

Coordination and Payment Mechanisms for Electric Vehicle Aggregators

Alvaro Perez-Diaz^{a,b,*}, Enrico Gerding^a, Frank McGroarty^c

^a*Electronics and Computer Science, University of Southampton, SO17 1BJ, United Kingdom*

^b*Engineering and the Environment, University of Southampton, SO17 1BJ, United Kingdom*

^c*Southampton Business School, University of Southampton, SO17 1BJ, United Kingdom*

Abstract

Motivated by the high electric vehicle (EV) penetration percentages foreseen for the near future, this paper studies the participation of large fleets of EVs in electricity day-ahead markets. Specifically, we consider a scenario where a number of independent and self-interested EV aggregators participate in the day-ahead market to purchase energy to satisfy their clients' driving needs. In this scenario, independent bidding can drive prices up unnecessarily, resulting in increased electricity costs for all participants. Inter-aggregator cooperation can mitigate this by producing coordinated bids. However, this is challenging due to the self-interested nature of the aggregators, who may try to manipulate the system in order to obtain personal benefit. In order to overcome this issue, we employ techniques from mechanism design to develop a coordination mechanism which incentivises self-interested EV aggregators to report their energy requirements truthfully to a third-party coordinator. This coordinator is then able to employ a day-ahead bidding algorithm to optimise the global bids on their behalf, extending the benefits of smart bidding to groups of competing EV aggregators. Importantly, the proposed coordination mechanism is straightforward to implement and does not require any additional infrastructure. To ensure scalability and computational tractability, a novel price-maker day-ahead bidding algorithm is proposed, which is formulated in terms of simple energy requirement constraints. The coordination mechanism substantially reduces bidding costs, as shown in a case study which uses real market and driver data from the Iberian Peninsula.

Keywords: electric vehicle charging, electric vehicle aggregation, inter-aggregator coordination, mechanism design, Vickrey-Clarke-Groves, day-ahead electricity market

1. Introduction

Climate change and environment conservation constitute two of the main challenges to address in the twenty-first century. In particular, fossil fuels account for a great proportion of the global contaminant emissions. Nowadays, around 29% of the total energy consumption in the US is attributable to the transportation sector, and fossil fuels power around 95% of this amount [1]. Similarly, in the UK, transportation is the biggest energy consumer accounting for 40% of the total energy consumption [2], where 96% of this consumption depends on fossil fuels [3].

Transportation is then a key participant in the modernisation and improvement of the energy generation-

demand duo. Specifically, the electrification of transportation has the potential to reduce the dependence on fossil fuels and allow the effective use of renewable electricity sources. However, it presents important challenges, such as accommodating the very large electricity requirements of a large electric vehicle (EV) fleet. If left unmanaged, uncoordinated operation can pose a novel and heavy strain on the existing electricity generation, transmission and pricing methods. As an example, currently, the UK has a fleet of nearly 100,000 EVs, combining purely electrical and hybrid vehicles [4]. Furthermore, the UK has a 10% target for electrification in the transportation sector by 2020 [5]. Similarly, at a global scale, there are targets to achieve 100 to 140 million of EVs by 2020 [6]. Hence, given the current and targeted EV penetration numbers, EV fleet management must be seen as a priority [7].

In order to address this issue, one of the proposed methods is to employ an EV aggregator. Originally in-

*Corresponding author

Email addresses: a.perez-diaz@soton.ac.uk (Alvaro Perez-Diaz), eg@soton.ac.uk (Enrico Gerding), f.j.mcgroarty@soton.ac.uk (Frank McGroarty)

troduced in [8], this entity acts as an intermediary between an EV fleet and the electricity grid and markets, and has control of the charging, and possibly discharging, of the fleet’s batteries. By harvesting the combined capacity of a number of EVs, the EV aggregator can participate in wholesale electricity markets and provide ancillary services. However, given the increasing numbers of EVs, we envision a scenario where independent and self-interested EV aggregators compete in the same electricity market, trying to maximise their own profit [9]. Each of these aggregators can optimise its own operation, but lack of coordination can cause global inefficiencies.

In more detail, we consider EV aggregators purchasing energy in day-ahead markets, in order to meet their clients’ electricity requirements. Without coordination, energy bidding could become concentrated in time, driving prices up and resulting in more expensive costs. This issue can be overcome by sharing information, *cooperating*, which can translate in cost decreases, a more stable grid and cleaner production methods, as demand peaks can be softened. However, manipulation can exist due to the self-interested nature of the aggregators, who could choose to cheat the system if greater personal benefit is perceived.

EV aggregator participation in day-ahead markets has been widely studied in the literature. Existing works consider either price-taker approaches [10, 11, 12, 13], where prices are exogenous and unaffected by the EV aggregator activity, or price-maker approaches [14, 15, 16, 17], where the aggregator’s bids affect electricity prices. However, all these works consider a single EV aggregator which is able to optimise its own bidding, but do not consider the realistic scenario of several aggregators coexisting. To extend these works to a multi-aggregator setting, and to address the inter-aggregator cooperation challenge presented above, we propose a novel coordination mechanism which uses techniques from mechanism design to allow extending the benefits of optimised bidding to groups of self-interested EV aggregators. Mechanism design studies agent interaction protocols which take into account the fact that agents are rational and self-interested [18]. Our approach is based on the well-known Vickrey-Clarke-Groves (VCG) mechanism [18], and we consider three approaches for computing the payments using well-known redistribution mechanisms to ensure the payments made by the aggregators are fair. Two of these payment approaches are guaranteed to be truthful in our setting, meaning that participating EV aggregators are incentivised to cooperate by truthfully reporting their requirements to the third-party coordinator. By doing so,

globally informed bidding decisions can be made, lowering total energy costs. This challenge has not been previously addressed for this setting.

In addition to the coordination mechanism, we introduce a novel price-maker bidding algorithm which is scalable and computationally tractable. This algorithm extends the price-taker algorithm from [13] by accounting for price impact through using residual supply curves [19]. We consider two approaches. The first utilises raw historical supply and demand data, while the second formulation employs a quadratic convex approximation. These approaches are in contrast with the linear price impact utilised in [14, 15, 16]. Moreover, unlike [17], who also consider non-linear price impact, our approach has only 72 linear constraints. Hence, it is simpler to implement and scales to arbitrary numbers of EVs. Another advantage of our convex approximation approach is that it guarantees finding the global minimum of the approximated function. This is important for using the VCG-type mechanisms, which require an optimal solution to ensure the truthfulness property.¹

More precisely, the contributions of this paper are three-fold.

1. The development of a novel price-maker day-ahead bidding algorithm for EV aggregators, formulated in terms of simple driver requirement constraints, which scales to very large fleets with very little computational burden.
2. The development of a novel EV aggregator coordination mechanism in which different self-interested EV aggregators can coordinate their bids in the day-ahead market to achieve reduced costs. This mechanism is generic and can be employed in combination with any underlying bidding algorithm. However, when employing our proposed bidding algorithm, the coordination mechanism scales linearly with number of aggregators and fleet size, allowing it to be applied to very large populations. Moreover, when combined with the convex approximation variant of the bidding algorithm, the mechanism incentivises truthful reporting.
3. Evaluation of the proposed bidding algorithm and coordination mechanism in a realistic scenario using real market and driver behaviour data from the Iberian Peninsula.

¹More precisely, even though the function is not optimal from the coordinator perspective, because the allocation and payments are both based on this function, from the individual aggregator perspective the allocation is perceived to be optimal and their utility cannot be improved by misreporting.

The rest of the paper is structured as follows. Section 2 introduces the novel price-maker day-ahead bidding algorithm. Section 3 details the proposed inter-aggregator coordination mechanisms, which utilise the bidding algorithm from Section 2. Then, Section 4 presents a realistic case study to evaluate the performance of the proposed algorithms. Finally, Section 5 presents the conclusions.

2. Participation of an EV Aggregator in the Day-Ahead Market

This section introduces the proposed day-ahead bidding framework. Firstly, the typical day-ahead market found in most countries is detailed, together with the mathematical formalism to quantify the impact of the aggregator's bids on electricity prices. Then, the proposed EV aggregator and novel bidding algorithm are described.

2.1. The Day-Ahead Market

Typical day-ahead markets are daily forward markets featuring a uniform-priced double-sided auction. In more detail, they run every day of the year and each day is divided into 24 one-hour slots. A separate auction is run for each hourly slot. All bids and offers for each hourly slot of day $D + 1$ need to be submitted before market closure time, noon on day D . Uniform price refers to marginal pricing, by which all the accepted orders have the same price, which coincides with the price of the intersection between supply and demand.

