

An Online Mechanism for Multi-Speed Electric Vehicle Charging

Valentin Robu¹, Sebastian Stein¹, Enrico H. Gerding¹,
David C. Parkes², Alex Rogers¹, and Nicholas R. Jennings¹

¹ University of Southampton, SO17 1BJ, Southampton, UK
{vr2, ss2, eg, acr, nrj}@ecs.soton.ac.uk

² Harvard University, Cambridge, MA 02138, USA
parkes@eecs.harvard.edu

Abstract. As plug-in electric vehicles become more widespread, their charging needs to be coordinated, in order to ensure that capacity constraints are not exceeded. This is becoming particularly critical as new fast-charging technologies are being developed that place additional burden on local transformers. To address this problem, we propose a novel online mechanism in which agents representing vehicle owners are incentivised to be truthful not only about their marginal valuations for electricity units, but also about their arrival, departure and maximum charging speeds. The work extends the state of the art in several ways. We develop an online, model-free mechanism that handles multi-unit demand per period, thus accommodating vehicles with heterogeneous and flexible charging speeds; we provide competitive worst-case bounds for our mechanism; finally, we simulate the proposed online mechanism using data from a real-world trial of electric vehicles in the UK, showing that using fast charging leads to significant cost savings.

1 Introduction

Recent advances in battery technology, and pressures to reduce the carbon emissions of transport, have stimulated renewed interest in electric vehicles (EVs). New hybrid designs, equipped with both an electric motor and an internal combustion engine (that can be used to drive or charge the battery), address common worries about the limited range of such vehicles, and EVs are expected to represent close to 10% of all vehicle sales by 2020, according to a recent Gartner report [7]. However, this potential growth has generated concerns that if many of these vehicles are plugged in, and charged simultaneously, they risk overloading local electricity distribution networks (Shao et al[10]).

To address this problem, a number of researchers have begun to investigate mechanisms to schedule the charging of EVs, such that the local constraints of the distribution network are not exceeded. For example, Clement, Haesen & Driesen [1] propose a centralised scheduler which makes optimal use of the network capacity when vehicle owners truthfully report their expected future vehicle use to the system. However, such approaches fail to address the fact that owners will likely misreport this information if it is in their interest to do so (for example, reporting that they require their vehicle earlier than is actually the case to receive preferential charging). To this end, in this paper we use *online mechanism design*, in order to engineer payment mechanisms that provide

incentives for these owners to report truthfully their value for receiving electricity, their willingness to wait and their maximum charging rate.

Online mechanism design is an important topic in distributed AI and economics. For example, Parkes and Singh [8] propose an online variant of the Vickrey-Clarke-Groves (VCG) mechanism, which uses an MDP-type framework to predict future arrivals. Unlike their work, the mechanism proposed here is model-free (i.e., assumes no knowledge about the future). Model-free online settings have been considered by Porter [9] and Hajiaghayi et al. [5], who study the problem of online scheduling of a single, reusable resource over a finite time period. We extend this work by considering multi-unit domains, with preferences described by a non-increasing vector of values.

A different approach for dynamic problems is considered by Juda and Parkes [6]. They consider a mechanism in which agents are allocated options (a right to buy) for the goods, instead of the goods themselves. The concept of options is promising, but would need modifications to apply to our setting with perishable electricity units. In addition to theoretical results, several applications have been suggested for online mechanisms, including: the allocation of Wi-Fi bandwidth (Friedman & Parkes [3]), scheduling of jobs on a server (Porter [9]) and the reservation of display space in online advertising (Constantin et al. [2]).

In recent work (Gerding et al. [4]), we propose the first online mechanism to deal with the problem of coordinating the charging of a set of plug-in hybrid electric vehicles (PHEVs) under limited supply. This earlier model, however, has several limitations that we address in this work. First, all vehicles participating in the system are assumed to have the same charging speed. In fact, given the large numbers of competing manufacturers entering this space, it is likely that domestic EV chargers with a wide range of maximum charging speeds will become available.¹ High performance chargers may pose additional burdens on the local transmission network and, for a real-world deployment, the allocation model and market design needs to be able to deal with such a challenge. Moreover, the presence of multiple, asymmetric charging speeds may considerably affect the dynamics of such a market, as they enable different allocations of the limited network capacity to become feasible. Finally, no theoretical worst-case bound for the mechanism was presented, whereas such a guarantee is highly desirable, especially as online allocation leaves some items unallocated to guarantee truthfulness. Against this background, this paper makes the following contributions to the state of the art:

- We develop a novel online mechanism that deals with multi-unit demands *per time step*, by extending the mechanism proposed by Gerding et al. [4] to accommodate heterogeneous and flexible charging speeds.
- We provide competitive bounds on allocative efficiency compared to the optimal off-line allocation (which assumes prior knowledge of future arrivals).
- We simulate the proposed mechanism using data from a real-world trial of electric vehicles in the UK. We show that the use of fast chargers can lead to significant savings in fuel consumption, and is beneficial both from the perspective of individual vehicle owners, and for the allocative efficiency of the whole market.

