Revenue Maximization for Electric Vehicle Charging Service Providers Using Sequential Dynamic Pricing

Jan Mrkos
Artificial Intelligence Center, Faculty of Electrical Engineering, Czech Technical University in Prague
Prague, Czech Republic
mrkos@agents.fel.cvut.cz

Antonín Komenda
Artificial Intelligence Center, Faculty of Electrical Engineering, Czech Technical University in Prague
Prague, Czech Republic
komenda@agents.fel.cvut.cz

Michal Jakob
Artificial Intelligence Center, Faculty of Electrical Engineering, Czech Technical University in Prague
Prague, Czech Republic
jakob@agents.fel.cvut.cz

ABSTRACT

With the increasing prevalence of electric vehicles (EVs), the provision of EV charging is becoming a standard commercial service. With this shift, EV charging service providers are looking for ways to make their business more profitable. Dynamic pricing is a proven technique to increase revenue in markets with time-variant, heterogeneous demand. In this paper, we propose a Markov Decision Process (MDP)-based approach to revenue-maximizing dynamic pricing for charging service providers. We implement the approach using an ensemble of policy iteration MDP solvers and evaluate it using a simulation based on real-world data. We show that our proposed method achieves significantly higher revenue than methods utilizing flat-based pricing. In addition to achieving higher revenue for charging service providers, the method also increases the efficiency of allocation measured in terms of the total utilization of the charging station.

KEYWORDS

Dynamic Pricing, Markov Decision Process, Innovative agents and multiagent applications, EV charging, Electric vehicle, Resource Allocation

ACM Reference Format:
Jan Mrkos, Antonín Komenda, and Michal Jakob. 2018. Revenue Maximization for Electric Vehicle Charging Service Providers Using Sequential Dynamic Pricing. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 9 pages.

1 INTRODUCTION

After more than hundred years of niche use, electric vehicles (EVs) seem on the cusp of displacing internal combustion engine (ICE) vehicles in personal transportation. Better fuel efficiency, environmental friendliness and lowering cost give EVs an edge over ICE vehicles. In Europe, EVs have been forecast to match ICE cars in the total cost of ownership as early as 2018 and by 2023 account for 30% of new car sales[1].

The rise of EVs drives interest from many different actors, including governments, cities, car manufacturers, environmental groups and electric utilities. Each is trying to prepare for the expected rise of EVs. For cities and electric utilities, the widespread use of EVs may require large investments into infrastructure because large numbers of EVs could increase the peak load on the grid up to threefold[2]. However, methods of load balancing – moving of the load "peaks" into the load "valleys" – can prevent the infrastructure costs from growing.

This concern for the future infrastructure investment is one of the primary motivations in the recent interest in dynamic pricing. Dynamic pricing [1, 17] is a technique to balance load in various domains. It is studied in economics, revenue management or supply chain management. In the field of smart mobility where the system cannot be controlled centrally, dynamic pricing was proposed to improve efficiency of taxi systems [8, 9] or power grid management for electromobility [5, 11, 20], balancing power load, power quality and other grid-related metrics. These fields recognize dynamic pricing as a critical lever for influencing the behavior of buyers.

Until recently, most of the research of charging for electromobility focused on the optimization of charging station placement [7, 12, 19, 21, 22]. Such approaches are only a seeming remedy in a changing environment where charging station placement is no longer optimal in the new environment. On the other hand, dynamic pricing of the charging services and its application to load balancing is robust to the dynamically changing situation in the infrastructure, demand and energy costs. This direction was recently taken by Xiong et al. [23]. The pricing problem proposed considered EV drivers' travel patterns and self-interested charging behavior. The problem can be seen as a variation on sequential posted pricing [3] for charging stations. The solution the authors proposed uses mixed integer non-convex optimization of social welfare in the model.

Dynamic pricing of charging services is a method that can potentially provide cheap and robust alternative to expensive upgrades of our current infrastructure. However, applications proposed above focus on the dynamic pricing primarily toward optimization of social welfare. Yet in real world situations, prospective charging station providers are often privately owned and as such not strongly incentivized to improve the social welfare.

Instead, private investors are concerned with the costs of installing and providing charging services and their financial returns[3].

[1]UBS report from 2017: Q-Series UBS Evidence Lab Electric Car Teardown – Disruption Ahead? Avaialable at http://www.iom3.org/sites/default/files/news-documents/Automotive_Materials_EV_UBS_May_2017.pdf

[2]Based on the IEEE Spectrum article: http://spectrum.ieee.org/transportation/advanced-cars/speed-bumps-ahead-for-electricvehicle-charging

[3]From report "An Industry Study on Electric Vehicle Adoption in Hong Kong" by the Hong Kong Productivity Council (2014): www.hkpc.org/images/stories/corp_info/hkpc_pub/evstudyreport.pdf

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

Modeling charging station providers as self-interested agents instead of cooperative group seems to be closer to reality in many places around the world.

For this reason, in this work we focus on maximizing the revenue of charging station provider and show that this can be done while still improving social welfare. To this end, we propose MDP based dynamic pricing strategy for a self-interested charging service provider that aims to maximize its revenue while improving utilization of the limited available grid resources.

2 RELATED WORK

To put dynamic pricing into perspective, we can see it as pricing of services that are in high demand or that each buyer values differently, such as hotel rooms [10] or airline tickets [18]. For airfares and hotel rooms, the price is changing based on the expected demand throughout the season, existing bookings, as well as the customers segment (business or tourist). Services such as airfares and hotel rooms have strict expiration deadline, that is, the departure of the airplane and the arrival of given booking day. Similarly, the EV charging resources in a given time window expire if there are no vehicles to use them.

With such a type of expiring services, the goal is to sell the available service capacity for a profit under the constraints given by their expiration and fluctuations in the demand. Unused service capacity is a wasted profit opportunity. Maximizing revenue from these expiring services is the topic of revenue management [17].