The supply side, consisting of electricity producers, offers volumes of electricity at different prices. All supply side offers are aggregated by low-price priority, resulting in a generation stack curve, which relates production size to price per energy unit, as seen in Figure 1a. Similarly, the demand side consists of large electricity consumers, such as industries, retailers or aggregators, who similarly submit their bids to the day-ahead auction, consisting of both desired electricity volumes and prices. All the bids are aggregated by high-price priority forming the aggregated demand curve, as shown in Fig 1a.

Based on supply and demand, a clearing price is determined. The exact clearing and pricing mechanisms depend on the particular market, but the most common approach is to find the intersection between the aggregated supply and demand curves. The resulting price will be the clearing price which applies to all accepted bids and offers (uniform pricing). Note that there is usually a maximum price for bids and offers, p_{\max} . The accepted bids and offers are the ones lying towards the

left-hand side of the intersection, *i.e.* cheaper offers and more expensive bids than the clearing price.

2.2. Price Impact

Price impact is an important characteristic associated with large market participants, such as large EV aggregators, bidding individually or as a cooperating group. In case of large buyers, their bids will significantly contribute towards the aggregated demand curve, and hence will impact the clearing price by pushing it up. In the power systems literature, residual demand techniques have been widely employed to quantify and study price impact, especially in day-ahead markets, and we describe the mathematical framework below.

Following standard notation from [19], for any given hour t , let $D_t(p)$ and $S_t(p)$ be the aggregated demand and supply curves respectively, as a function of price, p . Consider a new agent participating in the market as a buyer. The residual supply curve is defined as:

$$R_t(p) = S_t(p) - D_t(p) = E$$

and represents the amount of energy, E , the new agent can bid for while maintaining a clearing price p . Without loss of generality, we assume that the agent's bids are set at maximum price, p_{\max} , to guarantee execution.

More relevant here is the clearing price when requesting a quantity E , which can be obtained from the residual supply curve $p = R_t^{-1}(E)$. Introducing the notation $\mathcal{P}_t(E) = R_t^{-1}(E)$, the clearing price when the new agent bids an amount E is:

$$p = \mathcal{P}_t(E)$$

and the price impact Δp of this order is:

$$\Delta p = \mathcal{P}_t(E) - \mathcal{P}_t(0)$$

where $\mathcal{P}_t(0)$ represents the *base price* at hour t , *i.e.* the price without the agent's new bid. This formalism is depicted in Figures 1b and 1c.

Focusing on EVs, a single EV will have a negligible price impact, but for an aggregator managing a large number of EVs, the combined capacity of their batteries gets sufficiently large to have a non-negligible impact on the market. Due to the high degree of flexibility characteristic of EVs, there is ample room for coordination which will result in lowered costs and better grid functioning, which motivates the need for efficient energy buying algorithms under the price-maker approach considered in this paper.

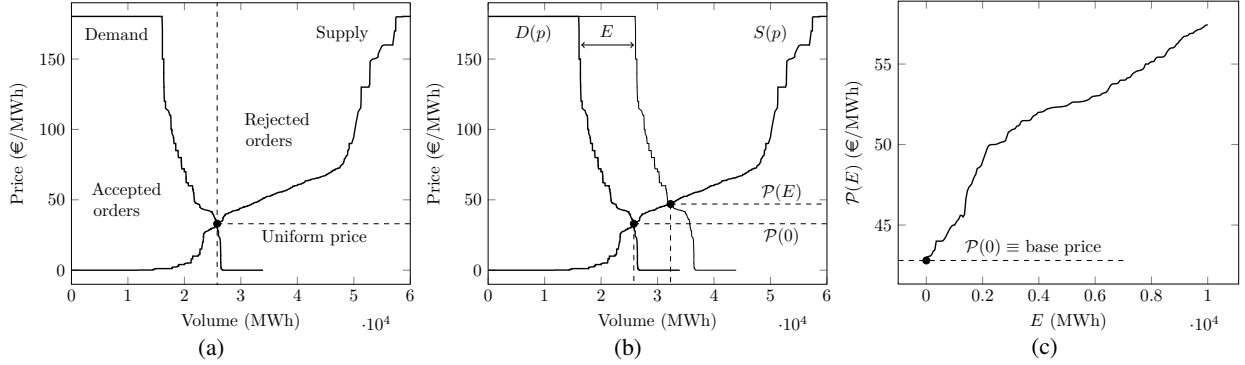


Figure 1: (a) Aggregated supply and demand curves, and market clearing mechanism. (b) Price impact of a buy order with volume E and maximum price p_{\max} . (c) Final price function $\mathcal{P}(E)$ for a volume range up to 10 MWh. Source: OMIE, 01/11/2016, 11th hour.

2.3. EV Aggregator Model

An EV aggregator is the entity in charge of coordinating fleets of EVs as described below. Specifically, their main purpose is to manage the charging of a number of EVs which use its services, acting as an intermediary between the fleet of EVs and the day-ahead electricity market, in order to minimize costs for itself and its customers/members [13, 20, 9, 17]. Examples of EV aggregators range from a private car park, to a residential micro-grid or a virtual cooperative with no physical proximity.

In our model, on arrival, each EV driver will input her requirements in the form of a departure time t_d^i , and a desired state of charge SoC_d^i at departure time. It is assumed that the aggregator can automatically record the arrival time t_0^i , the state of charge at arrival time SoC_0^i , and the battery size, $\text{SoC}_{\text{total}}^i$. Based on this information, the aggregator will make scheduling decisions, guaranteeing the desired SoC_d^i at departure time t_d^i ². The possibility to charge the EV's battery in an informed way (rather than immediately at arrival time, or randomly when charging resources are available) will result in cheaper charging costs which ultimately means cheaper tariffs for the EV drivers and/or increased profit for the aggregator. Each EV has a maximum charging speed, P_{\max}^i , in kW. This charging speed will depend on two factors: the available physical infrastructure (electricity grid) and the battery of each particular EV. Without loss of generality, we will assume that all EVs have the same charging speed, $P_{\max}^i = P_{\max}$. Note that the proposed bidding algorithm can be easily generalised to accommodate different charging speeds, which would have little impact in the results presented in Section 4. Lastly, in order to convert from quantity charged at hour

²We assume that the input requirements are feasible.

t to charging speed, $E = P\Delta t$, we employ the length of one time step, $\Delta t = 1$ hour.

Due to the nature of the day-ahead market, in which bids need to be submitted a number of hours prior to actual delivery time, bidding is based on several forecasts, including EV requirements and electricity prices, as detailed in the next section.

2.4. Forecasting of Energy Requirements

As day-ahead markets close around 12h prior to electricity delivery time, the EV aggregator needs to forecast its electricity requirements 12 to 36h in advance. The amount of purchased electricity needs to be enough to satisfy its clients' needs, at the cheapest possible price. Energy requirements are defined in a similar fashion as in [13, 14].

In more detail, the requirements of each EV driver, as described in the Section 2.3, translate into two requirement vectors, $\mathbf{r}^{\min,i}$ and $\mathbf{r}^{\max,i}$, each with 24 entries. Specifically, $r_t^{\min,i}$ is the amount of energy needed at hour t assuming charging has been left for the last possible moment and that the charging requirements need to be fulfilled. Conversely, $r_t^{\max,i}$ is the amount of energy needed at hour t assuming charging starts as soon as possible. For example, consider an EV arriving at the EV aggregator at 3pm and stating 9pm departure time and 8kWh charging needs with $P_{\max} = 3\text{kW}$. Then, $\mathbf{r}^{\min,i}$ would be as specified in Table 1. Specifically, charging can happen at any time, at a maximum rate of $P_{\max} = 3\text{kW}$, but if 6pm is reached with no charging done, at least 2kW of energy needs to be charged between 6-7pm in order to fulfil the EV driver requirements. The same applies with 3kW between 7-8pm and 8-9pm. Similarly, for the same scenario, the requirement vector $\mathbf{r}^{\max,i}$ would be as specified in Table 2.

$r_3^{\min,i}$	$r_4^{\min,i}$	$r_5^{\min,i}$	$r_6^{\min,i}$	$r_7^{\min,i}$	$r_8^{\min,i}$	$r_9^{\min,i}$
0	0	0	2	3	3	0

Table 1: Example of requirement vector $\mathbf{r}^{\min,i}$

$r_3^{\max,i}$	$r_4^{\max,i}$	$r_5^{\max,i}$	$r_6^{\max,i}$	$r_7^{\max,i}$	$r_8^{\max,i}$	$r_9^{\max,i}$
3	3	2	0	0	0	0

Table 2: Example of requirement vector $\mathbf{r}^{\max,i}$

Then, two global energy requirement vectors, \mathbf{R}^{\min} and \mathbf{R}^{\max} , can be obtained by adding together the hourly requirements of all the aggregated EVs, *i.e.* $R_t^{\min} = \sum_{i=1}^N r_t^{\min,i}$ and $R_t^{\max} = \sum_{i=1}^N r_t^{\max,i}$.