¹ <http://www.pod-point.com/pod-point-home> provides an example of a domestic charger that already offers a 2.5 times speed increase over standard chargers.

2 Multi-Speed EV Charging Model

We consider a setting in which multiple units of electricity are periodically sold at fixed time steps, t (e.g., once every hour). A *unit* of electricity is defined as the amount of kWh when charging at the lowest rate during that interval. Importantly, we assume that the charging rate can be flexible, and each vehicle has a maximum charging rate (which depends on the battery and its charger). We assume units to be indivisible, which means that the charging rate is a multiple of the lowest rate (this is not a limitation since units can be defined to be arbitrarily small). Moreover, we denote by $S(t)$ the supply of electricity available for EV charging, i.e. the number of units to be sold at time t . Note that, since our allocation is essentially greedy (i.e. units are allocated just before they are charged and there is no pre-allocation), this allows us to distribute electricity coming from uncertain sources such as a shared renewable generator (e.g. a shared neighbourhood wind turbine). For now, we consider a market for electricity for EV charging that is separate from that for household consumption, and so $S(t)$ can be considered the residual supply once household consumption has been removed. Beyond providing a manageable model, a practical reason for this separation of concern is to protect one neighbour from higher electricity prices for running regular jobs (lightning, domestic appliances) in the case that his neighbours have purchased EVs.

Let $I = \{1, 2, \dots, N\}$ denote the set of agents, each of which operates on behalf of a single EV and its owner. Vehicles come and go and are not always available for charging. Furthermore, they can have different maximum charging rates, and their owners have different valuations for the electricity. Given this, an agent i 's *type* is described by the tuple $\theta_i = \langle \mathbf{v}_i, a_i, d_i, r_i \rangle$, where $\mathbf{v}_i$ is the marginal valuation vector, a_i and d_i are the arrival and departure times (the earliest and latest times that the EV is available for charging) and r_i is the *maximum* charging rate (i.e., the maximum number of units agent i can charge at any time t).

Each element $v_{i,k}$ of $\mathbf{v}_i$ represents the agent's willingness to pay for the k^{th} unit of electricity. We assume non-increasing marginal valuations, i.e., $v_{i,k} \geq v_{i,k+1}$. This is a realistic assumption for PHEVs since the first few units of electricity are always more likely to be used [4]. Furthermore, a_i and d_i define when the agent is present in the market, where a_i is the agent's *arrival* or earliest time the vehicle can be charged, and d_i is the point of *departure*, after which the vehicle is required by the owner.

Given this, a *mechanism* asks the agents to report their types and decides on an appropriate allocation and payment. We denote the reported type by $\hat{\theta}_i = \langle \hat{\mathbf{v}}_i, \hat{a}_i, \hat{d}_i, \hat{r}_i \rangle$. In practice, the arrival report $\hat{a}_i$ is the time at which the owner plugs a vehicle into the electricity network. At the time of arrival, an agent is also required to report its marginal valuation vector $\hat{\mathbf{v}}_i$. The departure report $\hat{d}_i$ is *not* required in advance and it simply represents the time when the vehicle is unplugged from its unit. A vehicle owner may decide to change its (reported) departure time, by simply unplugging her vehicle.

