Dynamic Pricing Strategy for Electromobility using Markov Decision Processes

Jan Mrkos, Antonín Komenda and Michal Jakob
Artificial Intelligence Center, Faculty of Electrical Engineering, Czech Technical University in Prague, Czech Republic

Keywords: Electric Vehicles, Demand-response, Dynamic Pricing, Charging, MDP, Markov Decision Process.

Abstract: Efficient allocation of charging capacity to electric vehicle (EV) users is a key prerequisite for large-scale adaption of electric vehicles. Dynamic pricing represents a flexible framework for balancing the supply and demand for limited resources. In this paper, we show how dynamic pricing can be employed for allocation of EV charging capacity. Our approach uses Markov Decision Process (MDP) to implement demand-response pricing which can take into account both revenue maximization at the side of the charging station provider and the minimization of cost of charging on the side of the EV driver. We experimentally evaluate our method on a real-world data set. We compare our dynamic pricing method with the flat rate time-of-use pricing that is used today by most paid charging stations and show significant benefits of dynamically allocating charging station capacity through dynamic pricing.

1 INTRODUCTION

Electrification of personal transportation is commencing. There is a multitude of reasons, primarily environmental concerns, energy supply independence, and overall falling costs of production of both electric vehicles and the needed energy. Hand in hand with the clear benefits of a wide-spread deployment of electric vehicles (EVs) come many challenges. One of the most pressing problems is how to efficiently and cheaply distribute the energy from often unstable renewable sources to the EVs.

To illustrate the gravity of the situation, let us take the recent target to charge future EVs by no less than 300kW (Dyer et al., 2013). Provided a charging station with ten charging slots, we get to 3MW power intake if all the slots are charging EVs in parallel. For a comparison, an average instantaneous power consumption[1] of a U.S. household is about 1.2kW. The costs of upgrading the distribution network to cover such intakes would be extreme, on par with building the grid for additional three times the number of households[2].

Provided that the charging stations are not permanently fully occupied by EVs, an alternative to upgrading the grid is to charge stationary batteries at charging station, which are later used to fast charge EVs. In this approach, the initial costs of the upgrade of the grid are transferred to charging station owners in the cost of stationary batteries. Another approach is to use grid-centric methods ensuring fairness of charging such as the packetized charging management (Rezaei et al., 2014). However, such methods do not guarantee the charging service capacity, therefore the charging duration can not be guaranteed.

Existing regulations and laws design the overall mechanism for the allocation of charging resources. Within the rules of this mechanism, participants in this mechanism are free to act in a way that promotes their self-interests. These participants are power distribution service operators, charging service providers, and EV drivers responsible for charging their cars. However, self-interested strategies employed by the participants can be dangerous to the system as a whole. Multi-agent paradigm is suitable to efficiently balance the supply of the (renewable) power and EV charging demand. More precisely, the field of multi-agent resource allocation (Chevaleyre et al., 2006) provides techniques for solving problems

[1]Based on the 2015 statistics of the U.S. Energy Information Administration: https://www.eia.gov/tools/faqs/faq.php?id=97&t=3https://www.eia.gov/tools/faqs/faq.php?id=97&t=3

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

Mrkos, J., Komenda, A. and Jakob, M.
Dynamic Pricing Strategy for Electromobility using Markov Decision Processes.
DOI: 10.5220/0006601505070514
In *Proceedings of the 10th International Conference on Agents and Artificial Intelligence (ICAART 2018) - Volume 2*, pages 507-514
ISBN: 978-989-758-275-2

with self-interested agents.

We base our approach to dynamic pricing on the pricing of charging services as whole. This means that we do not consider the charging station to be selling electricity with price per kWh as is currently the norm. Instead, in our view, the charging station is selling a charging service that has multiple parameters that include time of day, volume of consumed electricity etc. This approach does not directly depend on the details of the low-level battery charging patterns and its optimization as proposed in (Cao et al., 2012). (Li et al., 2014) suggests an idea similar to our approach in the locational pricing; however, their method is based on the solution of nonlinear optimization towards the social welfare to get charging prices. Our approach solves the same problem as a solution to a set of decentralized Markov Decision Processes (MDPs), where the resulting decisions are prices of charging services at various times.

