

Effective Management of Electric Vehicle Storage Using Smart Charging

Konstantina Valogianni and Wolfgang Ketter

Erasmus University Rotterdam
{kvalogianni,wketter}@rsm.nl

John Collins

University of Minnesota
jcollins@cs.umn.edu

Dmitry Zhdanov

University of Connecticut
Dmitry.Zhdanov@business.uconn.edu

Abstract

The growing Electric Vehicles' (EVs) popularity among commuters creates new challenges for the smart grid. The most important of them is the uncoordinated EV charging that substantially increases the energy demand peaks, putting the smart grid under constant strain. In order to cope with these peaks the grid needs extra infrastructure, a costly solution. We propose an Adaptive Management of EV Storage (AMEVS) algorithm, implemented through a learning agent that acts on behalf of individual EV owners and schedules EV charging over a weekly horizon. It accounts for individual preferences so that mobility service is not violated but also individual benefit is maximized. We observe that it reshapes the energy demand making it less volatile so that fewer resources are needed to cover peaks. It assumes Vehicle-to-Grid discharging when the customer has excess capacity. Our agent uses Reinforcement Learning trained on real world data to learn individual household consumption behavior and to schedule EV charging. Unlike previous work, AMEVS is a fully distributed approach. We show that AMEVS achieves significant reshaping of the energy demand curve and peak reduction, which is correlated with customer preferences regarding perceived utility of energy availability. Additionally, we show that the average and peak energy prices are reduced as a result of smarter energy use.

Introduction

Energy markets are moving towards a new decentralized structure, where renewable sources and storage facilities have significant penetration and adoption by energy customers. This new decentralized formation is known as *smart grid* (Amin and Wollenberg 2005). Smart grid's major challenges are the increasing fuel prices (according to EU reports we had a quintupling of oil prices between 2002 and 2010 (Malby 2013)) and the need to become independent from nuclear power. Both of the challenges can be addressed with large scale integration of renewable sources that are less carbon intensive and provide sustainable solutions. All these changes require the energy customer an active participant in the energy production and consumption process. Now, the customer can produce energy (e.g. with rooftop photovoltaic panels) and consume part of this energy locally, offering the rest to the grid. This new role introduces new decisions for

the energy customers. Electric Vehicles (EVs), especially because of their energy storage capability, are valuable tools towards a sustainable solution. They can directly use wind and solar energy and substantially reduce the amount of primary energy used for transportation, since they are much more efficient than internal combustion vehicles. Additionally, when EVs are plugged in, their batteries can offset the volatility of wind and solar production. Massive EV integration in the Energy Grid is anticipated by the main players in the energy policy landscape (US Dept. of Energy (Department of Energy 2011), European Union (Commission and others 2011)). The uncoordinated use of EVs, though, will lead to high demand peaks during charging periods. Specifically, considering customers range anxiety (feeling that the battery capacity will not be sufficient (Franke et al. 2011)), this charging may threaten the grid's stability. Therefore, smart charging algorithms are needed to alleviate this strain.

We propose an adaptive smart charging algorithm, "Adaptive Management of Electric Vehicle Storage" (AMEVS), that adjusts the EV charging based on individual customer's utility. We use *Reinforcement Learning* (RL) to learn individual consumption behavior and schedule charging with respect to individual benefit maximization objective. Furthermore, we design statistical customer models to simulate the diverse EV customer behavior. In order to model the customers we use a *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. In our future work we plan to integrate AMEVS within Power TAC (Ketter, Peters, and Collins 2013), a realistic simulation of the Smart Grid. This will allow us to evaluate our algorithm on realistic conditions and enhance applicability on real world.

Related Work

Customer modeling within the Smart Grid's context has been addressed in the literature (Reddy and Veloso 2011; 2012) aiming to provide useful insights about energy customers' behavior. Apart from individual customer models, energy customer cooperatives have been examined under the prism of computational sustainability. In (Akasiadis and Chalkiadakis 2013) the authors propose an algorithm for

creating customer cooperatives so that peaks in energy demand are reduced and the grid is benefited by this formation. In a similar fashion (Veit et al. 2013) propose a coordination mechanism for energy cooperatives achieving balanced energy consumption and consequently reducing peak demand.

