A Computational Monetary Market for Plug-In Electric Vehicle Charging

Matteo Vasirani and Sascha Ossowski

Centre for Intelligent Information Technology,
University Rey Juan Carlos,
Madrid, Spain
{matteo.vasirani,sascha.ossowski}@urjc.es

Abstract. The penetration of plug-in electric vehicles (PEV) into near-future traffic and power infrastructure is expected to be large enough to have a serious impact on the grid. If each PEV arrives at home and charges immediately, the distribution network can incur in serious problems. Therefore, the charging process of the PEVs has to be coordinated, on the basis of the grid capabilities, generation and pricing. In this paper, we put forward a computational monetary market intended as an automatic scheduler for the charging problem. The market is designed so as i) to satisfy the constraints of the distribution network, ii) to guarantee a reasonable level of fairness and allocation efficiency, while at the same time iii) to give the possibility to each PEV to transiently increase its share of the charging capacity of the local distribution network when needed.

Keywords: Monetary markets, lottery scheduling, smart grid, plug-in electric vehicles, exchange rates.

1 Introduction

The rate at which plug-in electric vehicles (PEV) enter the vehicle stock depends on many factors, including battery cost and reliability, the price of gasoline, and government incentive programs. However, the penetration of these vehicles is expected to be large enough to have a serious impact on the grid. For instance, in [5] it has been estimated that for some scenarios where two PEVs are charging at the same time on the same distribution network, the addition of a hair dryer on the same network will seriously stress it. If each PEV arrives at home and charges immediately, the distribution network can incur in serious problems, since historically it has not been designed for that kind of electricity intensive loads. Therefore, the charging process of PEVs has to be coordinated, on the basis of the grid capabilities, generation and pricing.

One way to address this problem is using time-of-use (TOU) pricing plans to dissuade PEV owners from charging at peak times, when the distribution network the PEV is plugged to is already close to the maximum capacity. However, such an approach in inherently static and rely on the response of the human

P. Coles et al. (Eds.): AMMA 2011, LNICST 80, pp. 88–99, 2012.

owners, who should adapt their behaviour and manually change the time at which their PEVs have to be charged.

A more advanced way to address the charging problem is automatic scheduling. With automatic scheduling we envision a system where PEVs are equipped with intelligent agents that are able to negotiate with some infrastructure component that controls the local distribution network (e.g., a transformer) in order to charge their batteries, while satisfying the constraints of the distribution network.

The paper is structured as follows. In Section 2, we detail a list of properties that our resource management need to satisfy. In Section 3 we discuss some related works and their limitations. Section 4 formalises the PEV charging problem. In Section 5 we propose our resource management mechanism for PEV charging, which is experimentally evaluated in Section 6. Finally, we conclude in Section 7.

2 Desiderata

In designing a mechanism for the automatic scheduling of PEV charging, we aim to satisfy the following list of properties:

Starvation-free. The scheduling mechanism must guarantee that all PEVs have a non-zero probability of charging, at least partially, their batteries.

Fairness and competitiveness. Since different PEVs have different needs, related to the quantity of electricity they demand or when the electricity is needed, the scheduling mechanism should give the PEVs the possibility of increasing their share of the distribution network capacity when needed. However, it is desirable that no agent can unlimitedly increase this share and that a fair share of the distribution network capacity is statistically guaranteed to every PEV.

Environments. Power systems are nowadays highly regulated, and in general consumers pay electricity according to fixed per unit price plans. The proposed scheduling mechanism must work in regulated environments, where electricity is paid at a fixed price, as well as in deregulated environments, where dynamic pricing [3] can be used so as to influence the consumers decision making.

3 Related Work

The problems and challenges posed by the electrification of urban transportation are quite recent. Therefore, the application of agent-based techniques as the means of tackling these challenges is still in its infancy.

In [6], an agent-based solution for the automatic scheduling of PEV charging is presented. This work allows PEV owners to specify the time at which the PEV will be available for charging, the quantity of energy that must be stored in the battery, as well as the time at which the PEV is needed to be charged.