In this paper, we provide an experimental comparison of the MDP demand-response pricing strategy applicable in the context of multi-agent resource allocation for electromobility and today widespread flat rate time-of-use pricing (Versi and Allington, 2016).

2 DEMAND-RESPONSE PRICING

Economics, revenue management, and supply chain management have extensively studied demand-response pricing mechanisms of various kinds of services (Albadi and El-Saadany, 2008; McGill and van Ryzin, 1999). These fields recognize demand-response pricing as a critical lever for influencing the behavior of buyers. For this quality, we choose demand-response pricing as a way of dealing with increasing loads on the power grid caused by uptake of EVs.

To put the charging services pricing into context, we can view it as pricing of perishable goods, such as seasonal clothing, hotel rooms or airline tickets (Subramanian et al., 1999). These goods have value only until a certain point in time. For clothing, that is the end of the season, for airline tickets, it is the departure of the airplane. In the case of charging stations, the commodity is the charging resources available in a given time window. With perishable products, the goal is to sell the available stock for a profit before the stock expires. Same with the charging services, charging resources are a missed profit opportunity if they are left unused at the end of some time window.

Pricing of airline tickets has been extensively studied in different variations and with focus on various aspects of the problem (Chiang et al., 2007). The first step is usually the construction of an approximate model of user behavior (such as customer price sensitivity, seasonality of demand, no shows, etc.). Next, airlines need to determine how many tickets at which price to sell through the maximization of expected revenue. In brief, optimal rule for accepting or rejecting bookings is as follows: "Is the profit from this booking greater or smaller than the expected profit from this seat that we could get later? Confirm the booking now if the profit is bigger than the expected profit later. If it is smaller, deny this booking."

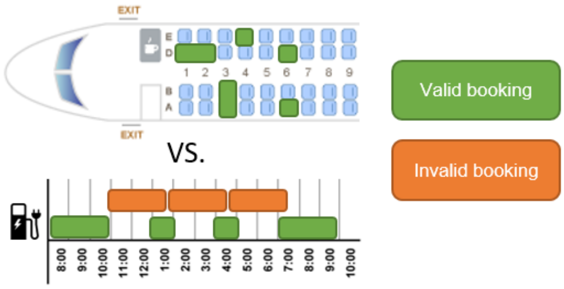

Figure 1: Difference between airline pricing and charging station pricing. Green rectangles show valid bookings. Seat bookings do not significantly affect bookings of other seats (except for large group bookings). On the other hand, a booking of short charging sessions and 1:00 and 4:00 blocks bookings of longer charging sessions shown in red.

However, each accepted or rejected booking can influence following bookings as consecutive customers are not able to book seats at the same price. If we include connecting flights and group bookings into the problem, pricing decisions can have a domino effect on the pricing in the whole network of an airline operator. Although the problem can look simple, due to this knock-down effect, the survey (McGill and van Ryzin, 1999) notes the complexity of this issue. For this reason, most work on this subject restricts the problem in some way.

In the following sections, we approach the charging station pricing strategy in a similar way to how airline revenue management approaches the pricing of airline tickets. Both problems focus on perishable goods where the knockdown effect plays a role as both bookings of airline tickets and bookings of charging station can affect following sales.

Similarly to airline revenue management, we do not directly consider competition between different service providers. We aggregate charging station customers (whose actions are based the options offered by different service providers as well as individual circumstances) in an environment model that adjusts demand as a response to changing price.

An important distinction between the airline pricing and charging station pricing is the interconnectedness of the bookings in the case of charging services; this is not present with the sales of the airline tickets. In the case of the airline tickets, it is not particularly important which seat (in a given class) was sold as booking of single seat does not block booking of surrounding seats. This distinction is illustrated in Figure 1.