EV charging attracts significant attention in the research community, since correct charging coordination can reduce peak loads on the Smart Grid and support sustainability. In (Vandael, Holvoet, and Deconinck 2013) the authors aggregate EV customer profiles with main objective to coordinate their charging. The charging coordination is fully performed by an aggregator (EV fleet aggregator). Furthermore, in (Gerding et al. 2013) the authors propose a two sided market approach to allocate charging timeslots among the EV customers and avoid charging congestion. In (Gerding et al. 2011) the authors present an online auction mechanism where the owners of EVs state their timeslots available for charging and also bid for power. In (Stein et al. 2012) the authors describe an online mechanism with pre-commitment for coordinating the EV charging. In (Valogianni, Ketter, and Collins 2014) an EV charging strategy is analyzed without assuming V2G capabilities and without formalizing the customers' preferences. Finally, in (Vandael et al. 2013) a three-step approach to coordinate EV charging is presented, being scalable and achieving demand shifting.

All previous works coordinate the charging from the point of view of an external party (charging coordinator) whereas in the current work we present a fully distributed approach without any external coordinator. We propose an algorithm (AMEVS) implemented through customer agent, that taking the stand point of the individual, schedules charging and discharging so that preferences are satisfied. We show, extending our previous work (Valogianni, Ketter, and Collins 2014), that if each individual adopts AMEVS instead of uncontrolled charging and discharging, not only the customer has savings, but also the grid is benefited in terms of *peak-to-average ratio* (PAR) and energy price reduction.

EV Customer Agents

Each customer agent consists of an *input module*, a *learning module* and an *optimization module*. The *input module* gets as inputs all the individual characteristics of each customer (gender, profession, daily routine etc.) as well as the customer's driving profile and the utility function. The *learning module* is responsible for learning the individual household consumption using RL. Finally the *optimization module* taking inputs from the other two modules, optimizes the charging and discharging ensuring the customer's mobility service at each point in time the customer may need the EV.

Driving Profiles To simulate large diverse populations of EV customers we create a statistical model reflecting the driving patterns of a whole population. We base our design on statistical commuting data coming from the Dutch Statistics Office (CBS)¹. The population is divided according to gender and the social groups: part-time employees, full-time employees, students and pupils, unemployed and retired per-

sons. For each group there are different activities together with the distance needed per day and per activity. Having determined the activities related to each group, we create driving profiles corresponding to the distance that each customer drives per day (assuming average driving speed). Additionally, we determine the customer's EV type and consequently the respective storage capacity (Equations (1) and (2)).

$$E\{Dist_t\} = LT(G_k, H_l) \quad (1)$$

where $Dist_t$ is the cumulative distance driven up to the timeslot t and G_k, H_l denote social group and activities and are drawn from the discrete set $\mathbf{G} \in \{part-time\ employees, full-time\ employees, students, unemployed, retired\}$ and $\mathbf{H} \in \{work, shopping, business\ trips, visits, leisure\ activities, school\}$ (each group has respective probabilities for each activity). LT is a look-up table function that has as inputs the social group and the activity and as output the expected average distance.

$$C_t = C_{t-1} - E\{Dist_t\} \cdot \rho \quad (2)$$

where C_t is the battery's state of charge up to timeslot t , and ρ is the capacity/distance rate given by the automotive companies.

We assume customers that own purely electric cars like Nissan Leaf and Tesla. With regard to the customer's charging and discharging availability we assume that the customers can charge the 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." The described EV statistical model is depicted in Figure 1. The model's output is the charging demand at each point in time (*timeslot*) according to the inputs given. The charging must be done within the *charging envelope* shown in Figure 2. It displays the feasible charging region bounded by the minimum and maximum state of charge. The charging should end at the time that the customer indicates that wants to use the car. Therefore if the EV charges at nominal charging rate, the battery will fill up to a certain capacity lower or equal to the maximum SoC (depending on the start and end time).

