

# Managing Overstaying Electric Vehicles in Park-and-Charge Facilities

Arpita Biswas, Ragavendran Gopalakrishnan, Partha Dutta

Xerox Research Centre India

{Arpita.Biswas, Ragavendran.Gopalakrishnan,Partha.Dutta}@xerox.com

## Abstract

With the increase in adoption of Electric Vehicles (EVs), proper utilization of the charging infrastructure is an emerging challenge for service providers. Overstaying of an EV after a charging event is a key contributor to low utilization. Since overstaying is easily detectable by monitoring the power drawn from the charger, managing this problem primarily involves designing an appropriate “penalty” during the overstaying period. Higher penalties do discourage overstaying; however, due to uncertainty in parking duration, less people would find such penalties acceptable, leading to decreased utilization (and revenue). To analyze this central trade-off, we develop a novel framework that integrates models for realistic user behavior into queueing dynamics to locate the optimal penalty from the points of view of utilization and revenue, for different values of the external charging demand. Next, when the model parameters are unknown, we show how an online learning algorithm, such as UCB, can be adapted to learn the optimal penalty. Our experimental validation, based on charging data from London, shows that an appropriate penalty can increase both utilization and revenue while significantly reducing overstaying.

## 1 Introduction

As the number of on-road Electric Vehicles (EVs) increases rapidly, lack of adequate charging infrastructure is an area of growing concern. For example, a recent New York Times article reports that in California, scarcity of park-and-charge spots is a chronic problem: “Electric-vehicle owners are unplugging one another’s cars, trading insults, and creating black markets and side deals to trade spots in corporate parking lots” [Richtel, 2015]. While investing in a wide-spread deployment of charging stations by both public and private service providers is needed to address the infrastructure problem [Gopalakrishnan *et al.*, 2016] in the long-term, effective use of the existing charging infrastructure can help in significantly reducing this problem. In particular, curbing improper utilization of park-and-charge spots can improve their availability. Two prominent causes of utilization degradation are

(i) the overstaying problem, where an EV continues to occupy a park-and-charge spot even after it is fully charged, and (ii) the “icing” problem, where a gas-powered car (Internal Combustion Engine or ICE vehicle) occupies a park-and-charge spot.

While icing and overstaying in park-and-charge spots are illegal in increasingly many jurisdictions, there is little or no enforcement [Loveday, 2013], as a result of which the frustrated users resort to ad-hoc measures ranging from mild (e.g., leaving courtesy notices) [Blink, 2012; 2015] to drastic (e.g., publicly shaming the violator by posting a picture of the violation showing the violator’s license plate on blogs and social media) [Tumblr, 2015; Zapatag, 2015]. Extension cables that can help charge a vehicle parked a few spots away from an occupied park-and-charge spot are available in the market, but are very expensive [Moloughney, 2015]. A longer term solution to these problems, however, will require both enforcement and an appropriate penalty for these events.

While enforcement with heavy penalty may help curb icing, a gentler approach is prudent to manage overstaying EVs. Depending on the demand for charging, overstaying EVs potentially block access to other EVs that might need charging, and so, it is important to discourage such behavior by imposing penalties. But, imposing too high a penalty might turn away EVs from using the park-and-charge facility altogether due to increased risk of a steep fine, since EV users may not exactly know their parking duration beforehand.

Our central contribution in this paper is a novel framework that combines a realistic user behavior model with traditional queueing dynamics to capture this trade-off and study the optimal penalty from the points of view of both utilization and revenue. Next, when the model parameters are unknown, we show how the well known UCB[Auer *et al.*, 2002] learning algorithm can be adapted to learn the optimal penalty over a period of time. Our experiments, based on charging data from London, show that an appropriate penalty results in increased utilization and significantly increased revenue. Also, perhaps surprisingly, we observe that the utilization achieved by imposing the revenue-maximizing penalty is very close to the maximum utilization.

### 1.1 Related Work

Dynamic pricing of parking has been an area of active study in the transportation literature. In [Zoeter *et al.*, 2014;

Rowe and Fiorucci, 2011], the authors present dynamic pricing schemes for regular parking based on estimated demand to reduce both congestion and underuse. On the other hand, significant work has also been done on dynamic pricing of electric vehicle charging. In [Tushar *et al.*, 2012], the authors model the problem as a Stackelberg game between the smart grid as the leader setting the price, and the vehicle owner as the follower deciding their charging strategies. In [Anglin *et al.*, 2013; Gadh *et al.*, 2012; Hafner *et al.*, 2009], dynamic pricing of electricity for charging EVs is proposed based on the usage data in a given location and time. None of the above work, however, consider the problem of overstaying EVs. To the best of our knowledge, this is the first paper to investigate designing penalty schemes for park-and-charge facilities to combat overstaying EVs.

To address the situation where the probability distributions required for modelling user behaviour remain unavailable to the system until the users are actually observed, and the user behaviour needs to be learned over time, we model our problem as a Multi-Armed Bandit (MAB) problem, which has been well studied [Bubeck and Cesa-Bianchi, 2012; Auer *et al.*, 2002; Cesa-Bianchi, 2006; Agrawal, 1995] and applied to a wide variety of domains such as crowdsourcing, online advertising, dynamic pricing, and smart grids. However, we are not aware of any application of MABs to the park-and-charge scenario.