3 PRICING STRATEGY

In the next section, we focus on the formalization and modeling of the pricing strategies as Markov Decision Processes (MDP)(Bellman, 1957).

3.1 Demand-response Pricing Strategy

We focus on the demand-response pricing strategy (Albadi and El-Saadany, 2008) for single charging station. As input, we use a discretization of possible charging parameters (time, duration and location of charging), current and historical utilization of single charging station and the expected price elasticity of the demand for charging services.

The goal of the pricing strategy described below is the maximization of charging station revenue within particular time horizon. However, other optimization criteria are possible to achieve different goals. For example, a publicly owned charging station that is not concerned with profits may attempt to maximize charging station utilization or minimize waiting times at the charging station.

3.2 Problem Formalization

In this section, we formalize the problem of dynamic pricing of charging station offers. In this formalization, our focus is on the offers that use the uniform discretization of time and some form of discretization of the other offer parameters. Regular structuring in time simplifies the formalization. We consider the set T of *times* $t_1, t_2, \dots, t_n$ for which the *prices* $p_1, p_2, \dots, p_n$ need to be determined.

The times in T denote the starting times of time intervals of the same length that start at time t_i and end at time t_{i+1}. For simplicity, we denote both the time interval and the associated start time with t_i. c_i is the expected free charging *capacity* of the charging station in each time interval t_i.

In this formalization, we use the symbol c_i, the *expected free* capacity to be the aggregate of all charging station constraints, such as power grid capacity or the number of available charging connectors.

Customers may book charging in any future time interval. Thus, we will denote price and capacity as functions $p(t_i, \tau)$ and $c(t_i, \tau)$, meaning the price or capacity of ith time interval at time $\tau \in T$. Fixing the time τ, both price and capacity functions are elements of the space of step functions over real numbers $\mathcal{L}$.

Given an offer, each *reservation* $r_j = (r_j^0, r_j^1, \tau_j) \in R$ is made for one or more consecutive time intervals, starting at time interval r_j^0 and ending in r_j^1. The arrival of a reservation is denoted τ_j. The *price of the reservation* $\pi(r_j)$ is the sum of prices associated with the time intervals at the time of the reservation:

$$\pi(r_j = (r_j^0, r_j^1, \tau_j)) = \sum_{t=r_j^0}^{r_j^1} p(t, \tau_j)$$

Reservations arrive randomly according to a demand distribution that is dependent on the price function as well as external factors. The set of reservations R depends on the pricing function because changes to the price influence demand. Thus, R is a function $R(p) : \mathcal{L} \times T \to 2T \times T \times T$, where $2T \times T \times T$ is the superset of the set of all possible reservations. The initial free capacity of the time interval relies on the state of the grid. We model this as stationary distribution.

The goal of the charging station is to maximize profits. In each time interval, charging station needs to cover the ground *cost of maintaining the infrastructure* denoted Γ_g. During charging, a charging station needs to pay for the electricity consumed from the grid $\gamma(r_j)$. This cost is unique to each charging session as it depends on the total charge delivered to the EV and possibly variable price of electricity and charging rate. Given that the price $\pi(r_j)$ of a reservation r_j, *profit or loss* at the end of the time horizon, after n time intervals, can be written as the sum across all reservations:

$$\Pi = -n\Gamma_g + \sum_{r_j \in R} \pi(r_j) - \gamma(r_j)$$

However, this point of view of the profit is not particularly useful for optimization. As each price $\pi(r_j)$ is calculated as the sum of prices of booked time intervals, we can rewrite the profit as a function of pricing of time intervals:

$$\Pi(p) = -n\Gamma_g + \sum_{r_j \in R(p)} \sum_{t=r_j^0}^{r_j^1} p(t, \tau_j) - \gamma(r_j)$$

The optimization goal of the pricing is then:

$$p' = \arg\max_{p \in \mathcal{L}} \Pi(p)$$

Finding optimal pricing function is not an easy task. The pricing function is part of the innermost sum in the calculation of Π. The sum itself is also dependent on p, as the set of reservations R is dependent on p. In fact, the set of reservations is dependent on p through the actions and responses of individual customers. However, it would be challenging to model behavior of each customer to get optimal pricing strategy.