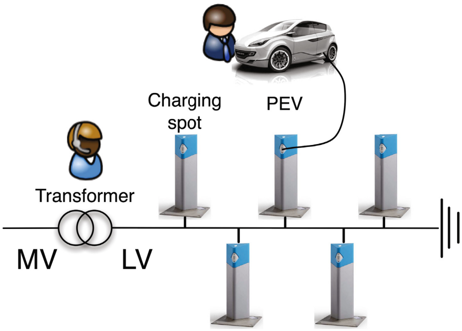

Fig. 1. Distribution network with charging spots

The PEV agent communicates this information to the agent that resides on the local transformer, which performs the automatic scheduling while satisfying the constraints of the distribution network. In general, such an approach is unable to dissuade owner to misreport information. For instance, an owner may indicate an earlier departure time or further travel distances in order to receive preferential charging. However, if all the agents are truthful, it implicitly guarantees a fair charging, since the scheduling is performed by the transformer agents that do not have preferences with respect to the PEVs.

A way to address the above shortcomings is allowing dynamic pricing [3] and defining incentive compatible market clearing mechanisms, such as in [4]. These mechanisms aim at providing PEV agent the monetary incentives to truthfully report their demand vector and the vehicle availability, regardless of the others' reports. However, this approach has two main shortcomings. The mechanism occasionally requires units of electricity to be "burned", either at the time of allocation or on departure of the agent. The former way of "burning" is more realistic but it makes the computation of the allocation hard. The latter is easier to compute, but requires the battery to be discharged on departure, which may not always be feasible in practice, since discharging is not instantaneous and the charging spot must physically impede the PEV owner to unplug the PEV. Furthermore, such an approach is strictly dependent on the possibility of applying dynamic pricing, while nowadays power systems are highly regulated and consumers typically pay electricity according to fixed per unit price plans.

4 PEV Charging

We assume a typical configuration of a local distribution network (see Fig. 1). A transformer device converts the voltage from medium to low. The low voltage feeds several charging spots (from a dozen to one hundred spots), where PEVs connect to charge their batteries. We assume that both the local distribution transformer and the PEVs are equipped with intelligent software agents.

4.1 Transformer Agent

The transformer agent of the local distribution network is in charge of activating the spots at which PEVs are connected to enable charging. As part of the infrastructure, we assume that the system designer is able to program the behaviour of this agent, as well as the scheduling mechanism that regulates which charging spot will be activated. We consider a model with discrete time steps $t \in \mathcal{T}$ with fixed duration Δ. For sake of simplicity, we assume that at each time step only one charging spot can be activated.

4.2 PEV Agent

Let $\mathcal{V} = \{1, 2, \ldots, n\}$ be a set of PEVs. For every PEV, let $c \in \mathbb{R}^+$ be the quantity of electricity demanded by the PEV, which for simplicity is equal to the capacity of the battery, $a \in \mathcal{T}$ the time slot when the PEV is plugged to the charging spot, and $d \in \mathcal{T}$ the departure time slot (i.e., the time slot at which the PEV need to have charged the demanded quantity in the battery). Given the charge rate γ, it is possible to compute the number of time steps $k = c/(\gamma \cdot \Delta)$ at which the PEV needs to be charging. The objective of each PEV is therefore having its charging spot activated for k time slots after its arrival and before its departure.

5 Resource Management for PEV Charging

The resource management mechanism we propose is inspired by *lottery scheduling* [7], a randomised resource allocation mechanism that has been developed for resource allocation in operating systems. In lottery scheduling resource rights are represented by lottery tickets, and an allocation is determined by holding a lottery. The agent with the winning ticket is granted the resource. This effectively allocates resources to competing agent in proportion to the number of tickets that they hold. Lottery tickets are denominated in different *currencies* (one for each PEV), and each currency is backed by a *base commodity*. Ultimately, lottery tickets are equivalent to monetary notes, which are issued in different denominations and are backed by a base commodity (e.g., gold).

