Smart Charging of Electric Vehicles Using Reinforcement Learning

Konstantina Valogianni
kvalogianni@rsm.nl
Rotterdam School of Management
Erasmus University

Wolfgang Ketter
wketter@rsm.nl
Rotterdam School of Management
Erasmus University

John Collins
jcollins@cs.umn.edu
Computer Science Department
University of Minnesota

Abstract

The introduction of Electric Vehicles (EVs) in the existing Energy Grid raises many issues regarding Grid stability and charging behavior. Uncontrolled charging on the customer's side may increase the already high peaks in the energy demand that lead to respective increase in the energy prices. We propose a novel smart charging algorithm that maximizes individual welfare and reduces the individual energy expenses. We use Reinforcement Learning trained on real world data to learn the individual household consumption behavior and propose a charging algorithm with respect to individual welfare maximization objective. Furthermore, we use statistical customer models to simulate the EV customer behavior. We show that the individual customers, represented by intelligent agents, using the proposed charging algorithm reduce their energy expenses. Additionally, we show that the average energy prices, on an aggregated level, are reduced as a result of smarter use of the energy available. Finally we prove that the presented algorithm achieves significant peak reduction and reshaping of the energy demand curve.

Introduction

Key energy initiatives worldwide have as their main focus to become fully independent from nuclear power and move towards massive integration of renewable energy sources (such as Germany to become totally independent from nuclear power by 2020). These sources are highly volatile and create instability in the energy flow process. Therefore, energy policy makers stress the importance of effective balancing between energy supply and demand to avoid outages.

Electric Vehicles (EVs) are important tools towards a sustainable solution, since they have storage features. Massive EV integration in the Energy Grid has been outlined by the main players in the energy policy landscape: according to US president's energy plan 1 million EVs are to be integrated in the US energy market by 2015 (Department of Energy 2011), similar aspirations have been expressed by the German, Dutch and UK governments by 2020. The uncoordinated use of EVs, though, will lead to high price peaks during the charging time. Specifically, considering customers *range anxiety* (Franke et al. 2011), this charging may threaten the grid's stability. This particular anxiety refers to the people's feeling that they *may run out battery capaciy* while they are driving. Therefore they select to charge their cars more often than needed. People plug in their EVs when the loads on the grid are already high and without controlled charging this might lead to problems.

Consequently, EVs are expected to change the logic behind individual power consumption. Uncoordinated charging of EVs will lead to critical stress test of the current grid. Therefore, there is need for smart charging algorithms that will alleviate this strain. We propose a smart charging algorithm that maximizes individual welfare and reduces the individual energy expenses. As a result, it reduces the peak consumption and supports grid's stability. We use Q-learning trained on real world data (described in the Data Description subsection) to learn the individual consumption behavior and propose a charging algorithm with respect to individual welfare maximization objective. Furthermore, we use statistical customer models to simulate the EV customer behavior. We show that the individual customers, represented by intelligent agents, using the proposed charging algorithm reduce their energy expenses. Additionally, we show that the average energy prices, on an aggregated level, are reduced as a result of smarter use of the energy available.

This paper is structured as follows. First we present related work with regard to EVs and their smart use in the Smart Grid. Secondly, we describe the designed EV customer model and explain the parameters used. Next, we propose a novel Smart Charging algorithm based on RL. Additionally, we present various results from our analysis and discuss their impact on the Energy Grid. And finally, we add conclude remarks and outline streams for future research.

Background and Related Work

EVs comprise an essential party for absorbing demand peaks on the grid, resulting from renewables' volatility or other causes. They have a prominent role because of the the storage features available in the vehicle's battery (Kempton and Letendre 1997). These batteries play two roles, apart from providing mobility to the owners. The first use is to store power when there is available surplus and return it to the grid when there is shortage. This can happen through remote-control mechanisms, or by observing price variations (Kahlen et al. 2012). The second role is short-term balancing, which requires less capacity, and is controlled by external signaling. EVs show different consumption behavior from the plug-in hybrid EVs which use conventional fuels, as well. In the presented research we aim at the less possible dependence from fossil fuels, thus we focus on pure EVs.

The main issues that arise with the large scale use of EVs in the grid are the charging availability of each individual and the charging coordination to avoid demand peaks caused by simultaneous charging. A number of scholar publications deal with the issue of proper coordination strategies for EVs with main objectives varying from making profit as individual energy providers (*brokers*) to increasing the social welfare by reducing the energy prices. In (Acha et al. 2011) the authors examine the optimal power flow using agent-based modeling or EVs. The work (Gerding et al. 2011) presents an online auction mechanism where the owners of EVs state their timeslots available for charging and also bid for power. The authors propose mechanism for charging coordination which achieves better allocative efficiency compared to some fixed benchmarks. In (Lopes et al. 2009) the authors focus on the Portuguese energy grid and proposed a smart charging coordination with the objective of maximizing the EVs integration. Finally, in (Vytelingum et al. 2010) the authors deal with storage batteries without specifying their use in EVs. However, they propose a charging coordination mechanism that leads to price reduction and increases the social welfare in a smart home.

