

UNIVERSITY OF SOUTHAMPTON

Faculty of Engineering and Physical Sciences
Electronics and Computer Science

**Coordination Mechanisms for Electric
Vehicle Aggregators**

by

Alvaro Perez-Diaz
ORCID ID: 0000-0001-8081-0772

A thesis submitted for the degree of
Doctor of Philosophy

August 2019

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING AND PHYSICAL SCIENCES
ELECTRONICS AND COMPUTER SCIENCE

by Alvaro Perez-Diaz

Motivated by the high electric vehicle (EV) penetration percentages foreseen for the near future, we study the integration of large fleets of EVs into the existing electricity market structure, with special emphasis on day-ahead markets. More specifically, given the large energy requirements characteristic of the widespread use of EVs, the operation of an EV fleet requires careful managing in order to limit environmental and infrastructural impacts. On the other hand, a key characteristic of an EVs is their flexibility, given that most vehicles are idle most of the day. This translates into coordination opportunities which allow for the temporal spreading of the energy consumption of the fleet, avoiding unnecessary demand peaks which translate into increased electricity prices and stress in the grid. Electric vehicle coordination has received a lot of attention in the literature in recent years, with many works studying this problem under the perspective of an aggregator controlling a fleet of electric vehicles and taking coordinated charging decisions. However, we envision a scenario where a number of independent and self-interested aggregators compete in the electricity markets in order to obtain the energy needed to meet their clients' needs. Extending the benefits of coordinated operation to this multi-aggregator setting can provide reduced prices and increased grid stability, but is challenging due to the self-interested nature of the aggregators. Indeed, lower costs translate into more profit for the aggregator or cheaper charging tariffs for their customers. Moreover, an additional challenge arises once the necessary energy has been purchased. In more detail, the aggregator needs to decide on how to allocate the available energy among the available EVs, without knowledge about future arrivals.

In order to address these challenges, the first part of this thesis proposes a novel price-maker day-ahead bidding algorithm for electric vehicle aggregators. Importantly, the price impact of the aggregator's bids is taken into account, providing greatly improved results for large aggregators, when compared to existing price-taker algorithms. With respect to existing price-maker algorithms, our algorithm is formulated in terms of a scalable and computationally tractable optimisation algorithm, which allows its application to very large problem sizes. Furthermore, the formulation of our proposed bidding algorithm in terms of simple linear constraints allows the application of the coordination mechanisms proposed next.

The second part of this thesis proposes the first inter-aggregator coordination mechanism, which incentivises competing aggregators to produce joint coordinated bidding in the day-ahead market. This cooperation mitigates the increased electricity costs resulting from uncoordinated independent bidding. Moreover, techniques from mechanism design are employed to design the energy cost distribution mechanism in a way that discourages strategic manipulation, in order to ensure inter-aggregator cooperation. We conclude this part with an empirical case study based on real market and driver data from the Iberian Peninsula, showing substantially reduced energy costs when our proposed coordination algorithm is employed.

The third part of this thesis studies the aforementioned multi-aggregator scenario from a different perspective. Specifically, while in the previous mechanism design approach we assume a central coordinator and a number of aggregators that can use its services, we now consider a scenario where the aggregators can form willing coalitions to perform joint bidding. Given that the aggregators are self-interested, this scenario naturally lies within the field of cooperative game theory. As a consequence, we model this scenario as a coalitional game and prove that it is superadditive and balanced, hence the grand coalition where all aggregators cooperate with each other has the incentive to form. Moreover, we propose a payment mechanism, namely the least-core, which lies in the core of the game and hence stabilises the grand coalition. Finally, we use the empirical evaluation presented in the previous chapter to evaluate the coalitional game and the proposed payment mechanism.

Next, in the fourth contribution, we address the trust and privacy problems arising from the centralised coordination approaches proposed above. In more detail, an aggregator participating in the aforementioned coordination approaches needs to disclose key business information to the centre, which then acts as a black box and decides energy allocations and payments. By reformulating these centralised mechanisms using the Alternating Method of Multipliers algorithm, we obtain a decentralised model which greatly limits the share of information. Moreover, the decentralised algorithm is well suited for implementation on a blockchain, allowing execution in a trustless environment and providing transparency and anti-tampering guarantees. Finally, we assess the performance of the new model in a realistic case study akin to the one used in the previous contributions.

However, despite the interesting properties of the decentralised coordination mechanism proposed above, it also introduces new challenges, which are addressed in the fifth contribution. In more detail, the fact that the self-interested aggregators are not responsible only for sharing their requirements, but also for computation, introduces new strategic manipulation opportunities. More specifically, an aggregator can choose to deviate from the decentralised algorithm and manipulate its computation in order to increase its utility (*i.e.* reduce its energy costs) in detriment of the other participating aggregators. To study this problem, we define several attack vectors that deviate from the vanilla algorithm, and study their effects on the attacker's utility. We find that the attacker is able to artificially reduce its energy costs and increase its competitors'. In order to address this issue, we propose a manipulation detection framework which monitors successive rounds of the decentralised algorithm. The performance of this algorithm is also tested empirically and we found that it significantly outperforms a naive benchmark, presenting high levels of accuracy in a variety of scenarios. Finally, we would like to note that the issues of computational manipulation described in this contribution are not limited to our considered scenario and are present in any decentralised optimisation mechanism. Furthermore, the proposed detection mechanism can also be used in general settings.

Next, turning our attention to EV charging scheduling, we study the problem of allocating energy units. In more detail, after a given aggregator has purchased the energy necessary to satisfy its clients' needs, it needs to decide how to allocate it to the available EVs, without knowledge about future arrivals. In order to address this issue, we study existing and novel online scheduling algorithms, both theoretically and empirically. Moreover, given that it is unrealistic to require EV drivers to report valuations, we consider fairness objectives such as envy-freeness and maximising the number of satisfied drivers instead. The results from the realistic empirical evaluation show

that the proposed algorithms are able to achieve good performance with respect to the considered fairness objectives. Moreover, we note that the proposed model and algorithms can be used in other settings with perishable goods, such as processor scheduling and cloud computing.

Finally, in order to support the operation of EV aggregators and, more specifically, the day-ahead bidding algorithms proposed above, we turn our attention to forecasting. In more detail, given the nature of the electricity markets, bidding happens days or hours in advance, and needs to be based on forecasts. As a consequence, accurate forecasts improve the performance of bidding and coordination algorithms. More specifically, in this contribution we focus on the forecasting of price impact curves, which relate bid volume with resulting price, and are essential for large market participants such as EV aggregators. To this end, we draw from the field of machine learning and present novel artificial neural network models which significantly outperform previous approaches in the literature. ¹

¹This work was supported by an EPSRC Doctoral Training Centre grant (EP/L015382/1) through the Next Generation Computational Modelling CDT in the University of Southampton.

Contents

List of Tables	xi
List of Figures	xiii
Declaration of Authorship	xvii
Acknowledgements	xix
1 Introduction	1
1.1 Research Research Challenges	3
1.2 Research Contributions	5
1.3 Report Outline	7
2 Background	9
2.1 Smart Grid and Electric Vehicles	9
2.2 Electricity Markets	10
2.3 The Day-Ahead Market	11
2.3.1 Order Aggregation	12
2.3.2 Market Clearing and Pricing Mechanism	12
2.3.3 Price Impact	14
2.4 EV Aggregation	15
2.5 Scenarios with Multiple EV Aggregators	18
2.6 Mechanism Design Applied to EV Management	19
2.7 Cooperative Game Theory in Smart Grid Scenarios	19
2.8 Decentralised Management of EVs	20
2.9 Smart Grid and Blockchain	21
2.10 Strategic Manipulation of Decentralised Optimisation Algorithms	22
2.11 Online Energy Allocation	23
2.12 Price Impact Forecasting in Day-Ahead Markets	24
2.13 Summary	25
3 A Novel Price-Maker Day-Ahead Bidding Algorithm	27
3.1 The EV Aggregator Model	27
3.1.1 Modelling Energy Requirements	28
3.1.2 Forecasting of Energy Requirements	28
3.2 A Novel Price-Maker Day-Ahead Bidding Algorithm	29
3.3 Case Study	32
3.3.1 EV Aggregator Characteristics	32
3.3.2 Day-Ahead Market Data	33
3.3.3 Driver Behavioural Data	34
3.3.4 Results: Performance of the Bidding Algorithm	34
3.3.4.1 Comparing the Raw and Convex Formulations	34
3.3.4.2 Comparison with Benchmark Algorithms	35

3.3.4.3	Comparison Between the Exact And Approximate Formulations	38
3.4	Summary	39
4	Inter-Aggregator Coordination: a Mechanism Design Approach	41
4.1	Stage 1: Coordinated Bidding	42
4.2	Stage 2: Hourly Energy Distribution	43
4.3	Stage 3: Payment Distribution	44
4.3.1	Pure VCG	44
4.3.2	VCG-Based Truthful Redistribution	45
4.3.3	VCG-Based Proportional Redistribution	46
4.4	The Coordination Mechanism in Practice	47
4.5	Case Study	47
4.6	Summary	49
5	Inter-Aggregator Coalition Formation	51
5.1	A Brief Introduction to Cooperative Game Theory	52
5.1.1	Characteristic Function Games	52
5.1.2	Payoff Allocations	52
5.1.3	The Shapley Value	53
5.1.4	The Nucleolus	54
5.2	Defining the Multi-EV Aggregator Coalitional Game	54
5.3	Value Function with Externalities	55
5.3.1	The outsider coalition (oc) conjecture	56
5.3.2	The γ -conjecture	56
5.4	Properties of the Coalitional Game	57
5.5	Payoff Allocations	58
5.6	Case Study	59
5.6.1	Approximating the Shapley Value	59
5.6.2	Experimental Results	59
5.7	Summary	62
6	Decentralised Inter-Aggregator Coordination	63
6.1	Decentralised Optimisation Algorithm	63
6.2	Proposed Blockchain Implementation	65
6.3	Empirical Evaluation	66
6.4	Summary	67
7	Strategic Manipulation of Decentralised Optimisation Algorithms	69
7.1	Strategic Manipulation of the ADMM Algorithm	70
7.1.1	Shift Attack	70
7.1.2	Proportional Attack	72
7.1.3	Freeze Attack	72
7.1.4	Adversarial Attack	73
7.2	Detecting Manipulation	73
7.2.1	Quantifying Manipulation	74
7.2.2	Detecting Manipulation	75
7.3	Empirical Evaluation	77
7.3.1	Attack Vectors: Utility and Convergence	77
7.3.2	Threshold-Based Detection Results	78
7.3.3	Performance of <i>Adversarial</i>	82
7.4	Summary	82
8	Online Allocation of Energy Units	85
8.1	Model	85