5.1 Registration

When a PEV is plugged in the charging spot for the first time, the corresponding PEV agent registers itself with the local transformer agent. The PEV agent sends a `register` message, providing its identification number, assumed to be unique. The transformer agent confirms the registration, providing the initial monetary base m of the PEV agent which is determined as follows. Let G be an arbitrary constant integer value. This value represents the overall quantity of base commodity that backs the different currencies[1]. Let q be the number

[1] In monetary terms, G is equivalent to the quantity of gold in circulation.

of charging spots in the local distribution network. The base commodity G is *equally split* among the PEVs, so as that each PEV owns a share $g = G/q$ of the base commodity. This share funds the monetary base m of the PEV agent, according to the exchange rate r. The exchange rate determines the quantity of base commodity a monetary note is worth, according to the equation $m = r \cdot g$

In the registration phase, the initial exchange rate r is set to 1, and therefore the monetary base m is equal to g. We remark that while G and g are constant, the monetary base m and the exchange rate r change with time[2]. In particular, the monetary base can be expanded by the PEV agent, which issues the monetary notes, while the exchange rate is set by the transformer agent. The transformer agent maintains a table with the actual values of the base commodity g, the monetary base m with its currency[3], and the exchange rate r.

5.2 Lottery

As said before, we assume a model with discrete time steps $t \in \mathcal{T}$ with fixed duration Δ. This means that at each time step one of the charging spots with a connected PEV must be activated to enable charging for the whole duration of the time step. Thus, at the beginning of each time step, the transformer agent sends a proposal, contained in a `call_for_lottery` message, to all the registered PEVs. A PEV that wants to participate in the lottery replies with an `accept` message and reports the monetary base m, which determines the probability that its charging spot will be activated. Since the PEV agent control the monetary base, it is possible that the new monetary base differs from the actual monetary base that is stored by the transformer agent. In fact, a PEV may try to increase its relative probability by *inflating* its monetary base.

The transformer agent then updates the exchange rate of all the PEVs. Let $M^t = \sum_{i \in \mathcal{V}} m_i^t$ be the sum of the PEV agents' monetary base submitted at time t, G the total base commodity, $\eta \in (0, 1]$ a tunable parameter, and $H(\cdot)$ the Heaviside step function, whose value is zero for negative argument and one for positive argument. The PEV exchange rates is updated according to Eq. 1, where $\alpha(m^t, m^{t-1}) : \mathbb{R}^+ \times \mathbb{R}^+ \to (0, 1]$ is an update step function, which depends on the new monetary base m^t, submitted by the PEV agent, and the old monetary base m^{t-1}, stored by the transformer agent. The rationale of this exchange rate update scheme will be detailed in section 5.3.

$$r \longleftarrow r + \alpha(m^t, m^{t-1}) \cdot \left(\frac{m^t}{g} - r\right)^{1+\eta H\left(\frac{m^t}{g} - r\right)\left(\frac{M^t}{G} - 1\right)} \tag{1}$$

After the update of the exchange rate, also the monetary base of each PEV that is stored by the transformer agent is updated to the new value m^t. Then the lottery takes place and the transformer agent draws the winner ticket according

[2] When necessary, we will use the notation m^t and r^t respectively to express the dependence of time of these quantities.

[3] For simplicity, the currency is set to the ID of the PEV.

to Eq. 2. The ratio between the monetary base and the exchange rate gives the probability of a charging spot to be activated.

$$p(i) = \frac{\dfrac{m_i}{r_i}}{\sum\limits_{j \in \mathcal{V}} \dfrac{m_j}{r_j}} \tag{2}$$

The transformer agent probabilistically selects the winner of the lottery and notifies the participants about the outcome. After that, the transformer agent physically activates the charging spot to which the winner PEV is connected. The PEV then starts charging the battery for the entire duration of the time slot, and it is charged for the electricity that it draws from the grid according to a fixed per unit price plan.

5.3 Exchange Rate Update

As said before, the monetary base m of a PEV is related to the base commodity g through the exchange rate r. For example, let's suppose that the exchange rate is 1. In this case, the value of one monetary note equals the value of one unit of the base commodity. If a PEV agent expands its monetary base by doubling the number of notes, the value of the monetary note should halve, since the base commodity g that funds the monetary base is constant. However, if the inflationary adjustment of the exchange rate were instantaneous, an expansion of the monetary base performed by a PEV agent would not have any positive effect on the probability of having its charging spot activated: if m is doubled and r is halved, the probability of winning the lottery remains constant (see Eq. 2).