We propose a smart charging algorithm that maximizes individual welfare and reduces the individual energy expenses. We use *Reinforcement Learning* (RL) to learn the individual consumption behavior and propose a charging algorithm with respect to individual welfare maximization objective. Furthermore, we design statistical customer models to simulate the EV customer behavior. In order to model the customers we use the *bottom-up* design. This approach (Christoph 1998; Valogianni et al. 2012) focuses on each individual household (or EV customer) and attempts to create a detailed user profile. We show that the individual customers, represented by intelligent agents, using the proposed charging algorithm reduce their energy expenses.

EV Customer Modeling

The proposed customer model focuses on EV owners and simulates their driving and charging behavior. More specifically, simulates each individual driving behavior and distance for various activities performed per day, as well as the household energy consumption behavior. Regarding charging, we assume initially only regular charging without having any *fast charging*. We plan to integrate *fast charging* capabilities later on. However, this addition will create higher challenges, since the demand peaks will be higher (i.e. higher demand at shorter time). We base our simulation on Smart Electricity Markets as discussed by (Bichler, Gupta, and Ketter 2010) (phase 1 and 2) and we are working towards integrating it in the Power TAC[1] environment (Ketter et al. 2012). An important factor in modeling the EV owners is their driving profile. This profile directly determines the battery capacity that a customer needs for driving and consequently the capacity available to offset supply-demand imbalances. For the precise creation of the customers' driving profiles we use mobility data from the Dutch Statistics Office(CBS)[2]. The population is divided according to gender and the social groups that comprise the total population. Those social groups with their special characteristics are: part-time employees, full-time employees, students and pupils, unemployed and retired persons. Here full-time employees are considered those who work 30 hours per week or more, whereas part-time employees are those with 12-30 hours of work per week. For each group there are different activities accompanied with the kilometers needed per day for each activity.

Table 1: Social groups and characteristics.

Social Group	Start work (time of the day)	Absence for work (hours)
Part-time employee	$[7am - 8.30am]$	4
Full-time employee	$[7am - 8.30am]$	8
Pupil/ Student	$[7am - 8.30am]$	7
Unemployed, Retired	-	-

Second step in the modeling process is the day determination (weekday or weekend). Having determined the activities related to each group considering the day, we create driving profiles corresponding to the distance that each customer drives per day (assuming average driving speed). Additionally, we determine the EV type that the customer owns and consequently the respective storage capacity. We assume that the customers in our population own purely electric cars like Nissan Leaf[3] and Tesla[4] (Table 2). With regard to the customer's charging and discharging availability we assume that the customers can charge their EV's battery when they are not only at home but also at work ("standard" charging with direct billing to the customer), which is nowadays implemented by large businesses in order to encourage their employees to drive "green." In this phase we do not assume that the customers discharge their remaining battery capacity for covering daily consumption needs. The *Vehicle to Grid* (V2G) concept will have a different effect on the energy consumption and we plan to examine it in the future.

The minimum charge level, the customer expects to have available for unplanned use of the vehicle, expresses customer's risk attitude towards range anxiety. Customers who are risk averse, want their EV fully charged as soon as possible after it's plugged in, and never want the charge to be less than 100% once it's charged. On the other hand totally risk seeking customers expect just the amount needed for planned driving at the times they plan to drive. In other words, they do not expect to use the vehicle for unanticipated driving. Thus, we experiment with populations expressing various risk attitudes.

The described model is depicted in Figure 1. The dark colored parameters denote the inputs in the model, while

[1]powertac.org

[2]www.cbs.nl

[3]www.nissanusa.com/leaf-electric-car/

[4]www.teslamotors.com

Table 2: Electric Vehicles specifications.

	Tesla S			Nissan Leaf
Battery Capacity (KWh)	40	60	85	24
Distance with full battery (km)	257.6	370.1	563.3	222.5
Charging Time for full battery (h)	12	17.5	23	7

the light colored ones are those calculated by each module (using look-up tables, and functions derived from the data). The model's output is the charging demand at each point of time (*timeslot*) according to the inputs given. This is the customer model used in the rest of the paper to simulate the individual driving behavior. It can represent various types of individuals according to the parameters given as input. In this case we parametrize it based on probabilities given from the Dutch Statistics office data (i.e the probabilities of each customer type, activities etc within the total population). Consequently, the demand coming from charging includes a stochasticity factor. This factor is interpreted as the uncertainty regarding which social group the customer belongs to, which kind of activities he/she has chosen for the coming day etc.

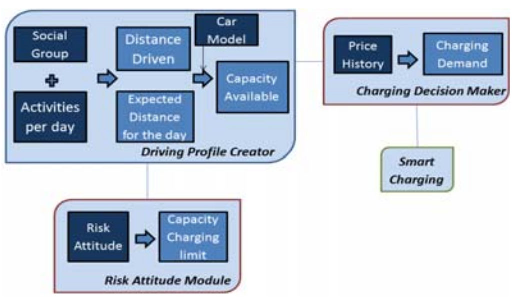

Figure 1: Electric Vehicle Customer Model.