8.2	Objectives	86
8.3	Offline Algorithms	88
8.4	Online Algorithms	89
8.5	Empirical Evaluation	92
	8.5.1 Supply Settings	92
	8.5.2 EV Settings	93
	8.5.3 Results	93
8.6	Summary	95
9	Forecasting Residual Supply Curves	97
9.1	The Data	97
9.2	Neural Network Models	98
	9.2.1 Targets	98
	9.2.2 Features	98
9.3	Methodology	99
	9.3.1 Data Stationarity	99
	9.3.2 Inter-hour Correlations	99
	9.3.3 Data Normalisation	100
	9.3.4 Sliding Window	101
	9.3.5 Validation and Test Sets	101
	9.3.6 Optimising the Number of Epochs	101
	9.3.7 Quantifying Prediction Errors	102
9.4	Results	102
	9.4.1 Hyperparameters and Architecture Tuning	103
	9.4.2 Model Results	103
	9.4.3 Single-hour versus Inter-hour Models	103
	9.4.4 Comparison with Results in the Literature	104
9.5	Summary	105
10	Conclusions and Future Work	107
10.1	Conclusions	107
10.2	Future Work	110
	Bibliography	113

List of Tables

3.1	Example of requirement vector $\mathbf{r}^{\min,i}$	28
3.2	Example of requirement vector $\mathbf{r}^{\max,i}$	28
3.3	Possible arrival times rounded to the nearest hour, with their respective probabilities.	34
3.4	Possible departure times rounded to the nearest hour, with their respective probabilities.	34
7.1	Summary of the proposed attack vectors.	70
8.1	Limits, averages and standard deviations of the Gaussian distributions of each of the scenario’s stochastic variables.	92
9.1	Optimal hyperparameters found for each model. The number of features for <i>NN Simple</i> is equal to the number of points considered for the price impact functions (50) multiplied by the number of hours (24). For <i>NN SingleHour</i> , only one hour is considered as input, so the number of features is just 50. The models <i>NN Wind</i> , <i>NN Demand</i> and <i>NN Both</i> augment <i>NN Simple</i> by considering 24 extra data points (one for each hour) for wind and/or demand forecasts.	103
9.2	MAE errors (MAE stds) for each model, evaluated on the test set (every day of 2017). Units: EUR/MWh.	104
9.3	Percentage improvements (using MAE) with respect to the <i>DayBefore</i> baseline. SCOMB is the best-performing model from existing literature [Aneiros et al., 2011, 2013].	105

List of Figures

2.1	Aggregated supply curve. Source: OMIE, 01/11/2016, 11 th hour.	13
2.2	Aggregated demand curve. Source: OMIE, 01/11/2016, 11 th hour.	13
2.3	Market clearing and uniform pricing mechanism. Source: OMIE, 01/11/2016, 11 th hour.	13
2.4	Price impact of a buy order with volume E and maximum price P_{\max} . Source: OMIE, 01/11/2016, 11 th hour.	15
2.5	Price impact function, $\mathcal{P}(E)$, for a volume range up to 10 MWh. Source: OMIE, 01/11/2016, 11 th hour.	15
3.1	Real and approximated price impact curves, showing a good (LHS) and a poor (RHS) approximation. Source: OMIE, 01/11/2016, LHS: 3 rd hour, RHS: 12 th hour.	31
3.2	(LHS) Raw and matched market orders, together with the hourly clearing price. (RHS) Appending the rejected supply orders (dashed) to the accepted matched supply curve (solid). Source: OMIE, 02/11/2016, 11 th hour.	33
3.3	Average daily payments per EV, for both bidding algorithm formulations and for different EV aggregator sizes, employing perfect forecasts.	35
3.4	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 OMIE, 09-10/11/2016.	36
3.5	Average daily payments per EV, for each different bidding strategy and for different EV aggregator sizes. Both with naive and perfect forecasts.	37
3.6	Percentage payment reduction when using the novel algorithm with price impact, compared to its counterpart without price impact.	37
3.7	(LHS) Runtimes for the approximate algorithm. (Center) Runtimes for the exact algorithm. (RHS) Normalised error between the solutions of the exact and approximate algorithm. Results are averaged across the first 15 days of November 2015 and error bars represent standard error.	38
4.1	Graphical depiction of the proposed coordination mechanism.	42
4.2	Average daily payments per EV when using a perfect forecast.	48
4.3	Average daily monetary surplus for the coordinator when using a perfect forecast.	48
4.4	Percentage payment reduction when using the proportional redistribution mechanism w.r.t. uncoordinated bidding, using both naive and perfect forecasts.	48
5.1	Comparison between uncoordinated and fully coordinated bidding when varying the number of aggregators.	60
5.2	Runtimes for the approx. SV and least-core payment mechanisms. Error bars represent standard deviations.	60
5.3	(LHS) Daily least-core payment allocations for four aggregators with different sizes. (RHS) Daily difference between the least-core and approximate Shapley value payments	60

5.4	(LHS) Daily least-core payment allocations for three aggregators of the same size, but with different flexibilities. (RHS) Daily difference between the least-core and approximate Shapley value payments.	60
6.1	Convergence of the ADMM decentralised algorithm to the optimal centralised solution, for different values of ρ . (Top) Simulations with two aggregators, each with 150 000 EVs. (Bottom) Simulations with ten aggregators, each with 150 000 EVs. Market data from 1/11/2016.	67
7.1	(a) Truthful allocation from aggregator i to aggregator j following Eq. 6.2b, $\mathbf{E}_{[k+1]}^{(i),j}$. (b) Attacked $\hat{\mathbf{E}}_{[k+1]}^{(i),j}$ employing a shift attack with $\mu = 2$ as given by Eq. 7.2. . .	71
7.2	Real example of the effects of <i>ShiftAll</i> on the resulting energy allocations. The scenario consists of three aggregators of the same size with the third aggregator being the attacker and employing $\mu = 1$. Results for both vanilla and attacked scenarios are presented.	71
7.3	Real example of the effects of <i>ProportionalAll</i> on the resulting energy allocations. The scenario consists of three aggregators of the same size with the third aggregator being the attacker and employing $\lambda = 0.66$. Results for both vanilla and attacked scenarios are presented.	72
7.4	<i>Difference matrix</i> (top) and <i>normalised difference matrix</i> (bottom), d and \bar{d} respectively, for a scenario with two aggregators of size 50 000 EVs (1 and 4) and two aggregators of size 150 000 EVs (2 and 3). One of the small aggregators (Aggregator 4) performs a <i>Shift</i> attack with $\mu = 1$ against the other small aggregator (Aggregator 1), displayed as light grey.	75
7.5	Cost and convergence analysis for each of the attack vectors described in Section 7.1. Scenarios with three aggregators (LHS) and four aggregators (RHS) and averaged over the ten first days of November 2016. All aggregators have capacity for 150 000 EVs.	78
7.6	Accuracy analyses for each of the proposed attack vectors, attack strengths and number of aggregators. Results averaged over every day of November 2016. All aggregators have capacity for 150 000 EVs. Dashed lines represent the naive benchmark for each scenario, which considers every aggregator to be benign. . .	79
7.7	Accuracy analyses for each of the proposed attack vectors, attack strengths and aggregator size. Results averaged over every day of November 2016. Dashed lines represent the naive benchmark for each scenario, which considers every aggregator to be benign.	80
7.8	Analysis of the efficacy of <i>Adversarial</i> for two attack strengths in the <i>Smaller</i> scenario, as described in Section 7.3.2. Dashed lines correspond to the percentage of times the attacked aggregator is incorrectly classified as deviator, in the case with no manipulation. Solid lines represent the same quantity when the aggregator is the target of <i>Adversarial</i> . Results averaged over every day of November 2016. .	81
8.1	Example allocation. For each EV and time step, each empty square represents one potential unit of energy, while a grey square represents an allocated unit. . .	86
8.2	Example showing that MaxDelivered does not imply MaxSatisfied.	87
8.3	Example showing that MaxSatisfied does not imply MaxDelivered.	87
8.4	Number of envious agents (LHS) and number of satisfied agents (RHS) when varying the number of agents present in the simulation. The stochastic available supply units follow the Gaussian distribution specified in Table 8.1. The dashed line corresponds to the optimal offline MaxSatisfied allocation. Error bars indicate standard deviations.	93

8.5	Number of envious agents (LHS) and number of satisfied agents (RHS) when varying the amount of supply available at each time step. Total number of agents is 45. The dashed line corresponds to the optimal offline MaxSatisfied allocation, averaged over 20 instances. Legend is shared with Figure 8.4. Error bars indicate standard deviations.	94
8.6	Runtimes for the proposed algorithms averaged over 50 instances.	94
9.1	(Top) Time series $\{\mathcal{P}_{h=0}^d(E_1)\}$ for 2008 to 2018. (Bottom) Differenced time series $\{\chi_{h=0}^d(E_1)\}$ for 2008 to 2018	100
9.2	Pairwise correlations between the time series $\{\mathcal{P}_h^d(E_1)\}$ for all d in 2016.	100
9.3	Illustration of the sliding window procedure to assign a training window to two given test points, d and d'	101
9.4	Results of the optimised architectures for each of the considered models.	104

Declaration of Authorship

I, Alvaro Perez-Diaz, declare that this thesis entitled Coordination Mechanisms for Electric Vehicle Aggregators and the work presented in it are my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as specified in Section 1.2.