Agents (or their owners) can misreport their availability, for example, by unplugging the vehicle early or plugging in the vehicle some time after arrival to try and get a better price. Also, they could report a lower maximum charging rate or higher valuations. Our aim is to develop a mechanism which is *dominant strategy incentive compatible* (DSIC), i.e., agents are best off reporting $\hat{\theta}_i = \theta_i$, no matter what other agents report. Formally, a mechanism is given by the allocation policy $\pi_i^{(t)}(\hat{\theta}_I | \mathbf{k}^{(t)})$, $i \in I$, which determines the number of units allocated to agent i at time t , and payment policy $x_i(\hat{\theta}_{-i} | k_i)$, $i \in I$, which calculates the total payment on (reported) departure of an agent. We denote by

$\theta_I = \{\theta_i | i \in I\}$ the types of all agents, and $\theta_{-i} = \{\theta_j | j \in I, j \neq i\}$ all agents except i . Payments are independent of an agent's own report. Here, k_i is agent i 's *endowment* (number of items allocated) on departure, and $\mathbf{k}^{(t)} = \langle k_1^{(t)}, \dots, k_N^{(t)} \rangle$ denotes the *endowment* of all agents at time t , where $k_i^{(t+1)} = k_i^{(t)} + \pi_i^{(t)}(\hat{\theta}_I | \mathbf{k}^{(t)})$.

3 The Online Mechanism

Essentially, our mechanism uses a greedy allocation policy, which allocates available units to the agents with the current highest marginal valuations (given their current endowments). While this is not always optimal (since it does not consider the fact that some agents leave earlier than others), it is generally applicable, as it does not require information about future arrivals or departures (i.e., it is “model free”).

A key problem, identified in our previous work [4], is that greedy allocation in online domains sometimes overallocates (given the prices it has to pay, the agent would prefer to get fewer units than it would be allocated). To address this problem, we correct the allocation by leaving some units unallocated under well specified conditions, which we refer to as ‘burning’ units. We consider two types of burning, discussed below: *immediate* burning, where units are simply left unallocated (i.e., none of the agents receive the unit, even if there is a demand for them), and *on-departure*, where units are initially allocated using the greedy approach (i.e., the battery is being charged), but then on departure of the agent, any overallocated units are discharged from the battery. Both of these have advantages/disadvantages. The model with on-departure burning is more efficient (i.e., generally burns fewer units) and is also more computationally efficient in computing the payments. However, it may not be realistic to expect that we can partially discharge a car's battery on the departure, so the model with immediate burning may be more realistic, given the application.

Formally, let the vector $\mathbf{b}_i^{(t)} = \langle \hat{v}_{i, \mathbf{k}_i^{(t)}+1}, \dots, \hat{v}_{i, \mathbf{k}_i^{(t)}+\hat{r}_i} \rangle$ denote agent i 's reported marginal values for the next $\hat{r}_i$ units, given its endowment $k_i^{(t)}$ at time t . It is convenient to think of this as the agent's *bids* for the units available at time t . Furthermore, let $B^{(t)}$ denote the multiset of such bids from all agents that are present in the market at time t , i.e., from all $i \in I$ s.t. $\hat{a}_i \leq t \leq \hat{d}_i$. The allocation rule is:

Definition 1 (Greedy Allocation). *At each time step t , allocate the $S(t)$ units of electricity to the highest bids in $B^{(t)}$.*

To ensure that the payment is report-independent, the mechanism reruns the market without agent i (from $\hat{a}_i$ onwards, since agent i did not affect the market prior to this time). Let $B_{-i}^{(t)}$ denote the multiset of the bids placed by all agents in the market at time t if agent i were removed and the market were rerun from $\hat{a}_i \leq t$ onwards. In case $|B_{-i}^{(t)}| < S(t)$, we add a number of zero-valued bids and refer to this enlarged set as $B_{-i \cup 0}^{(t)}$, to ensure that $|B_{-i \cup 0}^{(t)}| \geq S(t)$. Next, we define set operators $\max_k B$ and $\min_k B$ to return the highest and, respectively, lowest k elements of multiset B (or, if $|B| < k$, to return B). Then, we define the *externality* that agent i would impose on other agents if it won $\min(r_i, S(t))$ out of $S(t)$ units at time t as:

$$E_i^{(t)} = \min_{r_i} (\max_{S(t)} B_{-i \cup 0}^{(t)})$$

Intuitively, the multiset $E_i^{(t)}$ here contains the bids from other agents that would lose out if agent i were to win r_i units at time t . Note that the intuition here is the same as in the




Fig. 1. Example showing arrivals, departures, and valuation vectors of 2 time units and 3 agents.

regular Vickrey-Clarke-Groves (VCG) mechanism, as the total payment corresponds to the sum of the externalities. However, to compute the overall payments online, we need to combine these externalities across all time steps in the agent's active period up to current time t . To do this, we define an ordered vector of *prices*, $p_{-i}^{(t)}$, as follows:

$$\mathbf{p}_{-i}^{(t)} = \text{incr} \left(\bigcup_{t'=a_i}^t E_i^{(t')} \right),$$

where incr is an operator that orders elements from a multiset in increasing order, and we use the union symbol to denote the union of multisets (and so the same element can appear multiple times). Now, $p_{-i,k}^{(t)}$ is the price that agent i is charged for the k -th unit of electricity. Intuitively, this is the minimum valuation that the agent could report for winning this unit by time t . Note we use $p_{-i,k}$ to denote this price at time $\hat{d}_i$. Given this, the mechanism is as follows.