To make the problem tractable, we aggregate behavior of a multitude of customers into the probability distributions that describe the behavior of customers together. As such, we can no longer maximize the profit in absolute numbers. Instead, we maximize the expected profit (in the statistical sense):

$$p' = \arg\max_{p \in \mathcal{L}} E(\Pi(p)) \tag{1}$$

The framework that deals with problems posed this way is the framework of Markov Decision Processes.

3.3 Modeling as a Markov Decision Process

The described optimization problem is a complex one. To solve it, we choose to model the charging services pricing problem as a Markov Decision Process (MDP) (Bellman, 1957; Puterman., 1994). First, we decompose the optimization problem into a sequence of decisions, where at each time point τ, we need to select new pricing function $p \in \mathcal{L}$. Markov Decision Processes provide a framework for modeling a broad range of sequential decision problems, where an agent must submit a sequence of decisions as responses to the developing environment.

An MDP is a tuple $\langle S, A, R, P, s_0 \rangle$, where S is a finite set of *states*, A is a finite set of actions; $P : S \times A \times S \to [0,1]$ is the transition function forming the transition model giving a probability $P(s'|s,a)$ of getting to the state s' from the state s after application of the action a; and a reward function $R : S \times A \times S \to \mathbb{R}$. Starting in initial state s_0, any action from A can be chosen. Based on this action, the system develops and moves 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.

When solving the MDP, the goal is to select a sequence of actions that will in expectation lead to the highest accumulated reward.

For the implementation, we consider a charging station with integer capacity between 0 and c_{max} and possible prices being integers between 1 and p_{max}. We consider a single day of 24 time intervals, each 1 hour long. For computational feasibility reasons, we split the MDP into multiple MDPs, one for each time window. This splitting gives us 24 MDPs, each responsible for setting the price of the corresponding time window. Each MDP generates a decision policy for its own time window. While these policies are optimal in a sense that they maximize the reward in given time window, together they may not maximize the revenue in the whole day.

MDP-1 is in charge of setting the price between 00:00 and 01:00, MDP-2 for setting the price in the time window between 01:00 and 02:00 and so on. As the bookings arrive ahead of the charging, we include time t to the kth time window in the state of the kth MDP. For example, in our experiments, at the time between 13:00 and 14:00, time t to MDP-18 is 4.

State s of MDP-k is thus defined by the capacity, price and time; $s = (c, p, t)$. The actions are changes to the price in the kth time interval, that is, $a = p'$. The transition model for the kth MDP, $P^k(s'|s,a)$ then determines, given the current price, capacity and time to the kth time interval, whether somebody books charging (which reduces capacity in the time window.). We calculate the transition probabilities from the given price elasticity function $P_e(p)$ and discrete historical probability $D^k(t)$ of a booking arriving t ahead of kth time window:

$$\begin{aligned} P^k((c-1, p', t-1)|(c,p,t), p') &= P_e(p) D^k(t), \\ & \quad c > 0, t > 1 \\ P^k((c, p', t-1)|(c,p,t), p') &= 1 - P_e(p) D^k(t), \\ & \quad t > 0 \\ P^k((c', p', t')|(c,p,t), p') &= 0 \quad \text{otherwise} \end{aligned}$$

$D^k(t)$ give the probability of a booking for charging during kth time interval arriving t ahead of the kth time interval.Values of $D^k(t)$ are taken from the historical data. The use of $D^k(t)$ in the transition function forms a simplified demand model that models the distribution of demand within one day, but one that is independent of the absolute demand expected within this day. Price elasticity function $P_e(p)$ gives the probability that given price p a customer will accept the price of the booking.