For this reason, we perform a *delayed update* of the exchange rate r, which enables the PEV agents to *transiently increase* the probability of having its charging spot selected. In this way, a PEV that needs urgent charging may try to improve the probability of actually being allowed to charge. Analysing Eq. 1, we can distinguish three cases:

1. When the actual exchange rate r is lower then m/g, it means that the exchange rate is being adjusted towards the "true" exchange rate as a reaction to the expansion of the monetary base performed by the PEV agent. In this case, the update step function α plays a central role. A value of α close to 0 considerably delays the inflationary adjustment. On the other hand, a value of α close to 1 implies an almost instantaneous inflationary adjustment, which in practice impedes the inflationary agent to actually increase its probability of being able to charge. In this work, we implement a relative update step function. With this function, the update step increases with the expansion of the monetary base (see Eq. 3). The minimum update step is set to a constant value $\beta \in (0, 1]$, while the degree of increase is determined by another constant parameter $\rho \in \mathbb{R}^+$

$$\alpha(m^t, m^{t-1}) = \begin{cases} \beta & if\ 0 < \dfrac{m^t}{m^{t-1}} \leq 1 \\ 1-(1-\beta)e^{\rho-\rho\dfrac{m^t}{m^{t-1}}} & if\ \dfrac{m^t}{m^{t-1}} > 1 \end{cases} \tag{3}$$

Additionally, when $r < m/g$ the Heaviside function $H(\cdot)$ returns 1, which implies that the adjustment that is done towards the "true" exchange rate is raised to a power exponent that depends on the ratio between the overall monetary base of all the agents (M^t) and the total base commodity (G). This exponent penalises excessive inflation, which in turn makes the monetary expansion not effective[4]. The constant parameter η implicitly determines from which extent the overall expansion of the monetary base starts to be detrimental for the inflationary agents, actually reducing, rather than increasing, their probability of winning the lottery.

2. When the actual exchange rate r is greater then the "true" exchange rate m/g, it means that the exchange rate is being reduced towards the "true" exchange rate, as a reaction to an eventual deflation of the monetary base performed by the PEV agent. In this case, $\alpha(m^t, m^{t-1})$ is constant and equal to β, and the Heaviside function returns 0, so as the adjustment is raised to the power of 1.
3. When the actual exchange rate r is equal to the "true" exchange rate m/g no adjustment takes place.

6 Experimental Evaluation

In this section, we empirically evaluate our proposed mechanism. The objective of the experiments is assessing *allocative efficiency* of the resource management mechanism, as well as the *social welfare* of the PEV agents.

The allocative efficiency (ϕ) refers to how close to satisfy the aggregated demand is the resource management mechanism. The demanded electricity is represented by the overall capacity of the PEV batteries. Given that only one PEV can charge in a single time slot, satisfying the aggregated demand is not always possible, since it depends on the number of PEVs. Let $s^{\max}$ be the maximum feasible supply, which can be computed by a centralised scheduler with full knowledge of the aggregated demand and the arrival and departure time slots[5]. Finally, let s be the supply provided by the resource management mechanism, given by the summatory over the state of charge of the PEV batteries at the end of the time window. The allocative efficiency is therefore $\phi = s/s^{\max}$.

[4] In fact, expanding the monetary base is effective *only if* a certain subset of the PEV agents acts in that way.

[5] Due to lack of space, we omit the detail of the optimal computation of $s^{\max}$.

There are manifold measures of social welfare, given some utility function u_i, to assess the overall quality of an allocation [2]. Let b_i be the battery state of charge at the end of the time window, and c_i the demanded electricity, i.e., the capacity of the battery. We define the PEV utility as $u_i = b_i/c_i$. The social welfare metrics we use are the utilitarian social welfare, σ_u, and the Nash product, σ_N (Eq. 4). The utilitarian social welfare is simply the average of the utility gained by each PEV. The Nash product is the product of the utilities of each agent. This notion of social welfare favours both increases in overall utility and inequality-reducing redistributions.

$$\sigma_u = \frac{1}{|\mathcal{V}|} \cdot \sum_{i \in \mathcal{V}} u_i \qquad\qquad \sigma_N = \prod_{i \in \mathcal{V}} u_i \tag{4}$$