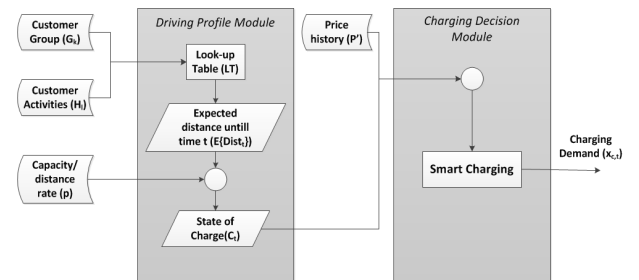


Figure 1: Electric Vehicle Statistical Model.

Individual Utility and Consumer Benefit After having the driving profile, the customer agent has to define the individual utility function and consequently the utility the customer gets from energy consumption. Assuming that the

¹www.cbs.nl




Figure 2: Electric Vehicle *Charging Envelope*.

total consumption consists of the two components: $x_{h,t}$, household demand (KWh) and $x_{c,t}$ demand from EV charging (KWh), we have the total utility $U(x_{h,t}, x_{c,t})$. We assume that in order to obtain a positive value, a customer must have non-zero energy consumption both for EV and household purpose. In the results section we will experiment with various utility function families. According to consumer theory (Mas-Colell, Whinston, and Green 1995), the individual consumer benefit is defined as in (3).

$$W(x_{h,t}, x_{c,t}) = U(x_{h,t}, x_{c,t}) - (x_{h,t} + x_{c,t}) \cdot P_t \quad (3)$$

where P_t stands for the price per energy consumption unit (€/KWh). We assume real time pricing, and as an example we use the EPEX spot price-trend over 24h horizons (Figure 3). This price curve can be substituted by any variable pricing tariff scheme.




Figure 3: Price Demand relationship (source: EPEX SPOT).

Consumption Learning Policy The *learning module* uses RL to learn the customers' energy consumption pattern. RL is based on a *reward mechanism* that provides the algorithm with positive and negative rewards for optimal or non-optimal decisions, respectively. In this particular problem 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):

$$\begin{aligned} \text{State Space} \quad \mathbf{S} &= \mathbf{T} \times \mathbf{L} = \{S_{t,j}\} \\ \text{Action Space} \quad \mathbf{A} &= \{A_i\} \\ \text{Rewards} \quad \mathbf{R} &= \mathbf{T} \times \mathbf{W} - \mathbf{T} \times \mathbf{E} = \{w_{t,j} - \varepsilon_{t,j}\} \end{aligned} \quad (4)$$

$\mathbf{T} = \{t\}$ is the set with the time intervals where $t \in [1, N]$ and N is the size of the horizon over which we want to learn

the consumption and $j \in [1, M]$ with M the maximum consumption level. \mathbf{L} is the set with the consumption levels discretized at the level of 1 KWh. The learning rewards are the individual welfare that the customer gets at each state reduced by a penalty factor $\varepsilon_{t,j}$ linearly dependent on the deviation from the actual consumption level per minute $L_{t,j}$. \mathbf{W} is the discrete set with the welfare per consumption level $L_{t,j}$ and \mathbf{E} is the set of the deviations from the actual consumption level. As $w_{t,j}$ we denote the welfare of each consumption level $L_{t,j}$, i.e. $w_{t,j} = W(L_{t,j}, x_{c,t})$. The Action Space, \mathbf{A} includes all the transitions from $S_{t,j}$ to $S_{m,n}$ under the hard constraint that $m > t$ (temporal constraint).

More specifically the *learning module* uses Q-learning as described by (5) and (6) (Mitchell 1997; Sutton and Barto 1998). The customer agent has to learn the individual household consumption pattern through rewards $R(S, A)$ that are offered to it for each (S,A). Each reward, $R(S, A)$, expresses the welfare for each state reduced by a penalty factor in sub optimal choices. After having learned the individual household consumption, the agent can adjust the EV charging, so that the total individual benefit obtained (from household consumption and EV charging) is maximum. The optimal valuation $Q(S, A)$ of each state is summarized as (Watkins and Dayan 1992):

$$Q(S, A) = R(S, A) + \gamma \cdot v^*(\delta(S, A)) \quad (5)$$

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. $S_{t+1} = \delta(S, A)$. The optimal evaluation of the states gives the learned household consumption (\mathbf{x}_h^* is a vector over temporal dimension):

$$\mathbf{x}_h^* = \text{argmax}_A \{Q(S, A)\} \quad (6)$$

where $\gamma \in [0, 1]$ is the discount factor and practically expresses the weight of the previous state rewards.