For household consumption we use real world data available here (Frank and Asuncion 2010). These data refer to household consumption including all the various household appliances and refer both to weekdays and weekends. For more information see the respective subsection. Firstly, we need to define the percentage of the battery that a customer can charge in a 24h horizon. From Table 2 we conclude that customers owning Nissan Leaf cars, are more likely to charge up to full battery's capacity, since the nominal capacity is smaller compared to the Tesla car models. In the latter case a fully charged battery is not needed for a 24h horizon, unless the customer is planning a long trip. These specifications refer to the most optimistic case without taking into the urban driving environment. Thus, in future work we take into account various driving attitudes depending on the urban environment (traffic lights, congestions, rush hours etc.), so that we achieve higher realism.

Individual Utility and Welfare Function

Customers' utility results both from power consumption within the household and from consuming power for charging the EV. Assuming that the total consumption consists of the two components: x_h, household demand (KW) and x_c demand from EV charging we have the total utility in equation (1). This gives us the utility that the customer gets both from consuming energy through household devices and through EV charging. It does not apply on a time scale. Instead, it gives a relation between consumption and received utility, whenever this consumption might happen. Furthermore, it expresses an additive relationship since EV can be considered as an extra device with different consumption features that adds utility to the existing utility from household consumption. In other words, the customer gets utility from using each household appliance (and consuming energy) and EV is considered one of this devices, which has no inter-dependency with other devices. Thus, given that the utility from all the household devices is $U_h(x_h, \omega)$, the utility from the EV is additional to that and expressed as $U_c(x_c, \lambda)$. There is no relation between household consumption and EV charging, since the customer decides based on different activity set for consuming household energy and on different activity set for EV charging. The user is charged for the sum of his/her consumption which in our case is the sum of the household consumption and the consumption coming form EV charging.

$$U((x, \omega, \lambda)) = \kappa \cdot U_h(x_h, \omega) + (1 - \kappa) \cdot U_c(x_c, \lambda) \quad (1)$$

where $U_h(x_h, \omega)$ the utility from household consumption and $U_c(x_c, \lambda)$ the utility from charging. The parameter $\kappa \in (0, 1)$ indicates the weight of those two utility components to the customers overall utility from power consumption. The parameters ω and λ characterize each individual customer and their role is explained in the next paragraph. This utility function does not apply to a timescale. Instead yields the utility that the individual receives from the consumption of each extra energy unit, whenever the customer consumes.

Regarding customer's utility function resulting from household consumption, $U_h(x_h, \omega)$ we use the quadratic utility function (2) which expresses the individual satisfaction level for individual power consumption. This particular family of functions is used for power consumption since the customer receives increasingly lower benefit for each extra energy unit he/she consumes.

$$U_h(x_h, \omega) = \begin{cases} \omega \cdot x_h - \frac{\alpha}{2} \cdot x_h^2 & \text{if } \ 0 \leq x_h \leq \frac{\omega}{\alpha} \\ \frac{\omega^2}{2 \cdot \alpha} & \text{if } \ \ x_h > \frac{\omega}{\alpha} \end{cases} \quad (2)$$

where ω, stands for the *level of satisfaction obtained by the user as a function of its power consumption*. It varies among customers and gets values $\omega \in [0, 1]$. The variable x_h (KW) stands for the individual power consumption and α is a predefined parameter (e.g. 0.5 used in (Fahrioglu and Alvarado 2000)). This particular function family is chosen for power consumption, since shows linear decreasing marginal benefit and represents sufficiently the power consumption behavior. The marginal benefit $V_h(x_h, w)$ is presented in equation (3).

$$V_h(x,\omega) = \begin{cases} \omega - \alpha \cdot x_h & \text{if } \ 0 \leq x_h \leq \frac{\omega}{\alpha} \\ 0 & \text{if } \ \ x_h > \frac{\omega}{\alpha} \end{cases} \tag{3}$$

Figure 2 presents some examples of utility functions coming from household demand (continuous lines) and respective marginal benefits (dashed lines) for various values of ω and with predefined $\alpha = 0.2$. It becomes clear that higher values of ω yield higher utility to the customer. We also observe, that the higher the value for ω, the higher the individual consumption that yields maximum utility to the energy customer. This parameter may vary not only among customers but also across time periods for the same customer. Thus it is an strong indication for the customer's elasticity of demand.

With regard to customer's utility function resulting from EV charging, $U_c(x_c, \lambda)$ we use sigmoid utility function (Equation (4)). Intuitively, this choice results from the customers satisfaction level which is little until the customer has charged *enough to cover his/her driving needs*, and after this point each extra consumption unit adds less utility to the customer. The point that is enough to cover the driving needs varies among customer groups and attitudes towards *range anxiety*. The more risk averse customers are, the more utility for each extra power unit they get.

