A Multiagent Approach to Variable-Rate Electric Vehicle Charging Coordination

Konstantina Valogianni
Erasmus University Rotterdam
The Netherlands
kvalogianni@rsm.nl

Wolfgang Ketter
Erasmus University Rotterdam
The Netherlands
wketter@rsm.nl

John Collins
University of Minnesota
Minneapolis, MN, USA
jcollins@cs.umn.edu

ABSTRACT

We present a coordination mechanism for variable-rate electric vehicle (EV) charging that combines the benefits of a decentralized decision making approach with a top-down control mechanism based on price functions. This combination allows independent decisions by self-interested EV agents, while ensuring that aggregate power demand from EV charging converges to a desired profile as determined by a control agent. This profile can yield reduced demand volatility, or a better match between the output of renewable energy sources and overall demand. We observe that price signals from the control agent are sufficient to motivate self-interested EV agents to adjust charging rate and produce the desired profile, even given a range of individual user preferences. We show that our hybrid coordination mechanism prevents *herding* in EV charging, which is typical in populations where all agents receive the same price signals and make similar charging decisions. Specifically, the control agent learns the responses of the EV agents in order to adjust its price signals, and eventually converges to coordinated charging, thereby producing the desired demand profile. We compare our approach with various benchmarks and show that EV charging congestion (*herding*) is reduced, while peaks and volatility of demand are mitigated.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

Keywords

Electric Vehicles; Hybrid Charging Coordination; Multiagent Systems; Optimization; Trading Agents; Smart Grid

1. INTRODUCTION

Electric vehicles (EVs) have the potential to significantly improve the energy efficiency and reduce the carbon intensity of our transportation system [6]. The hose at a filling station delivers energy from a local storage tank to a vehicle fuel tank at a rate of over 10 MW, while EV chargers draw energy from the shared electricity grid, typically at a maximum rate of 25 kW. But 25 kW is about half the total capacity of the electric service in most U.S. homes, far higher than the power draw of any common household device. Current electricity grids are not designed to support the load of large numbers of EVs charging their batteries during the early evening hours [21], at the same time electricity demand peaks due to energy-intensive household activities like cooking and cleaning [7]. What is needed is a way to coordinate the charging of large numbers of EVs in a way that minimizes stress on the grid, and perhaps makes the best use of available renewable energy.

EV charging coordination can be either centralized (top-down) or decentralized (bottom-up). The benefits of top-down coordination mechanisms are that they easily satisfy the constraints imposed by the coordinator (e.g. smart grid manager), leading to a balanced system. However, there are significant shortcomings in these type of approaches. The most important challenge is that often the coordinator must intervene and exogenously control the EV battery, violating the EV driver's comfort. Also, sometimes in auction based approaches [10], it becomes hard to practically implement such a mechanism, because EV charging relates to instantaneous decisions that cannot be handled properly by an auction mechanism. Specifically, it becomes difficult for EV agents to bid in every time instant for charging power and wait until the market is cleared to get the power allocated to them, because they are driving. Bottom-up approaches on the other hand, have as major benefit that customers' individual comfort is not violated and the agents have the freedom to schedule their EV charging based on their individual preferences. However, the main disadvantage is that since the same price signals are provided to all customer agents, the EV charging schedules coincide. Specifically, since all agents are cost minimizers, they tend to shift power demand to the cheaper time instants, creating new peaks.

We propose a multiagent method that aligns the objectives of smart grid managers or energy retailers with the objectives of EV owners. We are especially interested in using market-based mechanisms for coordination, because they support distributed decision making among self-interested agents. Therefore, we have designed a hybrid pricing mechanism to achieve charging coordination through the use of price functions. This hybrid pricing scheme combines the features of a decentralized approach with the top-down features that a smart grid operator needs in order to manage grid stability and achieve a desirable match between energy production and consumption. Our approach combines two types of agents:

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*

1. Energy providers or smart grid managers who want to minimize capacity investment by redistributing peak demand or who want to shape demand over time to follow the generation profile of renewable sources.

2. EV owners who receive price signals and modify their EV charging activities to satisfy their individual preferences, including cost minimization and risk (of running out battery [3]) reduction.

The key to our approach is the use of price *functions*, specifically prices that vary with charging rate, rather than simple price values. We show that simple time-varying prices, in the absence of other top-down control mechanisms, lead to herding behavior among self-interested EV agents. On the other hand, if prices vary not only by time but also by rate (measured in kW), then self-interested EV agents will adjust their charging rates over time to minimize their costs, allowing the price-setting agent to shape the overall demand profile of the EV agent population.

The remainder of this paper is organized as follows. First, we present relevant literature, addressing EV charging coordination. Secondly, we describe our algorithm and show how the decision processes of the two types of agents interplay in a multiagent simulation. Later on, we outline the assumptions of our multiagent simulation together with the data used to build it. Furthermore, we present the effect of our algorithm on the energy peak demand and we compare its performance with other commonly used benchmarks. Finally, we conclude by providing a summary of our results and describing future steps.

2. RELATED WORK

EV charging, if not controlled properly, is anticipated to bring extreme peak load on the electricity grid [15] and put the infrastructure under critical stress. Therefore, significant research has been dedicated to the EV charging coordination challenge. The related work can be divided into centralized (*top-down*) and decentralized (*bottom-up*) coordination mechanisms.

Kahlen et al. [8] present a centralized mechanism managed by a fleet operator that aims to coordinate EV charging and make profits. Vandael et al. [20] describe a three-step approach to coordinate the EV charging in a top-down fashion. Gerding et al. [4] tackle the coordination problem with an online mechanism that accounts for individual preferences in the form of time availability and bids for power and schedules EV charging accordingly. Building on this work, Stein et al. [18] introduce a pre-commitment mechanism for EV charging coordination. De Craemer et al. [2] present a dual implementation for shifting EV charging based on a central auctioneer. Kwak et al. [14] present a top-down coordination framework in the context of a household, where different appliances' functionality can be shifted. All these approaches have the potential to achieve balance on the grid but most of the times they do not satisfy individual comfort and require direct control, which might not be easy to implement in practice.