Instead of maximizing the profit of the whole day (in the sense of Equation 1), each MDP factor maximizes the profit in given time window. This adjustment is an important omission regarding optimality of the resulting pricing. As Figure 1 shows, sales of time windows affect the value of the neighboring time windows. However, as we will show in Section 4, even in this factored form with simple demand model, MDP demand-response pricing can bring some benefits. In the model, splitting of the MDP into MDP

Table 1: Summary statistics of the E-WALD data for the selected charging station with three charging points.

CS Dataset Statistics	Mean	Std
Charging session duration	0.726 h	0.794 h
Charge per charging session	6.72 kWh	5.19 kWh
# of daily charging sessions	2.53	1.49

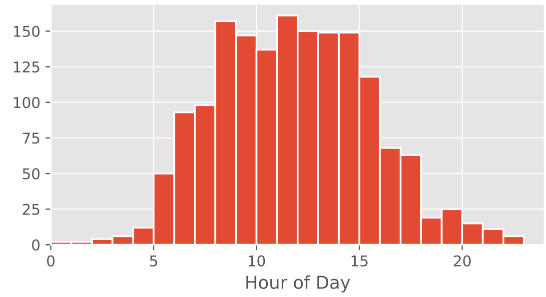

Figure 2: Histogram of the charging session start times in the E-WALD dataset for the selected charging station.

factors means that changes to capacity and price do not affect on the neighboring prices.

We find the optimal policies for MDP-1 to MDP-24 through policy iteration. We implemented all structures and algorithms in Python, using commonly used Python packages such as *NumPy* (van der Walt et al., 2011) and *Pandas* (McKinney, 2011). Policy iteration algorithm that we used is from the *pymdptoolbox*[3].

4 EXPERIMENTS

We evaluate the MDP dynamic pricing algorithm on real data provided by E-WALD[4], EV charging station provider in Germany. 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.

4.1 Dataset

The dataset contains information on charging sessions realized at one of the E-WALD charging stations. This information 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 the charging session and 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 2.

The particular charging station dataset does not contain any pricing information about the charging sessions. However, E-WALD uses only flat rate pricing in all their pricing stations.

4.2 Experimental Setup

In our experiments, we compare the performance of the flat rate pricing to the MDP based dynamic pricing. To compare their performance we use four metrics, charging station revenue, charging station utilization time, charge delivered by the charging station and price per unit of energy sold by the charging station. A detailed description of these metrics is given in Table 2.

We use the real E-WALD charging station data to simulate 24 hour period of the charging station operation. As the data was collected at charging station with three charging slots, in our experiments we consider our station to have three charging points. That is, we use $c_{max} = 3$ and $p_{max} = 5$. We consider the charging points to be capable of realizing any charging session recorded in our dataset. This means that at most three charging sessions can be realized at any point in time. In our simulation, customers book charging sessions ahead of time. Charging station rejects the booking if all three charging points are already booked for any portion of the requested time. If the station can realize the booking, pricing scheme is used to determine the price of the charging session which the station offers to the customer. Based on this price, the customer either accepts or rejects the offer. Price elasticity of demand described below determines whether the offer is accepted or rejected by the customer.

To make it possible for users to plan their trips in the environment where the charging capacity may not be readily available, we use ahead of time bookings in our simulation. We simulate how much ahead each customer books the charging session by drawing from the uniform distribution. The maximum period ahead of which customer can book a charging session is in our experiments set to 6 hours. The time of the bookings determines the order of the arrival of the bookings to the.

The simulation starts by drawing n charging ses-

[3] https://github.com/sawcordwell/pymdptoolbox

[4] We would like to thank E-WALD for providing us with the charging data for this study.

Table 2: Description of evaluation metrics.