Literature on the revenue management including dynamic pricing of hotel rooms and airline tickets is classically not primarily concerned with competition of the service providers nor with competition among customers [4, 14]. The primary concern is the profitable utilization of the sold resource (rooms, tickets). In case of the pricing of EV charging resources, there exist approaches that model competition between customers [6]. However, competition between different charging stations have not been investigated yet. In our wiev, this is because the deployment of the EVs is still limited and there is no data to validate competitive aspects of dynamic pricing of charging services. While competition between the charging stations might become important in the future, we focus on the problem of optimal pricing in an environment where other charging stations are not having an important effect. Thus, we adopt similar approach to the early work on the airline revenue management that disregards competitive aspect of pricing and focuses on ther factors deemed more important at the time.

The common approach to the pricing of airline tickets, extensively studied e.g., in [4], is to determine the number of tickets which can be sold for a particular price in order to maximize the expected revenue. These decisions require the modeling of customers, which is based on parameters such as price sensitivity, seasonality of demand and others. The optimal rule for accepting or rejecting air ticket bookings is as follows: *Is the profit from this booking higher than the expected profit from the booked seat that we could get later? Confirm the booking now if the profit is higher than the expected profit later. If it is lower, reject this booking.* Moreover each *accepted* booking can influence following bookings as customers booking later are potentially not able to book seats at the same price. Additionally, if we include connecting flights and group bookings,

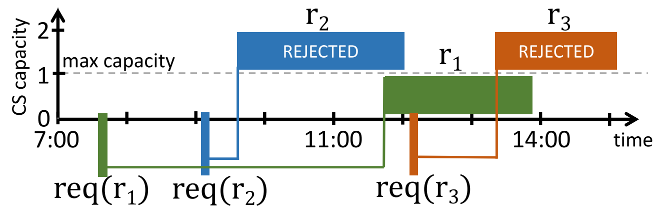

Figure 1: The requests for charging in the future arrive to the charging station in a sequence. Accepted charging requests and their duration have an effect on whether request that arrive later can be accomodated. If the charging station accepts requests as they arrive, r_1 arriving at $req(r_1)$ will block r_2 and r_3. If the charging station rejected r_1 instead, r_2 and r_3 could be accepted. Whether to accept or reject a charging request can be decided by comparing the value of the charging request with the expected value of not yet allocated available capacity.

pricing decisions can affect and be affected by other subsequent pricing processes. Because of this snowball effect, the complexity of the air ticket pricing problem (and similarly the charging service pricing problem) is in general intractable [17]. For this reason, the problem is often simplified (e.g., by assuming independence of demand between booking classes, allowing no cancellations, etc.) so that a solution can be found computationally.

An important distinction between airline pricing and charging station pricing is coupling of the bookings in the case of charging services, which is not present when selling airline tickets. In the case of airline tickets, it is not particularly important which seat (in a given class) was sold as the booking of a single seat does not block the booking of surrounding seats. However, booking of a single time window can influence time windows before and after it, as it can block other charging opportunities.

In the following sections, we approach the charging station pricing strategy in a way similar as revenue management deals with pricing of airline tickets. Both problems focus on expiring services where earlier transactions affect transactions that follow.

3 MODEL OF CHARGING IN THE ELECTROMOBILITY ENVIRONMENT

Our model describes a multi-agent system with a single charging station operated by one charging service provider and several EV drivers wanting to charge their vehicles. Within such system, we focus on the problem of dynamic pricing strategy [1]. The strategy is used by the single charging service provider to post prices, which are accepted or rejected by EV drivers. The goal of the pricing strategy described below is maximization of charging service provider revenue within a particular time horizon. As discussed earlier, we disregard the competition between charging service providers in our pricing model and we leave it for future work.

In our multi-agent model of charging in electromobility, we consider EV drivers as *n agents* and one charging station provider as one additional agent. The model is formally defined as a tuple $M = \langle D, c, \phi_1, \ldots, \phi_n, \Phi \rangle$. $D = (r_1, \ldots, r_m)$ is a demand expressed as a sequence of *charging services requests* r_i sent by the EV driver

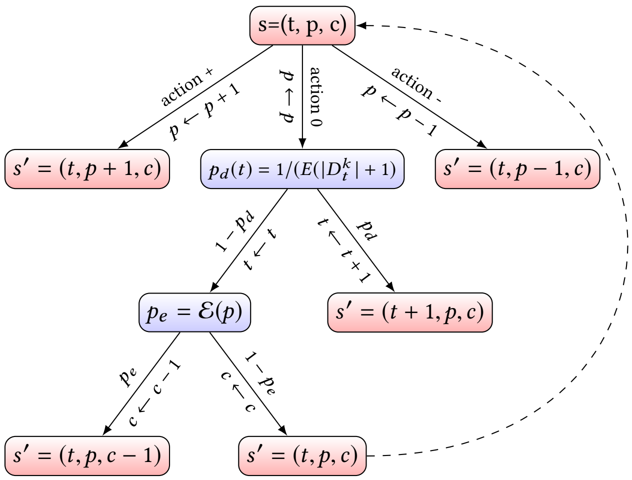

Figure 2: The structure of the transition function P. Given state s, the probability of getting to the next state s' is given by multiplying the probabilities along edges. States are the decision nodes (in red), chance states are in blue and contain the definition of the probability used on the edge.

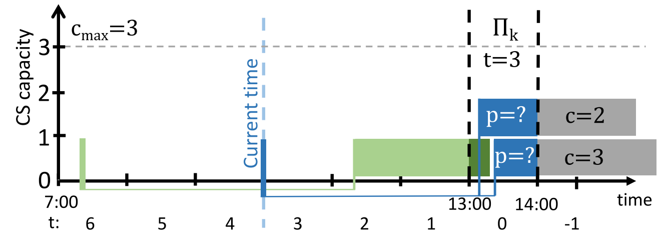

Figure 3: The figure shows state of MDP Π_k for k-th time interval (13:00 to 14:00) of a charging station with maximal capacity $c_{\max} = 3$ and expected demand for charginging in k-th $E(|D_t^k|)$ (curve in red). State of Π_k is given by $t = 3$, the *number of time intervals to charging* from current time, *current price* p and *free capacity* c that depends on the real-time overlap of the accepted charging requests (leftmost rectangle) with the incoming charging request (two rectangles with text show two possible incoming requests).

agents to the charging service provider agent in a sequence. Each request is defined by its $start(r_i)$ and $end(r_i)$ *times* of charging, time $req(r_i)$ when the request was issued and $char(r_i)$, the requested charge in kWh. We assume that $req(r_i) \leq start(r_i) < end(r_i)$ and $req(r_i) < req(r_{i+1})$. For each time we define *free capacity* $c : \mathbb{R} \rightarrow \mathbb{R}_0^+$ that determines the maximal charging capacity of the charging station in time.