Regarding the energy consumption data, we use the household consumption data from the Netherlands obtained in collaboration with a European Utility Company. This data set includes detailed consumption per 15 minutes aggregated in hourly intervals, for 24 different households. The measurements are gathered in 2010. Using this profile pool, we draw randomly household patterns to create 10^7 customers.

Adaptive Management of EV storage Taking as inputs the learned consumption and the individual driving and behavioral characteristics, the *optimization module* schedules EV charging with respect to maximum individual welfare. For each time horizon $N = 168h$ (one week) the customer agent calculates the charging vector \mathbf{x}_c^* based on (7).

$$\mathbf{x}_c^* = \text{argmax}_{x_{c,t}} \sum_{t=1}^N W(x_{h,t}^*, x_{c,t}) \quad (7)$$

subject to the constraints (8), (9), (10):

$$-ub_t \leq x_{c,t} \leq ub_t \quad \forall t = 1, \dots, N \quad (8)$$

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 characteristics of the residential connection.

$$x_{c,t} = C_t - C_{t-1} + E\{Dist_t\} \cdot \rho \quad \forall t = 1, \dots, N \quad (9)$$

$$C_0 = SoC_{min} \quad (10)$$

where C_t is the state of charge on timeslot t , and ρ is the capacity/distance rate given by the automotive industry specifications and SoC_{min} is the minimum allowed state of charge that does not destroy the battery's lifetime. This double constraint ensures that the agents do 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. The variable $x_{h,t}^*$ stands for the household consumption and comes from the Reinforcement Learning part. The prices \hat{P}_t at each hour are predicted by the intelligent agent using a moving window of the 7 past days, averaged over each hour respectively. Table 1 presents a general formalization of AMEVS algorithm.

Table 1: Adaptive Management of EV Storage - AMEVS

AMEVS	
1	Initialization
2	for each state $(S, A) \in \mathbf{S} \times \mathbf{A}$
3	Update state valuation as: $Q(S, A) = R(S, A) + \gamma \cdot v^*(\delta(S, A))$
4	end for
5	Calculate household demand: $\mathbf{x}_h^* = \text{argmax}_A \{Q(S, A)\}$
6	Calculate optimal charging vector as $\mathbf{x}_c^* = \text{argmax}_{x_{c,t}} \sum_{t=1}^N W(x_{h,t}^*, x_{c,t})$
7	return \mathbf{x}_c^*

Experimental Evaluation

We evaluate AMEVS in different populations and examine its effect on the individual demand curve but also on an aggregate level of peak demand and price reduction. We see that the adoption of AMEVS by all the customer agents leads to peak demand and price reduction on the market. This means that AMEVS achieves an implicit coordination of charging without the presence of an actual coordinator. Further, we examine how AMEVS influences the EV charging landscape as a function of the EV ownership penetration.

Simulation Environment

Our experimental setting consists of diverse EV customer populations (see Numerical Results) whose household consumption comes from data provided by a European Energy Utility. The dataset includes 15 minute household consumption information from the Netherlands, aggregated in hourly intervals. The customers driving profiles come from the statistical model trained on Dutch mobility data (Dutch Statistics Office (CBS)). Furthermore, we assume variable pricing on the market and use the example of EPEX SPOT prices, using price as a proxy for energy availability. However, all the models can be trained on different data sets (e.g. US mobility, pricing data) to examine effects on different populations. Basic assumption is that the EV customers interact

with the energy market through an energy provider (aggregator) (Peters et al. 2013) and buy energy from the market to cover both their household and their EV charging an aggregate level. Our simulation environment is based on Power TAC (Ketter, Collins, and Reddy 2013) and smart markets in general (Bichler, Gupta, and Ketter 2010).