Valogianni et al. [19] describe a decentralized EV charging mechanism that aims to reduce peak load. Similar bottom-up mechanisms are also implemented in [5] but in the smart home context. In these situations individual comfort is not violated but since all agents are cost minimizers, they tend to shift power demand to the cheaper time instants, creating new peaks in the power demand and thus *herding*.

Both top-down and bottom-up approaches do not address the herding issue in EV charging because they assume same signals offered to the agents for changing their behavior. Also, all these approaches price all charging speeds (slow or fast charging) in the same way or just ask for a premium in the fast charging case. So, energy policy makers do not know exactly how to price the different charging speeds. And certainly the prices should not be the same for all charging speeds because fast charging creates higher instantaneous peaks in the demand, stressing the grid infrastructure. Therefore, we propose a hybrid mechanism in which prices are a function of charging rate (kW), can mitigate herding and achieve a desired demand profile. This mechanism additionally to benefiting from its hybrid nature, provides an answer to pricing charging rates so that grid overload is reduced.

3. HYBRID COORDINATION

The proposed hybrid coordination mechanism combines distributed, independent decision making with a top-down control mechanism to shape aggregate power demand. We assume that each individual EV owner is represented by an intelligent agent responsible for EV charging, installed in the EV's charging controller. The agent interacts with the user by estimating arrival and departure preferences, risk tolerance and expected driving distances. This approach broadens the decision spectrum and overcomes bounded rationality barriers [17].

The control agent might represent a grid operator or energy portfolio manager [16]. It acts by broadcasting price signals to the EV agents and monitoring their aggregate consumption. This agent is given a desired aggregate demand profile over some time horizon, and uses a learning component to adjust price signals, adapting to the EV agent population it faces. The price function adjustment is made through a learning factor λ that varies among control agents. Our approach requires no vehicle-to-grid (V2G) [9] capability to achieve the desired demand curve, making it compatible with current grid infrastructure that does not support large scale V2G. Figure 1 provides an overview of the hybrid charging coordination mechanism.

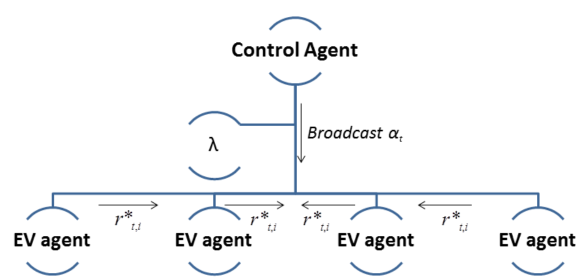

Figure 1: Hybrid coordination mechanism - multiagent implementation.

3.1 EV Driver's Agent Decision Problem

We assume the EV agents $i \in \mathbf{I}$ are self-interested (they represent their owners preferences) and wish to minimize energy cost over a time horizon T. The time horizon T is discretized to time intervals $t = 1 \dots T$. Energy cost over T is the sum of costs for each interval

$$\sum_{t=1}^{T} c_t = \sum_{t=1}^{T} e_t \cdot P_t(\cdot) \quad (1)$$

where c_t is the cost of energy during time t, e_t is the energy consumed in kWh during the interval, and $P_t(\cdot)$ is the (possibly rate-dependent) price of energy during this time. If we assume time intervals of one hour and charging at a constant rate r_t in kWh, then $e_t = r_t \cdot 1$. The decision function over T is then

$$min \sum_{t=1}^{T} r_t \cdot P_t(r_t) \quad (2)$$

subject to constraints (3)-(5):

$$0 \leq r_t \leq r_{max} \qquad \forall t \in \mathbf{T} \quad (3)$$

where r_{max} is the highest allowable rate, commonly 25 kW.

The choice of charging rate r_t by the EV agents may be influenced by the user's *range anxiety* [3]. Range anxiety is the fear that the battery's state of charge will be insufficient for unexpected driving needs. Typically, people with higher range anxiety prefer higher charging rates, allowing them to achieve a higher state of charge over a given time interval. This is an issue for EV owners due to long charging times and low density of charging facilities in most areas.

$$r_t = SoC_t - (SoC_{t-1} - E_t) \qquad \forall t \in \mathbf{T} \quad (4)$$

E_t determines how much energy the agent should charge to cover the driving needs for the next time instant t, so $E_t, \forall t \in \mathbf{T}$ accounts for the charging needs over the whole planning horizon T. This constraint ensures that the state of charge in the battery will be at least the amount satisfying the driving needs, without violating individual comfort. SoC_t indicates the battery's state of charge at each time instant t. We assume the SoC is at its minimum value at the beginning of the time period:

$$SoC_0 = SoC_{min} \quad (5)$$

Each EV agent i has also a set of preferences θ_i, including n arrival $t^n_{a,i}$ and departure times $t^n_{d,i}$over the horizon T: $\theta_i = \{t^n_{a,i}, t^n_{d,i}\}\ \forall n \in \mathbf{N}$, where $\mathbf{N}$ is the set of intervals during which the vehicle is connected to a charger based on the user's driving profile. These preferences should always be satisfied by the decision function (2) of a self-interested agent so that individual comfort is not violated. Therefore, equation (2) becomes:

$$min \sum_{n=1}^{N} \sum_{t=t^n_{a,i}}^{t^n_{d,i}} r_t \cdot P_t(r_t) \quad (6)$$

subject to constraints (3)-(5).

3.2 Smart Grid Manager's Decision Problem

The grid manager's agent (control agent) advertises prices for each time period over some time horizon to all EV agents. These prices can vary across time, and may also depend on the charging rate. Without this rate-dependent approach we observe herding, in which self-interested agents always charge at their maximum rates when price is lowest. One possible formulation is the linear function

$$P_t(r_t) = P_{0,t} + \alpha_t \cdot r_t \quad (7)$$

where r_t is the charging rate (power consumption) during timestep t and $P_{0,t}$ is the price for zero demand and can either be constant, or be exogenously determined as e.g. a) the wholesale price of electricity at time t or b) another variable price that is known ahead of time. The control agent's goal is to determine α_t at each timestep t that will produce the desired aggregate demand profile. The coefficient α_t determines the slope of the price curve with respect to charging rate (power).