The interaction of EV driver agent and the charging service provider agent is modeled via requests r_i from EV driver agent ϕ_j and prices p_i returned by the charging service provider agent. Formally, $p_i = \Phi(r_i)$, where Φ is the *charging service pricing function* of the charging service provider agent, r_i is the charging service request and p_i is the resulting proposed price. If the charging service provider agent can not accept a given request, it sets $p_i = \infty$.

$\phi_1, \ldots, \phi_n$ denote the *decision processes* of the EV driver agents that determine whether the proposed price p_i is accepted by the driver agent. We write $\phi_j(p) = \top$ iff the j-th EV driver agent is willing to pay price p for the requested charging service r_i and $\phi_j(p) = \bot$ otherwise ($\phi_j(\infty) = \bot$ always). Provided that $\phi_j(\Phi(r_i)) = \top$, the EV driver agent accepts the proposed price for charging request r_i and we assume both the charging and payment eventually happen in the system. An execution of the model M is a sequence of prices and decisions $A = (\langle p_1, \delta_1\rangle, \ldots, \langle p_n, \delta_n\rangle)$ for all agents and their charging requests $r_i \in D$, such that $p_i = \Phi(r_i)$ and $\delta_i = \phi_j(p_i)$.

The goal of the charging service provider is to maximize its revenue by optimally setting prices with Φ. Given that the price p_i of reservation r_i, the *revenue* ρ at the end of the time horizon can be written as the sum of prices across all realized reservations (Equation 1):

$$\rho(D, c, \phi_1, \ldots, \phi_n, \Phi) = \sum_{r_i \in D} \Phi(r_i) \mathbb{1}_{\phi_j(\Phi(r_i))=\top} \tag{1}$$

We are looking for revenue maximizing pricing function Φ^*:

$$\Phi^* = \arg\max_{\Phi} \rho(D, c, \phi_1, \ldots, \phi_n, \Phi) \tag{2}$$

The maximization is constrained by the free capacity:

$$\forall t \in \mathbb{R},\ c(t) \geq \sum_{r_i \in D} \mathbb{1}_{\phi_j(\Phi(r_i))=\top} \mathbb{1}_{start(r_i) \leq t < end(r_i)} \tag{3}$$

Here, $\mathbb{1}_{\phi_j(\Phi(r_i))=\top}$ is the indicator function that equals 1 for requests accepted by the EV driver agent and $\mathbb{1}_{start(r_i) \leq t < end(r_i)}$ is the indicator function of the charging interval of request r_i.

4 MDP-BASED DYNAMIC PRICING OF CHARGING SERVICES

The maximization problem given by Equation 2 is sequential in nature. The pricing function has to respond to the confirmed reservations as it can not exceed the capacity of the charging station (Equation 3). Additionally, the charging station provider agent will generally not have exact knowledge of the individual ϕ_j or of the number of requests to be made in one day. However, the charging station can have a probabilistic model of the EV driver behavior and of expected demand. Thus, from the perspective of the charging station, revenue maximization is a Markov Decision Process (MDP) [2].

We aggregate the charging station provider's understanding of the EV drivers decision processes $\phi_1, \ldots, \phi_n$ into the price elasticity function $\mathcal{E}$. Given a request r and a generated price p, price elasticity $\mathcal{E}(p)$ is the probability of an EV driver agent accepting the price p. Arrivals of requests are modeled using two random variables, $E(|D|)$, the number of requests arriving in some time interval, and probability distribution over $req(r_i)$, $start(r_i)$, $end(r_i)$, the parameters of the request r_i.

Using these probabilities, we reformulate the maximization goal from (2) as the maximization of the expected revenue across possible decision policies π:

$$\pi^* = \arg\max_{\pi} E(\rho(D, c, \phi_1, \ldots, \phi_n, \pi)) \tag{4}$$

The expectation is with respect to the joined probability distribution of arrivals and probability of acceptance and the maximization is subject to (3).

Based on the price elasticity function, we define the MDP pricing strategy that determines prices offered by the charging station provider throughout the day. We discretize time into n time intervals, price into n_p price levels and free capacity into n_c capacity levels. Since we assume in our model that all charging sessions use the same electrical power, the maximal *free capacity* in the time interval is the upper limit on the number of concurrent charging sessions. Free capacity incorporates all charging station provider constraints, including the power grid capacity or the number of available charging connectors, into one number.

4.1 MDP-based Pricing in Discrete Time Intervals

In each charging time window, there are n_p price and n_c capacity levels and up to n time windows prior to the charging. As such, the state space in each time window has size nn_pn_c and the branching factor is up to 5 (see Figure 2). Because price or capacity change in any time window can have an effect on any other time window, finding solution to (4) means finding solution for n time windows together; that is in a state space with $n(n_pn_c)^n$ states and branching factor at least 5^n.

$$\pi_k^* = \arg\max_{\pi_k} E(\rho(D, c, \phi_1, \ldots, \phi_n, \pi_k)),\ k \in 1, \ldots, l \tag{5}$$

To avoid this combinatorial explosion, we find optimal pricing policy for each time window independently through n independent MDPs that maximize expected revenue in each time interval (5). However, this means that the new problem is no longer optimal in the sense of (4) but only as a set of n independent solutions to (5). Combined together, summed maximums from (5) are bound above by maximum from (4). The tightness of this bound is data dependent. As such, we do not attempt to determine the tightness of this bound with an experiment using smaller, but full MDP as this would require us to oversimplify the data.

The pricing strategy uses one MDP for each time interval. The solution to MDP for the k-th interval is the pricing policy π_k.