Signed: Alvaro Perez-Diaz

Date: 10 August 2019

Acknowledgements

I would like to thank my partner, family, friends, supervisors Enrico Gerding and Frank McGroarty, colleagues, co-authors and progression report examiners for all their support during the last three years.

Chapter 1

Introduction

Climate change and environment conservation rise as two of the main challenges to address in the twenty-first century. Fossil fuels employed to satisfy a very large percentage of the global energy demands account for a great proportion of the pollution and emissions related to those. Nowadays, around 29% of the total energy consumption in the US corresponds to the transportation sector, and fossil fuels power around 95% of this consumption [EIA, 2016]. Similarly, in the UK transportation is the biggest energy consumer accounting for 40% of the total energy consumption [ONS, 2016], and 96% of this consumption depends on fossil fuels [POST, 2015].

In recent years there has been a shift towards renewable energy production to harvest non-polluting energy sources. This is seen as a fundamental aspect of the transition from fossil fuels to cleaner and more sustainable energy generation methods. For example, in the UK, renewable generation overtook coal in 2015 to consolidate as the second largest electricity generation method, accounting for 25% of the total UK electricity generation [Electricity Statistics, 2017]. A growing dependence on renewable sources will help lower the magnitude of greenhouse gas emissions and ensure the 2050 European targets, 80% reduction of emissions by 2050 [European Commission, 2011], are met.

Transportation is then a key participant in the modernisation and improvement of the energy generation-demand duo. The electrification of the transportation sector, which will stop the dependence on fossil fuels and allow working with renewable electricity sources, is seen as one of the main targets to achieve in the near future. In the UK, there exists a 10% target for electrification in the transportation sector by 2020 [ECCC, 2016]. Currently the UK has a fleet of nearly 100,000 electric vehicles (EVs), combining purely electrical and hybrid vehicles [SMMT, 2017]. In a global scale, there are targets to achieve 100 to 140 million of EVs by 2020 [International Energy Agency, 2016]. The transition from conventional fossil-fuelled to electric vehicles presents great opportunities but also large challenges.

Specifically, the main challenge is that the electrification of the transportation sector will present a heavy new load to existing electricity grids and power systems [Lopes et al., 2010]. More specifically, a single EV battery has a capacity of around 24 kWh [Gerding et al., 2016], depending on make and model. Comparing this capacity with the average daily electricity consumption per household in the UK, 10.8 kWh [World Energy Council, 2017], one can see that the widespread

use of EVs will cause a great increase in electricity demand. As a consequence, care needs to be taken to ensure a smooth transition and to accommodate the new EV fleet without the need of costly improvements in the electricity systems and extra generation infrastructure. At the same time, the new possibilities introduced by the penetration of EVs in the existing transportation ecology, *e.g.* the introduction of distributed storage systems (EV batteries), the ability of providing regulatory services to the electricity grid, *etc.*, could improve the behaviour of the electricity grid if used appropriately [Ramchurn et al., 2012].

To address this challenge, the academic literature has seen a high number of studies concerning a broad range of aspects related to the introduction of EVs in the power systems framework. Among those, of special relevance to this thesis is the literature on artificial intelligence applied to the management of a fleet of EVs [Rigas et al., 2015]. Typically, this involves an EV aggregator, an entity which can control the charging and discharging of a fleet of EVs, and thus make informed collective decisions [Bessa and Matos, 2010]. Moreover, different optimisation approaches have been proposed to address the aggregation of EVs and their participation on electricity markets, including energy scheduling and regulatory services [You et al., 2016]. However, a downside of most of the existing work is that it fails to account for the impact that a large fleet of EVs will have in electricity markets, as most works consider small EV penetration rate. In more detail, increased demand translates into increased electricity prices, the so-called price impact. Furthermore, the vast majority of the works in the literature consider the case of a single aggregator.

These two assumptions are unrealistic when considering the targeted levels of EV penetration in the near future (2030-2050) [International Energy Agency, 2016]. To this end, specialised algorithms for the participation of large EV fleets in the electricity markets, which incorporate the effect of the large load that future EV fleets will present, need to be devised. Moreover, the presence of multiple EV aggregators operating in the same markets and infrastructure presents exciting coordination opportunities. In more detail, given the high flexibility characteristic of EV charging, inter-aggregator cooperation can achieve reduced energy costs and limit physical impact on the grid. However, this coordination is not without challenges given the self-interested nature of the aggregators. Indeed, reduced electricity costs and better operation translates into increased profit for an aggregator and/or cheaper tariffs for their customers. As a consequence, multi-aggregator scenarios lie within the multi-agent systems field and coordination mechanisms need to be addressed carefully.

In addition, apart from participating on electricity markets, EV aggregators need to make real-time decisions such as how to best use the available energy. In more detail, the aggregator needs to devise appropriate charging schedules for each EV in a dynamic environment, where drivers arrive and depart over time and have different charging speeds and energy requirements. Moreover, the fact that energy is perishable, and needs to be used or wasted otherwise, adds complexity to the problem.

Furthermore, given that electricity markets are *future* markets, *i.e.* participation occurs days or hours before actual energy delivery, accurate forecasting is necessary in order to reduce imbalances and achieve optimised operation. This is true, not only for EV-related bidders, but for any market participant. To this end, modern machine learning techniques, which have achieved

impressive results in all research fields, are suitable candidates to provide precise forecasts, which translate into smoother electricity operation.

These issues, revolving around artificial intelligence techniques applied to EV management, are the main object of study of this thesis, and will be discussed in detail next

1.1 Research Research Challenges

As described in the previous section, EV penetration is expected to grow significantly in the near future, motivated by the offering of an environmentally friendly alternative to conventional fossil-fuelled transportation. However, this transition presents challenges to ensure the correct integration of the EVs into the existing power system and markets [Rigas et al., 2015]. Despite the large amount of literature published in the last decade, plenty of challenges remain open.

This thesis aims at addressing a subset of these, with a particular focus on large fleets of EVs and the presence of multiple self-interested EV aggregators. More specifically, we study the participation of multiple competing EV aggregators in day-ahead electricity markets. This includes the design of optimal bidding strategies and coordination mechanisms that enable joint-bidding, with the aim of reducing energy costs. Moreover, we consider both centralised and decentralised approaches, each with different strengths and weaknesses. In addition, we also study the forecasting of price impact curves in day-ahead markets, an important input for the aforementioned bidding algorithms. Finally, we turn our attention to real-time energy scheduling and to designing online algorithms for the fair allocation of energy resources among an EV fleet.

In more detail, the research challenges tackled in this thesis are as follows:

1. **Optimising the Charging of Large Fleets of EVs:** the management of a fleet of EVs has been studied in a variety of contexts [Rigas et al., 2015]. These include energy scheduling, operations management, integration with the smart grid and distribution networks, *etc.* (see Section 2.4 for a literature review). However, the vast majority of the literature considers low EV penetration, with fleets of hundreds of several thousands of EVs sharing the electricity markets and grids. Hence, the proposed algorithms and scenarios do not capture the scale of the targeted number of EVs in the near future [International Energy Agency, 2016]. This is important as the electricity demand characteristic of large EV penetration will have a significant impact on electricity prices. To this end, novel computationally tractable algorithms which take into account the large combined magnitude of a large fleet of EVs need to be developed, in order to ensure efficient power systems operation and reduce energy costs. More specifically, this can be achieved by utilising EV aggregators which are able to exploit the information about the fleet as a whole.
2. **Inter-Aggregator Coordination:** similarly to how an EV aggregator can make more informed decisions than individual EVs, it is reasonable to expect that inter-aggregator coordination can provide similar advantages. By sharing energy requirement information, optimised bidding can be performed, not only for the fleet of an individual aggregator, but for the whole group of aggregators. This is analogous to coordination approaches proposed in other smart grid scenarios, such as producer-consumer cooperatives [Akasiadis and

Chalkiadakis, 2016]. Considering now our multi-aggregator scenario, and with a focus on cost minimising, joint bidding will reduce unnecessary overlapping in the energy schedules of different aggregators, thus reducing price impact and energy costs. Despite the wide benefits arising from cooperation, the implementation of coordination mechanisms is not without challenges. More specifically, the main issue is the self-interested nature of EV aggregators. Indeed, lower electricity costs translate into more profit for the aggregator or lower charging tariffs for their clients. As a result, this coordination scenario naturally lies within the field of multi-agent systems [Wooldridge, 2009] and any coordination mechanism needs to be carefully designed in order to incentivise cooperation, rather than strategic manipulation.