$$U_c(x_c, \lambda) = \frac{1}{1 + e^{(-x_c + \lambda)}} \tag{4}$$

with $0 \leq x_c \leq x_{c,max}$. The parameter $\lambda \in [0, \lambda_{max}]$ could be assumed as an indication for range anxiety, since the higher λ is, the steeper the increase in the utility function for each extra power unit, that the customer "consumes" (charges in this case). The max value for λ, λ_{max} is the nominal capacity of the EV battery. The range anxiety and parameter λ association becomes more clear from Equation (5) which shows the marginal utility from EV charging. We see that the marginal utility increases up to a maximum point (where the customer feels that has enough battery capacity to drive) and then decreases up to a point where no extra utility is added.

$$V_c(x_c, \lambda) = \frac{e^{(-x_c + \lambda)}}{(1 + e^{(-x_c + \lambda)})^2} \tag{5}$$

with $0 \leq x_c \leq x_{c,max}$. Figure 3 displays the utility function coming from EV charging and the marginal utility for various values of λ. We see that for higher λ values, the customer needs more power, so that he feels that has enough battery capacity. Thus, the higher the λ, the more risk averse towards range anxiety the customer is.

According to consumer theory (Mas-Colell, Whinston, and Green 1995), the individual welfare is defined as shown in (6).

$$W(x, \omega, \lambda) = U(x, \omega, \lambda) - x \cdot P \tag{6}$$

where P stands for the price per power consumption unit (€/KW). Substituting the Equations (1)-(5) in (6), we have:

$$W(x,\omega,\lambda) = \begin{cases} \kappa \cdot (\omega \cdot x_h - \frac{\alpha}{2} \cdot x_h^2) + (1-\kappa) \cdot \frac{1}{1+e^{(-x_c+\lambda)}} \cdots \\ \quad .. - (x_h + x_c) \cdot P \\ \quad \text{if } \ 0 \leq x_h \leq \frac{\omega}{\alpha} \ \text{ and } \ 0 \leq x_c \leq x_{c,max} \\ \kappa \cdot \frac{\omega^2}{2 \cdot \alpha} + (1-\kappa) \cdot \frac{1}{1+e^{(-x_c+\lambda)}} - (x_h + x_c) \cdot P \\ \quad \text{if } \ \ x_h > \frac{\omega}{\alpha} \ \text{ and } \ 0 \leq x_c \leq x_{c,max} \end{cases} \tag{7}$$

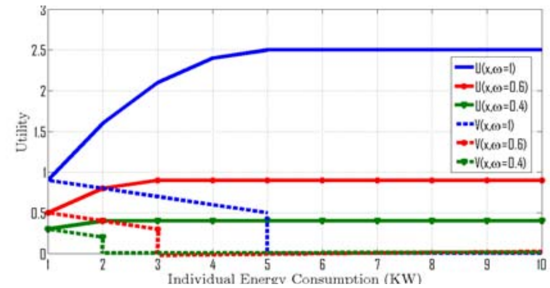

Figure 2: Utility function coming from household consumption and marginal benefit for various values of ω.

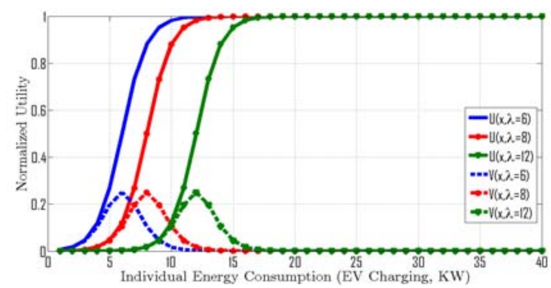

Figure 3: Utility function coming from EV charging and marginal benefit for various values of ω.

Given equations (1)-(7), we assume that each energy customer/EV owner is represented by an intelligent agent that acts towards the individual welfare maximization objective. More specifically, the agent is responsible for scheduling the EV charging so that individual welfare is maximum for each *timeslot*.

Learning Individual Power Consumption

Using the described customer model we simulate large populations of customers with different behavioral characteristics and socio-demographics. However, it is important for the customer agents to learn the individual household consumption before they decide on EV charging. If they have knowledge about each individual consumption they can apply charging algorithms that maximizes individual welfare.

We chose *Reinforcement Learning* (RL) to learn the customers' energy consumption. RL is based on a *reward mechanism* that provides the algorithm with positive and negative rewards for optimal or non-optimal decisions, respectively. We argue that this learning approach is suitable for the particular context, since the energy consumption shows a 24h periodicity and similarities in the trends throughout this 24h time frame. More specifically, we use the Q-learning mechanism as described by equations (8) and (9) (Mitchell 1997). The customer agent has to learn the optimal policy through rewards $r(s,a)$ that are offered to it for each (s,a). In the described problem, the learned policy corresponds to the individual household consumption over a daily horizon. The value of each reward is stated as $Q(s,a)$. The rewards are related to the utility function, giving higher utility, the closer the state is to the actual consumption. The (s,a) represents a state s to which the agent ends up as a result of an action a. The states in this problem are the various consumption levels discretized into levels of 1 Watt. Hence, the optimal policy is the one that yields the maximum reward evaluation $Q(s,a)$ for each action a. In other words, the agent selects the policy with the highest rewards which correspond to highest utility obtained from each state (consumption level). After having learned the individual household consumption, the agent can adjust the EV charging, so that the total individual welfare obtained (from household consumption and EV charging) is maximal.

$$Q(s,a) = r(s,a) + \gamma \cdot v^*(\delta(s,a)) \quad (8)$$

The function $v^*(\cdot)$ represents the discounted cumulative reward achieved by the policy starting from state s. The function $\delta(\cdot)$ is the one that determines the next state that the agent should proceed, i.e $s_{i+1} = \delta(s,a)$.