## 2 Park-and-Charge System Model

We adopt a simple three-part model for a parking area: (a) Section 2.1 introduces the pricing function (during charging) and penalty function (during overstaying) for an EV, (b) Section 2.2 models how EV users respond to the posted penalty scheme, i.e., whether they agree to its terms and enter the parking area, and if so, how long they stay, and (c) Section 2.3 models the flow of EV users in and out of the parking area using queueing dynamics. We end the section by defining some performance measures of interest in Section 2.4.

### 2.1 Pricing and Penalty Functions

Let the price function while actively charging be  $p_c(t)$  and the penalty function for overstaying at a park-and-charge spot after charging is complete be  $p_o(t)$ , where  $p_c(0) = p_o(0) = 0$ , and  $p_c(t)$  and  $p_o(t)$  are continuous, nondecreasing functions. In addition, we assume that  $p_o(t)$  is a strictly increasing function, since we will need its inverse function,  $p_o^{-1}$ , to be well defined.<sup>1</sup> To be realistically implementable (easily understandable by an average user), these functions should not be more complicated than simple piecewise linear functions; however, our framework is general and can accommodate arbitrary functions. These functions are illustrated in Figure 1.

### 2.2 EV User Behavior

We define the following parameters, for each EV user:

- $T_c$ , the duration of time it would take to fully charge their electric vehicle from the initial state,

<sup>1</sup>The framework can be extended to the case where  $p_o(t)$  is not strictly increasing, by defining  $p_o^{-1}(x) = \sup\{t : p_o(t) = x\}$ . This could accommodate, for example, an initial “grace period” (no penalty) after charging is complete.




Figure 1: Pricing and penalty functions for an example scenario.

- $T_a$ , the length of their appointment, i.e., the EV user’s preferred parking duration, and
- $C_{\max}$ , the *penalty threshold*, i.e., the maximum penalty due to overstaying that they would bear for the convenience of an uninterrupted appointment.

Let  $T_c, T_a, C_{\max}$  be independent, nonnegative random variables with cumulative distribution functions  $F_c, F_a, F_{\max}$ .

**Penalty Acceptance Probability:** When a user arrives at the parking area and inspects the posted penalty scheme, they know their  $T_c$  and  $C_{\max}$ , but they may not know exactly how long their appointment will last; hence  $T_a$  is not yet realized. However, the user can still estimate the likelihood that their penalty would not exceed  $C_{\max}$ , as follows:

$$q = Pr(p_o(T_a - T_c) \leq C_{\max}) = F_a(T_c + p_o^{-1}(C_{\max})). \quad (1)$$

Thus, we assume that an arriving user will enter the parking lot with probability  $q$ . Let  $\bar{q}$  denote its mean, given by:

$$\bar{q} = \int \int F_a(T_c + p_o^{-1}(C_{\max})) dF_c dF_{\max}. \quad (2)$$

**Actual Parking Duration:** If a user decides to enter the parking area, they park their EV at a park-and-charge spot and leave for their appointment.  $T_a$  is then realized, and if it exceeds  $T_c + p_o^{-1}(C_{\max})$ , we assume that the user reluctantly interrupts their appointment to avoid paying more than  $C_{\max}$  penalty. Thus, the actual parking duration is given by:<sup>2</sup>

$$T_{pc} = \min\{T_c + p_o^{-1}(C_{\max}), T_a\}, \quad (3)$$

and hence, the duration of time the EV overstays is given by:

$$T_o = (T_{pc} - T_c)^+ = \max\{T_{pc} - T_c, 0\}. \quad (4)$$

**Revenue Collected:** The revenue from an EV’s time at the park-and-charge spot (the cost to the EV user) is given by:

$$R = p_c(T_{pc} - T_o) + p_o(T_o). \quad (5)$$

### 2.3 Queueing Model

Since there is a finite number of park-and-charge spots in the parking area, the impact of overstaying depends on the arrival process to the parking lot, e.g., if the inter-arrival times are

<sup>2</sup>If the realized value of  $T_a$  is less than  $T_c$ , we assume that the user leaves immediately after  $T_a$  time units, without waiting for their EV to finish charging.

uniformly distributed with a low arrival rate, a limited amount of overstaying may be harmless, as opposed to a high arrival rate during weekends at a mall. We capture such phenomena using an  $M/G/N/N$  queueing model [Kleinrock, 1975], where the EVs that accept the posted penalty scheme and enter the parking lot are the “jobs”, and the  $N$  park-and-charge spots are the “servers”. The parameters and assumptions are:

- EVs arrive at the parking area according to a Poisson process with rate  $\lambda$ . However, only those EVs whose users accept the posted penalty scheme (with probability  $q$ ), and enter the parking area are counted for further analysis. Using the total probability theorem and the well known “Poisson splitting” property, it can be shown [Biswas *et al.*, 2016] that this “filtered” arrival process is still Poisson, with rate  $\lambda \bar{q}$ .
- An arriving EV that finds all  $N$  park-and-charge spots occupied leaves the system (there is no waiting). Otherwise, the EV stays at a park-and-charge spot for a duration of  $T_{pc}$ ; thus, the service times are i.i.d. according to the cumulative distribution function of  $T_{pc}$ .<sup>3</sup>

If  $N_{pc}$  is the random variable denoting the number of EVs in the parking area, then its steady state distribution is:

$$Prob(N_{pc} = i) = \frac{\rho^i / i!}{\sum_{j=0}^N \rho^j / j!}, \quad 0 \leq i \leq N, \quad (6)$$

where  $\rho = \lambda \bar{q} \mathbb{E}[T_{pc}]$  is the “offered load”. Thus, in steady state, the mean number of EVs in the parking area is:

$$\mathbb{E}[N_{pc}] = \rho \left( 1 - \frac{\rho^N / N!}{\sum_{j=0}^N \rho^j / j!} \right). \quad (7)$$

## 2.4 Performance Measures

We define the following performance measures of interest. Throughout,  $T > 0$  denotes an arbitrary duration of time.