Benchmark Algorithms

First we describe the benchmarks we use to evaluate AMEVS performance. We assume that the alternative to AMEVS is the Uncontrolled Charging resulting from the customer's behavior. In other words when the customer is available for charging and the battery capacity is not enough, he/she plugs the EV for charging (Table 2).

Table 2: Charging Benchmark 1

Uncontrolled Charging - UC	
1	Initialization
2	for t:N
3	Calculate $CA_t, E\{Dist_t\}$
4	if $CA_t == TRUE$ & $C_t < E\{Dist_t\} \cdot \rho$
5	$D_t = x_{h,t} + x_{c,t}$
6	endif
7	endfor
8	return \mathbf{D}

Here CA_t is charging availability vector ($t \in [1, N]$), $E\{Dist_t\} \cdot \rho$ is the expected capacity needed for driving up to timeslot t , \mathbf{D} – total demand vector, $\mathbf{x}_h, \mathbf{x}_c$ are the household and charging demand vectors over time, respectively.

We use as second benchmark a Heuristic Charging approach. Here the agent using a simple moving average model, predicts the prices over a time horizon consisting of N timeslots. Assuming \hat{P}_t stands for the energy price estimate per KWh and $x_{c,t}$ for the charging demand at the timeslot t , the customer agent acts based on the following heuristic: if $\hat{P}_t \leq \hat{P}_{t+1}$ charge the pre-scheduled amount, resulting from the behavioral model, otherwise split the charging demand (i.e. the respective charging time) evenly to the time horizon N . More formally, if $x_{c,t}^{max}$ stands from the maximum amount can be pulled from the network per timeslot, the Heuristic Benchmark is shown in Table 3. More specifically, here we use the myopic approach of this heuristic with $N=2$ to compare AMEVS with a totally myopic benchmark.

Numerical Results

We will examine the performance of AMEVS compared to the other Charging Benchmarks with respect to ability to reduce energy peaks and energy prices (peak and average).

Energy Consumption Utility Due to lack of actual utility functions for energy consumption, we experiment with some of the most commonly used utility function families and compare the results. These utility functions are approximations to help us derive concrete results. Firstly, we assume that the customer shows linear utility towards energy consumption ($x_{h,t}$) and EV charging consumption ($x_{c,t}$) both

Table 3: Charging Benchmark 2

Heuristic Charging - HC	
1	Initialization
2	for t:N
3	Calculate $CA_t, E\{Dist_t\}$
4	if $\widehat{P}_t \leq \widehat{P}_{t+1}$ & $C_t < E\{Dist_t\} \cdot \rho$
5	$x_{c,t} = x_{c,t}^{max}$
6	else
7	$x_{c,t} = \frac{x_{c,t}^{max}}{N}$ and $x_{c,t+N-1} = \frac{x_{c,t}^{max}}{N} + x_{c,t+N-1}^{max}$
8	endif
9	$D_t = x_{h,t} + x_{c,t}$
10	endfor
11	return \mathbf{D}

getting different weights depending on the customer but also depending on time (Equation 11).

$$U(x_{h,t}, x_{c,t}, \beta_t, \eta_t) = \beta_t \cdot x_{h,t} + \eta_t \cdot x_{c,t} \quad (11)$$

where $\beta_t + \eta_t = 1$. Using (3), (7), (11) and $\beta_t = 1 - \eta_t$ the charging vector becomes:

$$\mathbf{x}_c^* = \underset{x_{c,t}}{\operatorname{argmax}} \sum_{t=1}^N \{ (1 - \eta_t - \widehat{P}_t) \cdot x_{h,t}^* + (\eta_t - \widehat{P}_t) \cdot x_{c,t} \} \quad (12)$$

subject to constraints(8)-(10). We will call this variation of the algorithm LinAMEVS. Secondly, instead of linear relationship between utility and energy consumption, we assume quadratic:

$$U(x_{h,t}, x_{c,t}, \omega) = \begin{cases} \omega \cdot (x_{h,t} + x_{c,t}) - \frac{\alpha}{2} \cdot (x_{h,t} + x_{c,t})^2 & , 0 \leq (x_{h,t} + x_{c,t}) \leq \frac{\omega}{\alpha} \\ \frac{\omega^2}{2 \cdot \alpha} & , (x_{h,t} + x_{c,t}) > \frac{\omega}{\alpha} \end{cases} \quad (13)$$

where ω , stands for the *level of satisfaction obtained by the user as a function of its energy consumption* and varies among customers. The variable α is a predefined parameter (e.g. 0.5 used in (Fahrioglu and Alvarado 2000)). And the charging vector using (3), (7) and (13) now becomes:

$$\mathbf{x}_c^* = \underset{x_{c,t}}{\operatorname{argmax}} \sum_{t=1}^N \{ (\omega - \alpha \cdot x_{h,t}^* - \widehat{P}_t) \cdot x_{c,t} - \frac{\alpha}{2} \cdot x_{c,t}^2 \} \quad (14)$$

We call this variation of the algorithm QuadAMEVS. For our experiments we create populations with all the possible parameter combinations in Equations (11) and (13) so that all possible cases are represented in our sample.

Energy Peak Reduction Experiments showed that AMEVS reshapes the individual demand. Therefore, we examine its effect on peak demand in a population of 10^7 customer agents (all EV owners). The average individual demand curve is reshaped as shown in Figure 4. Here we observe that the customers adopting QuadAMEVS tend to consume more when the prices are low (e.g. 0 am-5am) and sell energy back to the grid when prices

are high (e.g. 4pm-8pm) (assuming the same price for buying and selling energy back to the grid). As a result, the individual curve is reshaped having less peaks and less volatility. Furthermore, we observe that using QuadAMEVS incurs peak demand reduction not only compared with the Uncontrolled Charging but also compared with the case where no EVs exist on the Smart Grid. This is attributed to the storage features of EVs that buffers part of the demand to lower demand periods, flattening the demand curve. On the contrary, LinAMEVS even though shifts the demand peaks to earlier timeslots, it creates the same volatility. Consequently, after some time applying LinAMEVS in the market, the energy and price peaks have been shifted to earlier time periods. Finally, MixedAMEVS assumes both customers with linear and quadratic utility with regards to energy. It shows less volatility compared to LinAMEVS reducing the peaks by a small amount, but compared to QuadAMEVS still performs worse. In Figure 5 we compare QuadAMEVS with Uncontrolled Charging not only based on the average steady state result, but also with regards to variability of results and the worst case scenario (outlined with the error bars in Figure 5). We observe that even the worst case scenario of QuadAMEVS flattens the demand curve compared to Uncontrolled Charging.




Figure 4: Individual Demand Curve: No EV charging, AMEVS and Uncontrolled Charging.




Figure 5: Individual Demand Curve Variability: QuadAMEVS and Uncontrolled Charging.

To provide a more complete comparison (Table 4) we use the peak-to-average ratio (PAR) reduction ($PAR = \frac{x_{peak}}{x_{rms}} = \frac{x_{peak}}{\sqrt{\frac{1}{N} \sum_{t=1}^N x_t^2}}$), the peak demand reduction and the demand volatility ($\frac{x_t - x_{t-1}}{x_t}$) reduction. PAR is also known as crest factor and indicates how extreme the peaks in a waveform are. PAR reduction is important because much of the cost of energy supply is driven by peak demand. The

demand volatility is rather important for Smart Grids stability, since a less volatile demand, requires less peak load resources. In other words, the demand can be covered by base load which is more sustainable and does not need extra resources for peak coverage.