3. **Decentralised Computation:** an essential requirement for any inter-aggregator coordination mechanism is the sharing of information which allows joint coordination. Given that EV aggregators are self-interested and, in general, private entities, they would be reluctant to sharing sensitive information. Moreover, the use of centralised approaches based on a *centre*, which receives information from the aggregators and acts as a black-box, requires trust and is not transparent. In order to address these issues, decentralised approaches in which each aggregator performs part of the coordination computation using its own private information, can be used instead [Munsing et al., 2017].
4. **Strategic Manipulation in Distributed Computation:** any distributed algorithm considering self-interested agents is susceptible to strategic manipulation, not only in the form of preference elicitation, as it is the case with centralised approaches, but also computationally [Shneidman and Parkes, 2003]. In more detail, an agent could deviate from the algorithm in order to increase its personal utility, in detriment of the other participants. More specifically, focusing on out multi-EV aggregator setting, an aggregator can deviate from the algorithm with the aim of reducing its energy costs or obtaining more advantageous energy allocations. As a consequence, in order to incentivise cooperation, the EV aggregators participating in the decentralised coordination mechanism need to be monitored in the look for manipulation, and appropriately penalised if deviation are found.
5. **Online Allocation of Energy Units:** after it has purchased a future electricity schedule in energy markets, an EV aggregator then needs to decide on how to allocate its available electricity units among its EV fleet in real-time. Interestingly, electricity is intrinsically different from other non-perishable resources or goods in that it needs to be used immediately or be wasted otherwise. Moreover, the highly dynamical and stochastic nature of the arrivals and departures of a fleet of EVs translates into the need of taking decisions online, *i.e.* without knowledge of the future. Furthermore, eliciting preferences and valuations from EV drivers is not practical or realistic, as in drivers do not generally have a technical background. In order to address these issues, online energy allocation algorithms need to be devised. Furthermore, in order to avoid the need for preference elicitation, efficiency and fairness objectives can be used instead.
6. **Price Impact Forecasting:** any effective bidding algorithm and coordination mechanism (see Research Challenges 1, 2 and 3) needs adequate forecasts such as EV fleet requirements, electricity prices, *etc.* With respect to the latter, the electricity market forecasting literature is mostly focused on prices and demand. However, forecasting price impact is

equally essential for large market participants and has received much less attention in previous works. From a technical perspective, it is different from demand and price forecasting, as price impact is a function (of bid volume) rather than a scalar quantity.

The work presented in this thesis will address the problems and challenges explained above. Against this background, we detail the precise contributions arising from our work in the next section.

1.2 Research Contributions

This thesis makes the following contributions to the state-of-the-art:

- We propose a novel price-maker day-ahead bidding algorithm for EV aggregators. This algorithm accounts for the price impact of the aggregator's orders and is formulated in terms of a non-linear optimisation algorithm with linear constraints, guaranteeing scalability to very large problem sizes and the applicability of the coordination mechanisms proposed below. The performance of the proposed bidding algorithm is tested in an empirical case study which employs real market and driver data. This addresses Research Challenge 1, as detailed in Section 1.1.
- We present the first inter-aggregator coordination mechanism for day-ahead markets. Specifically, we design a centralised approach and draw from the field of mechanism design to incentivise the participating aggregators to report their energy requirements to a coordinator, which then performs joint bidding. The performance of the proposed coordination mechanism is tested in an empirical case study which employs real market and driver data. This addresses Research Challenge 2, as detailed in Section 1.1.

These two previous contributions are described in the following paper [Perez-Diaz et al., 2018a]:

A. Perez-Diaz, E. Gerding and F. McGroarty. Coordination and payment mechanisms for electric vehicle aggregators. *Applied Energy*, 212, 185-195, 2018a.

- We present the first coalitional analysis of a multi-aggregator day-ahead scenario. In more detail, we study the theoretical properties of the coalitional game, and prove the stability of the grand coalition. A payment mechanism lying in the core is proposed in order to guarantee cooperation. Finally, we present an empirical study, also using real market and driver data. This addresses Research Challenge 2, as detailed in Section 1.1.

This contribution is described in the following paper [Perez-Diaz et al., 2018b]:

A. Perez-Diaz, E. Gerding and F. McGroarty. Coordination of Electric Vehicle Aggregators: A Coalitional Approach. *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)*.

The previous three contributions are described in the following paper [Perez-Diaz, 2018]:

A. Perez-Diaz. Coordination of Electric Vehicle Aggregator Participation in the Day-Ahead Market. *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018, Doctoral Consortium)*.

- We present the first decentralised coordination mechanism for EV aggregators participating in the day-ahead market. Specifically, we reformulate the centralised approach presented in the second and third contributions above using the Alternating Direction Method of Multipliers (ADMM) algorithm [Boyd et al., 2010]. In more detail, this well-known distributed optimisation algorithm presents very good convergence properties and has been widely used to tackle similar problems. Furthermore, in order to provide trustless execution, we discuss the potential of implementing the proposed coordination algorithm in a blockchain, providing transparency and anti-tampering guarantees. Finally, we test the proposed algorithm in a realistic scenario with real market and driver data. This addresses Research Challenge 4, as detailed in Section 1.1.

This contribution is described in the following paper [Perez-Diaz et al., 2018c]:

A. Perez-Diaz, E. Gerding and F. McGroarty. Decentralised Coordination of Electric Vehicle Aggregators. *International Workshop on Optimization in Multiagent Systems (OptMAS-18)*.

- We present the first study of strategic manipulation of decentralised ADMM algorithms. In more detail, focusing on the decentralised coordination mechanism proposed in the previous contribution, we describe and analyse different possible attack vectors and propose a mathematical framework to quantify and detect manipulation. Importantly, this detection framework is not limited to the considered EV scenario and can be applied to general ADMM algorithms. Finally, we empirically test the proposed attack vectors and manipulation detection algorithm in realistic scenarios. Our empirical results show that the decentralised algorithms convergence to the optimal solution can be effectively disrupted by manipulative attacks achieving convergence to a different non-optimal solution which benefits the attacker. With respect to the detection algorithm, results indicate that it achieves very high accuracies and significantly outperforms a naive benchmark. This addresses Research Challenge 4, as detailed in Section 1.1. This contribution is described in the following paper [Perez-Diaz et al., 2018d]:

A. Perez-Diaz, E. Gerding and F. McGroarty. Catching Cheats: Detecting Strategic Manipulation in Distributed Optimisation of Electric Vehicle Aggregators. *Under review in Journal of Artificial Intelligence Research*.

- We consider mechanisms for the online allocation of energy units. Unlike previous work, we consider mechanisms without money, and a range of objectives including fairness and efficiency. In doing so, we extend the concept of envy-freeness to online settings. Furthermore, we explore the trade-offs between different objectives and analyse their theoretical properties both in online and offline settings. We then introduce novel online scheduling algorithms and compare them in terms of both their theoretical properties and empirical performance. This addresses Research Challenge 5, as detailed in Section 1.1.

This contribution is described in the following paper [Gerding et al., 2019]:

E. Gerding, A. Perez-Diaz, H. Aziz, S. Gaspers, A. Marcu, N. Mattei and T. Walsh. Fair Online Allocation of Perishable Goods and its Application to Electric Vehicle Charging. *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI 2019)*¹

¹A. Perez-Diaz made the following contributions to this work: wrote the literature review (Section 2.11), designed the examples, proved Proposition 8.4, designed the experiments, ran the simulations, wrote the empirical evaluation (Section 8.5), produced all the figures and made significant overall editing and writing efforts.

- We study the prediction of residual supply curves in day-ahead electricity markets. To this end, we apply neural network models, more specifically multilayer perceptrons, and forecast the residual supply curves for each of the twenty-four hours of the next day. In more detail, we consider both intra- and inter-hour models, and also incorporate exogenous explanatory variables such as wind generation and total demand forecasts. We present empirical results using real data from the Spanish day-ahead market and show that our models outperform previous models in the literature, achieving up to 58.028% performance increase from the naive benchmark, compared to the previous 7.805% reported in the literature. Moreover, we find that inter-hour models achieve up to 6.028% performance increase when compared to intra-hour models. This addresses Research Challenge 6, as detailed in Section 1.1.

This contribution is described in the following paper [Perez-Diaz et al., 2019a]:

A. Perez-Diaz, A. Augustin, M. Nunes, E. Gerding and F. McGroarty. Forecasting Residual Supply Curves in Electricity Markets with Neural Networks. *Under review.*

1.3 Report Outline

The remainder of the thesis is structured as follows:

- Chapter 2 introduces the background needed to understand the challenges exposed above and compares our contributions with the state-of-the-art.
- Chapter 3 details our novel bidding algorithm, including its mathematical formalism and how to solve the optimisation problem. Its performance is analysed in an empirical case study with real data.
- Chapter 4 exposes our novel inter-aggregator coordination mechanism, from requirement aggregation, to collective bidding, energy distribution and payment handling. Its performance is analysed in an empirical case study with real data.
- Chapter 5 analyses the novel multi-aggregator coalitional game and proves several theoretical properties. This is also tested experimentally in an empirical case study with real data.
- Chapter 6 presents our novel decentralised coordination mechanism and studies its performance in a realistic empirical evaluation.
- Chapter 7 analyses the manipulation opportunities arising in decentralised optimisation problems, and presents our proposed detection algorithm. The effects of the proposed attack vectors and the detection algorithm are tested using real driver and market data.
- Chapter 8 studies the online allocation of energy units, or other perishable resources, and proposes novel scheduling algorithms. The performance of these algorithms, according to different fairness objectives, is tested in a realistic empirical evaluation.
- Chapter 9 presents our proposed artificial neural network models for the forecasting of residual supply curves in day-ahead markets.
- Chapter 10 serves as a conclusion for this thesis and outlines future work directions.

Chapter 2

Background

In this chapter, we present the background and literature review necessary to put the research objectives and challenges into context. The rest of the chapter is structured as follows. First, we discuss the smart grid paradigm in Section 2.1. Then, in Section 2.2, we proceed to detail the main characteristics of electricity markets, with a special focus on day-ahead markets and price impact. Next, we discuss the field of EV aggregation in Section 2.4, focusing on multi-aggregator scenarios in Section 2.5. Sections 2.6 and 2.7 describe the fields of mechanism design and cooperative game theory, respectively, with a focus on EV and smart grid applications. Next, we focus on decentralised algorithms for EV management in Section 2.8 and consider the application of blockchain technology in this type of scenario in Section 2.9. Moreover, we discuss the strategic manipulation of decentralised optimisation algorithms in Section 2.10. Next, we turn our attention to the issue energy scheduling, and to that end review the literature on the online scheduling of perishable resources, such as energy and computational power, in Section 2.11. Finally, we discuss the field of price impact forecasting, with a focus on day-ahead markets, in Section 2.12.