To achieve the desired aggregate power demand vector $\mathbf{D}$, the control agent sets prices so that summation of power demand over the EV agents comes as close as possible to the desired demand ($\mathbf{D} \approx \sum_{i=1}^{I} \mathbf{D_i}$). Since the EV drivers' preferences are unknown to the grid manager, it is unlikely to achieve an exact match of the desired aggregate demand and the summation of individual demands (i.e $\mathbf{D} = \sum_{i=1}^{I} \mathbf{D_i}$). Therefore, in Section 3.3 we present a learning component whereby the control agent observes the outcome of its actions on the EV driver population and adjusts its future actions accordingly.

In order to estimate initial values of α_t, the control agent takes the view of an EV agent. Substituting the price function (7) in (2) we have

$$min_{r_t} \sum_{t=1}^{T} r_t \cdot (P_{0,t} + \alpha_t \cdot r_t) \quad (8)$$

which has optimal solution: $\mathbf{r}^* = [r^*_1 ... r^*_t]$ for a time horizon T. Since the solution is bounded by constraint (3) we have

$$r^*_t \leq r_{t,max} \Rightarrow \sum_{t=1}^{T} r^*_t \leq \sum_{t=1}^{T} r_{t,max} \quad (9)$$

We now show that fixed (not rate-dependent) prices lead to herding, while rate-dependent price functions can spread demand over time.

THEOREM 1. *Assume a self-interested agent population who wishes to charge EV batteries by adding an amount of energy E over a time interval T, which is divided into a sequence of discrete intervals $t \in T$. We assume that such a self-interested agent will act to first minimize its cost c, and second to acquire its desired energy E sooner rather than later. Let $c_t = r_t \cdot P_t$ be a continuous cost function over a range of charging rates $[0, r_{max}]$. If P_t is constant, $P_t = \xi_t$, where ξ_t is a constant price (in monetary units/kWh) during a given time interval t, and P'_t is an increasing function of charging rate r_t, $P'_t(r_t) = P_{0,t} + \alpha_t \cdot r_t$, then P'_t reduces the "herding" of self-interested charging agents over multiple time intervals compared to P_t.*

PROOF. For price function $P_t = \xi_t$, the cost function is $c_t = r_t \cdot \xi_t$. The optimal charging rate r^*_t for a *self-interested* agent is always either zero or equal to the maximum charging rate $r_{t,max}$, since there is no price incentive for the agents to change their charging rate. If ξ_t is constant over time, then the agent's overall cost is $c = \xi \cdot E$ regardless of when the charging takes place. Therefore all such agents will immediately charge at r_{max} for E/r_{max} time periods. If large numbers of such vehicles are connected during the same time period, this will lead to herding. If ξ_t varies over time intervals, then such an agent will acquire as much energy as possible during the lowest-cost interval, followed by the next lowest-cost interval, and so on, until it has acquired E. This is illustrated in the top portion of Figure 2.

For price $P'_t(r_t) = P_{0,t} + \alpha_t \cdot r_t$, the cost function becomes: $c_t = r_t \cdot (P_{0,t} + \alpha_t \cdot r_t)$ and the optimal charging rate $r_t^{*'}$ for a *self-interested* agent is E/T. This is illustrated in the middle portion of Figure 2. The agent can arrive at this value incrementally as follows: divide E into an arbitrary number of small increments, and add each to the time period with the lowest price. If $P_{0,t} = \hat{P}$ and $\alpha_t = \hat{\alpha}$ are fixed, then this will always be the time period with the lowest allocated charge rate. The result will be constant-rate charging at a rate of E/T over the entire interval.

Furthermore, with price function $P'_t(r_t) = P_{0,t} + \alpha_t \cdot r_t$ the optimal charging can be exogenously determined by a central operating party (grid manager or energy retailer), through adjusting α_t across time, as shown in the lower portion of Figure 2. □

In Figure 2 we show an illustrative example of Theorem 1. Assume we have 5 hours to charge, max charge rate is

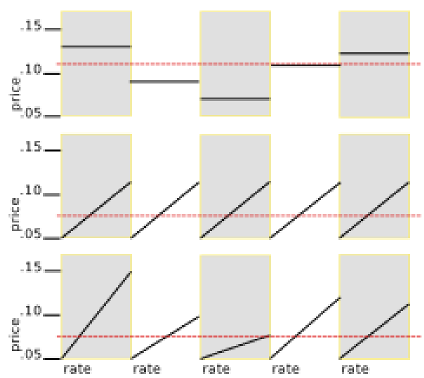

Figure 2: Illustration of Theorem 1.

10 kW, and we need 25 kWh during this charging period. With the flat-price scheme (top panel, $\alpha_t = 0\ \forall t \in \mathbf{T}$), we get 10kWh at 0.07/kWh, 10kWh at 0.09/kWh, and 5kWh at 0.11/kWh, for a total cost of 2.15. No charging is done during the first and last time periods.

With the linear price functions (bottom two panels), we set $P_{0,t} = 0.05\ \forall t \in \mathbf{T}$. The horizontal axis in each time period is in kW, from 0 to $10kW$. So if we charge at 5 kW during a period, we pay half the maximum price for the period. We can solve for the minimum cost by finding the price/kWh that gives us the total energy we need, in this case about $0.07/kWh$, for a total cost of 1.75. The charge rates in this example are (2.5, 5, 10, 3.5, 4). The total cost would be lower if the price we find is above the maximum price of one or more of the time periods.