The MDP is a tuple $\Pi_k = \langle S, A, R, P, s_0, S_e \rangle,\ k \in 1, \ldots, l$. In each Π_k, S is a finite set of *states*, A is a finite set of actions; $P : S \times A \times S \longrightarrow [0, 1]$ is the transition function forming the transition model giving a probability $P(s'|s, a)$ of getting to state s' from state s after action a; and a reward function $R : S \times A \times S \longrightarrow \mathbb{R}$.

Starting in the initial state s_0, any action from A can be chosen. Based on this action, the system develops as prescribed by P to the next state where another action can be applied. During the move, the reward can be received based on the $R(s, a, s')$ function.

A state s is defined by triplet (t, p, c). Here $t \in \{0, \ldots, t_{\max}\} \cup \{-1\}$ denotes *the number of time intervals* to the execution of charging (0 denotes the hour of charging, $t = -1$ marks the exit states S_e of the MDP). $c \in \{0, 1, \ldots, c_{\max}\}$ is the *current available capacity* in the time interval, i.e. how many more requests can be accommodated in the time interval, and $p \in \{0, 1, \ldots, p_{\max}\}$ is the current price level set in the time interval. The set of actions A contains three actions, price +1, price −1 and no change to the price.

Because each accepted request reduces capacity in the time windows by one, the reward function R generates reward p for any transition between states $s = (t, p, c)$ and $s' = (t, p, c - 1)$ for all $t \in \{0, \ldots, t_{\max}\}$, p and c in their domains.

The transition function P is based on the price elasticity function $\mathcal{E}(p)$ and the expected number of requests $E(|D_t^k|)$ for the charging in k-th interval t intervals prior execution of charging. Components of the transition function are given in (6) and (7). The way to combine these components into a transition function is shown in Figure 2.

$$p_d(t) = \frac{1}{E(|D_t^k|) + 1} \tag{6}$$

$$p_e(p) = \mathcal{E}(p) \tag{7}$$

The probability $p_d(t)$ of no more charging requests arriving in state (t, p, c) is calculated from the expected absolute demand for charging in k-th time interval t intervals prior to the start of charging $E(|D_t^k|)$ (red curve in Figure 3). Equation (6) is obtained by modeling $|D_t^k|$ as having geometric distribution, arrival of charging request as Bernoulli trial (failure in Bernoulli trial meaning no more charging request t intervals before charging) with failure probability $1 - p_d(t)$.

The geometric distribution is the probability distribution of the number Y of failures before first success in consecutive Bernoulli trials with success probability q (this means 0 is in the support). Y has expected value $\frac{1-q}{q}$. Identifying Y with $|D_t^k|$, we get $q = p_d(t)$. For each t, the MDP remains in some state $s \in \{(t, p, c)|c \in \{0, 1, \ldots, c_{\max}\}, p \in \{0, 1, \ldots, p_{\max}\}\}$ until successful Bernoulli trial with $p_d(t)$ being the probability of success.

Arriving charging request has probability of acceptance $\mathcal{E}(p)$. If the price offered by the pricing strategy is rejected by the customer, the MDP remains in the same state. This is illustrated in Figure 2 in the branch ending in $s' = (t, p, c)$.

There are few exceptions to the probabilities defined by Figure 2. One of them consists of the bounds of the domains of state variables. If the price is maximal resp. minimal, the action to increase resp. lower the price is not available. Similarly, when capacity is 0, $p_e(p) = 0$ as no additional request can be accommodated by the charging station provider in the given time interval. Finally, $t = -1$ denotes the set of exit states $S_e = \{(t, p, c) \in S|t = -1\}$ where the MDP terminates.

The pricing strategy uses n MDPs at once. For a charging request r, the free capacity c_r^k in each MDP is calculated from the overlap (in continous, not discretized, time) of accepted requests (green rectangle in Figure 3) and incoming request r (blue rectangles in Figure 3). In every *time window* I_k corresponding to Π_k, starting from the state (t^k, p^k, c_r^k), we apply pricing policy π_k repeatedly until the policy does not suggest change in price. The final price of the request r offered to the EV driver agent is the sum of the prices in time windows the request overlaps. In the case of partial overlap with some time window, the price is proportional to the size of the overlap:

$$\Phi(r) = \sum_{k=1}^{n} \frac{|I_k \cap (start(r), end(r))|}{|I_k|} \Pi_k(r) \tag{8}$$

Table 1: Summary statistics of the real-world charging data for the selected charging station with three charging points. The dataset contains charging sessions recorded over several weeks.

CS Dataset Statistics	Mean	Std
Charging session duration	0.726 h	0.794 h
Charge per charging session	6.72 kWh	5.19 kWh

Note that k-th MDP internally assumes that EV driver is making decision based on the price p^k (the state variable in Π_k). Depending on the length of the charging request, this price is only part of the price $\Phi(r)$ offered to the driver by the pricing strategy. Thus, the length of the time window should be chosen similar to the length of the average charging request.

5 EVALUATION OF THE MDP-BASED DYNAMIC PRICING METHOD

We evaluate the MDP dynamic pricing algorithm on real-world data. First, we provide summarizing statistics of the dataset and describe the preprocessing we performed on the data. Then we describe the experiments we conducted with the data and the results we obtained.

5.1 Real-world EV Charging Data

The dataset contains information on charging sessions realized at one of the E-WALD charging stations. E-WALD[4] is one of the biggest EV charging service providers in Germany. The E-WALD data includes timestamps of the beginning and the end of each charging session, the status of the electricity meter at the beginning and the end of each charging session and the anonymized identifier of a user who activated the charging session. In the preprocessing step, we remove clearly erroneous data points (such as charging sessions with negative duration) and merge some short charging sessions with following charging sessions if the same customer initiated both sessions.

The summary statistics of the dataset can be found in Table 1. Histogram of charging session start times can be seen in Figure 4. The dataset was collected at the charging station over the period of several weeks. In this period, the charging station averaged 2.53 charging sessions per day. With such low demand for charging, there were almost no conflicts in requested charging sessions. Thus, in our experiments, we randomly sample the dataset to generate single days with up to 60 daily charging sessions.

The particular charging station dataset does not contain any pricing information about the charging sessions for its three charging locations. However, E-WALD (similarly to the majority of other charging service providers) uses only flat rate pricing in all their charging stations.