Therefore, the optimal policy is summarized as (Watkins and Dayan 1992):

$$\pi^* = argmax_a Q(s,a) \quad (9)$$

where $\gamma \in [0,1]$ is the discount factor and practically expresses the weight of the previous state rewards. For our experiments we selected some representative values for gamma (e.g $\gamma = 0.7$), and we aim to explore the effect of the whole value range of this parameter in our future steps.

Markov Decision Process Representation

Based on the proposed Q-learning mechanism the customer agent has to decide on the customer's individual consumption value, based on training on previous consumption entries. More formally, the customer agents' decision making problem is outlined by the following Markov Decision Process (MDP) (Puterman 1994) representation.

$$\begin{array}{ll} State\ Space & S = \{S_0, ..., S_i\} \\ Action\ Space & A = \{a_0, ..., a_i\} \\ Probabilities & Pr(S_0, S_1) = \\ & Pr(S_{i+1} = S_1 | S_i = S_0, a_i = a_0) \\ Rewards & r_i = \{r_0, ..., r_i\} \end{array} \quad (10)$$

where $i \in [1, N]$ and N is the size of the horizon (minutes) over which we want to learn the consumption. The states here represent the energy consumption range, discretized at the level of 1 Watt. The learning rewards are related to the utility function and give lower utility when the state has consumption farther from the exact amount that the consumer needs. This choice is based on the fact that neither less nor more energy would give the appropriate policy. In summary, the learning algorithm learns the consumption trend through iterating over the states which represent consumption levels

Data Description

Regarding the energy consumption data, we use the household consumption data from University of California Irvine (UCI) machine learning repository (Frank and Asuncion 2010). This data set[5] includes detailed consumption per minute for a whole household. The measurements are gathered between December 2006 and November 2010 (47 months). The collected data come from France, thus we argue that represent the average European household consumption behavior. However, the data selection is not restrictive for the applicability of the algorithm in different consumption data patterns (e.g US energy consumption behavior). With regard to algorithm's training, we train it on data coming from one day (24h, 1440 observations) (randomly picked). This means that the algorithm each time is trained in random 24 hour time frame and then tested to the rest of the data set. We chose the 24h timeframe, since the periodicity of the learned policy is 24 hours. This training is repeated multiple times to achieve well trained algorithm. Part of our future work is to examine how much training is it needed to have a good learning performance without having over-fitting.

Smart Charging

After having learned the household consumption, the agents have to schedule EV charging in the optimal way, with respect to individual welfare maximization. For each timeslot, t (1 min time period) the customer agent calculates the amount that needs to be charged based on (11).

$$x_t^* = argmax_{x_t} W(x, \omega, \lambda) \quad (11)$$

subject to the constraint $lb_t < x_t < ub_t$.

The lower bound for daily charging, lb_t equals to the minimum capacity needed to cover the daily driving needs, as they result from the statistical customer model. The upper bound ub_t represents the maximum power that the customer agent can charge from the network per timeslot t. This represents the main network constraint and is dependent on the voltage the current characteristics of the residential connection. The selection of this double constraint lies on the fact that the agents must not violate the customer's comfort and have the EV always charged. Furthermore, the agents need to support network stability, therefore we decided on this particular upper bound.

Using (1)-(7), (11) becomes:

$$x_{c,t}^* = argmax_{x_{c,t}} \{ \frac{(1-\kappa)}{1 + e^{(-x_{c,t} + \lambda)}} - (x_{h,t} + x_{c,t}) \cdot P \} \quad (12)$$

which is a non-linear constrained maximization problem.

Table 3 presents a general formalization of Smart Charging Algorithm, as used by the customer agents.

[5]http://archive.ics.uci.edu/ml/machine-learning-databases/00235/

Table 3: Smart Charging Algorithm Formalization.

Algorithm:	**Smart Charging**
1	Initialization
2	Calculate optimal policy as
	$\pi^* = argmax_a Q(s, a)$
3	**for each timeslot** t
4	Learn household consumption $x_{h,t}$
5	Calculate charging amount as
	$x^*_{c,t} = argmax_{x_{c,t}}\{A\}$
with	$A = \{\frac{(1-\kappa)}{1+e^{(-x_{c,t}+\lambda)}} - (x_{h,t} + x_{c,t}) \cdot P\}$
6	**end for**
7	return $x^*_{c,t}$

Results

Using the designed customer model, we simulate populations of customers that own EVs. We assume that the energy customers are represented by an agent responsible for their EV charging. This agent uses the proposed Smart Charging based on RL. End goal is to examine how this smart charging algorithm affects the individual welfare and the general consumption behavior. We compare the results with customer agents that decide about EV charging based on purely behavioral characteristics, resulting from the designed EV customer model. Finally, we assume real time pricing, and more specifically we use the EEX price-trend over 24h horizons.