- (a) **Throughput:** The throughput of the park-and-charge system, denoted by  $\tau$ , is the average rate at which EV users leave the park-and-charge area, given by:

$$\tau = \frac{\mathbb{E}[N_{pc}]}{\mathbb{E}[T_{pc}]}. \quad (8)$$

- (b) **Overstay:** The fraction of time spent overstaying at park-and-charge spots, denoted by  $\gamma_o$ , is given by:

$$\gamma_o = \frac{\text{Overstaying Time}}{\text{Total Time}} = \frac{\tau T \mathbb{E}[T_o]}{N T} = \frac{\mathbb{E}[N_{pc}] \mathbb{E}[T_o]}{N \mathbb{E}[T_{pc}]}. \quad (9)$$

- (c) **Utilization:** The utilization, denoted by  $\gamma_u$ , is defined as the fraction of time the park-and-charge spots were used for charging, and is given by:

$$\gamma_u = \frac{\text{Charging Time}}{\text{Total Time}} = \frac{\tau T \mathbb{E}[T_{pc} - T_o]}{N T} = \frac{\mathbb{E}[N_{pc}]}{N} \left( 1 - \frac{\mathbb{E}[T_o]}{\mathbb{E}[T_{pc}]} \right). \quad (10)$$

- (d) **Revenue Rate:** The revenue rate  $\bar{R}$  is defined as the average rate at which revenue is accrued from the park-and-charge spots, and is given by:

$$\bar{R} = \frac{\text{Total Revenue}}{\text{Total Time}} = \frac{\tau T \mathbb{E}[R]}{T} = \frac{\mathbb{E}[N_{pc}] \mathbb{E}[R]}{\mathbb{E}[T_{pc}]}. \quad (11)$$

<sup>3</sup>This could be a general distribution; however, in an  $M/G/N/N$  system, the stationary distribution of its Markov chain (that tracks the number of EVs) depends on the service time distribution only through its mean.

When  $T_c$ ,  $T_a$ ,  $C_{\max}$  are generally distributed, and the price/penalty functions  $p_c(t)$ ,  $p_o(t)$  are general, the above quantities do not admit closed form expressions, and we defer their mathematical derivations to the full version [Biswas *et al.*, 2016]. However, in Section 3, we investigate a simple special case where the price/penalty functions are linear,  $T_c$  and  $T_a$  are Exponentially distributed, and  $C_{\max}$  is a constant, and derive closed form expressions for these performance measures.

**Benchmarking Performance:** In order to measure the effectiveness of our penalty scheme, we benchmark it against a system with *ideal* human behavior, one in which the EV users, on their own, do not overstay, i.e.,  $T_{pc} = \min\{T_c, T_a\}$ . Here,  $\bar{q} = 1$ , and the revenue from (5) is simply  $p_c(T_{pc})$ .

## 3 Example Scenario: Linear Functions, Exponential Distributions

In this section, we consider a simplified scenario for which we derive closed form expressions for the performance measures, which allows better analysis of the effectiveness of imposing simple, linear penalties for overstaying. In particular, we assume:

- **Linearity:** Let  $p_c(t) = \alpha_c t$  and  $p_o(t) = \alpha_o t$ , where  $\alpha_c > 0$  and  $\alpha_o \geq 0$  are the parameters.
- $T_c$  and  $T_a$  are Exponentially distributed with parameters  $\mu_c$  and  $\mu_a$  respectively.
- $C_{\max}$  is not a random variable, but a constant.

Due to limited space, we defer the often nontrivial intermediate steps in evaluating the expressions to the full version [Biswas *et al.*, 2016].

Under the above assumptions, the acceptance probability  $q$  from (1) and its mean  $\bar{q}$  from (2) can be evaluated to obtain  $q = 1 - \beta e^{-\mu_a T_c}$  and  $\bar{q} = 1 - \frac{\beta \mu_c}{\mu_a + \mu_c}$  respectively, where

$$\beta = e^{-\mu_a \frac{C_{\max}}{\alpha_o}}. \quad (12)$$

**Conditional Distribution of  $T_c$ :** While  $T_c$  is Exponentially distributed by assumption, given that (with probability  $q$ ), an EV accepts the posted penalty and enters the parking lot (call this event  $E$ ), the *conditional* random variable  $T_c|E$  need not be Exponential. ( $T_a$  remains unchanged, since  $E$  does not depend on it.) Thus, we first compute the conditional probability density function,  $f_{c|E}$ , using Bayes’s rule, to obtain:

$$f_{c|E}(T_c) = \frac{q f_c(T_c)}{\bar{q}} = \frac{1}{\bar{q}} \mu_c e^{-\mu_c T_c} (1 - \beta e^{-\mu_a T_c}).$$