[4]We would like to thank E-WALD (https://e-wald.eu/) for providing us with the charging data for this study.

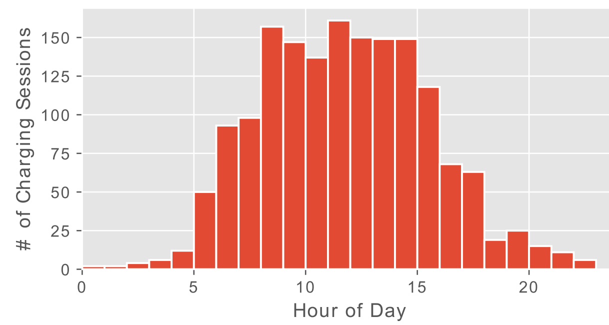

Figure 4: Histogram of the charging session start times in the real-world dataset for the selected charging location.

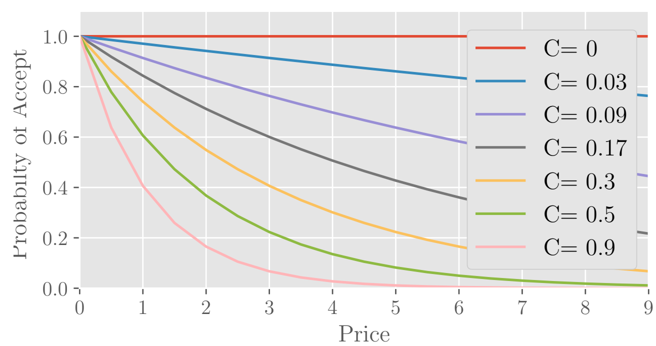

Figure 5: Price elasticity $\mathcal{E}(p)$ of demand curves for different values of the C parameter.

5.2 MDP-based Pricing Strategy Implementation

For the implementation of the MDP pricing strategy, we discretize single day into 24 time intervals, each 1 hour long. As the real-world data was collected at a charging station with three charging slots, in our experiments we consider our station to have three charging points. That is, we use $c_{max} = 3$. This means that at most three charging sessions can be realized at any point in time.

The dataset contains information only about realized charging sessions[5]. In our electromobility model, EV drivers can book charging sessions ahead of time. We model this by setting the request time $req(r)$ for each charging session r in the dataset randomly, with the request time drawn uniformly between 0 and 6 hours ahead of $start(r)$.

Values of $E(|D_t^k|)$ are estimated from the dataset. For each time interval k in the discretization of time, we calculate normalized histogram h_t^k of request times $req(r^k)$ of request r^k for which the charging interval $(start(r^k), end(r^k))$ is in k-th time interval. Bins of the histogram are the intervals of the time discretization and the normalization is done with the size of the used dataset. Given m, the *absolute number of requests in a day*, we set $E(|D_t^k|) = m \cdot h_t^k$.

Recall that $\mathcal{E}(p)$ is the probability of the EV driver agent accepting price p. Because we do not know the real price elasticity of

[5] E-WALD as most of the existing charging service providers does not yet allow booking of charging services ahead of time

Table 2: Description of the evaluation metrics.

Metric	Description
CS Revenue	Revenue of the charging station is the sum of prices of all charging sessions. Revenue is directly dependent on the selected pricing scheme Φ as given by Equation 1.
CS Utilization	Measured in hours. It is the added duration of all charging sessions realized by the charging station. This is a proxy of the social welfare of the EV drivers achieved through various pricing strategies. The higher the utilization, the more of the EV driver charging demand was satisfied by the charging station. Definition of CS utilization μ is given by Equation 9. $$\mu(D, c, \phi_1, \dots, \phi_n, \Phi) = \sum_{r_i \in D} (end(r_i) - (start(r_i))\mathbb{1}_{\phi_i(\Phi(r_i))=\top} \quad (9)$$

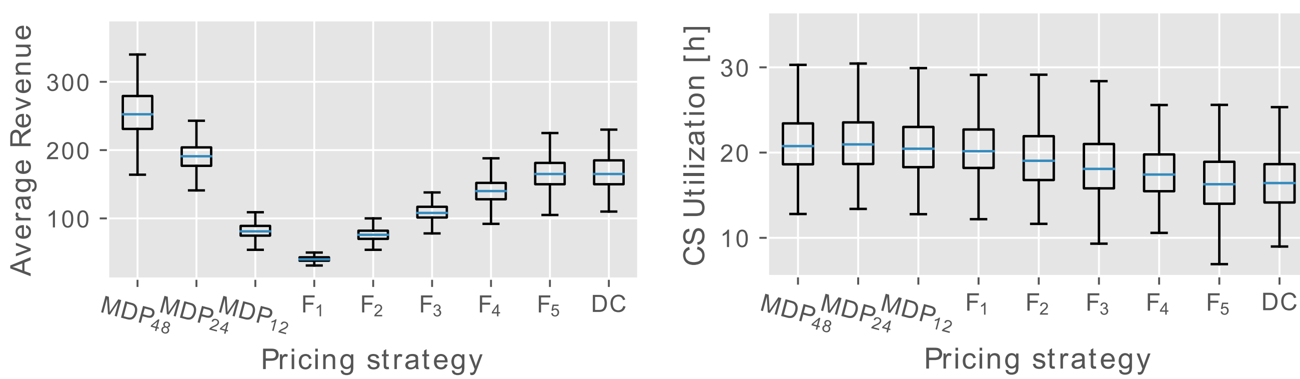

Figure 6: Experimental results using the evaluation metrics. Each boxplot is a result of 400 independent runs with $m = 40$ and $C = 0.03$. MDP lower index denotes number of time windows used for one day in the MDP dynamic pricing strategy. DC denotes demand correlated strategy and F_1 to F_5 the different flat rate strategies.

demand for EV charging services and we cannot estimate it from available data, we define the *price elasticity* function $\mathcal{E}(p)$ parametrically as $\mathcal{E}(p) = e^{-Cp}$. This is one of the standard demand functions in transport economic textbooks [16]. The different values of C and the corresponding shapes of price-elasticity curves are shown in Figure 5. The evaluation is done for multiple values of C. For $C = 0$, we talk about inelastic demand as customers will accept any price. At $C = 0.5$ the demand is highly elastic as small changes to the price have a big effect on customer's acceptance or refusal of the offer. For comparison, the price elasticity of demand for gas station services is usually described as relatively inelastic, meaning low values of C [15]. When data on customer price elasticity becomes available, parametric function $\mathcal{E}(p)$ can be replaced by a model learned from data.