6.1 Experimental Setup

In the experimental setup we assume that charging occurs within a fixed time window of 12 hours (e.g., from 8PM to 8AM of the next day), divided in time slots of $\Delta = 1$ min. Therefore, the set of time slots is $\mathcal{T} = \{1, 2, \ldots, 720\}$. The local distribution network is equipped with fast charging spots that provide a power of 40kW (400V@63A), which is a common standardised socket in European three-phase networks. A PEV is connected to each charging spot, with a battery capacity uniformly drawn from $15 - 25$kWh. This implies that each PEV needs $22.5 - 37.5$ min of charging. For each PEV, the arrival time slot a is drawn from a half-Gaussian probability distribution over the interval $[1, 720]$, with mean equal to 1 and standard deviation of 60 time slots. Similarly, the departure time slot d is drawn from a half-Gaussian probability distribution over the interval $[a, 720]$, with mean equal to 720 and standard deviation of 60 time slots. This means that 68.2% of the PEVs arrive in the first hour of the time window (e.g., from 8PM to 9PM) and departure in the last hour (e.g., from 7AM to 8AM).

As said in section 5.2, at the beginning of each time slot a PEV agent must accept or refuse to participate in the lottery. After receiving the `call_for_lottery` message, the PEV agent checks the state of charge of its battery and if the battery is not full it sends an `accept` message, which includes the reported monetary base. The real decision making of the PEV agent relies on the selection of the monetary base included in the `accept` message. In this evaluation, we conceive two different strategies: `Zero-Intelligence` and `Some-Intelligence`. With the `Zero-Intelligence` (ZI) strategy a PEV agent simply submits the monetary base that was set in the registration phase. With the `Some-Intelligence` (SI) strategy the PEV agent inflates its monetary base with a certain probability p_d:

$$p_d = \left(1 - \frac{d - t}{d}\right) \tag{5}$$

where t is the time slot at which the lottery is taking place and d is the departure time slot. The rationale of the probability p_d is that the closer the PEV is to the departure, the higher the probability of inflating the monetary base, trying to increase the likelihood of winning the lottery.

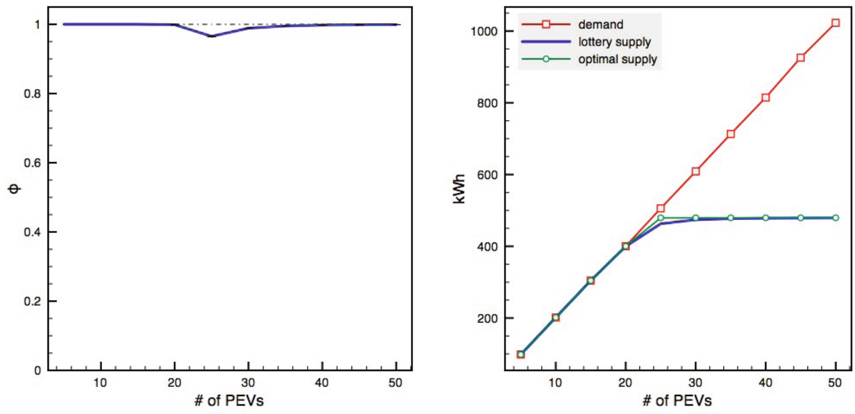

(a) Allocative efficiency (b) Demand and supply

Fig. 2. Efficiency metrics

We perform experiments for different values of β, ρ and η for the exchange rate update step function, and different compositions of the set of PEV agents, according to their strategies. For each experimental configuration, we run 100 trials and we compute average values of the metrics of interest. All error-bars denote 95% confidence intervals.

6.2 Results

Fig. 2(a) shows the allocative efficiency of the resource management mechanism. These results have been obtained for $\beta = 0.1$ and $\rho = 0.5$, although similar dynamics have been obtained with different combinations of these two values. It is noticeable how the mechanism guarantees a quite high allocative efficiency. Furthermore, as the number of PEVs increases, the allocative efficiency tends to 1, since the highest feasible supply is reached when more than 25 PEVs are connected to the distribution grid (see Fig. 2(b)). To evaluate the allocative efficiency we used a population entirely composed of ZI agents. We remark that the specific strategy followed by the PEV agent does not affect the value of the aggregated demand that is satisfied, but rather how the aggregated demand is individually satisfied among the agents.