In order to make informed bids in the day-ahead market, several quantities need to be forecasted by the EV aggregator, which are denoted by placing a hat over the variable name: hourly energy requirements, \hat{R}_t^{\min} and \hat{R}_t^{\max} , hourly number of available plugged-in EVs, \hat{N}_t , and hourly price impact functions, $\hat{\mathcal{P}}_t$.

Considering advanced forecasting is outside of the scope of this paper and simple forecasting is employed henceforth. Specifically, we consider two different forecasting approaches: a naive method, in which historical data from the day before is the forecast for the day after, and a perfect forecast, in which exact next day information is available [13, 17]. Furthermore, as shown in Section 4, forecasting is not essential for the correct working of our algorithms, apart from when considering very small fleet sizes.

2.5. Price-Maker Day-Ahead Bidding Algorithm

Now that the characteristics of the day-ahead market and the role of the EV aggregator have been defined, we proceed to detail our proposed price-maker day-ahead bidding algorithm. This algorithm is an extension of the optimisation algorithm proposed in [13], using residual supply curves to account for the price impact of the EV aggregator's energy bids.

The exact problem is as follows. Given an EV aggregator with capacity for N EVs participating in the day-ahead market, find the optimal distribution of energy quantities to bid across the 24 hourly slots of the next day, $\{E_0, \dots, E_{23}\}$, in order to satisfy its clients' charging needs while minimising the total cost of the purchased energy. We assume that the agent's bids are set at maximum price, p_{\max} , in order to guarantee execution.

Formally, the proposed optimization problem is defined as follows:

$$\min_{\{E_t\}} \sum_t \hat{\mathcal{P}}_t(E_t) \cdot E_t \quad (1)$$

$$\sum_{j=0}^t E_j \geq \sum_{j=0}^t \hat{R}_j^{\min}, \quad \forall t = 0, \dots, 23 \quad (2a)$$

$$\sum_{j=0}^t E_j \leq \sum_{j=0}^t \hat{R}_j^{\max}, \quad \forall t = 0, \dots, 23 \quad (2b)$$

$$E_t / \Delta t \leq \hat{N}_t P_{\max}, \quad \forall t = 0, \dots, 23 \quad (2c)$$

In more detail, the objective function (1) minimizes the total cost of the purchased energy. As explained in Sections 2.2 and 2.4, $\hat{\mathcal{P}}_t(E_t)$ is the forecasted price at hour t if the EV aggregator bids a quantity E_t . The constraints guarantee that the amount of purchased energy is enough to satisfy the forecasted demand (2a), that it is not purchased before the forecasted arrival of the EVs (2b) and that the energy purchased at each hour is not greater than the amount that the aggregator is able to charge at the given hour, based on the forecasted number of available vehicles (the aggregator cannot store energy). It is worth noting that the number of constraints is always 72, independent on the fleet size. In contrast with the original algorithm without price impact, the novel optimization problem given by Eqs. 1, 2a, 2b, 2c is no longer linear, making the problem more complex to solve.

Specifically, due to the arbitrary nature of real price impact functions obtained from real market historical data, the optimisation problem is non-convex and non-linear. Moreover, it presents a complex optimisation landscape with multiple local minima, which often prevents solving algorithms from obtaining the global minimum. In this paper, we compare two different optimisation approaches. First, we consider a basin-hopping algorithm [21], which sequentially solves the optimisation problem with different initial guesses, choosing the best solution. In addition, we consider an approximated convex bidding algorithm in which any local minima is actually the global minimum. The advantage of this formulation is that the global minimum is guaranteed to be found, which has attractive theoretical properties for the coordination mechanism proposed in Section 3.

In detail, we approximate the price impact functions, $\hat{\mathcal{P}}_t$ by fitting a convex quadratic curve, $\hat{\mathcal{P}}_t^{\text{convex}} = a_t E_t^2 + b_t E_t + \hat{\mathcal{P}}_t(0)$, where all the coefficients a_t and b_t are restricted to be positive. This guarantees the convexity of the optimisation problem, which in mathematical terms takes the following form:

$$\min_{\{E_t\}} \sum_t \hat{\mathcal{P}}_t^{\text{convex}}(E_t) \cdot E_t \quad (3)$$

$$\sum_{j=0}^t E_j \geq \sum_{j=0}^t \hat{R}_j^{\min}, \quad \forall t = 0, \dots, 23 \quad (4a)$$

$$\sum_{j=0}^t E_j \leq \sum_{j=0}^t \hat{R}_j^{\max}, \quad \forall t = 0, \dots, 23 \quad (4b)$$

$$E_t / \Delta t \leq \hat{N}_t P_{\max}, \quad \forall t = 0, \dots, 23 \quad (4c)$$

This resulting optimisation problem is then convex, where theoretical global optimality is guaranteed, and can be readily solved with little computational burden by using, for example, the SLSQP algorithm [22]. How-




Figure 2: Real and approximated price impact curves, showing a good (LHS) and a poor (RHS) approximation. Source: OMIE, 01/11/2016, LHS: 3rd hour, RHS: 12th hour.

ever, this convex approximation introduces a deviation from real data which will affect the accuracy of the solution. The degree of deviation depends on the considered day and hour, ranging from small to considerable in some cases, as shown in Figure 2. As shown in the case study (Section 4.4.1), solving the original algorithm with a SLSQP algorithm together with basin-hopping, provides lower costs than solving the convex approximation introduced in this section. However, these cost differences are small. We will henceforth refer to these two formulations of the bidding algorithm as *raw* and *convex*, respectively.

3. Inter-EV Aggregator Coordination

We are now ready to introduce the novel inter-EV aggregator coordination mechanism. We propose the use of an external entity, a third-party coordinator, which incentivises the participating aggregators to truthfully report their energy needs, and bids on their behalf in the day-ahead market. Specifically, it proceeds in three stages. Firstly, the individual requirements reported by the EV aggregators are aggregated, and the coordinator applies a bidding algorithm to obtain day-ahead energy in bulk for all the participants. Secondly, it distributes the obtained energy across the participants taking into account their individual constraints. Thirdly,

it computes payments in a way that promotes cooperation, instead of speculation to increase personal benefit in detriment of the other participants, by employing results from the field of mechanism design.

This proposed coordination mechanism is generic and independent on the particular bidding algorithm employed. However, the simple constraint formulation and little computational burden characteristic of the novel bidding algorithm introduced in Section 2 guarantee linear scaling with number of participating aggregators and fleet size. It is also worth noting that, in order to operationalise the coordination mechanism, there is no need for any additional infrastructure; it suffices to have a form of communication between each participating aggregator and the coordinator, such as an Internet connection. We will now proceed to detail each of the three stages, and the proposed implementation in practice.

3.1. Coordinated Bidding

In this first stage, the coordinator receives and combines the requirements reported by each of the participants. Let n be the number of EV aggregators, $\hat{R}_t^{\min,i}$ and $\hat{R}_t^{\max,i}$ aggregator i 's forecasted energy requirements for hour t , and \hat{N}_t^i the number of available EVs from aggregator i , as specified in Section 2.4. The combined requirements of all the aggregators are then:

$$\hat{R}_t^{\min} = \sum_{i=1}^n \hat{R}_t^{\min,i} \quad (5)$$

$$\hat{R}_t^{\max} = \sum_{i=1}^n \hat{R}_t^{\max,i} \quad (6)$$

$$\hat{N}_t = \sum_{i=1}^n \hat{N}_t^i \quad (7)$$

To find the optimal global energy bids, the bidding optimisation algorithm given by Eqs. 3, 4a, 4b, 4c can be applied with constraints given by the combined requirements (5), (6) and (7). This will result in obtaining a global day-ahead energy volume E_t^{global} for each hour t .