In each experiment, we use the same price elasticity function for all drivers. In reality each person will respond differently to changing price. However, using randomly selected C for each user does not give us different results when compared to using single C for all users. Using one C for all users means we are aggregating behavior of the population. Thus, for simplicity, we use one value of C for all users in the experiments described in Section 5.3.

We find the optimal policies for Π_k through the policy iteration algorithm [13].

5.3 Experiments and Results

In our experiments, we compare the performance of the MDP based dynamic pricing strategy to the flat rate strategies and demand-correlated pricing strategy. The flat rate pricing strategy calculates the charging request price in the same way as the MDP based dynamic pricing strategy, except it uses one price in all time windows (see (8)). The demand correlated strategy sets price based on the expected demand for each time interval. For time windows with smallest demand, it sets lowest price while for time windows with highest demand it sets the highest price.

In most experiments, we use $p_{max} = 5$ and five price levels, 1 to 5. The choice of the maximum price value is tied to the choice of price elasticity parameters. For given choice of price levels and given average length of charging session, we chose price elasticity parameters that generate price elasticity curves of varied shapes over the domain of probable charging session prices (see Figure 5). The number of price levels is selected at 5 as a tradeoff between the fidelity of pricing and size of the state space in the MDPs. We

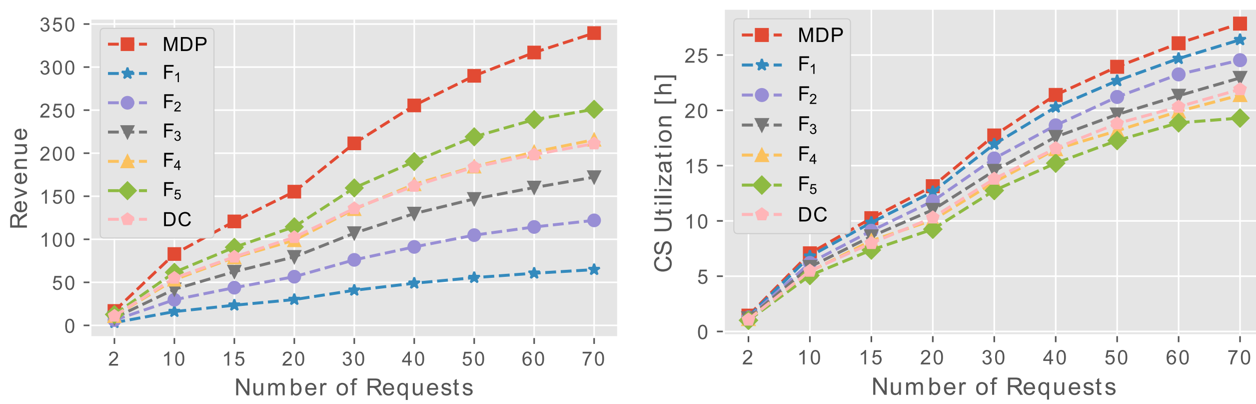

Figure 7: Performance of the MDP dynamic pricing compared to the performance of the flat rate (F_1 to F_5) and the demand correlated pricing strategy (DC) during one simulated day using prices between 1 and 5 when varying the number n of charging requests per day. The curves show the average values based on 400 runs with a random selection of n booking requests from the real-world dataset and with price elasticity parameter $C = 0.03$ in each run.

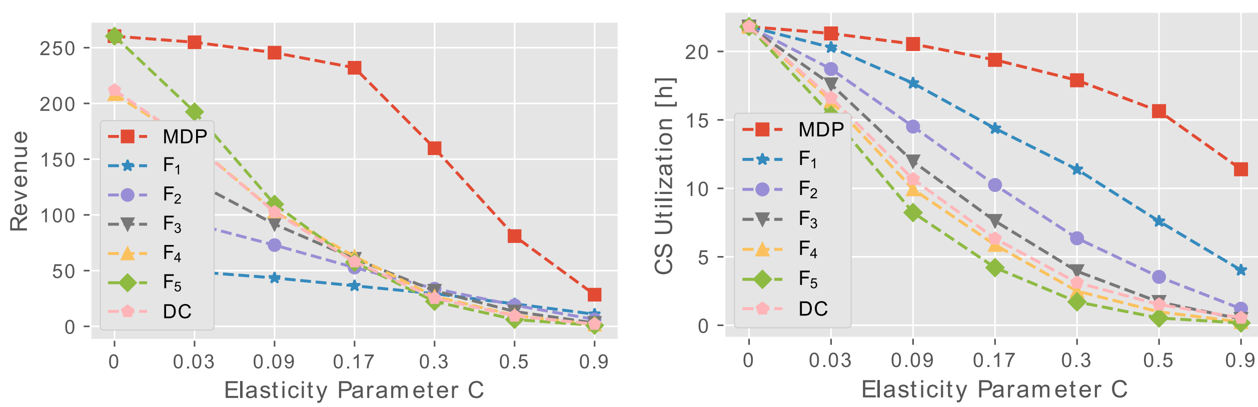

Figure 8: Performance of the MDP dynamic pricing compared to the performance of the flat rate (F_1 to F_5) and the demand correlated pricing strategy (DC) during one simulated day using prices between 1 and 5 when varying the price elasticity parameter C. The curves show the average value of the metric based on 400 runs with a random selection of 40 booking requests from the real-world dataset in each run.

use the same price levels for the flat rate strategies and the demand correlated strategy.

To compare the performance of different pricing strategies we use two metrics: charging station revenue and charging station utilization time. Detailed description of these metrics is given in Table 2.