To gain a deeper understanding of the presented algorithm we present some representative results using parameters $\omega = 1$ (we will experiment with the whole range of $\omega \in [0, 1]$ in coming subsection) and $\alpha = 0.5$ (commonly used in literature (Fahrioglu and Alvarado 2000), more experimentation with this parameter is part of our future work). Furthermore we use $\lambda = 8$ since it gives customer satisfaction for charging at around 50% of the nominal battery capacity. The minimum allowed State of Charge (SoC) for an EV is 20%, so $\lambda = 8$ represents the average customer. Next we use $\kappa = 0.5$ assuming equal split of the utility to household consumption utility and EV charging utility. End goal of this research project is to examine the effect of all the parameter across their whole range. At this level we present a general framework and some indicative results with some fixed parameters, assuming they represent the average consumer. Our results are based on the double bottom line: *individual* and *society*. Thus, we show the effect of the Smart Charging adoption to the individual energy expenditures and also to peak and price reduction for all the individuals in the market (societal perspective).

Energy Expenditures Reduction

To test the effect of the proposed Smart Charging on individuals, we examine the energy expenditures with and without the use of our algorithm. We observe that Smart Charging reduces significantly the average daily demand coming from charging by 23.5%.

In parallel, the expenditure reduction dependent on the parameter λ is shown in Figure 4. We chose to do sensitivity analysis for this parameter since it expresses the level of satisfaction that the customer receives from energy consumption, and thus it is indirectly related to the expenditures. This illustration shows that the expenditure reduction reaches a maximum for $\lambda = 6$ which corresponds to reduction of 30%. Then the reduction is stable for increasing values of λ. This illustration shows that there is no linear trend in the expenditure reduction and gives us the maximum of our algorithm performance. Also, it indicates that customers with $\lambda > 6$ get no extra expenditures reduction. This means that customers who are risk averse (higher λ) they do not get expenditure reduction. This is explained by the fact that the more risk averse a customer is the more he/she charges the EV battery, meaning that he/she consumes more and does not get benefits from Smart Charging which practically schedules the charging so that the individual is benefited from low price time intervals.

It becomes clear that the proposed algorithm is rather effective in terms of benefiting the individuals. A maximum of 30% decrease at the expenses for energy consumption is strong incentive for the customers to adopt Smart Charging for their EVs. In addition, it offers extra benefits to the energy policy makers since in supports stability and reduces the volatility within the network. In the following section we present analysis regarding this volatility and more specifically the peak reduction (*peak clipping*).

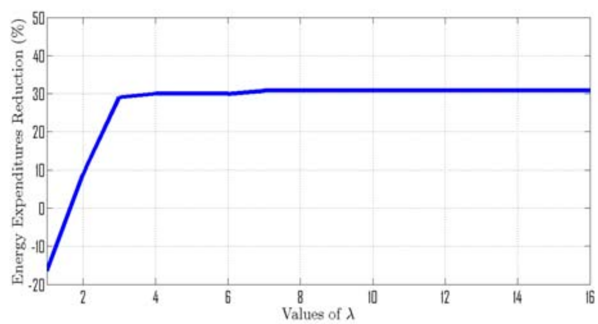

Figure 4: Energy Expenditures Reduction as a function of λ.

Energy Peak Reduction

In this section we illustrate the energy peak reduction as a function of the Smart Charging adoption in a population where all the customers own an EV and use either uncontrolled charging (results direction from the EV customer model) or Smart Charging. Specifically, we compare the Smart Charging approach with the Uncontrolled Charging, which results from the customers behavior. This practically means that the customer charges whenever he/she comes back from work and charges the EV according to the current state of charge. This summarized on Table 4. Herex is charging availability vector ($i \in [0, 23]$), $E\{c_i\}$ is expected capacity for driving, $\mathbf{D}$ – total demand vector, $\mathbf{d_h}$, $\mathbf{d_c}$ are the household and charging demand, respectively. We observe that the peak reduction is linear with maximum almost 25% of peak reduction. That means that there are

Table 4: Uncontrolled Charging Decision Algorithm.

Algorithm:	**Uncontrolled Charging**
1	Initialization
2	**for** i=0:23
	Calculate x_i
	Calculate $E\{c_i\}$
	endfor
3	**if** $x_i == TRUE$ & $SoC_i < E\{c_i\}$
	$D_i = d_{h,i} + d_{c,i}$
	endif
4	return D

strong incentives from the energy policy maker side to encourage the adoption of smart charging, since it will reduce the heavy loads from the network (Figure 5). The linear trend is straightforward, since the more agents move from uncontrolled charging to Smart Charging, the more the demand peaks are reduced. The maximum of almost 25% is comparable to many Demand Side Management practices that are used in the energy domain to shape the peaks towards a smoother demand curve. Part of our future work is to calibrate our algorithm so that achieves the desired demand shaping, given particular input parameters. Finally, this peak reduction leads to changes in the energy prices. The reason is that the high energy prices are mostly resulting from peaks in the energy demand.