3.3 Learning Component

The smart grid manager agent (control agent) needs to adapt to changes observed in the EV agent population, since we assume no prior knowledge related to the EV driver portfolio. Therefore, it needs to learn from observations related to EV agents' behavior and adapt the price signals accordingly, so that it achieves the desired aggregate demand profile **D**. We introduce a learning component in its decision algorithm that helps the hybrid coordination mechanism converge to the desired profile **D** without having knowledge about the EV agent population. This makes our coordination mechanism highly flexible since any potential additions of agents with different preferences or drop-outs of existing agents, can be observed online and the mechanism can adapt its behavior.

Specifically, a control agent observes and stores the deviations of the actual consumers profile and the intended profile that it wanted to achieve. Based on these observations it updates the error function over horizons T, $\sum_{t=1}^{T} \epsilon_t = \sum_{t=1}^{T} D_t - \sum_{t=1}^{T} \sum_{i=1}^{I} r_{i,t}^*$ and adjusts the value of α for the next period T based on the agent's *learning factor*, $\lambda > 0$, so that the aggregated demand profile created by the individuals approximates the intended demand profile. The learning factor λ varies across control agents and we experiment with different values in our simulation. Additionally, if $\sum_{t=1}^{T} \epsilon_t < 0$ it means that the produced aggregate result is higher because of higher charging rate of the individuals and using (7) we have to reduce charging rate $r_{i,t}^*$, and thus increase α_{t+T}: $\alpha_{t+T} = \lambda \cdot \alpha_t$. In the opposite case ($\sum_{t=1}^{T} \epsilon_t > 0$), the value of α_{t+T} needs to be decreased, so $\alpha_{t+T} = \frac{1}{\lambda} \cdot \alpha_t$. In summary, the learning component updates the next value of α_{t+T} based on the following rule:

$$\alpha_{t+T} = \begin{cases} \lambda \cdot \alpha_t & : \sum_{t=1}^{T} \epsilon_t < 0 \\ \frac{1}{\lambda} \cdot \alpha_t & : \sum_{t=1}^{T} \epsilon_t > 0 \end{cases} \tag{10}$$

This decision rule is repeated by the control agent until the error term $\sum_{t=1}^{T} \epsilon_t$ reduces to the desired error level $\sum_{t=1}^{T} \epsilon_{t,min}$. In Table 1 we summarize the proposed hybrid coordination mechanism in pseudo-code form.

Table 1: Hybrid coordination pseudo-code

Hybrid coordination

Initialization
Define desired aggregate demand profile $\mathbf{D} = \{D_t\},\ \forall\, t \in T$
Start with an initial value of parameter $\alpha_t,\ \forall\, t \in T$
Broadcast α_t to the self-interested agents $i \in I$
Observe aggregate charging rate $\sum_{i=1}^{I} r_{t,i}^*$ as calculated by each agent:
$r_{t,i}^* = argmin_{r_{t,i}} \sum_{t=1}^{T} r_{t,i} \cdot (P_{0,t} + \alpha_t \cdot r_{t,i})$
Calculate error:
$\sum_{t=1}^{T} \epsilon_t = \sum_{t=1}^{T} D_t - \sum_{t=1}^{T} \sum_{i=1}^{I} r_{t,i}^*$
while $\sum_{t=1}^{T} \epsilon_t \geq \sum_{t=1}^{T} \epsilon_{t,min}$ **do**:
  if $\sum_{t=1}^{T} \epsilon_t < 0$
    $\alpha_{t+T} = \lambda \cdot \alpha_t$
  else
    $\alpha_{t+T} = \frac{1}{\lambda} \cdot \alpha_t$
  endif
  Observe aggregate charging rate $\sum_{i=1}^{I} r_{t,i}^*$ as calculated by each agent:
  $r_{t,i}^* = argmin_{r_{t,i}} \sum_{t=1}^{T} r_{t,i} \cdot (P_{0,t} + \alpha_t \cdot r_{t,i})$
  Calculate error:
  $\sum_{t=1}^{T} \epsilon_t = \sum_{t=1}^{T} D_t - \sum_{t=1}^{T} \sum_{i=1}^{I} r_{t,i}^*$
end
return α_t

4. MULTIAGENT SIMULATION

To evaluate our coordination mechanism we create a multiagent simulation which consists of *self-interested* EV agents and a smart grid manager agent (control agent) who is responsible for keeping the aggregate demand closer to a stable level (desired aggregate profile, **D**). Our simulation environment is built according to the smart markets paradigm [1] and Power TAC's specifications [13, 11] since we aim to evaluate the mechanism within Power TAC's simulation platform [12]. Table 2 presents a summary of the notation used.

Table 2: Summary of notation

Symbol	Definition
c_t	electricity cost in time instant t
$\mathbf{D} = \{D_t\}$	desired aggregate power demand vector
E_t	estimated driving demand in time instant t
$\mathbf{I} = \{i\}$	discrete set of EV driver agents
M	upper bound of the learning factor λ
$\mathbf{N} = \{n\}$	set of intervals during which the EV is connected to a charger
$P_t(r_t)$	charging rate price in time instant t
$P_{0,t}$	constant factor of the linear price function per time instant t
r_{peak}	peak charging rate in a demand curve
r_{rms}	root mean square charging rate
r_t	charging rate per time instant t
$r_{t,max}$	maximum charging rate per time instant t
R_t	retail price of power per time instant t
SoC_t	EV battery's state of charge during time t
$t^n_{a,i}$	arrival time of agent i for activity n
$t^n_{d,i}$	arrival time of agent i for activity n
$\mathbf{T} = \{t\}$	discrete set with time instants
α_t	charging rate coefficient in function $P_t(r_t)$
ϵ_t	error factor in time instant t
$\epsilon_{t,min}$	error factor threshold in time instant t
θ_i	set of preferences for agent i
λ	learning factor in the hybrid coordination
ξ	constant value for price

4.1 Scenarios & Assumptions

In order to demonstrate the performance of the algorithm we will examine scenarios where the EV agents face prices given by $P_t(r_t) = P_{0,t} + \alpha_t \cdot r_t$. We create scenarios with both rate-independent ($\alpha_t = 0$) and linearly rate-dependent ($\alpha_t \neq 0$) prices. The constant factor of price function (7), $P_{0,t}$, may get either the average of wholesale price over a day or the corresponding retail price of each hour (R_t), representing the generation cost of this particular amount of charging power and taxes and network fees. We will use the latter option, since with this assumption price function (7) accounts for both the power generation cost, taxes and network fees and for an extra price factor $\alpha_t \cdot r_t$ which analogous to charging rate r_t. This factor can be interpreted as the premium the EV agents have to pay on top of the retail price to obtain a particular charging rate r_t. The scenarios examined are presented in Table 3. The following assumptions draw the boundaries of our simulation environment and determine our mechanism's goal:

Table 3: Simulation Scenarios.