In each run of the experiment, we randomly draw the set D of size m from the real-world E-WALD dataset. These requests are ordered by $req(r)$ and processed in parallel by each pricing strategy. Each strategy discards some requests due to capacity constraints, for other requests the issuing EV driver agents refuse the price. Metrics described in Table 2 are calculated from requests accepted by the charging stations with prices accepted by the EV driver agents. In each experiment, we perform 400 runs and average the resulting metrics. The runtime of the simulations and the solver implemented in Python is in the order of minutes on the Intel Core i7-3930K CPU @ 3.20GHz with 32 GB of RAM, with most of the time spent on pre-calculation of the policies for the MDPs.

For the first experiment, we fixed price elasticity parameter at $C = 0.03$ and $m = 40$ requests. In this experiment, we report the quartiles of the evaluation metrics in Figure 6. We can see that for given parameters, MDP dynamic pricing improves revenue. Furthermore, it also improves utilization. The same figure shows the effect of length of the MDP time window on the results. Using shorter time windows improves revenue but increases the variance of the results. For this reason, we use 24 time windows per day in the rest of the experiments. Figure 6 also shows that the results obtained for the MDP dynamic pricing can be achieved reliably, without increasing the variance of the observed metrics over the flat rate pricing.

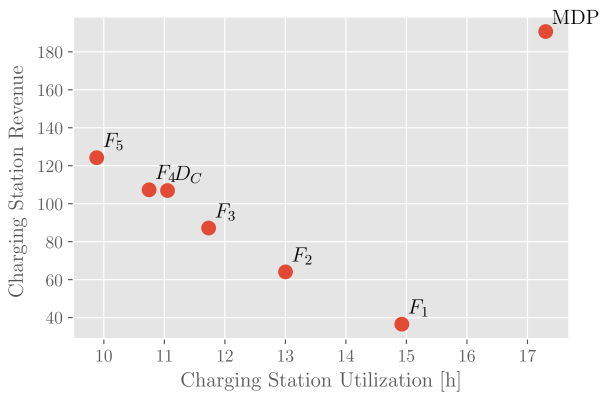

Figure 9: Performance of the MDP dynamic pricing strategy compared to he baselines (flat rates and demand correlated pricing) with demand and price elasticity parameters averaged out.

The selection of the number of price levels was done empirically; appropriate choice of price levels inherently depends on the concrete values of demand and price elasticity. However, any sufficiently high number of price levels is enough to demonstrate trends in the dependence on price elasticity and demand. Thus, for simplicity, we use $n_p = 5$ in the following experiments.

For the second experiment, we fixed the price elasticity parameter to $C = 0.03$ and varied m, the number of customers arriving per day from 2 to 70. Results of this experiment can be seen in Figure 7. As could be expected, increasing the number of booking requests increases revenue, utilization and delivered charge for all pricing schemes. Unsurprisingly, the revenue maximizing MDP strategy comes out ahead of other pricing strategies in terms of revenue. Moreover, MDP dynamic pricing outperforms the other pricing strategies in utilization across all numbers of booking requests. While the MDP dynamic pricing is closest to flat rate with price 1 in utilization and delivered charge, it is closest to flat rate with price 5 in revenue and average price per kWh.

For the third experiment, we varied the price elasticity parameter C through values given in Figure 5. We fixed m, the number of requests in D, to 40. Results of the second experiment are shown in Figure 8. With the increasing elasticity, revenue and utilization decrease. MDP dynamic pricing again comes on top in both metrics.

In the second and third experiment, the revenue of flat pricing strategy increases as we increase the flat rate. However, by increasing the number p_{max} and the number of price levels above 5, we found that the revenue of flat rate starts to decrease after certain threshold price level dependent on demand and elasticity. At all times, MDP pricing strategy that uses the same p_{max} achieves multiple times higher revenue.

Notice the utilization and delivered charge in the second experiment are the same for all pricing strategies when $C = 0$; when the demand is inelastic, customers always accept the offered price and the charging capacity is distributed solely on the first come, first serve basis. The downslope trend of the utilization and delivered charge with increasing elasticity were to be expected, given the fixed number of 40 booking requests at average duration 0.726 (the maximal theoretical utilization with three charging points would be $3 * 24$).

Figure 9 shows the aggregate results of our experiments with MDP pricing clearly outperforming the baseline pricing strategies. Note that the demand correlated strategy falls short of the MDP pricing strategy as it does not respond to the changes in capacity.

The results show that that in simulation, the MDP dynamic pricing will improve revenue compared to the baselines in all values of absolute daily demand and all values of price elasticity, with the exception of completely inelastic demand. The relative increase in revenue is greater if demand is higher or if it is more elastic. Moreover, the MDP dynamic pricing also improves the utilization and delivered charge when compared to the flat rate and demand correlated baselines.

6 CONCLUSION

We have shown how to use Markov Decision Process dynamic pricing of charging services for electric vehicles. The proposed method is focused on maximizing revenue, but it also improves the utilization of the charging station resources over the flat rate and demand correlated baselines through the improved allocation of charging services.

We have compared the proposed MDP dynamic pricing strategy with the baseline of currently most commonly used flat rate pricing across a range of system parameters, that is, the price elasticity of demand and volume of demand for charging services. The revenue generated by the proposed dynamic pricing strategy was up to 5 times higher than any flat rate pricing method with the relative revenue improvement increasing fast as the elasticity increases.

Moreover, while not directly optimizing the utilization of the charging station, our proposed method performed better than flat rate pricing and delivered energy across all considered scenarios. The improvement of our method in the utilization of the charging station over the flat rate pricing was up to 200%, depending on the price elasticity and demand.

The most obvious future work is to incorporate dependence of the consecutive time windows into the MDP model. Further, the model can be extended to a game theoretic setting. Such approach will, however, need substantial work to provide scalability for the solution to be practically usable. Another problem entirely is practicality of creating realistic instances of such games.

ACKNOWLEDGEMENTS

This research was funded by the European Union Horizon 2020 research and innovation programme under the grant agreement N°713864 and by the Grant Agency of the Czech Technical University in Prague, grant No. SGS16/235/OHK3/3T/13.