– **Decision Policy** The decision consists of two stages.

Stage 1 At each time point t , *pre-allocate* using Greedy Allocation (see Definition 1).

Stage 2 We consider two variations of when to decide to burn pre-allocated units:

– **Immediate Burning.** Leave any unit unallocated whenever the price for this unit is greater than the marginal value, i.e., whenever:

$$\hat{v}_{i,k} < p_{-i,k}^{(t)} \text{ for } k_i^{(t)} < k \leq k_i^{(t)} + \pi_i^{(t)}$$

– **On-Departure Burning.** For each departing agent, discharge any unit $k \leq k_i$ where $\hat{v}_{i,k} < p_{-i,k}$.

– **Payment Policy** Payment always occurs on reported departure (i.e., when the owner/agent unplugs the vehicle). Given that k_i units are allocated to agent i , the payment collected from i is:

$$x_i(\hat{\theta}_{-i}|k_i) = \sum_{k=1}^{k_i} p_{-i,k} \quad (1)$$

Figure 1 illustrates the mechanism through an example with 2 time steps and 3 agents, showing the agents' arrival, departure and valuations. Suppose furthermore, that supply is $S(t_1) = 2$ and $S(t_2) = 1$. Now, consider 2 distinct cases:

*The maximum charging speed of agent 1 is $r_1 = 1$.*² In this case, at most one marginal value is taken from each agent. At time t_1 , marginal valuations $v_{1,1} = 10$ of agent 1, and $v_{2,1} = 7$ of agent 2 are allocated, while at time t_2 , marginal value $v_{1,2} = 8$ of agent 1 is allocated. The prices charged to agent 1 are: $p_{-1} = \langle 0, 1 \rangle$, because without agent 1 in the market, there would be a free, spare unit at time t_1 and the available unit at t_2 would sell to agent 3 for 1. No units get burnt in this case, and the actual allocation is actually equivalent to the optimal offline allocation.

² Note that the other 2 agents only desire one unit, so their maximum charging speed is irrelevant in this example.

The maximum charging speed of agent 1 is $r_1 = 2$. Then, at time t_1 , the greedy mechanism described above allocates the 2 marginal values of agent 1: $v_{1,2} = 10$ and $v_{1,2} = 8$, as they are both higher than $v_{2,1} = 7$, and agent 2 drops out of the market. At time t_2 , the unit is again allocated to agent 1 (due to the marginal value of 3 being higher than 1). However, now the marginal payments vector required from agent 1 is $p_{-1} = \langle 0, 1, 7 \rangle$, while the marginal valuations are $v_1 = \langle 10, 8, 3 \rangle$. Given the prices, agent 1 prefers 2 units to 3 (because $10 + 8 - 1 > 10 + 8 + 3 - 1 - 7$), so the third is burnt. The overall market efficiency is lower, as the third available unit is now burnt, whereas with $r_1 = 1$ it was allocated to agent 2. Note, however, that even though the efficiency is much lower, agent 1 has an incentive to declare its true maximum charging speed $r_1 = 2$ as, in this case, its payment does not change.

4 Truthfulness Properties

In this section, we discuss the incentive compatibility properties of our mechanism under the following assumption:

Limited Misreports: *Agents cannot report an earlier arrival, a later departure, or a higher charging rate, i.e., $\hat{a}_i \geq a_i, \hat{d}_i \leq d_i, \hat{r}_i \leq r_i$ must hold.*

In our domain, this is a reasonable assumption since a vehicle owner cannot physically plug in her vehicle before it is available or unplug it after the actual departure. Note that the assumption $\hat{r}_i \leq r_i$ is natural for EV charging. While most electric batteries can be configured to charge at a slower rate, charging them at a faster rate than the one allowed by the manufacturer might destroy them. As the battery is an integral and expensive part of an EV, this by itself acts as a natural deterrent. An EV agent may strategise by reporting a $\hat{r}_i$ lower than its true maximal speed, but we show truthful reporting of r_i is a dominant strategy. Given this, we can state the following:

Theorem 1. *Assuming limited misreports, and given the decision and payment policy as described above, reporting $\hat{\theta}_i = \theta_i$ is a dominant strategy for $\forall i \in I$.*

Proof Sketch: The proof for this theorem has 3 parts:

1. Showing that, after fixing any tuple of $\langle \hat{a}_i, \hat{d}_i, \hat{r}_i \rangle$, reporting the valuation vector $\hat{\mathbf{v}}_i = \mathbf{v}_i$ truthfully is dominant strategy incentive compatible (DSIC).
2. Showing that for any fixed $\langle \hat{\mathbf{v}}_i, \hat{r}_i \rangle$ (and under the limited misreports assumption $\hat{a}_i \geq a_i, \hat{d}_i \leq d_i$), reporting truthfully $\hat{a}_i = a_i, \hat{d}_i = d_i$ is DSIC.
3. Showing that for any fixed $\langle \hat{\mathbf{v}}_i, \hat{a}_i, \hat{d}_i \rangle$, using the limited misreports assumption $\hat{r}_i \leq r_i$, reporting truthfully the maximum charging rate $\hat{r}_i = r_i$ is DSIC.

The first two parts can be shown by using the same proof techniques as in [4]. For the third part, we can show that the vector of marginal payments $\mathbf{p}^{\langle \hat{r}_i \rangle}$ (computed given the report $\hat{r}_i$) will always contain a *subset* of the elements from the vector $\mathbf{p}_i^{\langle r_i \rangle}$ (computed given the truthful report r_i), because $\hat{r}_i \leq r_i$. Thus, either an agent gets the same allocation, but pays less by reporting a higher speed (because the p_i vectors are increasingly ordered, and the payment is the sum of the first k_i units allocated), or the agent is allocated more units but then those additional units are needed, given the prices (otherwise, burning would apply).

5 Theoretical Bounds on Allocative Efficiency

An important question given the online allocation with burning proposed above is how the allocative efficiency compares to that of an optimal offline allocation (i.e., assuming

full knowledge of the future). To answer this, we consider the cases with immediate burning and on-departure burning separately. For the case of immediate burning, we can show that no lower bound can be established through an example:

Example 1: Let an agent A_{syn} with marginal valuation vector $\mathbf{v}_{\text{syn}} = \langle v_1, v_2, \dots, v_n \rangle$, where marginal valuations are strictly decreasing, i.e., $v_1 > v_2 > \dots > v_n$. Assume supply is $S(t) = 1$ for all t , and at each time step, one other “local” agent is present for that time step only and desires a single unit. The valuation of the first local bidder is $\mathbf{v}_1 = \langle v_1 - \epsilon \rangle$, the valuations of the next two are $\mathbf{v}_{2,3} = \langle v_2 - \epsilon \rangle$, the next three $\mathbf{v}_{3,4,5} = \langle v_3 - \epsilon \rangle$, and so on, where ϵ is an arbitrarily small quantity. In other words, each marginal value $v_k - \epsilon$ appears exactly k times.

In this example, agent A_{syn} is pre-allocated every unit, but for each valuation v_k , the first $k - 1$ pre-allocated units are burnt, with only the last unit being allocated (due to the expanding $\mathbf{p}_{-i}^{(t)}$ vector). Thus, in the limit, the fraction of units burnt goes to 1, while the efficiency goes to 0.

For on-departure burning, Theorem 2 provides the following worst-case guarantee:

Theorem 2. *The mechanism with on-departure burning is 2-competitive with the optimal offline allocation, for a setting with non-increasing marginal values.*

Proof Sketch: In order to establish a competitive bound with the optimal offline allocation, we use a “charging argument” similar to Hajiaghayi et al. [5].³ The basic idea is to “charge” (or match) all the marginal values of each agent that are allocated in the offline case with one or more values allocated online. Specifically, let $v_{i,k}^{\text{off}}$ denote the k -th marginal value of agent i that is allocated in the offline case. In the following we say that a unit is satisfied (or allocated) in the online case if it is actually allocated to agent i by the greedy mechanism and not burnt on departure.