### 3.1 Distribution and Mean of $T_{pc}$

The complementary cumulative distribution function of  $T_{pc}$ , defined as  $\bar{F}_{pc}(t) = Prob(T_{pc} > t)$ , can be evaluated as:

$$\bar{F}_{pc}(t) = \begin{cases} e^{-\mu_a t}, & t \leq \frac{C_{\max}}{\alpha_o} \\ \frac{e^{-\mu_a t}}{\bar{q}} e^{-\mu_c (t - \frac{C_{\max}}{\alpha_o})} (1 - \frac{\mu_c}{\mu_a + \mu_c} e^{-\mu_a t}), & t > \frac{C_{\max}}{\alpha_o} \end{cases}$$

Thus, the mean, given by  $\int_0^\infty \bar{F}_{pc}(t) dt$ , can be evaluated as:

$$\mathbb{E}[T_{pc}] = \frac{1}{\mu_a} - \frac{\beta}{2\mu_a + \mu_c} \left( \frac{\mu_a + \mu_c}{\mu_a} - \frac{\mu_a}{\mu_a + (1 - \beta)\mu_c} \right). \quad (13)$$

### 3.2 Distribution and Mean of $T_o$

The complementary cumulative distribution function of  $T_o$ , defined as  $\bar{F}_o(t) = \text{Prob}(T_o > t)$ , can be evaluated as:

$$\bar{F}_o(t) = \begin{cases} \frac{e^{-\mu_a t}}{q} \left( \frac{\mu_c}{\mu_a + \mu_c} - \frac{\beta \mu_c}{2\mu_a + \mu_c} \right), & t < \frac{C_{\max}}{\alpha_o} \\ 0, & t \geq \frac{C_{\max}}{\alpha_o} \end{cases}$$

Thus, the mean, given by  $\int_0^\infty \bar{F}_o(t) dt$ , can be evaluated as:

$$\mathbb{E}[T_o] = \frac{1 - \beta}{2\mu_a + \mu_c} \left( \frac{\mu_a + \mu_c}{\mu_a} - \frac{\mu_a}{\mu_a + (1 - \beta)\mu_c} \right). \quad (14)$$

### 3.3 Mean Revenue

When the pricing/penalty functions are linear, the revenue defined in (5) is given by  $R = \alpha_c(T_{pc} - T_o) + \alpha_o T_o$ , and by linearity of expectations, its mean can be evaluated to obtain:

$$\mathbb{E}[R] = \frac{\alpha_c}{2\mu_a + \mu_c} \left( 1 + \frac{\mu_a}{\mu_a + (1 - \beta)\mu_c} \right) + \alpha_o \mathbb{E}[T_o]. \quad (15)$$

### 3.4 Performance Measures

By inspecting the expressions, it can be seen that  $\beta$  is increasing in  $\alpha_o$ . Hence, with increasing penalty,  $\mathbb{E}[T_{pc}]$  increases and  $\mathbb{E}[T_o]$  decreases. The other quantities, including performance measures (8)-(11), are more complicated and their behavior is not obvious from inspection. We omit a detailed analysis with plots due to space constraints. As an example, when we set  $N = 10$  slots,  $\lambda = 8$  EVs per hour,  $\mu_c = \frac{60}{45}$  per hour,  $\mu_a = \frac{60}{105}$  per hour,  $C_{\max} = \$4$ ,  $\alpha_c = \$2$  per hour, we find that setting the penalty rate  $\alpha_o = \$2.37$  per hour provides the maximum value of utilization (30%), whereas the utilization with no penalty is about 26%. (The ideal utilization when there is no overstaying is 42%.) From the point of view of revenue, setting the penalty rate  $\alpha_o = \$3.07$  per hour provides the maximum revenue rate of \$15.36 per hour, but the decreased utilization at this penalty rate is only slightly less, at 29.5%. (The ideal revenue when there is no overstaying is merely \$8.34 per hour.) Thus, the revenue maximizing penalty rate provides increased utilization (by 4%), as well as significantly increased revenue (by 85%).

## 4 Dynamic Environment

In this section, we consider a real world setting where no prior information about the distributions  $F_c$ ,  $F_a$ ,  $F_{\max}$  is known, and thus, the expressions for the performance measures are also unknown to the framework. We assume that these distributions are unknown, but fixed for the day; hence, the utilization and revenue for the day depend only on  $\alpha_o$ .

In our model, on each day, a penalty rate  $\alpha_o$  is declared at the entrance to the parking area. Arriving customers behave according to their  $T_c$ ,  $T_a$ ,  $C_{max}$  values (which are unknown to the system, but known to the customers), and the posted penalty rate  $\alpha_o$ ; their behavior is as described in Section 2.2. The system observes the customer behavior only after they enter the parking area. Thus, the performance measures (such as revenue rate, utilization, etc.) corresponding to penalty rate  $\alpha_o$  are only observed at the end of the day. The task is to find

an optimal penalty rate  $\alpha_o^*$  for each day, such that the total revenue over all days is maximized.<sup>4</sup>

This problem is an example of sequential decision making in an unknown and dynamic environment, where the system seeks to optimize the average revenue over all days, while continuously gathering more information about the revenue obtained by using different penalty rates  $\alpha_o$ . This leads to a trade-off between exploration (imposing each penalty rate  $\alpha_o$  sufficiently often to obtain better estimates of the corresponding revenue accumulated) and exploitation (frequently imposing the optimal penalty for which the observed revenue is maximized). Such problems naturally fall into the category of stochastic multi-armed bandit (MAB) problems [Bubeck and Cesa-Bianchi, 2012]. Each penalty rate  $\alpha_o$  is considered as an arm, and the daily revenue obtained by imposing that penalty rate is analogous to the reward obtained by pulling the corresponding arm. The goal in the MAB setting is to determine the arm to be pulled each time in order to maximize the total reward obtained.