Metric	Description
CS Revenue	Revenue of the charging station is the sum of prices of all charging sessions. We do not express price in any currency. Instead, we use *unit price* as a basic unit. Revenue is directly dependent on the selected pricing scheme.
CS Utilization	Measured in hours, it is the added duration of all charging sessions realized by the charging station. This as a proxy of a social welfare of the EV drivers achieved through various pricing schemes. The higher the utilization, the more of the EV driver charging demand was satisfied by the charging station.
Delivered Charge	Measured in kWh, it is the charge delivered to all of the charging station customers. Delivered charge is another proxy of a social welfare of the EV drivers realized by various pricing schemes. The higher the delivered charge, the more of the EV charging demand was satisfied by the charging station. Because of each EV charging with different charging rate, this information complements the CS Utilization metric.
Energy price	Average price per unit of energy across all charging sessions realized by the CS.

sions from the dataset. As can be seen from Table 1, the mean number of customers at the charging station is quite small. Also, the dataset does not give us any information about the unsatisfied demand for charging services. Thus, in most of our experiments, we use higher values of n so that all demand cannot be satisfied by the given charging station.

Normalized histogram in Figure 2 and the simulated booking times are the basis of the historical probability $D_k(t)$ of a booking request arriving t ahead of kth time window.

The MDP dynamic pricing uses different price for every hour. To get the price of the charging session, we first split the charging session into segments that correspond to the various dynamic prices. The corresponding hourly rate then multiplies the length of each segment. Adding the partial prices together gets us the price of the charging session offered to the user. For the flat rate pricing, the duration of the charging session in hours is multiplied by the hourly rate.

When the customer receives the offered price, he can accept or reject the offer. We simulate this using the price elasticity of demand curve. The price elasticity function we use is $P_e(x) = e^{-Cx}$.

Because we do not know the real price elasticity of demand for EV charging services and we can not estimate it from data, we experiment with multiple values of C. The different values of C and the corresponding shapes of price-elasticity curves are shown in Figure 3. Having the price of the charging session, we apply the price elasticity function to this price. The resulting number is a probability that the customer accepts the offer. If the user accepts the offer, the charging session is added to the other already booked charging sessions. Rejected offer is discarded and no longer used by the system. For $C = 0$, we talk about inelastic demand as the customer will accept any price. At $C = 0.5$ the demand is highly elastic as small changes to the price have a big effect on users 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 (Lin Lawell and Prince, 2013).

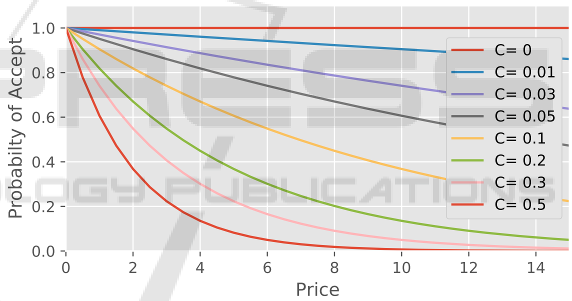

Figure 3: Price elasticity of demand curves for different values of the C parameter.

4.3 Results

To compare the performance of the MDP dynamic pricing and the flat rate pricing, we experimented with various numbers of customers and various parameters of price elasticity of these customers. In each experiment, we compare the MDP dynamic pricing that can set price in each time window to an integer value between 1 and 5. We compare it to the flat rate pricing that uses flat rates between 1 and 5.

For the first experiment, we fixed the price elasticity parameter to $C = 0.1$ and varied the number of customers arriving per day from 2 to 80. For the second experiment, we varied the number price elasticity parameter C through values given in Figure 3. We fixed the number of booking to 40.