3.2. Hourly Energy Distribution

In the second stage, the purchased energy needs to be distributed among the participating EV aggregators, based on their reported requirements. In particular, when employing the novel bidding algorithm from Section 2.5, the energy distribution problem can be formulated as follows. Letting E_t^i be the amount of energy allocated to EV aggregator i at time h , find E_t^i for

$t = 0, \dots, 23$ and $i = 1, \dots, n$ satisfying the following constraints:

$$\sum_{j=0}^t E_j \geq \sum_{j=0}^t \hat{R}_j^{\min}, \forall t = 0, \dots, 23; \forall i = 1, \dots, n \quad (8a)$$

$$\sum_{j=0}^t E_j \leq \sum_{j=0}^t \hat{R}_j^{\max}, \forall t = 0, \dots, 23; \forall i = 1, \dots, n \quad (8b)$$

$$E_t / \Delta t \leq \hat{N}_t^i P_{\max}, \forall t = 0, \dots, 23; \forall i = 1, \dots, n \quad (8c)$$

$$\sum_{i=1}^n E_t^i = E_t^{\text{global}}, \forall t = 0, \dots, 23 \quad (8d)$$

This is a so-called constraint satisfaction problem (CSP), in which the allocation of the existing resources that satisfies the imposed constraints needs to be found [23, 24]. In more detail, Eqs. 8a, 8b, 8c ensure that each EV aggregator has enough energy to satisfy its requirements (no more, no less) for each hour. Eq. 8d makes sure the sums of the allocated hourly energies add up to that available global energy.

Existence of a solution is guaranteed by definition of the optimization problem, Eqs. 3, 4a, 4b, 4c. However, uniqueness is not guaranteed. For example, consider the case with two EV aggregators who report the same preferences to the coordinator; energy purchasing will be spread along a number of different hours by the bidding algorithm, and given that the two aggregators have the same requirements, hourly energy can be distributed in an large number of ways, up to the Wh, or the smallest energy unit considered in the system. More precisely, if the coordinator's non-zero bids are $E_1 = 200$ and $E_2 = 300$, we could do an even distribution $E_1^1 = E_1^2 = 100$ and $E_2^1 = E_2^2 = 150$, or any other combination $E_1^1 = 100 + e$; $E_1^2 = 100 - e$ and $E_2^1 = 150 - e$; $E_2^2 = 150 + e$ for each $e \in [0, 100]$. Nonetheless, every possible distribution is entirely valid and will satisfy each EV aggregator's reported requirements.

This distribution problem given by Eqs. 8a, 8b, 8c and 8d grows linearly with the number of participating EV aggregators, n , being the number of constraints $72n + 24$. With this particular choice of model constraints, given by the proposed bidding algorithm, computation time scales linearly with the number of EV aggregators and fleet size, and presents very little computational burden. This constraint problem can be adapted to other bidding models by employing their requirements definition, and their constraints. However, the performance and scalability will depend on the particular model considered.

3.3. Payment Distribution

So far, the coordinator bids in the day-ahead market and distributes the energy to the individual EV aggregators. The third stage decides how to compute payments appropriately, so that each aggregator pays an appropriate price for the electricity it has obtained, and strategical manipulation is prevented or minimised. We propose using mechanisms from the Vickrey-Clarke-Groves (VCG) family [25, Ch. 9] which incentivise cooperation across the participating aggregators. In these mechanisms, the payments reflect the marginal cost that each EV aggregator incurs on the total overall cost. However, in our experiments (discussed in Section 4) we find that the typical VCG mechanism, known as VCG with Clarke pivot payments [25] and henceforth referred to here as *pure* VCG, results in too high payments in this setting. To alleviate this problem, we consider two so-called redistribution mechanisms, which attempt to redistribute some of the pure VCG payments back to the aggregators. The three payment mechanisms are detailed next.

3.3.1. Pure VCG

The VCG family of mechanisms are a classical approach from the field of mechanism design and have the desirable property of being *truthful* under certain conditions, which in our case means that reporting true requirements will yield each aggregator the best benefit and there is no rational incentive for cheating [18, 25]. This is essential, as otherwise participants could develop strategies to *cheat* the system, by reporting false preferences to the coordinator if a greater personal benefit is foreseen. Now, VCG requires global optimality³ in order to guarantee truthfulness which, from the perspective of the EV aggregators, is achieved when using the convex formulation of the proposed bidding algorithm. It is worth noting that solving the convex approximated problem does not actually yield a globally optimal solution to the bidding problem, due to the deviations produced by the approximation. However, both the allocation and the payments are calculated based on the approximated function, which EV aggregators are unable to influence. Hence, they have no opportunity to change the outcome to their benefit by being untruthful about their requirements.

Formally, let $\hat{\mathbf{R}}^{\min} = (\hat{R}_0^{\min}, \dots, \hat{R}_{23}^{\min})$, $\hat{\mathbf{R}}^{\max} = (\hat{R}_0^{\max}, \dots, \hat{R}_{23}^{\max})$ and $\hat{\mathbf{N}} = (\hat{N}_0, \dots, \hat{N}_{23})$ be the vectors of aggregated requirements. Let $E(\hat{\mathbf{R}}, \hat{\mathbf{N}}) =$

³More precisely, the allocation needs to be *optimal in range* [?].

$(E_0^{\text{global}}, \dots, E_{23}^{\text{global}}) = \mathbf{E}^{\text{global}}$ be the energy schedule provided by the bidding algorithm (3) with requirements $\hat{\mathbf{R}}^{\min}$, $\hat{\mathbf{R}}^{\max}$ and $\hat{\mathbf{N}}$. Let $\mathcal{P}^{\text{convex}}(\mathbf{E}) = (P_0^{\text{convex}}, \dots, P_{23}^{\text{convex}}) = \mathbf{P}^{\text{convex}}$ be the forecasted clearing prices using the convex approximation. Let $\mathcal{P}(\mathbf{E}) = (P_0, \dots, P_{23}) = \mathbf{P}$ be the real hourly clearing prices after market closure. Then, the total aggregated forecasted cost incurred by the coordinator is given by:

$$\text{Cost}^{\text{convex}}(\hat{\mathbf{R}}^{\min}, \hat{\mathbf{R}}^{\max}, \hat{\mathbf{N}}) = E(\hat{\mathbf{R}}^{\min}, \hat{\mathbf{R}}^{\max}, \hat{\mathbf{N}}) \cdot \mathcal{P}^{\text{convex}}(E(\hat{\mathbf{R}}^{\min}, \hat{\mathbf{R}}^{\max}, \hat{\mathbf{N}}))$$

and the real total aggregated cost incurred by the coordinator is given by:

$$\text{Cost}(\hat{\mathbf{R}}^{\min}, \hat{\mathbf{R}}^{\max}, \hat{\mathbf{N}}) = E(\hat{\mathbf{R}}^{\min}, \hat{\mathbf{R}}^{\max}, \hat{\mathbf{N}}) \cdot \mathcal{P}(E(\hat{\mathbf{R}}^{\min}, \hat{\mathbf{R}}^{\max}, \hat{\mathbf{N}}))$$

Given the use of the convex approximation, in general these two cost functions are not equal:

$$\text{Cost}(\hat{\mathbf{R}}^{\min}, \hat{\mathbf{R}}^{\max}, \hat{\mathbf{N}}) \neq \text{Cost}^{\text{convex}}(\hat{\mathbf{R}}^{\min}, \hat{\mathbf{R}}^{\max}, \hat{\mathbf{N}})$$

Furthermore, let $\hat{\mathbf{R}}_{-i}^{\min}$, $\hat{\mathbf{R}}_{-i}^{\max}$ and $\hat{\mathbf{N}}_{-i}$ be the vectors of requirements without the contribution from the i -th EV aggregator. Then, the VCG payment of EV aggregator i to the coordinator, p_i , is given by:

$$p_i = \text{Cost}^{\text{convex}}(\hat{\mathbf{R}}^{\min}, \hat{\mathbf{R}}^{\max}, \hat{\mathbf{N}}) - \text{Cost}^{\text{convex}}(\hat{\mathbf{R}}_{-i}^{\min}, \hat{\mathbf{R}}_{-i}^{\max}, \hat{\mathbf{N}}_{-i}) \quad (9)$$

Note that, to compute VCG payments, in case of n EV aggregators, the coordinator needs to compute the optimal solution $n + 1$ times (once to compute the total costs when all EV aggregators are present, and then once for each EV aggregator when they are removed from the market). Fortunately, due to the scalability of our bidding strategy, we can easily compute VCG payments for very large settings.

The VCG payments effectively mean that the price paid by each participant reflects the impact that its requirements have on the overall constraints, *i.e.* an aggregator which reports more flexible requirements will pay less than another with tighter constraints, which is a very desirable property when considering fairness of the payments. However, there are other properties to consider. First, the question is whether the coordinator will run into a *deficit*, meaning that the sum of money received from the EV aggregators is insufficient to cover the payments for the electricity incurred by the coordinator. In order to maintain the desirable properties of VCG mechanisms, *i.e.* truthfulness, the allocations and

payments received from the EV aggregators are based on the estimated price impact curves. These estimated are due to both the approximated convex function and because it uses a *predicted* price curve. As a result, the coordinator can, depending on the accuracy of the prediction, incur a loss.