### 4.1 Learning Algorithm UCB-PC

Let  $A$  denote a finite, ordered set of penalty rates, and let  $\mathcal{R}_i$  denote the expected daily revenue corresponding to penalty rate  $A[i]$ , which is unknown. The (unknown) optimal penalty rate is thus given by  $A[i^*]$ , where  $i^* = \arg \max_i \mathcal{R}_i$ . In order to estimate  $\{\mathcal{R}_i\}$ , on each day, we impose a penalty rate and observe the daily revenue, which is the sum of the payments made by the customers who choose to enter the parking area that day.<sup>5</sup> In doing so, we keep track of the observed average daily revenues  $\{\hat{\mathcal{R}}_i\}$ , which are then used to gradually learn the optimal penalty rate  $A[i^*]$  for the parking area. Algorithm 1 is based on the techniques used in the UCB1 algorithm [Auer *et al.*, 2002], which is a well known tool for solving stochastic multi-armed bandit problems. However, UCB1 assumes that the system is allowed to run only for a finite number of trials  $T$  ("finite horizon multi-armed bandit"), whereas we operate under the assumption that the system runs for infinite time, and must therefore tackle the challenges involved in adapting UCB1 to the infinite horizon setting.<sup>6</sup>

---

#### Algorithm 1 UCB-PC

---

- 1: **Input:** A finite set  $A$  of penalty rates
  - 2:  $\text{TR}_i$  stores the total observed revenue for  $A[i]$
  - 3:  $K_i$  stores the number of days  $A[i]$  is imposed
  - 4: **for**  $t \leftarrow 1$  **to**  $|A|$  **do**
  - 5:   Choose penalty rate  $A[t]$  on  $t^{\text{th}}$  day; Observe revenue earned  $r$ ;
  - 6:   Set  $\text{TR}_t \leftarrow r$ ; Set  $K_t \leftarrow 1$ ;
  - 7: Find  $i^* = \arg \max_i \left( \text{TR}_i + \sqrt{2 \ln |A|} \right)$ ;
  - 8: **for**  $t \leftarrow |A| + 1, |A| + 2, \dots$  **do**
  - 9:   Choose penalty rate  $A[i^*]$  on  $t^{\text{th}}$  day; Observe revenue earned  $r$ ;
  - 10:   Update  $\text{TR}_{i^*} \leftarrow \text{TR}_{i^*} + r$ ; Update  $K_{i^*} \leftarrow K_{i^*} + 1$ ;
  - 11:   Update  $\hat{\mathcal{R}}_{i^*} \leftarrow \frac{\text{TR}_{i^*}}{K_{i^*}}$
  - 12: Find  $i^* = \arg \max_i \left( \hat{\mathcal{R}}_i + \sqrt{\frac{2 \ln t}{K_i}} \right)$  for  $(t + 1)^{\text{th}}$  day;
- 

<sup>4</sup>Variants of the technique presented here can be used for other performance measures such as utilization or overstay.

<sup>5</sup>We observe, for each customer, whether they choose to enter the parking area or not, and if they do, their charging and overstaying times, as well as their final payment.

<sup>6</sup>While UCB1 was our choice to illustrate the technique, we believe that our framework can also accommodate modified versions of other stochastic multi-armed bandit algorithms (such as EXP3).

## 4.2 Performance Analysis

In this section, we compare the performance of our learning algorithm with an optimal algorithm that knows the expected revenues  $\{\mathcal{R}_i\}$  beforehand, and can therefore choose the optimal penalty  $A[i^*]$  every day. The loss incurred by the learning algorithm over  $K$  days is termed as the *regret*  $\mathcal{L}_K$ , and can be written as

$$\mathcal{L}_K = K \cdot \mathcal{R}_{i^*} - \left( \sum_{i=1}^{|A|} \mathbb{E}[K_i(K)] \mathcal{R}_i \right) = \sum_{i=1}^{|A|} \mathbb{E}[K_i(K)] (\mathcal{R}_{i^*} - \mathcal{R}_i) \quad (16)$$

where  $\mathbb{E}[K_i(K)]$  denotes the expected number of days (out of  $K$ ) that penalty rate  $A[i]$  is chosen by our learning algorithm. In the literature, the performance of multi-armed bandit algorithms is measured by obtaining an upper bound on this regret. In [Auer *et al.*, 2002], the authors show that the upper bound on the regret using UCB1 is sublinear within a finite, fixed horizon  $T$ ; in particular, they show that the regret is bounded above by a term that is  $\tilde{O}(\log T)$ . Next, in Theorem 1, we show that the regret for UCB-PC is upper bounded by  $\tilde{O}(\log K)$  after the algorithm runs for  $K$  days, for all  $K$ .