Each data point in Figures 4 and 5 is an average of

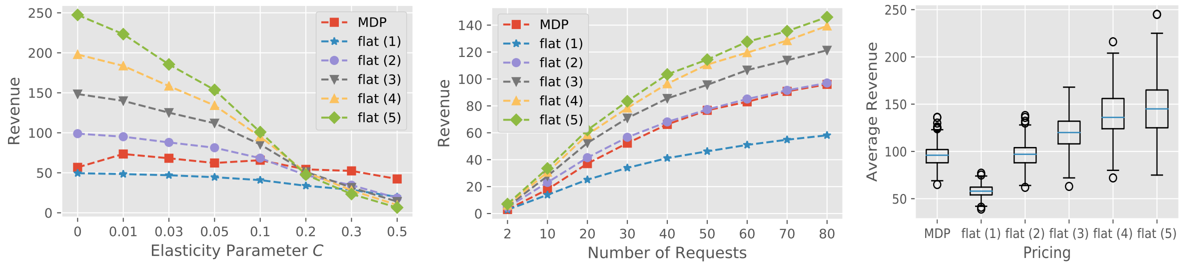

Figure 4: Performance in terms of revenue of the MDP dynamic pricing compared to the performance of the flat rate during one simulated day. The graphs are based on 400 runs with a random selection of booking requests from the E-WALD dataset. Price elasticity parameter $C = 0.1$and $n = 40$ booking requests in plots where these parameters are not on axis.

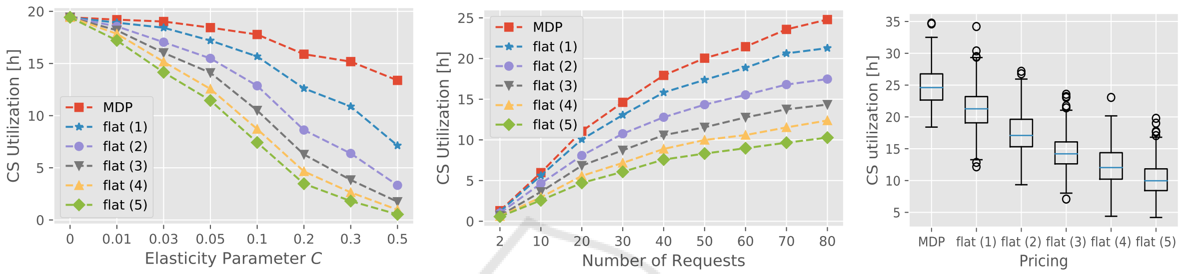

Figure 5: Performance in terms of CS utilization of the MDP dynamic pricing compared to the performance of the flat rate during one simulated day. The graphs are based on 400 runs with a random selection of booking requests from the E-WALD dataset. Price elasticity parameter $C = 0.1$and $n = 40$ booking requests in plots where these parameters are not on axis.

400 runs. In each run, we picked the booking requests randomly from the full E-WALD dataset. For price elasticity parameter $C = 0.1$ and 40 requests we give the quartiles of the evaluation metrics.

As could be expected, increasing the number of booking requests increases revenue, utilization and delivered charge for all pricing schemes. Figure 4 for revenue and Figure 5 for CS utilization illustrate this, curves for delivered charge display the same trends as figures for utilization.

Note that while the revenue is lower for the MDP demand-response pricing for the lower price elasticity curves ($C < 0.2$), the charging station utilization and delivered charge are better across all values of C. The utilization and delivered charge are same for all pricing schemes when $C = 0$; that is, 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.

For the experiment with variable elasticity parameter C, the downslope trend of the utilization and delivered charge with increasing elasticity are 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 $3 * 24$). As the price elasticity increases, the likelihood of any given customer booking for given price becomes lower.

Another notable result is that while the MDP price per kWh is for most values of C comparable to the flat rate of price 1, the revenue of MDP is consistently higher than the revenue of the flat rate of 1.

The results show that that in simulation, the MDP dynamic pricing will return greater revenue than flat rates with a price higher than one only if the demand for EV charging is somewhat elastic (elasticity parameter $C \geq 0.2$, Figure 3). However, dynamic pricing improves the utilization and energy delivered by the charging station across all values of the elasticity parameter C and any number of booking requests, while keeping the average price per kWh to the customer comparable to the flat rate pricing with the lowest price. Additionaly, these results for the demand-response MDP pricing are achieved reliably, without increasing the variance of the observed metrics over the flat rate pricing.