Consider each marginal value $v_{i,k}^{\text{off}}$ from agents i that are allocated in the *offline* case, and to “charge” the values as follows: 1. Marginal values allocated both in the online and offline case are charged to themselves. 2. Marginal values that are allocated in the offline case, but not in the online case will always be displaced by some higher value $v_{j,p}^{\text{off}}$. Note that, since these values could be allocated online and are higher, they must necessarily also be allocated in the offline case (but this may occur at a different time than in the online case, when they would displace a lower valued unit). To complete the charging argument, we need to show what happens to these units w.r.t. burning.

Lemma 1. A marginal utility value $v_{i,k}^{\text{off}}$ that is allocated to agent i *both* offline and online cannot be burnt on the departure of agent i in the online market. Moreover, if a unit $v_{i,k}^{\text{off}}$ is displaced by another unit $v_{j,p}^{\text{off}}$, the displacing unit $v_{j,p}^{\text{off}}$ cannot be burnt on departure of agent j .

Proof Sketch: The proof for this lemma relies on comparing the vectors of decreasing marginal values $\mathbf{v}_i$ and the payment vector $\mathbf{p}_{-i}$ as well as the condition for burning specified by the on-departure burning decision policy. In summary, one can show by contradiction that, if a unit $v_{i,k}^{\text{off}}$ is allocated both offline and online, there could not be enough higher valued marginal units in agent i ’s active interval $[a_i, d_i]$ for this unit (or a higher valued one) to be burnt.

Given Lemma 1, for all agents i , each marginal value unit that is allocated offline can be charged *at most* twice:

³ Here the term “charging” does not refer to electricity charging, but represents the name of a proof device used in online MD.

- At most once to itself (if it is also allocated online).
- At most once by another unit allocated offline, with a value less than its own, that it displaced at some time t when that unit was allocated in the offline case (recall that in the mechanism with on-departure burning, each unit $v_{i,j}$ can be the winning bid at most once, so it can displace at most one other unit). Thus each unit allocated offline is charged at most twice, giving the online allocation with on-departure burning and non-increasing marginal values a competitive bound of 2 with the optimal offline allocation.

Theorem 2 shows the multi-unit demand case with on-departure burning is no worse in terms of worst-case competitive bound than the case with single unit demand (c.f. Hajiaghayi et al. [5]) which may seem surprising, given the burning. However, the only units allocated online which could be burnt online on departure of an agent i are units of such low value that would not be allocated in the corresponding offline case.

6 Experimental Evaluation

In this section, we apply our mechanism to a range of settings that are based on real data collected during the first large-scale trial of EVs in the UK. Our experiments examine two main questions. First, we study how our online allocation mechanism performs compared to a range of benchmarks as demand for electricity increases in a neighbourhood with limited supply. Second, we look at how the gradual introduction of fast-charging EVs would affect the neighbourhood, both in terms of social welfare (i.e., overall fuel savings) and the financial savings of individuals.

6.1 Experimental Setup

We base our experiments on data gathered by the CABLED (Coventry And Birmingham Low Emissions Demonstration) project,⁴ which is the first large-scale endeavour in the UK to record and study the driving and charging behaviours of EV owners. The arrival and departure times of vehicle owners are generated based on this data. We also construct the agents' valuation vectors using the expected travel distances reported by the project, as well as the fuel and electric efficiencies and battery capacities of typical plug-in hybrid EVs, such as the Chevrolet Volt or Toyota Plug-In Prius. Such hybrid electric vehicles have a dual engine, and a per mile driving efficiency, which can be expressed either in units of electricity or litres of fuel (assuming a standard driving speed). As such, an agent's marginal valuation of a unit of electricity corresponds to the savings in fuel costs that the agent expects to make due to consuming this unit instead of regular fuel (based on a fuel price of £1.3/litre). For each configuration, 500 or 1000 runs were performed, with the expected driving distance for each vehicle for the next day being and its entry point into the market being randomized in each run. For full details, see Gerding et al. [4], where we use the same setup and dataset.

For each experiment, we consider a single day and, to calculate the capacity constraints (the supply function $S(t)$), we divide the day into 24 hourly time steps. For each time step, we first obtain the overall neighbourhood consumption based on real data.⁵ We then consider two possible scenarios. First, when supply is low, electricity is highly constrained, and we set the capacity limit to 90% of the peak overall consumption, i.e., about 0.9 kW per household. This scenario represents neighbourhoods where the local

⁴ See <http://cabled.org.uk/>.

⁵ We use the average data for domestic households in June 2010, as reported by SCE (<http://www.sce.com/>).