2.1 Smart Grid and Electric Vehicles

Energy and power systems are experiencing a radical change, transitioning from the traditional carbon-based structure into a smarter, cleaner, dynamic and bi-directional model, the so-called smart grid [Department of Energy and Climate Change, 2009]. This new paradigm envisions a future with decentralised production, transmission and distribution, widely employing renewable production methods and distributed storage systems. This will ensure a low-carbon and modern electricity grid with lower costs and improved efficiency [Ofgem and Department of Energy and Climate Change, 2014].

One of the key ingredients for this low-carbon and low emission future is the advent of electric vehicles. By electrifying transportation, dependency on fossil fuels is eliminated, and clean energy sources can be utilised to power the transportation sector, which accounts for a large proportion of the global consumed energy [Yong et al., 2015]. However, the introduction of EVs also presents a problem for the existing electricity grid, as this new heavy load will present a

heavy strain, which the current generation, transmission and distribution infrastructure may not be ready to accommodate [Pieltain Fernández et al., 2011, Lopes et al., 2010].

On the other hand, EVs can also be seen as distributed energy sources and storage systems, and could support the functioning of the electricity grid if managed appropriately [Ramchurn et al., 2012]. By storing energy when surplus is available and injecting it back to the grid when it is scarce, they could help regulate supply and demand. Given their high degree of flexibility [Pasaoglu et al., 2012, 2013], EVs can participate in demand response frameworks and provide auxiliary services that ensure stable grid operation [Shafie-khah et al., 2016].

In recent years, a lot of work has been done in the context of smart grid and EVs, as detailed in Section 2.4. However, further work is still necessary in order to ensure that the integration and operation of fleets of EVs is done in the best possible way. This challenge comprises several key aspects, such as the interaction with electricity markets, or with distribution and transmission grids. In this thesis, we will focus on former. To this end, we detail the structure and working principles of modern electricity markets next.

2.2 Electricity Markets

Electricity is produced in generation plants of different types and consumed in industries and households. With the ongoing electrification of the transportation sector, EVs are acquiring an important role as electricity consumers by obtaining the energy needed to charge their batteries. After the deregulation process which happened in most countries in the last few of decades, wholesale electricity is traded in energy markets as any other commodity: suppliers and consumers make offers and bids, the market is cleared, and electricity is delivered accordingly from sellers to buyers. While different countries and regions may be structured in different ways, the underlying structure and working principles are similar [Murray, 2009].

In the following, we present an overview of the common wholesale market structure, where several markets with different purposes and time horizons are available. Among these, energy scheduling markets are the ones where electricity is actually sold and bought. The purpose of these is to negotiate the future delivery of electricity, from generators to consumers, usually in an auction style. Different markets exist, ranging from long-term to very short-term time scales [Herranz et al., 2012, Ugedo et al., 2006]. Each market participant is free to take part in any combination of those, according to its needs and preferences.

On the long-term side one can find bilateral contracts and futures, in which energy delivery can be arranged and agreed upon months or years in advance. For example, energy retailers often employ this kind of contract to offer fixed-price energy to consumers.

Considering shorter time scales, *i.e.* days to hours in advance, two different markets are usually available: day-ahead and intra-day. Day-ahead markets present an auction style mechanism in which energy delivery deals are struck in an hourly fashion for the next day. Market participants forecast their needs or generation capabilities and make bids and offers to the market. When the market is cleared, energy delivery contracts for the next day are formalised. An important characteristic of electricity markets, different from other commodities, is the fact that energy is

difficult and expensive to store, hence supply and demand need to be constantly balanced to prevent losses. This causes the need for extra markets close to delivery time.

After the day-ahead market is cleared, intra-day markets become available. There are usually several instances of daily intra-day markets, each negotiating energy for 4 to 6-hour slots each day in an auction style. These markets allow market participants to adjust to deviations in demand or supply capabilities with respect to their orders in the earlier day-ahead market. The shortest-term includes balancing markets, which act as last-minute auctions in which any unbalanced supply and demand can be matched together to ensure best possible allocation. This sequence of energy scheduling markets is employed by wholesale buyers and sellers to agree on electricity delivery across different time-scales. EV aggregators, as large electricity consumers, participate in these scheduling markets as buyers to purchase electricity for their clients, as shown in Section 2.4.

Apart from scheduling markets, and due to the particular nature of electricity and the physical grid across which it is distributed, a second type of markets are available to ensure a correct and stable behaviour, namely the ancillary and regulatory markets. As mentioned above, electricity generation and consumption need to be balanced at all times. Indeed, generation and demand unbalances cause physical problems such as instability in alternate current frequency and voltage. Moreover, electricity needs to be distributed along power lines from generators and distributors to consumers. This imposes physical constraints to electricity delivery, and malfunction or congestion can affect the deals struck in advance in electricity markets. These and other issues create the need for regulatory markets, in which generators and consumers can participate to guarantee a stable electricity grid. The usual market structure consists of three reserve markets: primary, secondary and tertiary reserve. The exact working of these markets may again depend on the considered country, but they usually share the underlying structure and consider different time-scales, from one hour to several minutes ahead [Murray, 2009]. They are usually auction-based, where suppliers and consumers can submit orders to increase or decrease generation or consumption, respectively, based on the state of the grid and their own. All the orders are then aggregated and the market cleared, assigning a correction term to the generation or demand of each market participant, and some monetary compensation or penalty. Reserve markets can also be employed as last-minute electricity sources to trade in, for profit or to obtain energy if needed [Bessa and Matos, 2013c]. Ancillary markets have been employed in the EV literature to account for forecast deviations and to trade vehicle-to-grid (V2G) electricity, both for profit and for grid stability, as shown in Section 2.4.

This concludes the overview of the typical electricity market structure, providing a broad picture of the different wholesale markets available to suppliers and consumers, their particularities, and how they interact with each other. Of particular interest to this thesis is the day-ahead market, which will be discussed in detail next.

2.3 The Day-Ahead Market

Day-ahead markets are the main source of short-term wholesale energy [Stoft, 2002, Murray, 2009] and have been widely employed in the EV literature to obtain energy for battery charging.

This type of electricity market are available all across Europe (EPEXSPOT, GME, NORD-POOL, OMIE, OPCOM, OTE, TGE) and the US (CAISO, MISO, ISO-NE, NYISO, NWPP, WECC, PJM, SPP, ERCOT). Typical day-ahead markets are daily forward markets featuring a uniformly-priced double-sided auction. In more detail, they run everyday 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, usually noon on day D . Double-sided auction means that market participants can simultaneously submit orders to both the demand and the supply sides of the auction, and uniformly-priced means that each hourly auction uses uniform marginal pricing (see Section 2.3.2). We will now describe this type of market in detail.

2.3.1 Order Aggregation

As described above, participants in the day-ahead market submit their buy and sell orders to the auction before closure time. The usual day-ahead market order has three main characteristics: order side, volume and price. Order side determines whether it is a supply or demand order, referred to as an offer and bid respectively. The volume specifies the amount of electricity, usually in MWh, that the participant is willing to trade. The price specifies the lowest (highest) price the agent is willing to accept for that volume of offered (bided) electricity. Note that there is usually a maximum price for bids and offers, P_{\max} , which is far above the usual cleared prices.

Different markets have different order types, ranging from the basic order with three components we just described, to more sophisticated orders which can include a step-wise function of different volumes and prices, or conditions on acceptance based on other orders or particular market details. Henceforth, we will consider simple orders, consisting on volume-price pairs on both sides of the auction.

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 depicted in Figure 2.1. This generation stack curve will vary hourly and daily as a result of a myriad factors, such as the weather, the infrastructure available in the considered region, fossil fuel price fluctuations, new technologies (such as renewable energy) becoming available, and commercial interests of the market participants.

The demand side consists of large electricity consumers, such as industries, retailers or aggregators, which also submit their bids to the day-ahead auction, consisting as well of desired electricity volumes and prices. All the bids are aggregated by high-price priority forming the aggregated demand curve, as shown in Figure 2.2, based on real data from the Iberian electricity market OMIE.

Once all the orders have been aggregated, the generation stack and demand curves are employed to clear the market, as described in the next section.




FIGURE 2.1: Aggregated supply curve.
Source: OMIE, 01/11/2016, 11th hour.




FIGURE 2.2: Aggregated demand curve.
Source: OMIE, 01/11/2016, 11th hour.

2.3.2 Market Clearing and Pricing Mechanism

In order to find the clearing price and accepted orders of a given market hour, the aggregated generation and demand curves described in the previous section are employed. The most common approach is to find their intersection, as shown in Figure 2.3.

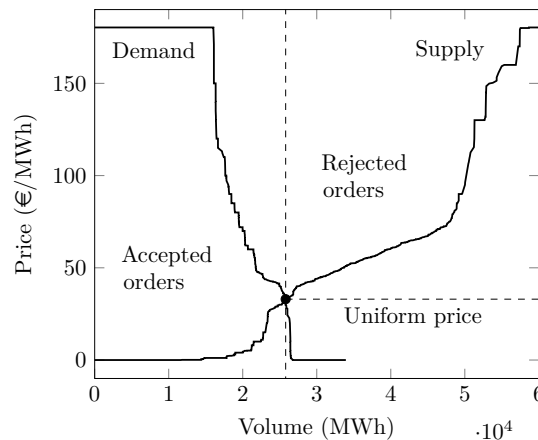


FIGURE 2.3: Market clearing and uniform pricing mechanism. Source: OMIE, 01/11/2016, 11th hour.

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. The clearing price will be the price at the intersection, which applies to all accepted bids and offers (uniform pricing). The exact clearing and pricing mechanisms vary across countries, as do the available types of orders. It is common to employ a two-stage algorithm, in which aggregated curves intersection and uniform pricing is applied in the first stage, determining the accepted bids and orders and a clearing price. In the second stage, the complex requirements of non-basic accepted order types are taken into account, moving the resulting price up or down according to the particular realisation. This is the case all across Europe, in which the EUPHEMIA algorithm is employed [EUPHEMIA, 2015]. This concludes a detailed exposition of the behaviour and structure of day-ahead markets, which conform the framework of our proposed bidding and coordination mechanisms, introduced in Chapters 3, 4 and 5. An important feature of financial and electricity markets, which describes how orders affect the market price, will be explained in the next section.