The runtime of the simulations 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.

5 CONCLUSION

We have shown how to use the Markov Decision Processes to model the problem of demand-response pricing of charging services for electric vehicles. Using the factored MDP demand-response pricing, we aimed at the core objectives of electromobility: distribution of cost between the grid and EV owners, signaling of power scarcity or abundance and incentivization of behavior change of the EV drivers.

Experimentally, we have compared the demand-response pricing strategy with the baseline of currently most commonly used time-of-use flat rate pricing across a wide range of environmental parameters, that is, the price elasticity of demand and volume of demand for charging services. While the revenue generated by the proposed demand-response pricing method was higher than the flat rate pricing methods only for specific values of the environmental parameters, our method performed better than any considered flat rate pricing in the achieved utilization of the charging station and delivered energy across all considered scenarios. The improvement of our method in the utilization of the charging station and delivered energy over the flat rate pricing of comparable revenue was up to 300%, depending on the price elasticity and the demand.

As we mentioned in the paper, the most obvious future work is to incorporate dependence of the consecutive time windows in the factored MDP model and improve the demand model. Further, the model is extendable to a game theoretic setting. Such approach will, however, need substantial work to provide scalability for practical use of the approach.

ACKNOWLEDGMENTS

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

REFERENCES

Albadi, M. and El-Saadany, E. (2008). A summary of demand response in electricity markets. *Electric Power Systems Research*, 78(11):1989 – 1996.

Bellman, R. (1957). A markovian decision process. *Journal of Mathematics and Mechanics*, 6:679–684.

Cao, Y., Tang, S., Li, C., Zhang, P., Tan, Y., Zhang, Z., and Li, J. (2012). An optimized EV charging model considering TOU price and SOC curve. *IEEE Trans. Smart Grid*, 3(1):388–393.

Chevaleyre, Y., Dunne, P. E., Endriss, U., Lang, J., Lemaître, M., Maudet, N., Padget, J. A., Phelps, S., Rodríguez-Aguilar, J. A., and Sousa, P. (2006). Issues in multiagent resource allocation. *Informatica (Slovenia)*, 30(1):3–31.

Chiang, W.-C., Chen, J. C., and Xu, X. (2007). An overview of research on revenue management: current issues and future research. *International Journal of Revenue Management*, 1(1):97–128.

Dyer, C., Epstein, M., and Culver, D. (2013). Station for rapidly charging an electric vehicle battery. US Patent 8,350,526.

Li, R., Wu, Q., and Oren, S. S. (2014). Distribution locational marginal pricing for optimal electric vehicle charging management. *IEEE Transactions on Power Systems*, 29(1):203–211.

Lin Lawell, C.-Y. C. and Prince, L. (2013). Gasoline price volatility and the elasticity of demand for gasoline. *Energy Economics*, 38(C):111–117.

McGill, J. I. and van Ryzin, G. J. (1999). Revenue management: Research overview and prospects. *Transportation Science*, 33(2):233–256.

McKinney, W. (2011). Pandas: a foundational python library for data analysis and statistics.

Puterman., M. L. (1994). *Markov Decision Processes*.

Rezaei, P., Frolik, J., and Hines, P. D. H. (2014). Packetized plug-in electric vehicle charge management. *IEEE Transactions on Smart Grid*, 5(2):642–650.

Subramanian, J., Jr., S. S., and Lautenbacher, C. J. (1999). Airline yield management with overbooking, cancellations, and no-shows. *Transportation Science*, 33(2):147–167.

van der Walt, S., Colbert, S. C., and Varoquaux, G. (2011). The numpy array: A structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2):22–30.

Versi, T. and Allington, M. (2016). Overview of the Electric Vehicle market and the potential of charge points for demand response. Technical report, ICF Consulting Services.