However, as shown in Section 4, we find that the payments are actually often too high, even higher than the costs the EV aggregators would incur by not participating in the mechanism. Technically, the mechanism is said to violate so-called *individual rationality* [25], also known as the participation constraint (meaning EV aggregators have no incentive to participate if given the choice). Hence, to address this, we consider so-called payment redistribution mechanisms which, in the case of a not-for-profit coordinator, should ideally result in a surplus close to zero.

3.3.2. VCG-Based Truthful Redistribution

In recent years, the problem of redistributing some or all the surplus of pure VCG payments, while preserving truthfulness, has been extensively studied [26, 27]. Formally, given the payments defined by Eq. 9, the redistributed amount to participant i is given by:

$$r_i = \frac{1}{n} \min_{(\hat{\mathbf{R}}_{-i}^{\min}, \hat{\mathbf{R}}_{-i}^{\max}, \hat{\mathbf{N}}_{-i})} \text{VCG} \left[(\hat{\mathbf{R}}^{\min}, \hat{\mathbf{R}}^{\max}, \hat{\mathbf{N}}_i), (\hat{\mathbf{R}}_{-i}^{\min}, \hat{\mathbf{R}}_{-i}^{\max}, \hat{\mathbf{N}}_{-i}) \right] = \frac{1}{n} \text{VCG}(\hat{\mathbf{R}}_{-i}^{\min}, \hat{\mathbf{R}}_{-i}^{\max}, \hat{\mathbf{N}}_{-i})$$

where $\text{VCG}()$ is the total VCG payment collected from those reports, *i.e.*:

$$\text{VCG}(\hat{\mathbf{R}}_{-i}^{\min}, \hat{\mathbf{R}}_{-i}^{\max}, \hat{\mathbf{N}}_{-i}) = \sum_{\substack{j=1 \\ j \neq i}}^n \left[\text{Cost}^{\text{convex}}(\hat{\mathbf{R}}_{-i}^{\min}, \hat{\mathbf{R}}_{-i}^{\max}, \hat{\mathbf{N}}_{-i}) - \text{Cost}^{\text{convex}}(\hat{\mathbf{R}}_{-i-j}^{\min}, \hat{\mathbf{R}}_{-i-j}^{\max}, \hat{\mathbf{N}}_{-i-j}) \right]$$

This way, the total payment of each EV aggregator i is given by:

$$p_i = \left[\text{Cost}^{\text{convex}}(\hat{\mathbf{R}}^{\min}, \hat{\mathbf{R}}^{\max}, \hat{\mathbf{N}}) - \text{Cost}^{\text{convex}}(\hat{\mathbf{R}}_{-i}^{\min}, \hat{\mathbf{R}}_{-i}^{\max}, \hat{\mathbf{N}}_{-i}) \right] - \frac{1}{n} \sum_{\substack{j=1 \\ j \neq i}}^n \left[\text{Cost}^{\text{convex}}(\hat{\mathbf{R}}_{-i}^{\min}, \hat{\mathbf{R}}_{-i}^{\max}, \hat{\mathbf{N}}_{-i}) - \text{Cost}^{\text{convex}}(\hat{\mathbf{R}}_{-i-j}^{\min}, \hat{\mathbf{R}}_{-i-j}^{\max}, \hat{\mathbf{N}}_{-i-j}) \right] \quad (10)$$

Note that, for this redistribution mechanism, optimal bids need to be computed in the order of n^2 times,

requiring significant additional computation. Furthermore, depending on the particular scenario, the redistribution provided by this technique can range from full to small proportions [27]. In our experiments (see Section 4), we find that the redistribution is too large, so the money paid by the aggregators is not enough to cover the costs for their purchased energy, and the coordinator incurs large losses. Therefore, we explore another redistribution mechanism, which achieves zero surplus and losses for the coordinator by sacrificing theoretical truthfulness.

3.3.3. VCG-Based Proportional Redistribution

One such non-truthful redistribution mechanisms is to redistribute the real monetary surplus among the group of EV aggregators proportionally to their size. This real monetary surplus is given by $\sum_{i=1}^n p_i - \text{Cost}(\hat{\mathbf{R}}^{\min}, \hat{\mathbf{R}}^{\max}, \hat{\mathbf{N}})$, the sum of all payments minus the real cost of the energy paid by the coordinator to the day-ahead market. In more detail, letting N_i be the vehicle capacity of EV aggregator i , the resulting payments are given by a two-stage algorithm, in which first VCG payments are computed, p_i^1 , and then the proportional redistribution takes place, providing the final payments assigned to each aggregator, p_i^2 . Specifically:

$$p_i^1 = \text{Cost}^{\text{convex}}(\hat{\mathbf{R}}^{\min}, \hat{\mathbf{R}}^{\max}, \hat{\mathbf{N}}) - \text{Cost}^{\text{convex}}(\hat{\mathbf{R}}_{-i}^{\min}, \hat{\mathbf{R}}_{-i}^{\max}, \hat{\mathbf{N}}_{-i})$$

$$p_i^2 = p_i^1 - \frac{N_i}{\sum_{j=1}^n N_j} \left[\sum_{j=1}^n p_j^1 - \text{Cost}(\hat{\mathbf{R}}^{\min}, \hat{\mathbf{R}}^{\max}, \hat{\mathbf{N}}) \right] \quad (11)$$

It is worth noting that, in contrast to the redistribution approach discussed in Section 3.3.2, this payment mechanism scales linearly with the number of participating EV aggregators. The loss of truthfulness means that strategic misreporting of requirements by an EV aggregator to the coordinator could improve an aggregator's performance. For example, it could be the case that misreporting less flexible energy requirements, resulting in earlier energy allocation at a more expensive price, provides the aggregator with earlier energy while the cost excess will be partially absorbed by other aggregators. In practise, however, the bulk of each aggregator's payments are computed using VCG, which is truthful. This means that the tighter the reported requirements, the more expensive the payment gets, so the room for manipulation is slim. Hence, in the case study presented in Section 4, we assume aggregators are truthful when using this mechanism. Analysis on the extent of strategic manipulation possible by this payment mechanism is left for future work.

3.4. The Coordination Mechanism in Practice

Now that the three stages of the coordination mechanism have been described, we proceed to discuss how it could be implemented in practice. As noted earlier, the coordination mechanism does not need any additional infrastructure, relying only on the aggregator's existing electrical infrastructure and an Internet connection. In more detail, the proposed implementation of the three-stage coordination mechanism works as follows:

1. Coordinated bidding: the participating EV aggregators securely and privately submit their electricity requirements to the coordinator over the Internet. Given the low computational cost of the proposed bidding algorithm and coordination mechanism, the coordinator is then able to use a standard computer to calculate the optimal bidding schedule, as shown in Section 3.1. The bids are submitted to the day-ahead market online.
2. Hourly energy distribution: once the market is cleared and an electricity schedule has been allocated to the coordinator, the hourly energy distribution algorithm detailed in Section 3.2 can be applied, using the same standard computer. Energy delivery to the aggregators is then managed by the distribution system operator (DSO) [13] in the same way as in the individual bidding case.
3. Payment distribution: payments are computed by the coordinator (see Section 3.3), and reported to the participating aggregators, which can then readily process them online.

By using this simple implementation, which requires no infrastructure investment, an EV aggregator can seamlessly transition from individual to cooperative operation. Similarly, the role of the coordinator can be easily assumed by any trusted third-party, such as a government agency or a private company.

4. Case Study

This section describes a case study based on real data⁴, in order to test the performance of the bidding strategy, and the coordination and payment mechanisms introduced in Sections 2 and 3, respectively. We will describe the simulation characteristics first, detailing the considered EV aggregators, market data and real driver behavioural data. Then, the results concerning the performance of the proposed algorithms are presented and described.

⁴All data supporting this case study are publicly available at [28].

4.1. EV Aggregator Characteristics

This case study considers a night-time residential scenario in which EVs arrive in the evening and need to be charged by the next morning. With respect to the EV fleet, we consider medium-sizes electric vehicles with battery capacities of 24kWh [29, 20]. Charging speed is considered to be the same for all EVs and set to $P_{\max} = 3.7\text{kW}$ [20, 13]. Charging efficiency is considered to be 90%, meaning that 10% of the consumed electricity is lost and does not contribute to the charging of the battery [30, 31].