REFERENCES

[1] M.H. Albadi and E.F. El-Saadany. 2008. A summary of demand response in electricity markets. *Electric Power Systems Research* 78, 11 (2008), 1989 – 1996. https://doi.org/10.1016/j.epsr.2008.04.002

[2] Richard Bellman. 1957. A Markovian decision process. *Journal of Mathematics and Mechanics* 6 (1957), 679–684.

[3] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010,*

Cambridge, Massachusetts, USA, 5-8 June 2010. 311–320. https://doi.org/10.1145/1806689.1806733

[4] Wen-Chyuan Chiang, Jason C.H. Chen, and Xiaojing Xu. 2007. An overview of research on revenue management: current issues and future research. *International Journal of Revenue Management* 1, 1 (2007), 97–128. https://ideas.repec.org/a/ids/ijrevm/v1y2007i1p97-128.html

[5] S. Deilami, A. S. Masoum, P. S. Moses, and M. A. S. Masoum. 2011. Real-Time Coordination of Plug-In Electric Vehicle Charging in Smart Grids to Minimize Power Losses and Improve Voltage Profile. *IEEE Transactions on Smart Grid* 2, 3 (Sept 2011), 456–467. https://doi.org/10.1109/TSG.2011.2159816

[6] S. Rasoul Etesami, Walid Saad, Narayan Mandayam, and H. Vincent Poor. 2017. Smart routing in smart grids. In *Decision and Control (CDC), 2017 IEEE 56th Annual Conference on*. IEEE, 2599–2604.

[7] Ines Frade, Anabela Ribeiro, Goncalo Goncalves, and Antonio Antunes. 2011. Optimal Location of Charging Stations for Electric Vehicles in a Neighborhood in Lisbon, Portugal. *Transportation Research Record: Journal of the Transportation Research Board* 2252 (2011), 91–98. https://doi.org/10.3141/2252-12 arXiv:https://doi.org/10.3141/2252-12

[8] Jiarui Gan, Bo An, and Chunyan Miao. 2015. Optimizing Efficiency of Taxi Systems: Scaling-up and Handling Arbitrary Constraints. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS '15)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 523–531. http://dl.acm.org/citation.cfm?id=2772879.2772946

[9] Jiarui Gan, Bo An, Haizhong Wang, Xiaoming Sun, and Zhongzhi Shi. 2013. Optimal Pricing for Improving Efficiency of Taxi Systems. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI '13)*. AAAI Press, 2811–2818. http://dl.acm.org/citation.cfm?id=2540128.2540534

[10] Zheng Gu. 1997. Proposing a room pricing model for optimizing profitability. *International Journal of Hospitality Management* 16, 3 (1997), 273 – 277. https://doi.org/10.1016/S0278-4319(97)00015-7

[11] Keiichiro Hayakawa, Enrico H. Gerding, Sebastian Stein, and Takahiro Shiga. 2015. Online Mechanisms for Charging Electric Vehicles in Settings with Varying Marginal Electricity Costs. In *Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI'15)*. AAAI Press, 2610–2616. http://dl.acm.org/citation.cfm?id=2832581.2832614

[12] Fang He, Di Wu, Yafeng Yin, and Yongpei Guan. 2013. Optimal deployment of public charging stations for plug-in hybrid electric vehicles. *Transportation Research Part B: Methodological* 47, Supplement C (2013), 87 – 101. https://doi.org/10.1016/j.trb.2012.09.007

[13] R.A. Howard. 1960. *Dynamic Programming and Markov Processes*. The M.I.T. Press. https://books.google.cz/books?id=fXJEAAAAIAAJ

[14] Shripad Kulkarni and Pushkar H. Joshi. 2017. Passenger Airline Revenue Management: Research Overview and Emerging Literature. *International Journal of Engineering and Management Research (IJEMR)* 7, 1 (2017), 387–389.

[15] C.-Y. Cynthia Lin Lawell and Lea Prince. 2013. Gasoline price volatility and the elasticity of demand for gasoline. *Energy Economics* 38, C (2013), 111–117. http://EconPapers.repec.org/RePEc:eee:eneeco:v:38:y:2013:i:c:p:111-117

[16] P.S. McCarthy. 2001. *Transportation Economics: Theory and Practice*. Blackwell Publishers. https://books.google.com.au/books?id=86PAtgAACAAJ

[17] Jeffrey I. McGill and Garrett J. van Ryzin. 1999. Revenue Management: Research Overview and Prospects. *Transportation Science* 33, 2 (1999), 233–256. https://doi.org/10.1287/trsc.33.2.233 arXiv:https://doi.org/10.1287/trsc.33.2.233

[18] Janakiram Subramanian, Shaler Stidham Jr., and Conrad J. Lautenbacher. 1999. Airline Yield Management with Overbooking, Cancellations, and No-Shows. *Transportation Science* 33, 2 (1999), 147–167. https://doi.org/10.1287/trsc.33.2.147

[19] Moby Khan T. Donna Chen, Kara M. Kockelman. 2013. The electric vehicle charging station location problem: A parking-based assignment method for Seattle. In *Proceedings of 92nd Annual Meeting of the Transportation Research Board (TRB)*.

[20] Theoni Versi and Mark Allington. 2016. *Overview of the Electric Vehicle market and the potential of charge points for demand response.* Technical Report. ICF Consulting Services.

[21] Hengsong Wang, Qi Huang, Changhua Zhang, and Aihua Xia. 2010. A novel approach for the layout of electric vehicle charging station. In *The 2010 International Conference on Apperceiving Computing and Intelligence Analysis Proceeding*. 64–70. https://doi.org/10.1109/ICACIA.2010.5709852

[22] Yanhai Xiong, Jiarui Gan, Bo An, Chunyan Miao, and Ana L. C. Bazzan. 2015. Optimal Electric Vehicle Charging Station Placement. In *Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI'15)*. AAAI Press, 2662–2668. http://dl.acm.org/citation.cfm?id=2832581.2832621

[23] Yanhai Xiong, Jiarui Gan, Bo An, Chunyan Miao, and Yeng Chai Soh. 2016. Optimal pricing for efficient electric vehicle charging station management. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 749–757. http://dl.acm.org/citation.cfm?id=2937035