Table 4: Energy Peak Reduction

	PAR red.(%)	Peak red.(%)	Volatility red.(%)
LinAMEVS vs. UC	-1.6	-25.5	-24.0
QuadAMEVS vs. UC	9.0	12.5	66.5
MixedAMEVS vs. UC	5.9	0.3	-2.1
HC vs. UC	0.3%	1.0	11.0
LinAMEVS vs. no EVs	2.2	-34.0	-17.1
QuadAMEVS vs. no EVs	11.3	6.9	75.2
MixedAMEVS vs. no EVs	9.4	-6.7	-1.2

Energy Price Reduction An immediate result of the previous figure is that apart from the peak demand, also the average energy price is reduced. Consequently, this price reduction is diffused in the market because of the demand shift and peak reduction. In Figure 6 we show this reduction for a scenario with values in the whole spectrum of ω , showing a maximum of 20% at 100% QuadAMEVS adoption (against UC). Also we show that above 40% and 75% penetration of LinAMEVS and MixedAMEVS respectively, there is a negative effect on the grid. So assuming a population with consumers with linear function penetration higher than 40% should not be encouraged. Similarly, in a mixed population an EV penetration above 75% necessitates the construction of additional infrastructure to accommodate the new increased peaks.




Figure 6: Average Price Reduction in a mixed sensitivity population (AMEVS vs. Uncontrolled Charging).

Regarding the effect on the prices we consider apart from the average price reduction, the peak price reduction and the price volatility reduction as shown in Table 5. From Tables 4 and 5 we conclude that LinAMEVS does not yield any benefits to the individuals and the market. This comes from its linear behavior which is more driven by price changes. On the other hand QuadAMEVS reshapes the demand curve reducing the peaks and the prices. This results from the customer's decreasing utility for each extra unit of energy he/she consumes. Intuitively the quadratic behavior is more realistic since naturally customers consume until one saturation point above which they get no extra utility.

Furthermore, quadratic approximation of energy consumption has been used in the literature (Samadi et al. 2010; Fahrioglu and Alvarado 2000; Hall and Mishkin 1982). Finally, linear utility function is not a good approximation of customer behavior since it represents very limited range of behavior. In order to be realistic we must have a functional form of utility that allows for a wide range of possible solutions, which is not attainable with a linear utility on an interval. Therefore, QuadAMEVS is better approximation of actual behavior. It is interesting, though, that the shape of the utility function for individual customers should have such dramatic welfare effects on the whole grid. Therefore, a fully distributed approach is suitable for this problem.

Table 5: Energy Price Reduction

	Avg. price red.(%)	Peak price red.(%)	Price volat. red.(%)
LinAMEVS vs. UC	-99.5	-96.1	-79.1
QuadAMEVS vs. UC	19.6	32.8	74.2
MixedAMEVS vs. UC	-13.9	1.0	2.0
HC vs. UC	1.5	1.7	11.2
LinAMEVS vs. no EVs	-160.4	-140.1	-55.9
QuadAMEVS vs. no EVs	9.0	6.0	88.7
MixedAMEVS vs. no EVs	-50.1	-21.4	-1.5

Another interesting result of the proposed distributed approach is that it does not encourage herding behavior or conflicts of usage. Generally, in a large population the probability of destructive collisions is small, assuming that not every agent responds immediately and in the same way to every change in price. Even if all the individuals had identical behavior (worst case scenario, Figure 5), the peaks would still be lower than the ones created by uncoordinated charging. Supporting evidence is that since our system is using real-time prices, simultaneous increase in demand from many users would lead to spot price increase. Since we generally observe price reductions, it indicates that the system is getting more balanced rather than becoming unbalanced.

Conclusions & Future Work

Electric Vehicles are undoubtedly one important part of the Smart Grid. If they are properly integrated in the market, they may yield significant benefits for the network and the energy users. However, the uncontrolled use of EVs may lead to energy debacles, due to spikes in the energy demand. Thus, we propose an Adaptive Management of EV Storage (AMEVS) algorithm to mitigate the negative influence on social benefit and enhance the robustness and reliability of the grid. The average energy prices are reduced for all customers in the market with the use of AMEVS against the Uncontrolled Charging. Consequently, using AMEVS EVs support grid's sustainability as peaks are significantly mitigated. Moreover, the unnecessary charging stemming from customers range anxiety is significantly reduced promoting EVs adoption. In the future we are planning to examine the effect of AMEVS on micro-grids and smart neighborhoods and explore how the consumer social network (smart neighborhood) influences the individual consumption behavior.