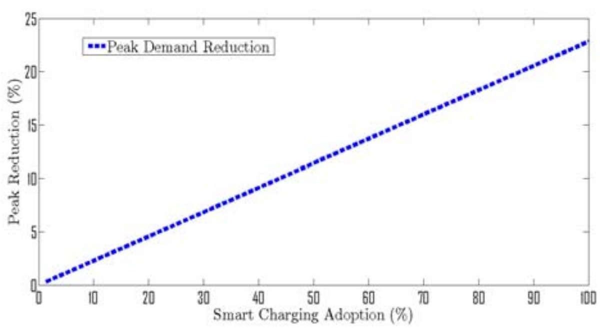

Figure 5: Aggregate Peak Reduction as a function of Smart Charging Adoption.

Energy Price Reduction

To examine the effect of the proposed Smart Charging to the average prices (societal persecutive), we simulate populations of individual customers that use the proposed algorithm and aggregate the consumption. This aggregation is done by an intermediary party or *broker* (Ketter, Peters, and Collins 2013). Making use of the $Price = f(Demand)$ relationship obtained by European Energy Exchange (EEX) (Figure 6), we present the average prices after the use of the Smart Charging algorithm (Figures 7 and 8) within a population of 10 million energy customers EV owners. Those two graphs outline a clear increase in the energy savings and price reduction as the Smart Charging Adoption increases.

In Figure 8 we observe that our algorithm leads to overall energy price reduction up to 16%. The same trend is observed in Figure 7, where we present the savings for each individual in week time period. In the latter case, the individuals may be either Smart Charging adopters or not. This means, that when there are Smart Charging adopters in the market, side benefits arise for the non-adopters, as well (as a result of the overall price reduction). Detailed analysis of this effect is presented in the following subsection.

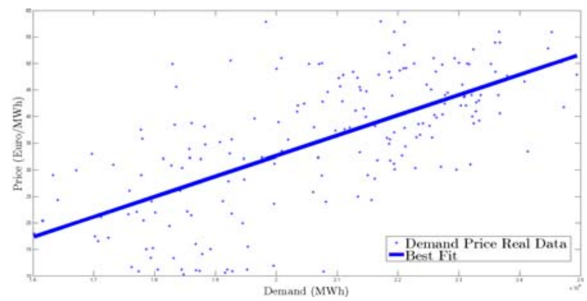

Figure 6: Price Demand relationship (source: EEX).

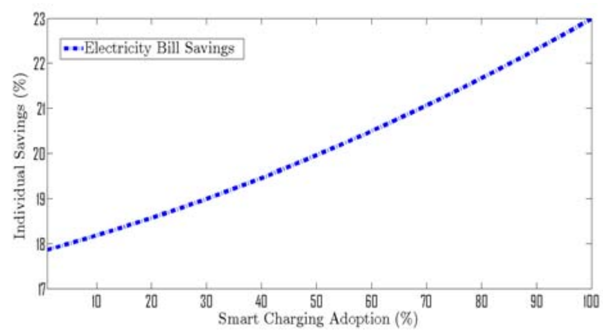

Figure 7: Individual savings as a function of Smart Charging Adoption.

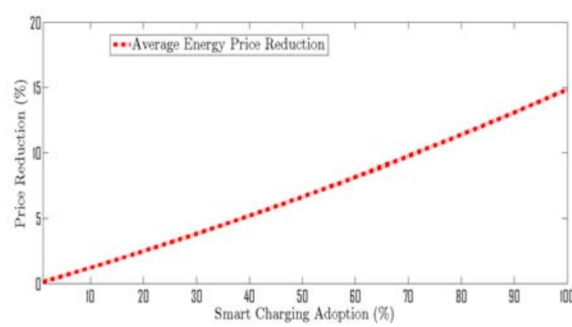

Figure 8: Average Price Reduction as a function of Smart Charging Adoption.

Savings for Smart Charging non-adopters and adopters

As explained previously, the presence of Smart Charging adopters in the market, brings average price reduction on an aggregate level. This reduction is spread in the market equally among energy customers. Therefore, there are benefits for non adopters as well. As depicted in Figure 9, the savings for non-adopters are in fact larger than the Smart Charging adopters. This results from the relative difference between the charging demand. In the case of non Smart Charging adopters the demand is higher in absolute terms, and when the prices are decreased, the benefit is higher. In contrast the Smart Charging adopters have already lower energy demand in absolute terms and thus the price decrease does not yield very large benefits. However, the Smart Charging Adopters pay less for energy compared to non-adopters, even though they receive lower benefits. Therefore, there is clear incentive for the EV owners to adopt Smart Charging. The linearity in the graphs comes from the linearity of Smart Charging adoption among the population. It is only influenced by the percentage of the population that adopts Smart Charging.