The considered fleet sizes range from 100 000 to 3 000 000 EVs. These values correspond to EV penetration rates of 0.27% and 8.14% in the Iberian Peninsula, whose electricity market is considered in the simulations. Note that European targets for EV penetration include, for example, 10% in the UK by 2020 [5].

4.2. Day-Ahead Market Data

Real market data from the Iberian day-ahead market, OMIE, is employed in the simulations [32]. Detailed order data is available online, containing hourly supply and demand data down to individual order level. All simulations utilize OMIE weekday market data from November and December 2016, and January and February 2017. Weekends are removed in order to eliminate weekly seasonality [19].

Residual supply curves are employed to build price impact functions as described in Section 2.2. In order to obtain the convex approximation of price impact functions described in Section 2.5, we proceed as follows. The interval 0-10 GWh is discretised in 50 equally spaced points, and the clearing price at each point is computed. Then, the corresponding quantity-clearing price pairs can be interpolated by using a quadratic curve as described in Section 2.5. A more detailed description about the pre-processing of the historical market data can be found in Appendix A.

4.3. Driver Behavioural Data

To model the behaviour of the EVs in our simulations, real data from the Spanish driver behaviour survey MOVILIA [33] is employed. This survey studies a plethora of driving patterns, among which are the average number of daily commutes, length and time of this trips, *etc.* Specifically, driver behaviour for our simulations is based on [34], which provides a convenient analysis of MOVILIA data which determines the distribution of times for the first and last trip from and to home. These distributions are given in Tables 3 and 4. To account for driver mobility, each EV will make use

of the aggregator’s services with 80% probability every day.

With respect to energy requirements, the state of charge of an EV at arrival and departure times is drawn from a uniform distribution as follows: $\text{SoC}_0 \in [\text{SoC}_{\text{total}}/4, \text{SoC}_{\text{total}}/2]$ and $\text{SoC}_f \in [2 \cdot \text{SoC}_{\text{total}}/3, \text{SoC}_{\text{total}}]$. Consequently, the EV charging requirements range between a large percentage of the battery (up to 75%), to a small percentage (down to 16%), accounting for long and short trips home, a choice consistent with the literature [30].

Time	19h	20h	21h	22h	23h
Probability	0.16	0.25	0.32	0.12	0.15

Table 3: Possible arrival times rounded to the nearest hour, with their respective probabilities.

Time	6h	7h	8h	9h	10h
Probability	0.04	0.02	0.34	0.5	0.1

Table 4: Possible departure times rounded to the nearest hour, with their respective probabilities.

The results from the simulations are described next, starting with the performance of the bidding algorithm proposed in Section 2 and then looking at the coordination mechanism proposed in Section 3.

4.4. Results: Performance of the Bidding Algorithm

In this section we first compare the two formulations of the novel price-maker day-ahead bidding algorithm proposed in Section 2, namely the raw problem with basin-hopping, and the convex approximation. Our experiments confirm the appropriateness of the convex approximation, and hence we focus on this bidding strategy in subsequent experiments. Then, the algorithm’s performance is assessed by direct comparison with two existing algorithms, in the night-time residential scenario detailed in Section 4.1.

4.4.1. Comparing the Raw and Convex Formulations

We now compare the performance of the two formulations of our proposed bidding algorithm: the raw formulation given by Eqs. 1, 2a, 2b, 2c and the convex approximation given by Eqs. 3, 4a, 4b, 4c. Both formulations are solved using the Sequential Least Squares Programming (SLSQP) algorithm for solving quadratic programmes, and we used the implementation from [35]. In the raw case, we present two separate results, by solving the optimisation problems with and without employing a basin-hopping algorithm (see Section 2.5).

In both cases, the initial guess is set to the optimal solution provided by the algorithm without price impact (which has a linear objective and can easily be solved optimally [13]).

Specifically, we present a comparison of EV aggregators of different sizes employing both formulations and the naive forecasting technique, during four months of day-ahead trading, as specified in Section 4.2. The results are shown in Figure 4a. As seen for small EV fleet sizes, both formulations achieve very similar results. However, as the price impact of the EV aggregator grows, the raw formulation without basin-hopping struggles to find the global minimum, resulting in worse solutions than the convex approximation. The raw formulation with basin-hopping explores the solution space and consistently achieves better solutions than the convex approximation. However, these cost differences are very small. This slight performance loss justifies employing the convex approximation in order to make use of its desirable theoretical properties for the VCG mechanism. Thus, we will focus on the convex formulation of the bidding algorithm from now on.

4.4.2. Comparison with Benchmark Algorithms

We now compare the proposed bidding algorithm using convex approximation with two existing bidding algorithms. The first of these algorithms is a simple zero-intelligence, *dumb*, algorithm which acts as a lower performance bound and simply charges each EV as soon as it becomes available, without any further consideration to future prices or requirements [13, 34]. The second is the price-taker algorithm proposed by [13], which inspired our proposed price-maker algorithm. We will henceforth refer to these strategies as *dumb* and *noPI* respectively, and to our strategy with price impact as *PI*. Moreover, depending on the employed forecast, naive or perfect (see Section 2.4), we will refer to the PI strategy as *PI naive* and *PI perfect*.

To demonstrate the behaviour of each algorithm, an example of the bidding decisions of the three compared strategies for a given day is presented in Figure 3. Here, we consider an EV aggregator managing a fleet of one million EVs, corresponding to an EV penetration of 2.71% in the Iberian Peninsula. For this magnitude of EV penetration, price impact is already very pronounced and plays an essential role in the bidding process. Figure 3 provides important insight into the behaviour of each of the three strategies. Specifically, the dumb strategy forecasts that EVs come home in the evening and tries to charge their batteries as soon as they are available. At this time of the day, prices are highest and large orders from the EV aggregator push them even

higher. The second strategy, noPI, is able to take advantage of the fact that EVs stay idle during the night and only need to be charged by the morning. By forecasting the EV requirements as explained in Section 2.4, it obtains the bidding schedule that provides the cheapest hourly prices. However, by neglecting price impact, the bids are mainly concentrated in three hours, and price impact is high. Lastly, our proposed strategy is able to mitigate price impact by spreading its bids in time, achieving lower costs.

Next, we consider the average results over the entire four-month time period. Specifically, we compare the average daily payments per EV for each strategy as we vary the size of the EV fleet. Results are shown in Figure 4b. In more detail, the dumb bidding strategy provides the highest payments, due to its inability to consider hourly electricity prices. The results of the noPI and PI strategies are comparable when the size of the EV fleet is moderate, but as the fleet gets larger, the noPI strategy incurs significantly increased payments, approaching the results from the dumb strategy for large fleets. In contrast, the PI strategy is able to maintain a sub-linear payment increase as the fleet grows, providing consistent payment reduction. In more detail, the payment reduction percentage between noPI and PI strategies is shown in Figure 4c. We can see that, once the price impact of the EV fleet becomes appreciable, with a size of around 500 000 vehicles, the improvement percentage grows linearly with fleet size, achieving around 10% improvement for EV penetration values around 8%, and higher percentages for larger fleets.

Also, it is important to note the performance difference between the perfect and the naive forecasts (see Section 2.4). The average payment improvement comparing the results of the proposed strategy with price impact with perfect and naive forecasts is around 2%. The importance of the forecasting technique is more crucial for small fleet sizes, where a poor forecast can cause the PI strategy to perform slightly worse than the noPI one. This is due to the PI strategy spreading bidding to hours with forecasted mid-range prices which, with a poor forecast, can have high prices in reality. In contrast, when employing the perfect forecast, the PI strategy consistently outperforms noPI.

4.5. Performance of the Coordination Mechanism

In this section we consider the setting where a number of self-interested and independent EV aggregators participate in the same day-ahead market, as described in Section 4.1. To this end, we compare the setting with *uncoordinated* bidding, where each EV aggregator does independent bidding based on past historical data,




Figure 3: Comparison of the three bidding strategies for a single day: dumb, without price impact and with price impact for the same scenario where the aggregators manage one million EVs. Curves represent hourly prices, both with and without the EV aggregator bids, and correspond to the LHS axis. Bars represent the amount of energy purchased at each hourly slot by the EV aggregator and correspond to the RHS axis. Market data corresponds to OME, 09-10/11/2016.