Fig. 2. Units of electricity available for EV charging.

transformer can only just support domestic demand with no spare capacity for vehicle charging during peak hours. Second, for *high* supply with significant spare capacity, we set the limit to 150% of peak consumption, i.e., approximately 1.5 kW per household.

In addition to our mechanism with *On-Departure* and *Immediate* burning, we consider a set of benchmark mechanisms.

(i): *Random* allocates units randomly to agents that have some positive marginal valuation for them.

(ii): *Fixed Price* allocates units at random to agents that have valuations greater than a fixed price and collects a payment equal to this fixed price. The price is chosen a priori to optimise the expected social welfare (i.e., total fuel savings), given full knowledge of the agent type distributions.

(iv): *Heuristic* is a common earliest-deadline-first scheduling heuristic that allocates units to agents to maximise a weighted sum of the agent's valuation for the unit and its deadline. The weights are chosen as for *Fixed Price*.

(v): *Optimal* allocates units to optimise the overall social welfare, assuming full knowledge of all future arrivals.

Note that apart from the *Random* baseline, these benchmarks assume additional knowledge either about agent type distributions or even future arrivals, and so they should be seen as upper bounds for our approach rather than actual realistic alternatives. Furthermore, only *Random* and *Fixed Price* are truthful (DSIC) like our mechanism.

6.2 Result 1: General Trends with Increasing Demand

First we consider the effect of rising electricity demand within a neighbourhood. To this end, Figure 3 shows the performance of our mechanism in a neighbourhood of 25 households with increasing numbers of EVs.⁶ These results are for the low supply setting only, as the variation of EV numbers also changes the supply/demand balance. We assume that electricity is allocated in units of 3 kWh (the charging rate of a standard UK power socket), and that maximum rates are chosen at random from $\{1,2,3\}$ units, corresponding to the rates of currently available fast chargers for domestic use. Due to the computational complexity of the *Optimal* and *Immediate* mechanisms, we only plot those up to 30 and 15 agents, respectively.⁷

⁶ Note that as this exceeds 25, some households will have multiple cars. We show this to indicate the effect of very high demand and assume there is no collusion within a household.

⁷ Briefly, *Immediate* is computationally hard, because the price vector $\mathbf{p}_{-i}^{(t)}$ has to be computed at every time step by simulating the market without agent i , and, recursively, for every agent within that market, as it is needed to compute when burning takes place.




Fig. 3. Effect of increasing demand for electricity in a small 25-household neighbourhood with low supply.

First, the graph on the left shows the daily fuel saving (excluding payments) for EV owners under the different mechanisms (all results reported are averaged over at least 500 samples and plotted with 95% confidence intervals). This shows that there is little difference between the various mechanisms when competition is low, as most agents are allocated; but as demand increases, our mechanism starts to outperform the other truthful benchmarks (by almost 50%) and achieves 95% or more of the optimal throughput. Interestingly, our mechanism sustains higher savings per agent for neighbourhoods with far higher EV penetration than the truthful benchmarks. For example, to guarantee a saving of at least £0.7 (just over half a litre of fuel), the *Random* benchmark can support up to 35 EVs, while our mechanism can cope with 50. The graph on the right shows the overall allocative efficiency achieved in the same setting, normalised to the optimal (beyond 30 EVs, this is normalised to the *Heuristic* as a close approximation).

Our principal finding is that online allocation with *On-Departure* burning achieves almost the same performance as the *Optimal*, calculated using full prior information about departures (the difference is not statistically significant). Moreover, the difference between our online allocation with *Immediate* burning and with *On-Departure* burning is also not significant. Surprising, given the different theoretical performance bounds of these policies, and indicates that burning is not a large problem in realistic settings.

6.3 Result 2: Proliferation of Fast-Charging EVs

We now consider in more detail how the presence of fast-charging vehicles affects market performance. To this end, we examine a larger neighbourhood of 50 households with one EV each. We choose this size, because it is realistic and leads to more interesting results, due to increasing likelihood of competition at each time step. Due to their computational cost, we no longer run the *Optimal* and *Immediate* mechanisms (but their performance is similar to the *Heuristic* and *On-Departure*, respectively, as shown previously). To investigate the impact of fast-charging, we assume there are two agent types — the first, *normal*, can charge a single unit of 3 kWh per time step, while the second, *fast*, are equipped with fast chargers that can charge up to three such units per time step. Throughout the experiments, we vary the number of fast-charging EVs (out of the total 50).