	Attributes	
Rate-independent scenario	$\alpha_t = 0$	$P_{0,t} = R_t$
Linear scenario	$\alpha_t \neq 0$	$P_{0,t} = R_t$

- The simulation includes *self-interested* EV agents that do not exchange information among each other. They only interact with the grid (via the control agent).
- The interaction of the EV agents with the grid is limited to receiving price signals (retail prices) and decide on EV charging rate and duration, based on these prices.
- The *self-interested* EV agents have preferences regarding departure and arrival times, which are derived by the data set described in Section 4.2.
- All EV agents are located and driving under the same distribution network, to avoid procurement of charging power from other distribution networks.
- The granularity of the designed simulation is 1 hour, since typically the EV charging rate is calculated in hourly intervals.
- The planning horizon is 1 week (T=168h) because there seems to be repetition of driving habits and overall consumer behavior within weekly intervals. However, the algorithm can be adjusted to produce results for different planning horizons.
- The grid manager is in control of steering the aggregate EV charging consumption towards a desired profile which might be either a less volatile demand curve or a demand curve that follows the production pattern of a renewable production unit (e.g. wind turbine).
- Goal of the hybrid coordination mechanism is to reach to a targeted aggregate charging profile.

4.2 Data Description

In this section we present the data sets used to calibrate our simulation and evaluate the performance of the hybrid coordination mechanism. All data refers to same region in the Netherlands and is collected during 2012-2013. Our simulation environment can be calibrated with data from other areas without affecting the mechanism's functionality.

4.2.1 Individual Preferences

We bootstrap our simulation with arrival and departure preferences obtained by the Central Bureau of Statistics (CBS)[1] in the Netherlands. This data includes different population clusters (full time employees, part-time employees, students, retired persons, etc.) with a variation of habits and driving behaviors (business commuting, leisure time driving, vacation, visits to relatives, shopping etc.). For each individual we get a driving profile with certain activities and driving demand for each activity, combined with arrival and departure times. The aggregate driving demand in (kms) of our population is displayed in Figure 3.

4.2.2 Energy Prices

The results for the Rate-independent Scenario, where the prices are fixed over a time period, are produced using as an example of wholesale prices offered by the European Power Exchange (EPEX) adjusted to account for network fees, taxes and VAT for the Netherlands (44% of the retail price[2]). These prices are the values of $P_{0,t} = R_t$ in both Linear and Rate-independent Scenarios. In Figure 4 we show 3 weeks of retail price data, as it is used in the simulation.

[1] www.cbs.nl

[2] http://www.nuon.nl/energie/energieprijzen-vergelijken/opbouw-energieprijs.jsp [Date Accessed: 02/05/2014]

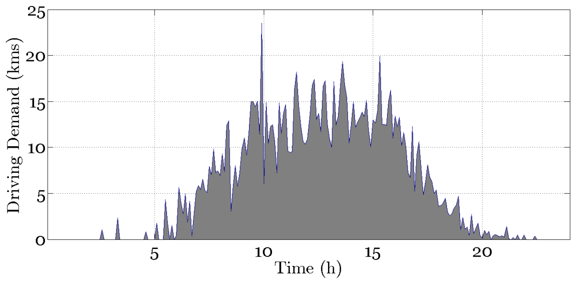

Figure 3: Aggregate driving demand of the EV agents population.

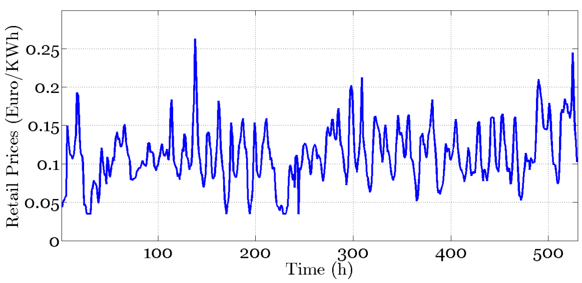

Figure 4: Retail prices (€/KWh) over 3 weeks.

4.2.3 *Learning Factor*

For the learning factor of the control agents we use values in the spectrum of $\lambda \in (0, M)$, where M is a sufficiently large number. Agents with $\lambda = 0$ are considered zero-intelligence agents that show no learning ability, whereas agents with $\lambda = M$ show the highest learning ability and are the most desirable control agents. High learning factor λ, indicates higher adaptability of the control agent to the EV agent portfolio changes and thus potentially quicker convergence to the desired profile $\mathbf{D}$. However, there is no proven direct analogy between learning factor and convergence, since the behavior of the EV agents includes stochasticity that cannot always be accounted for.

4.3 Benchmarks

The output of the hybrid coordination mechanism is the Linear Scenario, calibrated with the population's preferences. To evaluate its performance we compare it with the following benchmarks.