2.3.3 Price Impact

Consider now a market participant, an electricity buyer, who bids in the day-ahead market to obtain electricity to satisfy its needs. This market participant's bids will contribute towards the aggregated demand curve, and will affect the clearing price by pushing it up. Conversely, if the market participant is an energy seller, its ask orders will affect the supply curve and push the clearing price down. This is the so-called price impact, a term borrowed from the financial literature, where it has been extensively studied [Bouchaud, 2010]. The market impact of small participants can be negligible, but the larger the order size, the more pronounced effect it will have on the clearing price. Price impact is an important characteristic associated with large market participants, such as large EV aggregators, and needs to be taken into account when developing bidding schedules.

Specifically, price impact has also been applied to electricity markets, under the name of *price-making* or *market effect*, and several frameworks have been proposed for its mathematical treatment. Baillo et al. [2006] present a review of commonly used methods which are described as follows. The first technique is a two-level problem, in which firstly market participants submit their orders trying to maximise individual profit, and secondly the market operator determines optimal dispatch and market prices [Pereira et al., 2005]. The second method employs the so-called market distribution functions. This probabilistic approach considers the probability of quantity-price pairs being accepted in market closure and runs an optimisation problem on this two-dimensional space Anderson and Philpott [2002]. The last method utilises residual curves to represent the available levels of demand as a function of market clearing prices Baillo et al. [2004].

In more detail, residual curves were originally introduced to model the effects of electricity suppliers in electricity markets [Garcia-Gonzalez et al., 1999]. This model was extended to account for a portfolio of different generators by Berzal et al. [2001], by solving the optimisation problem for the global offer and distributing the scheduling among the generators. Conejo et al. [2002], De La Torre et al. [2002] considered a more general model, including the case of oligopolistic companies. Baillo et al. [2004] extended these works by including long-term contracts and a multi-unit scenario. A recent study extends the residual demand curves framework to considered more sophisticated day-ahead market structures, introducing zonal pricing and more complex order types [Portela González et al., 2017]. The similarities between generation and demand from the perspective of market participation are high, so the concepts introduced in the previously described works can be translated to an energy buying perspective. To this end, some works focus on the demand side of the market. Specifically, the existing literature focuses on an electricity retailer which participates in the day-ahead market to obtain energy to distribute to end-users, under different market scenarios [Ya'an Liu and Xiaohong Guan, 2003, Fleten and Pettersen, 2005, Philpott and Pettersen, 2006, Herranz et al., 2012]. Hence, residual curves seem the appropriate choice for modelling the price impact of an EV aggregator in our considered scenario. We will now proceed to formalise the concept of price impact in the context of uniform double-sided auctions, such as day-ahead markets, utilizing residual supply curves. We focus on electricity buyers, and the notation employed closely follows the literature [Baillo et al., 2004, Herranz et al., 2012].

Specifically, let $D(p)$ and $S(p)$ 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:

$$R(p) = S(p) - D(p) = E$$

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

More relevant to us is the clearing price when bidding an amount E , which can be obtained from the residual supply curve:

$$p = R^{-1}(E)$$

Introducing the notation $\mathcal{P}(E) = R^{-1}(E)$, the clearing price when the new agent bids an amount E is:

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

and the price impact Δp of this order is:

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

as $\mathcal{P}(0)$ represents the *base price*, that is the price without the agent's new bid. This formalism is depicted in Figs. 2.4 and 2.5.




FIGURE 2.4: Price impact of a buy order with volume E and maximum price P_{\max} . Source: OMIE, 01/11/2016, 11th hour.

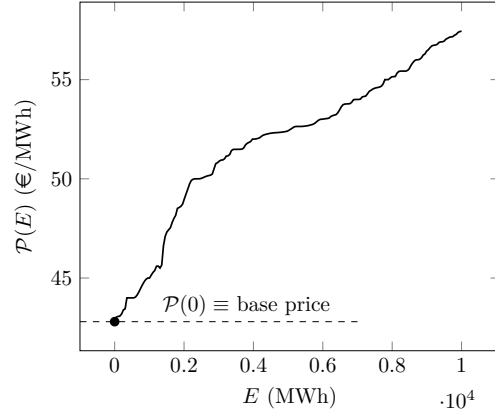


FIGURE 2.5: Price impact function, $\mathcal{P}(E)$, for a volume range up to 10 MWh. Source: OMIE, 01/11/2016, 11th hour.

Focusing on EVs, a single unit will have a negligible price impact as its battery size is too small compared to the size of the market. For an aggregator managing a large number of EVs, the combined capacity of their batteries gets large enough to have a non-negligible impact on the market, and efficient energy buying algorithms need to account for this. This will become clear when considering the novel price-maker bidding algorithm presented in Chapter 3.

Next, we will review the literature on EV aggregation, with a special focus on day-ahead market participation.

2.4 EV Aggregation

So far, we have explored the smart grid paradigm and the structure of current electricity markets, with a focus on market clearing and price impact in day-ahead markets. Intimately related to these is the transition to the widespread use of EVs. The quickly increasing EV penetration numbers will trigger the growth of EV aggregators, which will become an important part of the electricity markets participant ecosystem. As described in detail in previous sections, the heavy new load introduced by the electrification of the transportation sector carries considerable market impact, and may require large investments and new infrastructure if not carefully managed. However, this load is very flexible and offers a novel distributed storage system with very different behaviour compared to conventional electricity consuming demand. This motivates the need to study the interaction of a fleet of EVs with the existing power systems and to develop mechanisms to ensure this integration happens in the best possible way. This section reviews the existing literature working towards this goal.

From the perspective of an EV driver, their main concern is their vehicle's battery being charged on time to satisfy their needs. It is unreasonable to assume that the average EV user will have expertise on electricity markets, or that they will adjust their demand accordingly to hourly electricity prices or market needs. This motivates the idea of introducing an EV aggregator, first described in [Kempton et al., 2001] as an entity which acts as an intermediary between a number of EV users and electricity markets and infrastructure. By collecting information from its EV fleet, and by managing their batteries, the aggregator is able to make better informed decisions and has better market visibility than individual EVs. Since their introduction, EV aggregators have been extensively studied in a variety of settings. A thorough review of EV aggregator structure and its relation to existing electricity infrastructure and markets can be found in [Bessa and Matos, 2010]. In addition, detailed reviews of economical and technical aspects of EV aggregation can be found in [Bessa and Matos, 2012] and [You et al., 2016]. Hu et al. [2016] review the literature on service provision and EV fleet control strategies. A review of scheduling and optimisation methods can be found in [Yang et al., 2015]. A detailed overview of artificial intelligence techniques applied to the managing of fleets of EVs, including routing problems and congestion managing, is presented in [Rigas et al., 2015]. These works present a comprehensive review of the EV aggregation literature from different viewpoints. We will now focus on the most relevant works for this thesis.

As exposed in previous sections, one of the main concerns of growing EV penetration rates is how they integrate with the existing electric power systems and markets. Early studies investigate the main aspects necessary to ensure a smooth integration. Specifically, Galus et al. [2010] identify the actors in power systems planning and operation which will be impacted by the introduction of EVs, and present potential changes to mitigate this impact. They describe an EV aggregation framework which can be utilised for activities such as load management and service provision. A similar study presented by Lopes et al. [2010] focuses on both the grid technical operation and electricity market operation. The authors describe a list of overall challenges for large deployments of EVs, which include: (i) evaluating the impact that battery charging has in system operation; (ii) developing control strategies for battery charging periods; (iii) assessing the potential of EVs to provide power systems services. These challenges have thenceforth been studied in the literature in different scenarios.

Of special relevance to this thesis is the study of the impact of EV charging on electricity markets and the design of scheduling algorithms which balance the overall EV demand to mitigate demand peaks which can cause grid instability, increased costs, and require polluting and expensive generation methods. There are a number of works in the literature that study these issues, as detailed next. Bessa et al. [2012] consider the algorithmic participation of an EV aggregator in the day-ahead market, in order to obtain energy to charge its clients' batteries. Their linear programming algorithm accounts for the large degree of flexibility characteristic of the EV fleet and optimises the day-ahead bidding to lower monetary costs. Subsequent work considers different forecasting techniques, and compares the performance of modelling EV requirements individually and globally for the whole fleet [Bessa and Matos, 2013a,b]. Lastly, the same authors extend their model further by considering participation in the day-ahead and manual reserve markets, in which the EV aggregator can purchase energy in the day-ahead market but also bid in the reserve market, both for up and downward reserve [Bessa and Matos, 2013c]. Along slightly different lines, the participation of an EV aggregator in the day-ahead market is also considered in [Sousa et al., 2012]. In this study, V2G is included in the framework by considering a virtual power plant that negotiates V2G contracts with drivers and also bids in the day-ahead market. V2G is also considered in [Sarker et al., 2016], where battery degradation compensations are included in their model of participation in day-ahead and ancillary markets. A different aspect of EV aggregation, namely the demand response capabilities offered by the fleet, is considered in [Shafie-khah et al., 2016]. Their work includes a multi-market scenario with a day-ahead and several instances of intra-day markets. Moreover, the integration of EV aggregation with wind production is considered in Honarmand et al. [2014]. They propose an aggregation agent which considers the interaction between wind generators and a small number of EVs in a micro-grid setting. This work is extended in [Heydarian-Forushani et al., 2016], where more attention is paid to the uncertainty characteristic of EV fleets, both concerning arrival and departure times, and charging requirements.