**Theorem 1.** *The expected regret of UCB-PC, after the algorithm runs for  $K$  days, is*

$$\sum_{\substack{i=1 \\ i \neq i^*}}^{|A|} \left( \left\lceil \frac{8 \ln K}{(\mathcal{R}_{i^*} - \mathcal{R}_i)^2} \right\rceil + 1 + \frac{\pi^2}{3} \right) (\mathcal{R}_{i^*} - \mathcal{R}_i).$$

*Proof Sketch:* The proof follows from [Auer *et al.*, 2002], where the upper bound for regret is obtained by bounding the expected number of times an arm is pulled. In our case, the upper bound on  $\mathbb{E}[K_i(K)]$  for each  $i \neq i^*$  is given by:

$$\mathbb{E}[K_i(K)] \leq \tau + \sum_{t=\tau}^{K-1} \mathbb{P}(A[i] \text{ is imposed on day } (t+1) \wedge K_i(t) \geq \tau) \quad (17)$$

for any finite positive integer  $\tau$ . Next, in Lemma 1, we show that, for any suboptimal penalty index  $i \neq i^*$ , the term  $\mathbb{P}(A[i] \text{ is imposed on day } (t+1) \wedge K_i(t) \geq \tau)$  is decreasing in  $t$  after sufficient exploration.

**Lemma 1.** *For any  $i \neq i^*$ , if penalty rate  $A[i]$  is chosen for at least  $\left\lceil \frac{8 \ln t}{(\mathcal{R}_{i^*} - \mathcal{R}_i)^2} \right\rceil$  days among the first  $t$  days, then,  $\mathbb{P}(A[i] \text{ is imposed on day } (t+1) \wedge K_i(t) \geq \tau) \leq \frac{2}{t^2}$ .*

*Proof.* We begin by writing,

$$\begin{aligned} & \mathbb{P}(A[i] \text{ is imposed on day } (t+1) \wedge K_i(t) \geq \tau) \\ &= \sum_{s'=1}^t \sum_{s=\tau+1}^t \mathbb{P}(A[i] \text{ is imposed on day } (t+1) \\ & \quad \wedge K_i(t) = s \wedge K_{i^*}(t) = s'). \end{aligned} \quad (18)$$

After running UCB-PC for  $t$  days, a suboptimal penalty index  $i \neq i^*$  is chosen only when  $\hat{\mathcal{R}}_{i^*} + \sqrt{\frac{2 \ln t}{s'}} \leq \hat{\mathcal{R}}_i + \sqrt{\frac{2 \ln t}{s}}$ . Let this event be denoted by  $W$ . Also, let  $U$  denote the event  $(\hat{\mathcal{R}}_{i^*} \leq \mathcal{R}_{i^*} - \sqrt{\frac{2 \ln t}{s'}})$ , and  $V$  denote the event

$(\hat{\mathcal{R}}_i \geq \mathcal{R}_i + \sqrt{\frac{2 \ln t}{s}})$ . Now, the probability of choosing penalty rate  $A[i]$  on the  $(t+1)^{th}$  day can be written as:

$$\mathbb{P}(W) \leq \mathbb{P}(W|U)\mathbb{P}(U) + \mathbb{P}(W|V)\mathbb{P}(V) + \mathbb{P}(W|\bar{U} \wedge \bar{V})\mathbb{P}(\bar{U} \wedge \bar{V}). \quad (19)$$

Using the Chernoff-Hoeffding's inequality, it can be shown that  $\mathbb{P}(U) = \mathbb{P}(\hat{\mathcal{R}}_{i^*} \leq \mathcal{R}_{i^*} - \sqrt{\frac{2 \ln t}{s'}}) \leq e^{-2 \left( \sqrt{\frac{2 \ln t}{s'}} \right)^2 s'} = \frac{1}{t^4}$ , and  $\mathbb{P}(V) = \mathbb{P}(\hat{\mathcal{R}}_i \geq \mathcal{R}_i + \sqrt{\frac{2 \ln t}{s}}) \leq \frac{1}{t^4}$ . Next, we show that  $\mathbb{P}(W|\bar{U} \wedge \bar{V}) = 0$  if  $s \geq \frac{8 \ln t}{(\mathcal{R}_{i^*} - \mathcal{R}_i)^2}$ .

First, we observe that the events  $\bar{U}$  and  $\bar{V}$  can be equivalently written as:

$$\begin{aligned} \bar{U} &\Leftrightarrow \hat{\mathcal{R}}_{i^*} > \mathcal{R}_{i^*} - \sqrt{\frac{2 \ln t}{s'}} \Leftrightarrow \mathcal{R}_{i^*} < \hat{\mathcal{R}}_{i^*} + \sqrt{\frac{2 \ln t}{s'}} \\ \bar{V} &\Leftrightarrow \hat{\mathcal{R}}_i < \mathcal{R}_i + \sqrt{\frac{2 \ln t}{s}} \Leftrightarrow \hat{\mathcal{R}}_i + \sqrt{\frac{2 \ln t}{s}} < \mathcal{R}_i + 2\sqrt{\frac{2 \ln t}{s}} \end{aligned} \quad (20)$$