4.3.1 *Benchmark 1- Rate-independent Scenario*

This is the baseline scenario of our analysis since it assumes *self-interested agents* that minimize their costs based on a given variable retail price signal, which does not depend on the charging rate (Section 4.2.2). The household demand combined with the power demand of this scenario is depicted in Figure 5. We present the EV charging demand combined with the household demand, because this is the demand that the grid faces from each household. Additionally, on this graph it is more clear that EV charging fills the valleys created by the household demand during early morning hours. Firstly, here we observe that the EV agents are price sensitive and are solely driven by the high variations in prices. This makes them consume significant portion of the daily power demand during low price periods (early morning and

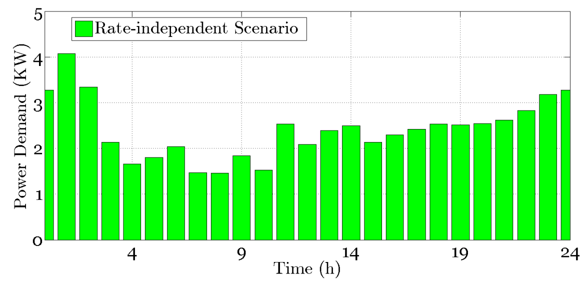

Figure 5: Rate-independent Scenario - Combination of EV charging and household demand.

late night) whereas, they mitigate EV charging when prices are high (noon and evening hours). Secondly, we observe that herding of charging is present because every agent gets the same price signals and besides small differences in preferences, the power demand of all agents coincides, creating new peaks during low price periods. This herding is exactly what our algorithm aims to prevent by adjusting the price signals and partially redistribute the peaks across the whole time horizon. This redistribution mitigates volatility of aggregate demand, which is highly beneficial for the smart grid's infrastructure.

4.3.2 *Benchmark 2- Real-world Charging*

As a second benchmark we use real-world EV charging data obtained in collaboration with EV charging infrastructure company in the Netherlands. The data set accounts for EV charging during 2012-2013 across the whole country. The steady state curve of this data, combined with the household demand is presented in Figure 6. From this graph we verify our initial assumption, that most of the people, without any control in EV charging, just plug their EV once they return home from work (around 6 pm) increasing peak demand.

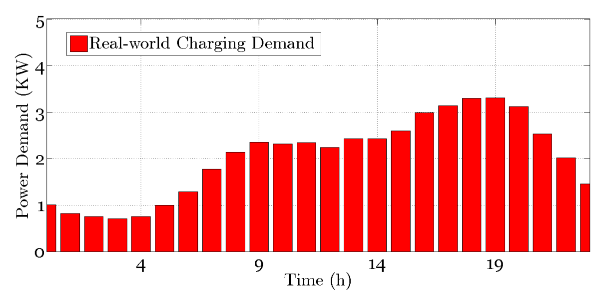

Figure 6: Combination of real-world EV charging and household demand.

4.3.3 *Benchmark 3- Aggregate Demand without EVs*

This benchmark represents the total power demand of our population, assuming that there is no EV charging involved. It is crucial to compare the performance of the algorithm with this benchmark because goal of a successful EV integration policy is not to create extra peaks on the already volatile aggregate household demand. Therefore, we want to see how close the algorithm's results are compared to this benchmark. We do not expect the algorithm to reduce peak demand since extra power demand is added, coming from EV charging. It is desirable though, to show that during peak hours the EV charging does not create higher volatility. The data for this benchmark comes from households in the Netherlands obtained in collaboration with European energy utility. In Figure 7 we show some individual power demand curves (anonymized for privacy reasons) and in Figure 8 the aggregate household demand of the population.

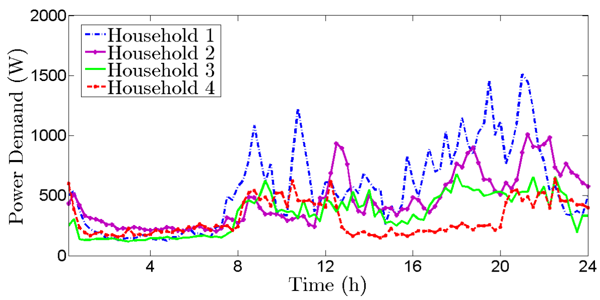

Figure 7: Typical household power demand of the EV driver agents population.

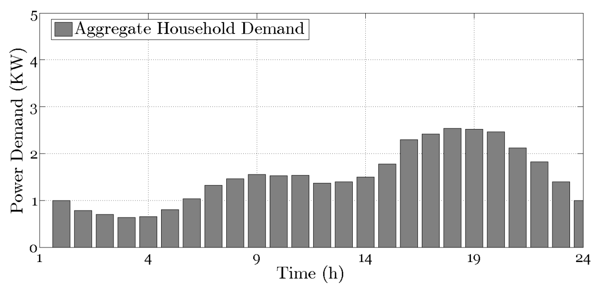

Figure 8: Aggregate household demand of the population over 4 month period.

5. NUMERICAL RESULTS

In this section we implement the hybrid coordination mechanism in the multiagent simulation described in Section 4 and present indicative performance results. We are mostly interested in the impact of the algorithm on the aggregate demand curve. Specifically, the peak demand of this curve is the determinant of installing extra capacity on the network. Therefore, the grid managers using coordination mechanisms strive to mitigate this peak demand. A second important factor is the demand's volatility. Reduced volatility of this curve protects the grid from critical strains.

5.1 Impact on Power Demand

Applying the algorithm on our EV agent population for a typical price function ($\alpha_t = 1, \forall t \in \mathbf{T}$): $P_t(r_t) = P_{0,t} + r_t$ we get the steady state power demand displayed in Figure 9. We observe that the EV charging demand is more evenly distributed on top of the household demand without having significant herding during low price periods.

Comparing the Linear Scenario with the Rate-independent Scenario and the real-world charging demand we get Figure 10. In this graph we notice first that the Rate-independent scenario shifts most of the charging during low price time

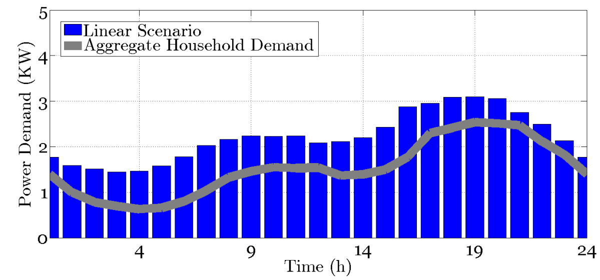

Figure 9: Power demand after applying the hybrid coordination mechanism - Linear Scenario ($\alpha_t = 1 \forall t \in \mathbf{T}$).

intervals (early in the morning or late at night), whereas during the day and specifically during high price periods, it does not charge at all. Consequently, significant herding is present because all of the *self-interested* EV agents congest to charge during the low price periods. That explains the high peak of 4KW around 2am-3am. This outcome is aligned with the results of Gottwalt et al. [5] where they use bottom-up cost minimization in the smart home context. They also observe significant herding during these time periods.