Furthermore, a different body of literature focuses on physical grid aspects, rather than being focused on electricity markets. For example, Sousa et al. [2014] study the issue of EV scheduling with the aim of flattening EV demand and successfully integrating the fleet into a smart grid with large penetration of distributed generation units. Similarly, Alonso et al. [2014] study the participation in the day-ahead market focusing on the physical aspects of EV penetration, and considering physical constraints such as thermal line limits, load level on transformers and voltage limits to find an optimal EV charging schedule.

However, all the works presented above consider price-taker approaches, in which the effect of the EV fleet in the market is neglected. In more detail, considering the large EV penetration levels forecasted for the near future (see Section 1), these models will not present an accurate analysis of the impact of large EV fleets on our existing power systems and market structure.

In order to address this issue, a small body of literature has recently started addressing the EV aggregator algorithmic participation in day-ahead markets under a price-maker perspective. In more detail, Kristoffersen et al. [2011] present the first of such studies. Specifically, they consider an EV aggregator able to charge and discharge its EVs batteries bidding in the day-ahead market, where price impact is approximated to be linear. Similarly, Sundström and Binding [2012] also consider linear price impact, but focus on physical grid distribution constraints. A related market scenario in which the market clearing process is explicitly solved is presented in

[Liu et al., 2017]. Finally, Gonzalez Vaya and Andersson [2015] consider non-linear price impact, and propose a bidding algorithm which relies on a bi-level optimisation problem in which the lower level is explicit market clearing. All these works paved the way towards large EV aggregator participation in day-ahead markets, but have several limitations. Firstly, price impact is often modelled in a linear fashion, which results in a loss of accuracy. Secondly, some of the models explicitly solve the market clearing process, introducing extra computational complexity. Thirdly, all the works are formulated in terms of complex optimisation algorithms with cumbersome constraints. This limits the scalability to large problem sizes and the applicability of coordination mechanisms (see Chapters 4, 5 and 6).

2.5 Scenarios with Multiple EV Aggregators

In the previous section, we have highlighted the higher performance potential achieved by replacing individual driver decisions by collective actions managed by an EV aggregator. Indeed, an EV aggregator, who has aggregate information from the EV fleet and superior market visibility, can obtain reduced costs and ensure better functioning of the electricity grid. Similarly, a higher level of inter-aggregator coordination would present similar advantages when compared to individual aggregator decisions without considering other aggregators which share the same market or operational area. Specifically, we envision a scenario where the high electrification levels in the transportation sector cause the apparition of many different competing EV aggregators which want to profit from the services provided, or to lower charging costs for their clients. These private and independent EV aggregators participate in the same day-ahead market, and can optimise their own bids via smart bidding algorithms. However, the lack of coordination among different aggregators can cause inefficient bidding, where hourly bids unnecessarily overlap driving prices up. The problem of coordination resides in the fact that EV aggregators are self-interested, the precise setting explored in the field of multi-agent systems [Wooldridge, 2009].

Similar multi-aggregator problems are studied in the literature. In more detail, Qi et al. [2014], Shao et al. [2016] study the hierarchical control of EV fleets where different aggregators are coordinated by a high-level coordinator. However, they do not consider self-interested aggregators and focus accommodating grid constraints and ensuring driver satisfaction instead. More related to our considered scenario, Yu et al. [2016] study a setup where a number of EV aggregators can trade energy among them in order to fix forecasting deviations, instead of purchasing the energy from the grid, and show that this improves the aggregators' energy costs. However, the authors do not consider price impact in their model and each aggregator performs independently. Another related study can be found in [Mukherjee and Gupta, 2017]. Their work considers a scenario where several private aggregators are present in a given city, and negotiate with each other in order to balance charging in the different limitedly available charging stations. The aim is to maximise the total number of EVs charged and the profit of the EV aggregators and their results indicate that coordinated operation improves the profit of the EV aggregators and the services offered to the drivers. Finally, in a similar vein to our work, Wu et al. [2016] study a multi-aggregator day-ahead bidding scenario and apply game theory to find Nash equilibria. In more detail, each aggregator tries to categorise the other aggregators and thus forecast their day-ahead bids, and adjust their bidding accordingly. However, after introducing several approximations in order to simplify the model's structure, the proposed model depends on a complex

optimisation algorithm which does not guarantee existence of Nash equilibria. Summarising, even though all these works study multi-aggregator scenarios, none focuses on the interaction of self-interested aggregators and their coordination possibilities.

2.6 Mechanism Design Applied to EV Management

Mechanism design studies agent interaction protocols which take into account the fact that agents are rational and self-interested [Vlassis, 2007]. One of the key contributions of the field is the concept of truthful or incentive compatible mechanisms. Such mechanisms are designed in a way that strategical manipulation is discouraged, so that the participating agents cannot obtain a personal benefit by lying or cheating the system. For more detail, the exact definitions and mathematical framework can be found in, for example, the works by Nisan et al. [2007] and Vlassis [2007]. Specially relevant to this thesis is the application of mechanism design to the context of EV aggregation and management, of which there are several studies in the literature. An early work by Gerding et al. [2011] considers a scenario where a limited amount of electricity is available throughout the day, and EV drivers negotiate charging times based on their needs and preferences. Mechanism design is employed to *truthfully* elicit EV driver requirements and to distribute the available energy accordingly to these requirements. This work is extended in [Stein et al., 2012] by introducing the idea of pre-commitment, in which energy resources are reserved for EVs but the particular charging time and rate are flexible. Incentive compatibility is maintained in this study to ensure a fair mechanism. Furthermore, these two works are extended in [Hayakawa et al., 2015] by accounting for multi-dimensional preferences and varying marginal prices. In a slightly different vein, Gerding et al. [2013] consider a market where EV drivers and charging points can negotiate charging. Different charging locations and times are available, and their mechanism truthfully elicits buyer requirements. Lastly, Gerding et al. [2016] consider a car park with V2G capabilities which truthfully elicits drivers' preferences and employs the parked vehicles to provide V2G services to obtain a profit. All these works employ mechanism design in the context of EV management to ensure fair competition among self-interested EV drivers and energy providing companies. However, the possibilities of cooperation among different EV aggregators which share the same infrastructure, particularly the same day-ahead market, have not been addressed in the literature.

2.7 Cooperative Game Theory in Smart Grid Scenarios

Cooperative game theory studies competitive multi-agent scenarios, in which the agents are allowed to form coalitions in order to increase their performance by cooperation [Peleg and Sudhölter, 2007]. More specifically, this field is concerned with the challenges of coalition structure optimisation and the distribution of value among the members of a given coalition. These issues are far from trivial given that the agents are rational and self-interested, so incentives need to be adequately designed, in a related vein to the field of mechanism design described in the previous section. Techniques from cooperative game theory have been widely applied in many fields, including smart grid scenarios as described below.

The cooperation among small power producers is studied in [Chalkiadakis et al., 2011, Bremer and Sonnenschein, 2013]. More specifically, Chalkiadakis et al. [2011] consider small distributed energy producers which are able to form coalitions to make joint offers in electricity markets. This is shown to be beneficial to gain market visibility, and to reduce the uncertainty related to their bids, resulting in increased profit. Bremer and Sonnenschein [2013] consider a similar scenario, and focus on profit distribution among the coalition members. Similarly, cooperation among wind producers is studied in Baeyens et al. [2013]. Their results show that cooperation increases their profit and that a coalition including all participants provides everyone the best benefit. Moreover, cooperative game theory has been employed to aggregate demand response providers, in order to improve their performance and grid stability [O'Brien et al., 2015, Chapman et al., 2017]. Lastly, cooperation among independent households with distributed generation and storage capabilities is studied in [Alam et al., 2013]. Their results show that battery degradation can be greatly reduced, while obtaining significant better energy efficiency, by cooperation. However, none of these works study EV charging and the interaction among self-interested EV aggregators.

2.8 Decentralised Management of EVs

Typical centralised algorithms have a *centre*, which collects information from every agent and makes decisions as a black box. In contrast, decentralised algorithms present important advantages, such as reduced information sharing, more transparent operation and trust-less execution, achieved by outsourcing computation from the centre to the participating agents. Moreover, decentralised algorithms are well-suited for implementation on a blockchain (see Section 2.9). Importantly, apart from monetary transactions, blockchains can be used as general computing systems by using smart contracts. As a consequence, blockchains are a suitable architecture for implementing the consensus or aggregation step present in decentralised optimisation algorithms (see Section 6.1). Furthermore, decentralised mechanisms can also be fully implemented in the blockchain by using cryptocurrencies.

Focusing now on EV related scenarios, these techniques have been applied in different settings in the literature. In more detail, several works study the decentralised scheduling of EV charging [Ardakanian et al., 2014, Wen et al., 2012, Gan et al., 2013, Ma et al., 2013, Le Floch et al., 2015, 2016]. Overall, these works consider the problem of scheduling the charging of EVs in different decentralised scenarios, considering each EV as an individual node in their respectively proposed algorithms. Specifically, Ardakanian et al. [2014] focus on physical grid constraints, considering an electricity network managed by different access points, and its interaction with a fleet of EVs. Similarly, Gan et al. [2013], Ma et al. [2013] consider decentralised valley-filling algorithms, where the aim is to flatten demand over time and each EV sequentially updates its own charging schedule by iterative interaction with a central utility company. In a related vein, Wen et al. [2012] consider a decentralised algorithm which employs discrete time intervals and selects subsets of EVs to be charged at each time interval, by iterative communication between each EV and their aggregator. Finally, Le Floch et al. [2015, 2016] consider V2G scenarios, where the EVs are able to inject energy back to the grid when needed. However, although all these works study different aspects of EV charging scheduling under decentralised algorithms, they do not consider the interaction among different self-interested aggregators.

Furthermore, decentralised algorithms have been employed in many non-EV related smart grid publications. Among these, an algorithm that has acquired great popularity in recent years due to its versatility and good convergence properties is the Alternating Direction Method of Multipliers (ADMM) algorithm [Boyd et al., 2010]. Examples of its application in the smart grid field include power flow [Wang et al., 2017, Sulc et al., 2014, Peng and Low, 2014, Scott and Thiébaux, 2014] and micro-grid [Munsing et al., 2017] scenarios. However, as mentioned above, the interaction among different self-interested EV aggregators has not been considered in the literature. We will now turn our attention to blockchain technologies and present a literature review on smart grid related applications in the next section.