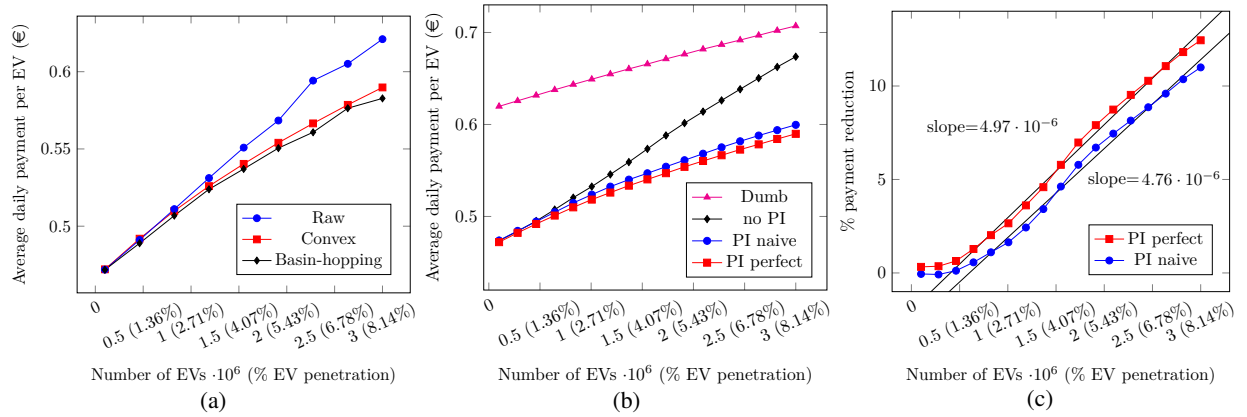


Figure 4: (a) Average daily payments per EV, for both bidding algorithm formulations and for different EV aggregator sizes, employing perfect forecasts. (b) Average daily payments per EV, for each different bidding strategy and for different EV aggregator sizes. Both with naive and perfect forecasts. (c) Percentage payment reduction when using the novel algorithm with price impact, compared to its counterpart without price impact.




Figure 5: (a) Average daily payments per EV when using a perfect forecast. (b) Average daily monetary surplus for the coordinator when using a perfect forecast. (c) Percentage payment reduction when using the proportional redistribution mechanism w.r.t. uncoordinated bidding, using both naive and perfect forecasts.

and *coordinated* bidding, where all the EV aggregators participate in the coordination mechanism presented in Section 3. Furthermore, for the coordinated setting, we

compare the three different payment mechanisms, *i.e.*, the VCG mechanism (Eq. 9), the truthful redistribution (Eq. 10) and proportional redistribution (Eq. 11).

Specifically, Figures 5a and 5b show the average EV payments and coordinator surplus respectively for the various mechanisms. We can see that the VCG payments are higher than the uncoordinated ones, and hence the coordination mechanism with VCG payments does not offer any advantage to the aggregators. Conversely, the coordination mechanism with truthful redistribution generates too low payments, causing the coordinator to incur very large monetary losses. These results indicate the inappropriateness of these two payment mechanisms, which incur too large and too little payments, respectively. At the same time, the proportional redistribution mechanism demonstrates a significant reduction in payments compared to the uncoordinated setting.

Considering the proportional redistribution mechanism in more detail, Figure 5c shows the average daily payment improvement compared to uncoordinated bidding, employing both the naive and perfect forecasts. We can see that the payment reductions w.r.t. uncoordinated bidding grow linearly with fleet size and number of participating EV aggregators. Again, when employing the naive forecast, and for small fleet sizes, the coordination mechanism presents a very small performance reduction when compared to uncoordinated bidding. However, for around 8% EV penetration values (see last data points in Figure 5c), we already see that the coordination mechanism achieves around a 10% payment reduction.

5. Conclusion and Discussion

This paper introduces a novel inter-EV aggregator coordination mechanism which allows cooperation among groups of self-interested aggregators bidding in the same day-ahead market. The mechanism incentivises the participating EV aggregators to truthfully report their energy requirements to a third-party coordinator, which is then able to employ a bidding algorithm to bid on their behalf. The purchased energy and the corresponding payments are then distributed among the participants. Techniques from mechanism design are employed to determine each aggregator’s payment, in a way that cooperation rather than strategic manipulation is encouraged. To this end, three different payment mechanisms are proposed.

In addition, to ensure scalability whilst taking into consideration price impact, a novel price-maker day-ahead bidding algorithm is proposed. This is formulated in terms of convex non-linear optimisation problem with simple constraints, presenting very little computational burden and scaling to very large fleet sizes.

By employing this bidding algorithm, the proposed coordination mechanism scales linearly with fleet size and number of participating aggregators, making it suitable for arbitrary large systems.

The performance of both the coordination and the bidding algorithms is assessed in a case study which uses real market and driver data from the Iberian Peninsula. Results indicate that the coordination mechanism (employing the proposed bidding algorithm) achieves substantial cost reductions, which grow linearly with fleet size and the number of participating EV aggregators. Specifically, the coordination mechanism achieves around 10% cost reduction for a scenario with around 8% EV penetration.

Lastly, it is worth noting that the proposed coordination mechanism can be adapted to consider participation in different electricity markets, such as intra-day or ancillary services. Similarly to the day-ahead scenario, the participants’ actions in this markets affect prices, and benefits could be obtained from cooperation. A relevant example can be found in a vehicle-to-grid (V2G) setup, where the energy stored in the vehicles’ batteries can be sold back to the grid in times of scarcity. V2G is seen as one of the main ambitions in the EV literature, and the application of our mechanism to this scenario is left for future work.

Acknowledgements

The authors would like to thank S. Stein and T.-D. Nguyen for helpful comments and discussions. This work was supported by an EPSRC Doctoral Training Centre grant (EP/L015382/1). All data supporting this study are openly available at [28].

Appendix A. Historical Market Data

All European markets employ the EUPHEMIA [36] clearing algorithm and allow different types of orders, some with complex requirements. Roughly, the clearing algorithm proceeds as follows. First, it aggregates all the supply and demand orders and finds the intersection point. This determines the accepted orders, and a preliminary clearing price. Secondly, the particularities of the accepted orders are taken into account, which effectively shifts the previously obtained clearing price, determining the final hourly price.

The historical data available from OMIE has two types of data. Firstly, the *offered* orders coming from all market participants for every hour and day of the year, corresponding to the first stage as defined above. Secondly, the *matched* accepted orders resulting from the

second stage of the clearing algorithm. An example of these two types of data is presented in Figure A.6a.

The procedure employed to obtain base prices and price impact functions in the case study (see Section 4) is as follows. Firstly, for each considered day and hour, the *matched* orders are aggregated together to build supply and demand curves and to determine the clearing price. This clearing price corresponds to the real hourly price. Secondly, in order to consider the price impact of the exogenous buy orders from our EV aggregators, we look at the unaccepted supply orders. Specifically, the unaccepted tail of the *offered* supply curve is appended to the intersection of the *matched* curves described above, as shown in Figure A.6b. Then, the bids from the external aggregators can be integrated in the *extended* supply and demand curves, and the intersection of the two will determine the final clearing price with the aggregators' bids.

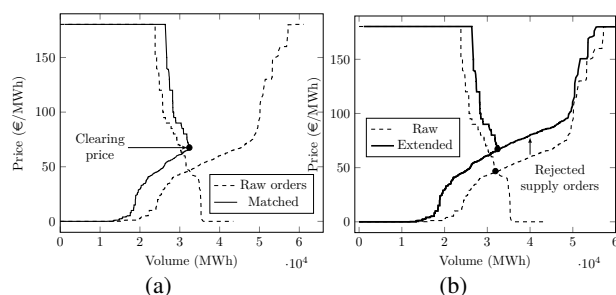


Figure A.6: (a) Raw and matched market orders, together with the hourly clearing price. (b) Appending the rejected supply orders (dashed) to the accepted matched supply curve (solid). Source: OMIE, 02/11/2016, 11th hour.

References

- [1] EIA, Monthly Energy Review, Tech. rep. (2016). doi:D0E/EIA-0035(2011/02). URL http://www.eia.gov/totalenergy/data/monthly/pdf/sec11_{_}5.pdf
- [2] ONS, UK Perspectives 2016: Energy and emissions in the UK, Tech. rep. (2016). URL <http://visual.ons.gov.uk/uk-perspectives-2016-energy-and-emissions-in-the-uk/>
- [3] POST, Trends in Energy Production., Tech. rep. (2015). URL <http://researchbriefings.parliament.uk/ResearchBriefing/Summary/POST-PN-0503>
- [4] SMMT, Electric & Alternatively-Fuelled Vehicle Registrations (2017). URL <https://www.smmt.co.uk/2017/05/april-2017-ev-registrations/>
- [5] ECCC, 2020 Renewable Heat and Transport Targets, Tech. rep. (2016). URL <https://www.publications.parliament.uk/pa/cm201617/cmselect/cmenergy/173/173.pdf>
- [6] International Energy Agency, Global EV Outlook 2016: Beyond one million electric cars, Tech. rep. (2016).