Additional interesting insights can be derived from the analysis of the social welfare metrics. In these experiments, we are interested in determining how the utility of an allocation is distributed among the agents, depending on the strategies followed by PEV agents. Therefore, we run experiments with different percentages of SI agents of the total population of PEVs, while the rest of the agents follow the ZI strategy.

Fig. 3(a) to 3(d) plots the utilitarian social welfare for different compositions of the set of PEV agents. These results have been obtained for $\beta = 0.1$, $\rho = 0.5$

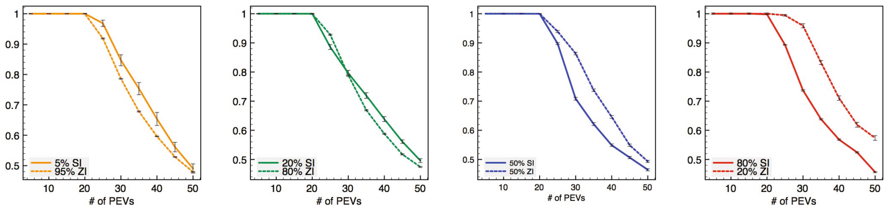

Fig. 3. Utilitarian social welfare

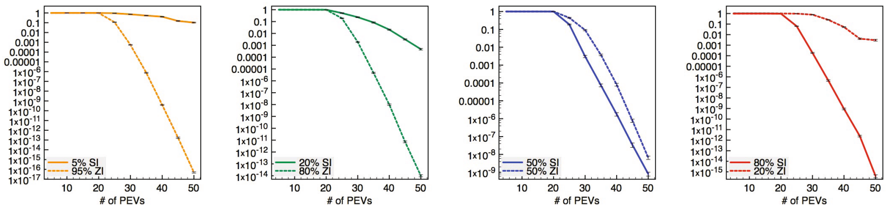

Fig. 4. Nash product

and $\eta = 10^{-27}$. It is possible to notice how the SI strategy is not beneficial when 50% of the agents use that strategy. If too many agents opt to play the SI strategy, the excessive inflation of the overall monetary base makes playing that strategy detrimental, since inflating the monetary base actually decreases the probability of having the PEV's spot selected for charging.

Fig. 4(a) to 4(d) plots the Nash product for the same experimental configuration. Again, the SI strategy does not perform well when 50% of the agents (or more) play that strategy. Therefore, given the parametrisation used in these experiments, there exists a particular splitting of the agent into two groups, ZI agents and SI agents, which represents an equilibrium condition, under which the utility gained by the SI agents must equal the utility gained by the ZI agents.

6.3 Learning to Select the Best Strategy

As highlighted by the social welfare analysis of the previous section, it is clear that PEV agents are faced with the problem of selecting a suitable strategy. Let's suppose the restricted case in which each agent must pick on of the two strategy, ZI or SI, when the PEV is plugged to the charging spot. If we consider a one-shot encounter (i.e., a single charging period, from 8PM to 8AM of the next day), it is hard for a PEV agent to speculate about the best strategy. In fact, playing the SI strategy may not be beneficial if many other PEV agents do the same, as a sort of minority game [1].

Therefore, we setup a learning scheme that is used by the agents to select, at the beginning of each charging time window, the strategy to follow. Let $\overline{u}(y)$ be the current estimate of the agent utility that is obtained if strategy $y \in \mathcal{Y}$

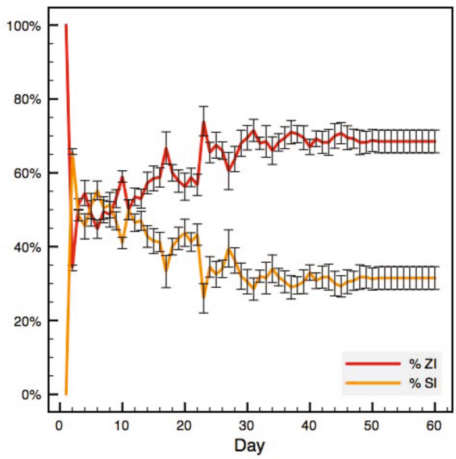

Fig. 5. Learning to pick the best strategy

is followed, where $\mathcal{Y}$ is the set of possible strategies. When the PEV is plugged to the charging spot, the PEV agent selects the charging strategy according to the Boltzmann probability distribution π (see Eq. 6), where $T \in (0, 1]$ is the "temperature" parameter used to enforce convergence over time.