Secondly, in Figure 10 we observe that the real-world charging mostly shows up during business hours despite the high prices. This happens because the current situation in Europe allows EV drivers to plug their car in their employers premises and charge it there while working. Other EV drivers leave their EV charging the whole night to cover their range anxiety. This situation is undesirable because the daily peaks around 6pm-8pm increase even more with EV charging.

Finally, we observe that in the Linear Scenario where we put a price function on the charging rate (charging speed) the *self-interested* EV agents schedule their charging in a way that prevents extreme peaks but also covers the driving needs. This happens because increasing charging speed leads to increasing costs. To measure the impact of our

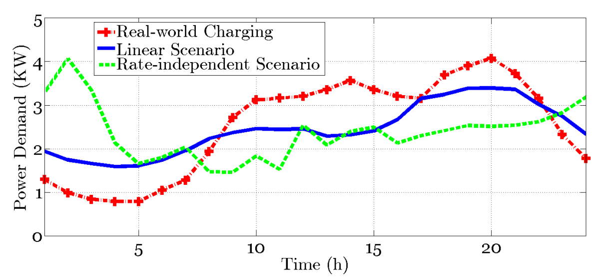

Figure 10: Comparison of Rate-independent, Linear and Real-world scenarios.

mechanism on the smart grid we use the peak-to-average ratio (PAR) metric: ($PAR = \frac{r_{peak}}{r_{rms}} = \frac{r_{peak}}{\sqrt{\frac{1}{T}\sum_{t=1}^{T} r_t^2}}$), which is also known as *crest factor* and measures the intensity of peaks or valleys in a curve. Secondly, we will measure the peak reduction incurred by our algorithm in comparison to the other benchmarks. Table 4 summarizes these metrics (negative reduction indicates increase). We observe that the Linear Scenario which uses our mechanism reduces the peak

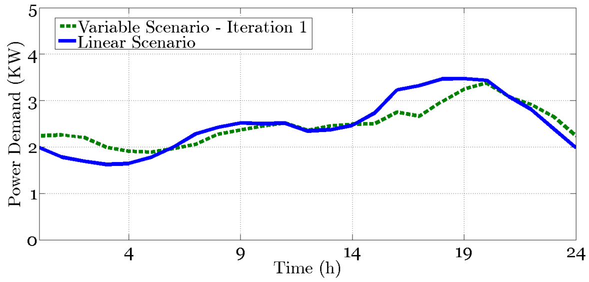

Figure 11: Comparison of Linear Scenario with constant $\alpha_t, \forall t \in \mathbf{T}$ and Variable Scenario with $\alpha_t = R_t, \forall t \in \mathbf{T}$.

demand compared to all the other benchmarks. Of course compared to the household demand it is not possible to reduce peak demand, because we add extra demand which is attributed to the EV charging. Similar are the results for the PAR reduction (volatility reduction). It is interesting to note that the Linear Scenario reduces PAR compared to the plain household demand, resulting in a less volatile curve. Therefore, there are strong incentives for the energy policy makers to introduce such kind of coordination mechanisms.

Table 4: Energy Peak and PAR Reduction

	PAR red. (%)	Peak red. (%)
Linear vs. Rate-independent	16.00	15.02
Linear vs. Real-world	9.61	16.73
Linear vs. Household	11.40	−36.40

5.2 Shaping Aggregate Power Demand

Besides the scenarios shown before where α_t has a fixed price $\forall t \in \mathbf{T}$, we present here a scenario where we give variable values to α_t. In this Variable Scenario we set $\alpha_t = R_t$. In Figure 11 we display the first iteration of the algorithm. This iteration is practically the first observation the control agent gets from the EV agent population. Depending on this observation it will adjust the α_t values to reach the desired profile $\mathbf{D}$. From Figure 11 we can also observe that by changing α_t over time the aggregate demand curve becomes smoother compared to the Linear Scenario where α_t had a constant value $\forall t \in \mathbf{T}$. After the first iteration we will see how the control agent adjusts the α_t values to reach the goal. We assume a learning factor $\lambda = 10$ since we want the algorithm to converge quickly. We set the error threshold $\epsilon_{t,min} = 0.2, \forall t \in \mathbf{T}$ since lower than this level cannot be achieved by the agents. This happens because they have as hard constraints to satisfy the EV drivers's needs and therefore, they have to deviate from the desired profile to have the battery charged for their owners. The algorithm converges after 21 iterations and in Figure 12 we show how the weekly charging demand changes after 21 iterations.

5.3 Sensitivity Analysis

Since price coefficient α_t drives the outcome, we provide some indicative results for this parameter's sensitivity. In Figure 13, we present results in the spectrum of $\alpha_t \in [1, 4]$.

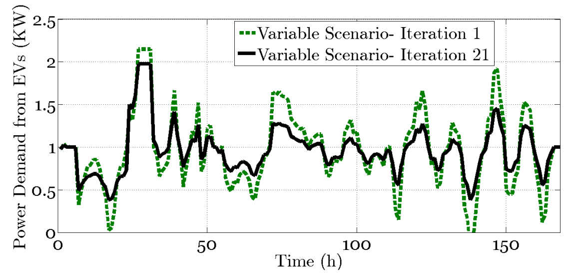

Figure 12: Demand after learning EV agents behavior - Variable Scenario with $\alpha_t = R_t, \forall t \in \mathbf{T}$.