- URL [http://www.eia.gov/forecasts/aeo/pdf/0383\(2016\).pdf](http://www.eia.gov/forecasts/aeo/pdf/0383(2016).pdf)
- [7] E. S. Rigas, S. Member, S. D. Ramchurn, N. Bassiliades, Managing Electric Vehicles in the Smart Grid Using Artificial Intelligence : A Survey, Ieee Transactions on Intelligent Transportation Systems 16 (4) (2015) 1619–1635.
 - [8] W. Kempton, J. Tomic, S. Letendre, A. Brooks, T. Lipman, Vehicle-to-Grid Power: Battery, Hybrid, and Fuel Cell Vehicles as Resources for Distributed Electric Power in California, Fuel Cell IUCD-ITS-R (June) (2001) 95. URL <http://escholarship.org/uc/item/0qp6s4mb.pdf>
 - [9] M. Shafie-khah, E. Heydarian-Forushani, M. E. H. Golshan, P. Siano, M. P. Moghaddam, M. K. Sheikh-El-Eslami, J. P. S. Catalão, Optimal trading of plug-in electric vehicle aggregation agents in a market environment for sustainability, Applied Energy 162 (2016) 601–612. doi:10.1016/j.apenergy.2015.10.134. URL <http://dx.doi.org/10.1016/j.apenergy.2015.10.134>
 - [10] Y. Xu, L. Xie, C. Singh, Optimal scheduling and operation of load aggregators with electric energy storage facing price and demand uncertainties, NAPS 2011 - 43rd North American Power Symposiumdoi:10.1109/NAPS.2011.6024888.
 - [11] S. Vagropoulos, A. G. Bakirtzis, Optimal bidding strategy for electric vehicle aggregators in electricity markets, IEEE Transactions on Power Systems 28 (4) (2013) 4031–4041. doi:10.1109/TPWRS.2013.2274673.
 - [12] D. Wu, D. C. Aliprantis, L. Ying, Load scheduling and dispatch for aggregators of plug-in electric vehicles, IEEE Transactions on Smart Grid 3 (1) (2012) 368–376. doi:10.1109/TSG.2011.2163174.
 - [13] R. J. Bessa, M. A. Matos, F. J. Soares, J. A. P. Lopes, Optimized bidding of a EV aggregation agent in the electricity market, IEEE Transactions on Smart Grid 3 (1) (2012) 443–452. doi:10.1109/TSG.2011.2159632.
 - [14] Z. Liu, Q. Wu, S. Huang, L. Wang, M. Shahidehpour, Y. Xue, Optimal Day-ahead Charging Scheduling of Electric Vehicles through an Aggregative Game Model, IEEE Transactions on Smart Grid 3053 (c) (2017) 1–1. doi:10.1109/TSG.2017.2682340. URL <http://ieeexplore.ieee.org/document/7879192/>
 - [15] O. Sundström, C. Binding, Flexible charging optimization for electric vehicles considering distribution grid constraints, IEEE Transactions on Smart Grid 3 (1) (2012) 26–37. doi:10.1109/TSG.2011.2168431.
 - [16] T. K. Kristoffersen, K. Capion, P. Meibom, Optimal charging of electric drive vehicles in a market environment, Applied Energy 88 (5) (2011) 1940–1948. doi:10.1016/j.apenergy.2010.12.015. URL <http://dx.doi.org/10.1016/j.apenergy.2010.12.015>
 - [17] M. Gonzalez Vaya, G. Andersson, Optimal Bidding Strategy of a Plug-In Electric Vehicle Aggregator in Day-Ahead Electricity Markets under Uncertainty, IEEE Transactions on Power Systems 30 (5) (2015) 2375–2385. doi:10.1109/TPWRS.2014.2363159.
 - [18] N. Vlassis, A Concise Introduction to Multiagent Systems and Distributed Artificial Intelligence, 2007. doi:http://dx.doi.org/10.2200/S00091ED1V01Y200705AIM002.
 - [19] R. Herranz, A. Muñoz San Roque, J. Villar, F. A. Campos, Optimal demand-side bidding strategies in electricity spot markets, IEEE Transactions on Power Systems 27 (3) (2012) 1204–1213. doi:10.1109/TPWRS.2012.2185960.

- [20] E. H. Gerding, S. Stein, S. Ceppi, V. Robu, Online mechanism design for vehicle-to-grid car parks, IJCAI International Joint Conference on Artificial Intelligence 2016-Janua (2016) 286–293.
- [21] D. J. Wales, Energy Landscapes, Cambridge University Press, Cambridge, UK., 2003.
- [22] J. Nocedal, S. J. Wright, Numerical Optimization, Springer, 2006. doi:10.5194/bgd-11-12733-2014.
- [23] E. Tsang, Foundations of constraint satisfaction: the classic text, Books on Demand, 1996.
URL <https://books.google.com/books?hl=en&lr={&}id=UFmRAwAAQBAJ{&}oi=fnd{&}pg=PP20{&}dq=Foundations+of+constraint+satisfaction:+the+classic+text.+BoDBooks+on+Demand.{&}ots=WKEG5MR3uW{&}sig=pcR62MaPIgPUgTRGofYOCemFP0c>
- [24] I. P. Gent, C. Jefferson, I. Miguel, Minion: A Fast Scalable Constraint Solver, The 17th European Conference on Artificial Intelligence 141 (2006) 98–102. doi:http://dl.acm.org/citation.cfm?id=1567016.1567043.
- [25] N. Nisan, E. Tardos, T. Wexler, V. Vazirani, Algorithmic Game Theory, 2007. arXiv:0907.4385, doi:10.1145/1785414.1785439.
- [26] R. Cavallo, Optimal Decision-Making With Minimal Waste: Strategyproof Redistribution of VCG Payments, Aamas-06 (2006) 603–607doi:10.1145/1160633.1160790.
URL [cavallo-aamas-06.pdf](#)
- [27] M. Guo, V. Conitzer, Undominated VCG Redistribution Mechanisms, Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS) (2008) 1039–1046.
- [28] A. Perez-Diaz, E. Gerding, F. McGroarty, Dataset for “Coordination and payment mechanisms for electric vehicle aggregators” article (dec 2017). doi:10.5258/S0T0N/D0339.
URL <https://eprints.soton.ac.uk/416291/>
- [29] E. Heydarian-Forushani, M. E. H. Golshan, M. Shafie-khah, Flexible interaction of plug-in electric vehicle parking lots for efficient wind integration, Applied Energy 179 (2016) 338–349. doi:10.1016/j.apenergy.2016.06.145.
URL <http://dx.doi.org/10.1016/j.apenergy.2016.06.145>
- [30] L. Yao, W. Lim, T. Tsai, A Real-Time Charging Scheme for Demand Response in Electric Vehicle Parking Station, IEEE Transactions on Smart Grid 3053 (c) (2016) 1–1. doi:10.1109/TSG.2016.2582749.
URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7496956>
- [31] W. Su, M. Y. Chow, Computational intelligence-based energy management for a large-scale PHEV/PEV enabled municipal parking deck, Applied Energy 96 (2012) 171–182. doi:10.1016/j.apenergy.2011.11.088.
URL <http://dx.doi.org/10.1016/j.apenergy.2011.11.088>
- [32] OMIE, Historical data from the Iberian day-ahead market (2017).
URL <http://www.omie.es/en/inicio>
- [33] J. M. Roperro Ortega, Encuesta Movilidad Cotidiana, Ministerio de Fomento. (1) (2014) 1–5. doi:10.1007/s13398-014-0173-7.2.
- [34] M. Alonso, H. Amaris, J. G. Germain, J. M. Galan, Optimal charging scheduling of electric vehicles in smart grids by heuristic algorithms, Energies 7 (4) (2014) 2449–2475. doi:10.3390/en7042449.
- [35] E. Jones, E. Oliphant, P. Peterson, SciPy: Open Source Scientific Tools for Python (2001).
URL <http://www.scipy.org/>
- [36] EPEX SPOT Nord Pool OMIE OPCOM GME OTE TGE, EUPHEMIA Public Description. PCR Market Coupling Algorithm, Tech. rep. (2015).
URL <https://www.nordpoolspot.com/globalassets/download-center/pcr/euphemia-public-documentation.pdf>