$$\pi(y) = \frac{e^{\overline{u}(y)/T}}{\sum_{y' \in \mathcal{Y}} e^{\overline{u}(y')/T}} \tag{6}$$

After having selected strategy y, we assume that for the entire duration of the time window the PEV agent will adhere to this strategy. At the end of the time window, the PEV agent computes the utility that it obtained and updates the estimate using the formula:

$$\overline{u}(y) = \overline{u}(y) + \frac{1}{w+1}\left(\frac{b}{c} - \overline{u}(y)\right) \tag{7}$$

where w is the number of times the strategy y has been selected so far, b is the battery state of charge at the end of the time window, and c is the available capacity.

To evaluate the outcome of such a learning scheme, we simulate a scenario with 40 PEVs, whose strategy space is $\mathcal{Y} = \{\text{ZI,SI}\}$. At the beginning, each PEV agent is assumed to play the ZI strategy. We also assume that for each PEV agent the initial estimate is $\overline{u}(y) = 1 \; \forall y \in \mathcal{Y}$. As before, we set $\beta = 0.1$, $\rho = 0.5$ and $\eta = 10^{-27}$. The temperature T is initialised to 1 and reduced geometrically by 0.9 every time window.

The result of the experiments are shown in Fig. 5. Using the simple learning scheme described above, the PEV agents selfishly learn to pick the best strategy.

The result is that they split into two groups, with approximately 70% of them playing the ZI strategy, and the remaining 30% playing the SI strategy. This result is in line with the results of the social welfare analysis. In fact, we noticed than when 50% of agents start to play the SI strategy, the utility that those agents receive is lower than that received by the ZI agents. However, for every agent that switches from ZI to SI, there is a reduction of the utility surplus. For the particular parametrisation of the system that we used in the experiments, the splitting 30%-70% between SI agent and ZI agents represents the equilibrium condition.

7 Conclusions

The forecast penetration of plug-in electric vehicles into our cities is expected to have a serious impact on the grid. In this context, we put forward a computational monetary market for the management of the charging process of PEVs. The experimental evaluation showed how the monetary market worked as an automatic scheduler that satisfied the constraints of the distribution network, guaranteeing fairness and allocation efficiency. Furthermore, since it offers the possibility of forcing a desired equilibrium between inflationary and not inflationary agents, it preserved the feature of giving the PEV agents the possibility of transiently increase their share of the disputed resource.

Acknowledgments. This research was partially supported by the Spanish Ministry of Science and Education through the project “AT” (CONSOLIDER CSD2007-0022, INGENIO 2010) and “OVAMAH” (TIN2009-13839-C03-02).

References

1. Arthur, W.B.: Inductive reasoning and bounded rationality. American Economic Review 84(2), 406–411 (1994)
2. Chevaleyre, Y., Dunne, P.E., Endriss, U., Lang, J., Lematre, M., Maudet, N., Padget, J., Phelps, S., Rodríguez-Aguilar, J.A., Sousa, P.: Issues in multiagent resource allocation. Informatica 30(1), 3–31 (2006)
3. Faruqui, A., Hledik, R., Tsoukalis, J.: The power of dynamic pricing. The Electricity Journal 22(3), 42–56 (2009)
4. Gerding, E., Robu, V., Stein, S., Parkes, D., Rogers, A., Jennings, N.: Online mechanism design for electric vehicle charging. In: Proceedings of 10th International Joint Conference on Autonomous Agents and Multi-Agent Systems (2011)
5. Rahman, S.: Analysis of the impact of phevs on the electric power distribution network. In: Transmission and Distribution Conference and Exposition (2010)
6. Vandael, S., Boucke, N., Holvoet, T., Deconinck, G.: Decentralized demand side management of plug-in hybrid vehicles in a smart grid. In: Proceedings of 1st International Workshop on Agent Technologies for Energy Systems, pp. 67–74 (2010)
7. Waldspurger, C., Weihl, W.H.: Lottery scheduling: flexible proportional-share resource management. In: Proceedings of the 1st USENIX conference on Operating Systems Design and Implementation, pp. 1–11. USENIX Association (1994)