Results for this setting are shown in Figure 4. First, we note that the trends for the two scenarios are different — when supply is low, the introduction of more fast charging




Fig. 4. Social welfare, individual utilities and units burnt in scenarios with low supply (left) and high supply (right), as more fast-charging EVs are introduced.

vehicles results only in a small overall improvement in social welfare. However, when it is high, there is a very marked improvement, with an additional 13–14 litres of fuel per day being saved by our mechanism as more fast-chargers are present (corresponding to a saving of around 5000 litres of fuel per year over the whole neighbourhood). This is because, at low demand times, there is some spare capacity that remains unallocated, unless vehicles have a high charging speed.

With respect to the utility of individual agents (including payments to the mechanism), we see that agents *in both settings* always have an incentive to switch to fast-charging EVs (e.g., by purchasing a domestic fast charger). With low supply, the expected daily saving when switching to a fast-charging EV is approximately £0.05, while with high supply, this is around £0.45. In both cases, this benefit is the result of increasing available supply per time step, as well as increasing the size of the price vector.

Finally, another benefit of introducing fast-charging vehicles is the reduced proportion of units of electricity that are burnt by our mechanism. In the low supply setting, the percentage of burnt units is more than halved to less than 1% of all units allocated

(on average 1.5 kWh) when all cars are fast-charging. With high supply, this is even more significant as only about a third of units are burnt. This reduction is a result of reduced competition and a larger price vector.

7 Conclusions and Further Work

This work makes several key contributions to the existing literature. On the theoretical side, we extend the multi-unit, online mechanism proposed in Gerding et al. [4] to also allow for modeling multi-unit demand *per time step*, in conjunction with multi-unit demand over time. For our EV application, this allows us to model more realistic markets, which include vehicles with heterogeneous charging capabilities. Moreover, we provide the first theoretical worst-case bounds for multi-unit online mechanisms. On the practical side, we build a detailed simulation of an EV charging market and show that our online mechanism performs close to the offline optimal in a variety of realistic settings. We also show that faster-charging batteries lead to savings in fuel consumption and increased allocative efficiency.

In future work, we plan to compare the model-free mechanism presented in this paper with online mechanisms that use a model of expected future arrivals. Moreover, we are interested in studying whether there are truthful mechanisms for this problem that do not require monetary payments.

Acknowledgements

This work was carried out as part of the IDEaS project (<http://www.ideasproject.info>), funded by a company active within the energy domain.

References

1. K. Clement, E. Haesen, and J. Driesen. Coordinated charging of multiple plug-in hybrid electric vehicles in residential distribution grids. In *Proc. of the IEEE/PES Power Systems Conference and Exposition*, pages 1–7, 2009.
2. F. Constantin, J. Feldman, S. Muthukrishnan, and M. Pal. An online mechanism for ad slot reservations with cancellations. In *Proc. SODA'09*, pages 1265–1274, 2009.
3. E. Friedman and D. C. Parkes. Pricing WiFi at Starbucks— Issues in online mechanism design. In *4th ACM Conference on Electronic Commerce*, pages 240–241, 2003.
4. E. H. Gerding, V. Robu, S. Stein, D. C. Parkes, A. Rogers, and N. R. Jennings. Online mechanism design for electric vehicle charging. In *Proc. of the 10th International Conference Autonomous Agents and Multi-Agent Systems (AAMAS'11)*, 2011.
5. M. Hajiaghayi, R. Kleinberg, M. Mahdian, and D. C. Parkes. Online auctions with re-usable goods. In *6th ACM Conference on Electronic Commerce (EC'05)*, pages 165–174, 2005.
6. A.I. Juda and D. C. Parkes. An options-based solution to the sequential auction problem. *Artificial Intelligence*, 173:876–899, 2009.
7. T. Koslowski. The electric vehicle's value chain and technology evolution. Technical report, Gartner Inc., 2009.
8. D. C. Parkes and S. Singh. An MDP-Based approach to Online Mechanism Design. In *Proc. of NIPS'03*, pages 791–798, 2003.
9. R. Porter. Mechanism design for online real-time scheduling. In *Proc. 5th ACM Conference on Electronic Commerce (EC'04)*, pages 61–70, 2004.
10. S. Shao, M. Pipattanasomporn, and S. Rahman. Challenges of PHEV penetration to the residential distribution network. In *IEEE Energy Society General Meeting*, 2009.