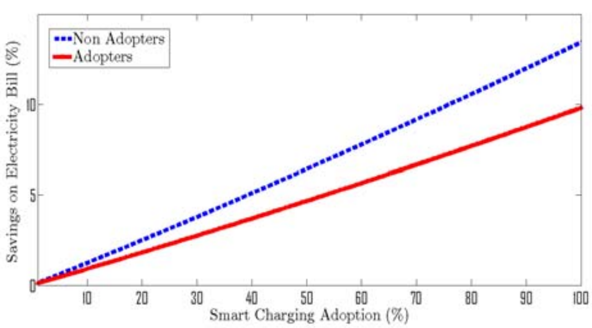

Figure 9: Savings for Smart Charging non-adopters and adopters.

Conclusions & Future Work

Electric Vehicles are undoubtedly one important part of the Smart Grid. If they are properly integrated in the grid, they may yield significant benefits for the network and the energy users. However, there is present the danger of uncontrolled use of EVs that may easily lead to energy debacles, due to ”spikes” in the energy demand. Thus, we propose our Smart Charging Algorithm to mitigate this danger and enhance a stable energy network. At the same time we achieve price reduction and individual savings on the electricity bill.

We designed a novel charging algorithm for EV customers based on individual welfare maximization objective. We propose a statistical customer model to simulate EV customers behavior and use Reinforcement Learning to learn the individual household consumption. Combining those components, we achieved reduction in the customers energy expenditures. Additionally, we proved that the proposed smart charging algorithm leads to peak reduction on an aggregate level that supports Grid stability. Finally, we show that the average energy prices are reduced for all customers in the market (social welfare improvement) with the use of smart charging against the uncontrolled charging.

In our future work we aim to increase the learning algorithm's accuracy and perform sensitivity analysis for the various parameters. Furthermore, we plan to include the V2G concept and examine the effect on individual and social welfare. Finally, we will incorporate the urban driving behavior in the customer model and take into account the electrification of transport.

References

Acha, S.; Van Dam, K. H.; Keirstead, J.; and Shah, N. 2011. Integrated modelling of agent-based electric vehicles into optimal power flow studies. In *21st International Conference on Electricity Distribution, Frankfurt*, 6–9.

Bichler, M.; Gupta, A.; and Ketter, W. 2010. Designing smart markets. *Information Systems Research* 21(4):688–699.

Christoph, B. 1998. The synthesis of bottom-up and top-down in energy policy modeling. *Energy Economics* 20(3):233 – 248.

Department of Energy, U. 2011. One million electric vehicles by 2015.

Fahrioglu, M., and Alvarado, F. L. 2000. Designing incentive compatible contracts for effective demand management. *Power Systems, IEEE Transactions on* 15(4):1255–1260.

Frank, A., and Asuncion, A. 2010. UCI machine learning repository.

Franke, T.; Neumann, I.; Bühler, F.; Cocron, P.; and Krems, J. 2011. Experiencing range in an electric vehicle: Understanding psychological barriers. *Applied Psychology*.

Gerding, E.; Robu, V.; Stein, S.; Parkes, D.; Rogers, A.; and Jennings, N. 2011. Online mechanism design for electric vehicle charging. In *The Tenth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2011)*, 811–818.

Kahlen, M.; Valogianni, K.; Ketter, W.; and van Dalen, J. 2012. A profitable business model for electric vehicle fleet owners. In *IEEE Conference on Smart Grid - Technologies, Economics, and Policies*, 1–5. Nuremberg: IEEE.

Kempton, W., and Letendre, S. 1997. Electric vehicles as a new power source for electric utilities. *Transportation Research Part D: Transport and Environment* 2(3):157–175.

Ketter, W.; Collins, J.; Reddy, P.; and de Weerdt, M. 2012. The 2012 Power Trading Agent Competition. Technical Report ERS-2012-010-LIS, RSM Erasmus University, Rotterdam, The Netherlands.

Ketter, W.; Peters, M.; and Collins, J. 2013. Autonomous agents in future energy markets: The 2012 power trading agent competition. In *Association for the Advancement of Artificial Intelligence (AAAI) Conference*, Forthcoming.

Lopes, J.; Soares, F.; Almeida, P.; and da Silva, M. 2009. Smart charging strategies for electric vehicles: Enhancing grid performance and maximizing the use of variable renewable energy resources. In *EVS24-The 24th International Battery, Hybrid and Fuel Cell Electric Vehicle Symposium & Exhibition, Stavanger, Norway*.

Mas-Colell, A.; Whinston, M.; and Green, J. 1995. *Microeconomic Theory*. Oxford University Press New York.

Mitchell, T. M. 1997. *Machine learning. WCB*. McGraw-Hill Boston, MA:.

Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. New York, NY, USA: John Wiley & Sons, Inc., 1st edition.

Valogianni, K.; Ketter, W.; de Weerdt, M.; and Collins, J. 2012. Analysis of smart grid balancing using realistic customer models. In *Conference on Information Systems and Technology*, 1–24.

Vytelingum, P.; Voice, T. D.; Ramchurn, S. D.; Rogers, A.; and Jennings, N. R. 2010. Agent-based micro-storage management for the smart grid. In *Proceedings of 9th International Conference on Autonomous Agents and Multi Agent Systems*, 10–14.

Watkins, C. J., and Dayan, P. 1992. Q-learning. *Machine learning* 8(3-4):279–292.