As expected increasing α_t decreases average charging rate. The interesting result of this graph is that increasing the

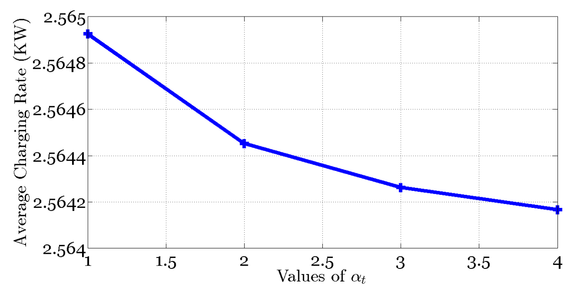

Figure 13: Sensitivity Analysis.

value of α_t by increments of 1, yields small changes in the average charging rate. Therefore, we can confirm the assumption that higher learning factors λ on the control agent's side are crucial for achieving the desired convergence.

6. CONCLUSIONS & FUTURE WORK

We presented a hybrid mechanism that coordinates EV charging. It combines the decentralized decision making on the EV agents' side with a central coordination party that ensures convergence of the aggregate EV charging to the desired (coordinated) outcome. Our mechanism is based on *price functions* for EV charging rates that create incentives for charging in low rates (low speed charging) when the prices are high and in high rates (high speed charging) when the prices are lower. The control agent does not require any prior knowledge of the EV agents portfolio to set the right prices since it learns their behavior online. Therefore, the mechanism is highly dynamic and can adjust quickly to exogenous shocks or portfolio changes. We show that the proposed mechanism prevents *herding* in EV charging, which is present in many coordination mechanisms and also distributes the EV charging demand in a way that peaks and volatility are reduced.

In future, we plan to investigate the integration of vehicle-to-grid(V2G) [9] in our mechanism. Furthermore, we plan to extend the price functions to other forms and evaluate their performance. Finally, we aim to test this mechanism in a real-world experiment using a mobile application.

7. ACKNOWLEDGEMENTS

This research is supported by the *Vereniging Trustfonds Erasmus University Rotterdam* (Ref. No 97090.16/14.0574/evt).

REFERENCES

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

[2] K. De Craemer, S. Vandael, B. Claessens, and G. Deconinck. An event-driven dual coordination mechanism for demand side management of phevs. *IEEE Transactions on Smart Grid*, 5(2):751–760, March 2014.

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

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

[5] S. Gottwalt, W. Ketter, C. Block, J. Collins, and C. Weinhardt. Demand side management - a simulation of household behavior under variable prices. *Energy Policy*, 39:8163–8174, 2011.

[6] International Energy Agency. *Global EV Outlook*. Organisation for Economic Co-operation and Development, Paris, 2013.

[7] A. Ipakchi and F. Albuyeh. Grid of the future. *IEEE Power and Energy Magazine*, 7(2):52–62, März 2009.

[8] M. Kahlen, W. Ketter, and J. van Dalen. Balancing with electric vehicles: A profitable business model. In *Proceedings of the 22nd European Conference on Information Systems*, pages 1–16, Tel Aviv, Israel, June 2014.

[9] W. Kempton, V. Udo, K. Huber, K. Komara, S. Letendre, S. Baker, D. Brunner, and N. Pearre. A test of vehicle-to-grid (v2g) for energy storage and frequency regulation in the pjm system. *Mid-Atlantic Grid Interactive Cars Consortium Jan*, 2009.

[10] W. Ketter, J. Collins, M. Gini, A. Gupta, and P. Schrater. Real-time tactical and strategic sales management for intelligent agents guided by economic regimes. *Information Systems Research*, 23(4):1263–1283, Dec. 2012.

[11] W. Ketter, J. Collins, and P. Reddy. Power TAC: A competitive economic simulation of the smart grid. *Energy Economics*, 39:262–270, 2013.

[12] W. Ketter, J. Collins, P. Reddy, and M. de Weerdt. The 2015 Power Trading Agent Competition. Technical report, RSM Erasmus University, 2015.

[13] W. Ketter, M. Peters, and J. Collins. Autonomous agents in future energy markets: The 2012 Power Trading Agent Competition. In *Association for the Advancement of Artificial Intelligence (AAAI) Conference Proceedings*, pages 1298–1304, Bellevue, WA, July 2013.

[14] J.-y. Kwak, P. Varakantham, R. Maheswaran, Y.-H. Chang, M. Tambe, B. Becerik-Gerber, and W. Wood. Tesla: an extended study of an energy-saving agent that leverages schedule flexibility. *Autonomous agents and multi-agent systems*, 28(4):605–636, 2014.

[15] A. S. Masoum, S. Deilami, P. S. Moses, and A. Abu-Siada. Impacts of battery charging rates of plug-in electric vehicle on smart grid distribution systems. In *Innovative Smart Grid Technologies Conference Europe (ISGT Europe), 2010 IEEE PES*, pages 1–6. IEEE, 2010.

[16] M. Peters, W. Ketter, M. Saar-Tsechansky, and J. E. Collins. A reinforcement learning approach to autonomous decision-making in smart electricity markets. *Machine Learning*, 92:5–39, 2013.

[17] H. A. Simon. Rational decision making in business organizations. *The American Economic Review*, 69(4):493–513, 1979.

[18] S. Stein, E. Gerding, V. Robu, and N. R. Jennings. 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*, pages 669–676. International Foundation for Autonomous Agents and Multiagent Systems, 2012.

[19] K. Valogianni, W. Ketter, J. Collins, and D. Zhdanov. Effective management of electric vehicle storage using smart charging. In *Proceedings of 28th AAAI Conference on Artificial Intelligence*, pages 472–478, 2014.

[20] S. Vandael, B. Claessens, M. Hommelberg, T. Holvoet, and G. Deconinck. A scalable three-step approach for demand side management of plug-in hybrid vehicles. *IEEE Transactions on Smart Grid*, 4(2):720 – 728, 2013.

[21] R. A. Verzijlbergh, M. O. Grond, Z. Lukszo, J. G. Slootweg, and M. D. Ilic. Network impacts and cost savings of controlled ev charging. *Smart Grid, IEEE Transactions on*, 3(3):1203–1212, 2012.