Then, we show that if (20) holds and  $s \geq \frac{8 \ln t}{(\mathcal{R}_{i^*} - \mathcal{R}_i)^2}$ , then the event  $W$  is a contradiction:

$$\begin{aligned} W &\Leftrightarrow \hat{\mathcal{R}}_{i^*} + \sqrt{\frac{2 \ln t}{s'}} \leq \hat{\mathcal{R}}_i + \sqrt{\frac{2 \ln t}{s}} \Rightarrow \mathcal{R}_{i^*} < \mathcal{R}_i + 2\sqrt{\frac{2 \ln t}{s}} \\ &\Rightarrow \mathcal{R}_{i^*} < \mathcal{R}_i + \sqrt{(\mathcal{R}_{i^*} - \mathcal{R}_i)^2} \Rightarrow \mathcal{R}_{i^*} < \mathcal{R}_{i^*}, \end{aligned}$$

which is a contradiction. Finally, substituting the results of the above evaluations in (19), we get  $\mathbb{P}(W) \leq \frac{2}{t^4}$ , which in turn, can be substituted in (18) to obtain

$$\mathbb{P}(A[i] \text{ imposed on } (t+1)^{th} \text{ day}) \leq \sum_{s'=1}^t \sum_{s=\tau+1}^t \frac{2}{t^4} \leq \frac{2}{t^2}.$$

This completes the proof.  $\square$

Using Lemma 1 and substituting  $\tau = \left\lceil \frac{8 \ln K}{(\mathcal{R}_{i^*} - \mathcal{R}_i)^2} \right\rceil$  in (17), we obtain the required upper bound on regret that is specified by Theorem 1. It can be shown that  $\mathcal{L}_K = \tilde{O}(\ln K)$  for a fixed set of penalty rates. Thus, the average daily regret, given by  $\frac{\mathcal{L}_K}{K}$ , vanishes as  $K \rightarrow \infty$ .

## 5 Experimental Results

For simulating the behavior of electric vehicles we study real-world data for the city of London, obtained from [www.gov.uk, 2013], which consists of 9961 charging events, including usage data of charge points and the duration of stay for EVs. We find it difficult to fit a single parametric model to these data without any restrictions.<sup>7</sup> Therefore, we only look at data entries that correspond to the ‘‘standard’’ charging point type and a specific duration of parking time (between 30 minutes and 180 minutes), in order to fit them in a single parametric model. We assume that the charging duration data corresponds to  $T_c$  and the parking duration data corresponds to  $T_a$ , the duration for which customers would park their car in the absence of any penalties. The histogram for the restricted parking duration data and charging duration data are shown in Figures 2 and 3.

<sup>7</sup>Mixture models can be effective, since they reveal the hidden heterogeneity that arises from a latent categorical variable such as charging point type, parking duration (very short, standard, long), time of arrival (peak hours and non-peak hours), etc.




Figure 2: Histogram of parking duration data




Figure 3: Histogram of charging duration data

The distributions that best fit these histograms (using Mathematica) correspond to a uniform distribution between 30 minutes and 180 minutes for  $T_a$ , and a generalized Gamma distribution with parameters  $\mu = -1.35188$ , scale parameter 33.7831, and shape parameters 1.44212 and 1.19403, for  $T_c$ .

We assume that the penalty threshold  $C_{max}$  is a discrete random variable that takes values \$4, \$8, \$10, \$20 with probabilities 0.4, 0.3, 0.2, 0.1 respectively. We also assume a linear pricing function during charging, with rate  $\alpha_c = \$2$  per hour. We simulate vehicles arriving to a parking area with 10 charging slots for a 6 hour time period, assuming that the arrivals follow a Poisson distribution with rate 10 per hour. We also assume that when all the slots of a parking lot are occupied, arriving users do not wait for a slot to be free and immediately leave the parking area. We perform two sets of experiments, discussed next.

**Experiment 1:** We simulate the parking area for a 6 hour time period per day for 100 days. Then, we observe the utilization and revenue obtained for each day by employing a linear penalty function  $p_o(t) = \alpha_o * t$  with  $\alpha_o \in \{0, 1, 2, 3, 4, 5, 6\}$ . Note that  $\alpha_o = 0$  corresponds to the “no-penalty” scenario. We compare these results against an ideal benchmark where EV users do not overstay at all. Figure 4 shows the results averaged over the 100 days.




Figure 4: Comparing the utilization and total revenue obtained by imposing penalty, no penalty and when no customer overstays

**Utilization:** The utilization of the parking area increases (dark blue line) after imposing a small penalty as compared to when no penalty (red dotted line) is charged. However, when a high penalty is imposed, most of the EV users decide not to enter the parking area, resulting in a steep drop in utilization. With a penalty rate \$4 per hour, the utilization is maximum, and also very close to the ideal utilization when there is no over-staying by the EV users (green line). Thus, imposing just a small penalty helps improve the utilization of charging spots in the parking area significantly.

**Revenue:** The daily revenue from the parking area quickly increases even after imposing a small penalty for overstaying. The daily revenue obtained in the ideal situation of no overstaying is more than when no penalty is imposed (albeit only slightly so), since more customers are served in the former situation than the latter. As the penalty rate increases, the daily revenue follows the same pattern as the utilization, and for the same reasons. It should be noted that the penalty rate \$4 per hour also maximizes the daily revenue.

**Experiment 2:** We simulate the parking area for two different settings: (a) the optimal penalty rate  $\alpha_o^*$  is chosen on all days, where  $\alpha_o^*$  is calculated by approximately maximizing the total expected revenue with complete knowledge of all the underlying distributions that are used, and (b) our learning algorithm UCB-PC determines the best  $\alpha_o$  for each day, based on the observed parameters for all previous days. We compare the convergence of the average revenue obtained by our learning method to the average revenue obtained by the (approximately) optimal method. Figure 5 shows the results.




