2\theta} du$$
$$= e^{-\mu\sigma^*}\left(\frac{1}{2} + \frac{1 - e^{-\mu\theta}}{2\mu\theta}\right),$$

and we can solve for the threshold in the same way as before:

$$\sigma^* = -\frac{1}{\mu}\log\left[\left(\frac{1}{2} + \frac{1 - e^{-\mu\theta}}{2\mu\theta}\right)^{-1}\frac{C}{\rho}\right]. \quad (23)$$

We also have the analogous result to Proposition 4 for the conditional attained work under LLF:

$$E[S_a \mid U] = \frac{e^{-\mu(U\mathbf{1}_{\{U>0\}} + \sigma^*)}}{\mu}, \quad (24)$$

where σ^* is given by (23). Again, the above expression is non-increasing in U and constant for $U < 0$, meaning there is no gain on average by under-reporting the deadline.

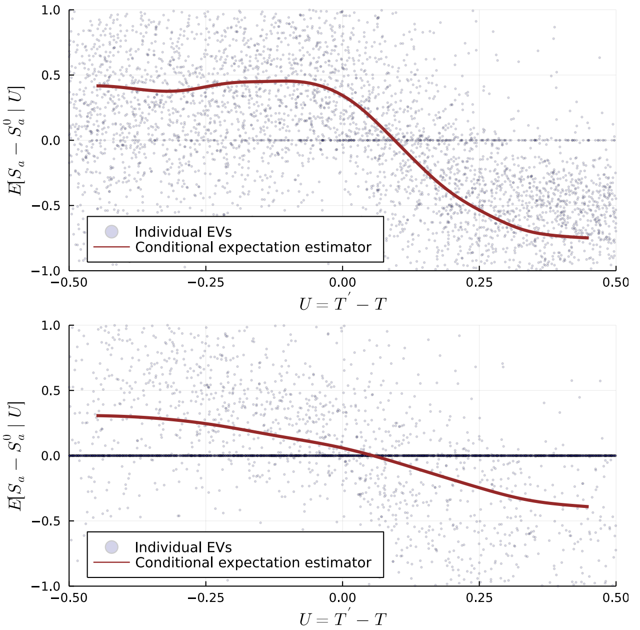

Figure 5: Attained energy difference in a real-world trace under curtailed EDF (above) and LLF (below).

Combining all of the above, we obtain an explicit expression for the relative gain in attained service for the curtailed LLF policy:

$$RG_{LLFc}(U) = \frac{C}{\rho}\left[\left(\frac{1}{2} + \frac{1-e^{-\mu\theta}}{2\mu\theta}\right)^{-1} e^{-\mu U \mathbf{1}_{\{U>0\}}} - 1\right]. \tag{25}$$

In Fig. 4, we show the results of the simulation experiment as before but applying the new curtailed policies. We can see that the average gain is curbed whenever $U < 0$, and coincides with the expressions (20), (25) based on mean field.

5. SIMULATIONS WITH REAL TRACES

Finally, we simulate our algorithms using real world traces from a parking lot at a Silicon Valley firm, kindly provided by the authors of [5]. For this purpose, we build a set with arrival times, departures, demands and power requested by EVs in a multi-day period with time-varying demand. The charging capacity is set at $C = 30$ charging stations in order to analyze the system in an overload situation. The average sojourn time in our set is 2.25 hours, and the average charging time is 1.77 hours. We introduce an uniform uncertainty in deadline reporting with $\theta = 0.5$ hours. For the particular dataset analyzed, the parking lot is overloaded 73% of the total simulation time.

Fig. 5 shows the results for both the EDF and LLF curtailed policies applied to this scenario, as well as the Nadaraya-Watson estimator for the conditional expectation. We observe similar patterns as the theoretic results, most importantly, the effect of enforcing an incentive compatible policy with the curtailment.

6. CONCLUSIONS

We analyzed the performance in overload of the EDF and LLF policies for scheduling EV charging at a parking lot, when departure times are not known exactly. In a suitable mean-field model [13] these policies are characterized by a service threshold; we extended the analysis to include deadline uncertainty. A relevant performance metric is the relative gain in service conditioned on the deadline reporting error; explicit formulas were obtained for a parametric case. We further proposed and analyzed a modification of both policies to curb incentives to under-report the deadline, and validated its performance with real world traces. In future work we plan to address the problem where users' decisions to join the system and their sojourn times may depend on the current congestion state.

7. REFERENCES

[1] A. Aveklouris, Y. Nakahira, M. Vlasiou, and B. Zwart. Electric vehicle charging: a queueing approach. *ACM SIGMETRICS Performance Evaluation Review*, 45(2):33–35, 2017.

[2] A. Aveklouris, M. Vlasiou, and B. Zwart. A stochastic resource-sharing network for electric vehicle charging. *IEEE Transactions on Control of Network Systems*, 6(3):1050–1061, 2019.

[3] `EVQueues.jl`. A Julia language library for simulating Electrical Vehicle charging policies. `https://github.com/Grupo-MATE/EVQueues.jl`, 2021.

[4] O. Hafez and K. Bhattacharya. Integrating ev charging stations as smart loads for demand response provisions in distribution systems. *IEEE Transactions on Smart Grid*, 9(2):1096–1106, 2018.

[5] Z. Lee, T. Li, and S. H. Low. ACN-Data: Analysis and applications of an open EV charging dataset. In *Proc. of ACM e-Energy 2019*.

[6] Y. Nakahira, N. Chen, L. Chen, and S. H. Low. Smoothed least-laxity-first algorithm for EV charging. In *Proc. of ACM e-Energy 2017*.

[7] B. Sun, Z. Huang, X. Tan, and D. H. Tsang. Optimal scheduling for electric vehicle charging with discrete charging levels in distribution grid. *IEEE Transactions on Smart Grid*, 9(2):624–634, 2018.

[8] W. Tang, S. Bi, and Y. J. Zhang. Online charging scheduling algorithms of electric vehicles in smart grid: An overview. *IEEE communications Magazine*, 54(12):76–83, 2016.

[9] D. Towsley and S. S. Panwar. On the optimality of minimum laxity and earliest deadline scheduling for real-time multiprocessors. In *Proc. of Euromicro'90*, pages 17–24, 1990.

[10] M. P. Wand and M. C. Jones. *Kernel smoothing*. CRC press, 1994.

[11] Y. Xu, F. Pan, and L. Tong. Dynamic scheduling for charging electric vehicles: A priority rule. *IEEE Transactions on Automatic Control*, 61(12):4094–4099, 2016.

[12] 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.

[13] M. Zeballos, A. Ferragut, and F. Paganini. Proportional fairness for EV charging in overload. *IEEE Transactions on Smart Grid*, 10(6):6792–6801, 2019.