References

- Akasiadis, C., and Chalkiadakis, G. 2013. Agent cooperatives for effective power consumption shifting. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI-13)*.
- Amin, M., and Wollenberg, B. 2005. Toward a smart grid: Power delivery for the 21st century. *IEEE Power & Energy Magazine* 3(5):34–41.
- 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.
- Commission, E., et al. 2011. Roadmap to a single european transport area—towards a competitive and resource efficient transport system (white paper). *Brussels (COM (2011) 144 final)*. Available online at <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do>.
- Department of Energy, U. 2011. One million electric vehicles by 2015.
- Fabrioglu, M., and Alvarado, F. L. 2000. Designing incentive compatible contracts for effective demand management. *IEEE Transactions on Power Systems* 15(4):1255–1260.
- 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.
- Gerding, E.; Stein, S.; Robu, V.; Zhao, D.; and Jennings, N. R. 2013. Two-sided online markets for electric vehicle charging. *The Twelfth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2013)* 989–996.
- Hall, R. E., and Mishkin, F. S. 1982. The sensitivity of consumption to transitory income: estimates from panel data on households.
- Ketter, W.; Collins, J.; and Reddy, P. 2013. Power TAC: A competitive economic simulation of the smart grid. *Energy Economics* 39:262–270.
- 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 Proceedings*, 1298–1304.
- Maltby, T. 2013. European union energy policy integration: A case of european commission policy entrepreneurship and increasing supranationalism. *Energy Policy* 55:435–444.
- 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.
- Peters, M.; Ketter, W.; Saar-Tsechansky, M.; and Collins, J. E. 2013. A reinforcement learning approach to autonomous decision-making in smart electricity markets. *Machine Learning* 92:5–39.
- Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. New York, NY, USA: John Wiley & Sons, Inc., 1st edition.
- Reddy, P., and Veloso, M. 2011. Learned Behaviors of Multiple Autonomous Agents in Smart Grid Markets. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence (AAAI-11)*.
- Reddy, P., and Veloso, M. 2012. Factored Models for Multiscale Decision Making in Smart Grid Customers. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI-12)*.
- Samadi, P.; Mohsenian-Rad, A.; Schober, R.; Wong, V. W.; and Jatskevich, J. 2010. Optimal real-time pricing algorithm based on utility maximization for smart grid. In *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, 415–420. IEEE.
- Stein, S.; Gerding, E.; Robu, V.; and Jennings, N. R. 2012. A model-based online mechanism with pre-commitment and its application to electric vehicle charging. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, 669–676. International Foundation for Autonomous Agents and Multiagent Systems.
- Sutton, R. S., and Barto, A. G. 1998. *Reinforcement Learning: An introduction*, volume 116. Cambridge Univ Press.
- Valogianni, K.; Ketter, W.; de Weerd, M.; and Collins, J. 2012. Analysis of smart grid balancing using realistic customer models. In *Conference on Information Systems and Technology*, 1–24.
- Valogianni, K.; Ketter, W.; and Collins, J. 2014. Learning to schedule electric vehicle charging given individual customer preferences. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS14)*.
- Vandael, S.; Claessens, B.; Hommelberg, M.; Holvoet, T.; and Deconinck, G. 2013. A scalable three-step approach for demand side management of plug-in hybrid vehicles. *IEEE Transactions on Smart Grid* 4(2):720 – 728.
- Vandael, S.; Holvoet, T.; and Deconinck, G. 2013. Aggregate demand models for electric vehicles in a smart electricity grid. In *AAMAS Workshop on Agent Technologies and Energy Systems (ATES)*, 1–4.
- Veit, A.; Xu, Y.; Zheng, R.; Chakraborty, N.; and Sycara, K. 2013. Multiagent coordination for energy consumption scheduling in consumer cooperatives. In *Proceedings of 27th AAAI Conference on Artificial Intelligence*.
- Watkins, C. J., and Dayan, P. 1992. Q-learning. *Machine learning* 8(3-4):279–292.