Figure 5: Comparing the performance of UCB-PC with respect to the optimal algorithm

We observe that the penalty rates chosen by our learning algorithm UCB-PC are almost equivalent to that of the optimal algorithm after 15 days in terms of the daily revenue obtained. We also show the difference between the total revenue by the optimal algorithm and UCB-PC averaged over the number of days, along with the theoretical upper bound given in Theorem 1. As expected, it can be seen that the average difference, that is,  $\frac{C_k}{k}$ , decreases with increase in the number of days after sufficient exploration (15 days).

## 6 Concluding Remarks

In this paper, we undertake a formal study, for the first time, of the problem of overstaying EVs in park-and-charge spots. We establish a novel framework in which we bring together an interdisciplinary mix of models and techniques: probabilistic user behaviour, queueing dynamics, online learning. This framework can be extended to accommodate different user behaviour models, queueing dynamics, and other learning techniques. One can imagine “instantiating” it for multiple parking areas in a city with a model for the population of EV users and study the interaction at a higher level, e.g., competition between parking lots. When viewed as a comprehensive model for park-and-charge, our framework could serve as a useful tool for future research.

## References

- [Agrawal, 1995] Rajeev Agrawal. Sample mean based index policies with  $O(\log n)$  regret for the multi-armed bandit problem. *Advances in Applied Probability*, pages 1054–1078, 1995.
- [Anglin *et al.*, 2013] Howard N. Anglin, Irgelkha Mejia, Nicholas J. Ruegger, Young, and Yvonne M. Recharging of battery electric vehicles on a smart electrical grid system. *United States Patent Application US 20130173326 A1*, 2013.
- [Auer *et al.*, 2002] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- [Biswas *et al.*, 2016] Arpita Biswas, Ragavendran Gopalakrishnan, and Partha Dutta. Managing overstaying electric vehicles in park-and-charge facilities, 2016. arXiv preprint arXiv:1604.05471.
- [Blink, 2012] Blink. The Blink Courtesy Notice, 2012. <http://www.blinknetwork.com/file/7741/Blink+Courtesy+Notices.pdf>.
- [Blink, 2015] Blink. The Blink Courtesy Notice, 2015. [http://www.blinknetwork.com/file/18760/Blink\\_Courtesy\\_Notices.pdf](http://www.blinknetwork.com/file/18760/Blink_Courtesy_Notices.pdf).
- [Bubeck and Cesa-Bianchi, 2012] Sébastien Bubeck and Nicolo Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems, 2012. arXiv preprint arXiv:1204.5721.
- [Cesa-Bianchi, 2006] Nicolo Cesa-Bianchi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- [Gadh *et al.*, 2012] Rajit Gadh *et al.* Smart electric vehicle (ev) charging and grid integration apparatus and methods. *United States Patent US 9026347 B2*, 2012.
- [Gopalakrishnan *et al.*, 2016] Ragavendran Gopalakrishnan, Arpita Biswas, Alefiya Lightwala, Skanda Vasudevan, Partha Dutta, and Abhishek Tripathi. Demand prediction and placement optimization for electric vehicle charging stations, 2016. arXiv preprint arXiv:1604.05472.
- [Hafner *et al.*, 2009] James Lee Hafner, Melissa Wiltsey O’mara, and Paul Stuart Williamson. Managing incentives for electric vehicle charging transactions. *United States Patent Application 20090313104*, 2009.
- [Kleinrock, 1975] Leonard Kleinrock, editor. *Queueing Systems Volume 1: Theory*. Wiley–Interscience, 1975.
- [Loveday, 2013] Eric Loveday. Electric Vehicle Parking Law Coming Under Fire for Not Being Enforced in Hawaii, 2013. <http://insideevs.com/electric-vehicle-parking-law-coming-under-fire-for-not-being-enforced-in-hawaii/>.
- [Moloughney, 2015] Tom Moloughney. Featured EV Product: The JLong, 2015. <http://bmw13.blogspot.co.uk/2015/03/featured-ev-product-jlong.html>.
- [Richtel, 2015] Matt Richtel. In California, Electric Cars Outpace Plugs, and Sparks Fly, 2015. <http://www.nytimes.com/2015/10/11/science/in-california-electric-cars-outpace-plugs-and-sparks-fly.html>.
- [Rowe and Fiorucci, 2011] Rick Rowe and Jean-louis Fiorucci. Parking locator system providing variably priced parking fees. *United States Patent 8688509*, 2011.
- [Tumblr, 2015] Tumblr. ICE Shaming, 2015. <http://iceshaming.tumblr.com/>.
- [Tushar *et al.*, 2012] W. Tushar, W. Saad, H.V. Poor, and D.B. Smith. Economics of electric vehicle charging: A game theoretic approach. *Smart Grid, IEEE Transactions on*, 3(4):1767–1778, 2012.
- [www.gov.uk, 2013] www.gov.uk. Plugged in places data, 2013. [https://www.gov.uk/government/uploads/system/uploads/attachment\\_data/file/236776/plugged-in-places.csv/preview1](https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/236776/plugged-in-places.csv/preview1).
- [Zapatag, 2015] Zapatag. Zapatag, 2015. <http://www.zapatag.com/>.
- [Zoeter *et al.*, 2014] Onno Zoeter, Christopher Dance, Stéphane Clinchant, and Jean-Marc Andreoli. New algorithms for parking demand management and a city-scale deployment. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD, 2014.