2.9 Smart Grid and Blockchain

Following the discussion in the previous section, blockchains and smart contracts present desirable advantages such as fully decentralised operation, transparency and anti-tampering guarantees. In more detail, a blockchain is a decentralised ledger and computation environment protected by cryptographic techniques, which allows the participants to agree on the state of the system at all times [Christidis and Devetsikiotis, 2016]. Moreover, a smart contract is simply a piece of code hosted publicly and immutably in the blockchain, which can receive messages from other blockchain users and send its own following its internal logic.

Due to these interesting properties, they have been widely applied in recent smart grid related works, both in academia and in industry, as detailed next. Specifically, Mylrea and Gourisetti [2017], Lombardi et al. [2018] study how the smart grid infrastructure can be improved by using blockchains and the internet-of-things. In more detail, they propose the use of smart contracts to transparently record energy transactions, which together with connected sensors (such as smart meters) would present a more robust, secure and efficient smart grid. Similarly, Horta et al. [2017] focus on distribution networks and study how blockchain technologies can be utilised in order to support distributed energy generation, EVs and demand response. In a slightly different vein, a number of works study micro-grid scenarios [Mattila et al., 2016, Munsing et al., 2017, Baroche et al., 2018, Zizzo et al., 2018]. More specifically, these works study the peer-to-peer (P2P) interchange of energy in local markets, where households can trade energy with their neighbours. In this scenarios, blockchain can introduce enhanced security and transparency, and removes the need for a centralised operator. Finally, we turn our attention to demand response [Akasiadis and Chalkiadakis, 2016, Pop et al., 2018]. In these works, consumers coordinate with each other in order to provide flexible energy consumption, which is shown to yield increased benefits than individual operation. Again, blockchains enable decentralised cooperation and also profit distribution by using specially designed cryptocurrencies. However, despite the wide application of blockchains and smart contracts to different smart grid settings, as described above, the problem of decentralised EV aggregator coordination has not been addressed in the literature.

Furthermore, we would like to note that these blockchain applications are not limited to academic research but have already been tested in real scenarios worldwide. We will now provide several illustrative examples. In the Netherlands, Spectral¹ manages two blockchain projects. Firstly, a

¹<https://spectral.energy/>

residential P2P micro-grid in Amsterdam, where neighbour can trade energy, including surplus from photovoltaic using a custom built token called Jouliette. Secondly, a similar project in the city of Groningen, which has the aim of being carbon-neutral by 2025. Currently, the project is scaling from small neighbourhoods to city-level P2P energy transactions. Similarly, PowerLedger ² provides similar services and participates in the largest commercial P2P trial in the world, based in Bangkok. Finally, in a slightly different vein, WePower ³ allows individuals to buy renewable electricity straight from the produces. Although this would not be possible by using conventional energy markets, given the small size of households compared to wholesale market participants, the use of blockchain and energy tokens allows coalitions of consumers to join together to purchase energy directly from generators in a transparent and secure way.

2.10 Strategic Manipulation of Decentralised Optimisation Algorithms

In the previous two sections we have discussed the use of decentralised algorithms to support the coordination of multiple EV aggregators. However, although this type of approach has significant advantages with respect to centralised mechanisms, such as removing the need of reporting private requirements or place trust in the coordinator, it introduces a new important problem, namely, the possibility of manipulating computation. More specifically, considering that each aggregator is self-interested, they can decide to deviate from the algorithm in order to increase their personal utility in detriment of the other participants. We would like to note that this is a general issue in multi-agent systems works using decentralised optimisation algorithms, and not restricted to EV charging.

As discussed in the previous two sections, the ADMM algorithm is extremely popular and has been applied in many fields in the last decade [Boyd et al., 2010]. There is, however, an large gap between the introduction of such algorithms and the study of their robustness to potential manipulative or malicious attacks [Munsing and Moura, 2018]. These important issues are studied in several recent works which are reviewed next. More specifically, following [Munsing and Moura, 2018], we can classify the existing literature based on the technique utilised:

- *Round-robin techniques* [Liao and Chakraborty, 2016, 2017]: these techniques seek to identify compromised nodes by replacing the coordination step of the ADMM algorithm by a round-robin detection algorithm which compares the proposals of different subsets of nodes in order to identify discrepancies. Once corrupted nodes have been identified, the coordinator switches back to the ADMM algorithm.
- *Filtering techniques* [Liao and Chakraborty, 2018]: these techniques do not try to identify compromised nodes, but to employ robust statistics and outlier detection techniques in order to accurately compute a solution even in the presence of malignant data.
- *Non-linear weighting techniques* [Chen et al., 2018]: similarly to filtering techniques, these techniques also do not try to identify compromised nodes. Instead, they employ data from all nodes, but introduce weights to scale down the impact of suspicious nodes.

²<https://www.powerledger.io>

³<https://wepower.network>

- *Convexity techniques* [Musing and Moura, 2018]: these techniques detect compromised nodes and false-data injection in convex algorithms by checking for convexity violations.

Overall, these works focus on cyber-security, *i.e.* the effects of external attacks which compromise a participant in the ADMM algorithm. Moreover, the papers discussed above focus on random noise injection by a malignant agent who tries to disrupt the algorithm’s convergence. In contrast, in our considered scenario we are concerned with the potential strategic manipulation arising from internal self-interested agents (Research Challenge 4). In more detail, rather than considering external malignant attackers, we consider algorithm participants that want to artificially increase their utility in detriment of the other participants, and deviate from the vanilla ADMM algorithm in order to do so. This differs from these existing works in two important aspects: (i) the manipulated algorithm should still converge to a stable outcome; (ii) the manipulating agents will use clever cheating techniques which stir the algorithm’s solution, as opposed to injecting random noise.

Furthermore, we would like to note the connections of this challenge with the field of decentralised mechanism design (DMD). Intimately related to the area of classical mechanism design (see Section 2.6), DMD is concerned with the design of rules governing decentralised multi-agent games [Dash et al., 2003]. In contrast to centralised games, where mechanism design is mainly concerned with preference elicitation, in decentralised scenarios the participating agents have additional important roles that include *message passing* and *computation* [Feigenbaum and Shenker, 2002]. As a consequence, in such a setting, not only truthful preference elicitation but also faithful computation are required in order to obtain strategy-proof cooperation. However, there exists no works in the literature addressing the issue of faithful computation in decentralised optimisation algorithms, such as ADMM and variants.

2.11 Online Energy Allocation

So far, we have considered the participation of EV aggregators in day-ahead electricity markets, and the issues optimal bidding and cooperation. Recall that day-ahead bidding happens one day in advance, so the aggregators forecast their requirements and purchase enough energy to satisfy their forecasted demand. Then, an intrinsically different problem arises when the time comes to schedule the available energy to the EV fleet in real-time. More specifically, the available energy is limited (indeed, it has been purchased in advance hours earlier), and a number of EVs arrive and depart over time, each with different energy requirements. Given these characteristics, online scheduling algorithms, rather than offline, are appropriate. Moreover, eliciting preferences and valuations from EV drivers is not practical or realistic, given that charging can happen daily or several times per day, and most drivers do not have a technical background. As a consequence, efficiency and fairness objectives are a suitable alternative.

These types of scenarios are extensively studied in various fields. Firstly, works that consider scheduling specifically for EV applications include [Mehta et al., 2018, Sun et al., 2018, Liu et al., 2018]. These works are concerned with cost minimisation, but are not interested in fair energy allocation and propose offline rather than online algorithms. Other works, more in line with our purpose, consider online EV scheduling, including [Gerding et al., 2011, 2016, Stein

et al., 2012, Bilh et al., 2018]. However, all of these works consider mechanisms with money instead of fairness. A recent closely related work is [De Weerd et al., 2018]. However, while they also analyse the computational complexity of battery charging algorithms, they focus on offline algorithms and monetary payments instead. Also, they consider binary valuation functions (demand fulfilled/non-fulfilled), rather than partial fulfilment.

Secondly, the extensive field of scheduling has studied related problems (see [Conway et al., 2003, Pinedo, 2012] for overviews). However, none of the existing models considers the features demanded by our scenario. More specifically, a model with dynamic arrivals and departures, soft completion times, and where each element of capacity is perishable. Furthermore, the area of power balancing and bandwidth allocation [Shah, 2012, Shi et al., 2014] has a similar flavour to our scenario, but focuses on the coordination of multiple decentralised actors instead. Finally, Porter [2004] investigates strategic aspects of maximizing weighted completion in online hard real-time scheduling where tasks have weights, release times, deadlines, and durations. However, this work does not allow partial fulfilment and does not consider metrics around fairness.

Thirdly, another relevant field is online fair division [Aleksandrov et al., 2015, Aleksandrov and Walsh, 2017]. In more detail, these works consider the problem of food allocation from food banks. While this literature presents some similarities, such as using an online setting and indivisible goods, their scenarios are very different and do not consider agents that dynamically arrive and depart over time.

The fourth and final similar area of research is that of energy efficient processor scheduling [Yao et al., 1995, Albers, 2010]. In these works, a number of jobs need to be processed on a variable-speed processor where each job has a release time, a deadline, and a processing volume. Certain processing power can be allocated to each job, and this needs to be suitably scheduled. In comparison, the scenario considered in this work is unique in that it has a multi-processor flavour, as each EV is independent.

2.12 Price Impact Forecasting in Day-Ahead Markets

Last, but not least, we will now turn our attention to Research Challenge 6, as described in Section 1.1. Recall that due to the nature of day-ahead markets, and electricity markets in general, bidding takes place hours or days in advance and needs to be based on forecasts. Among these, of special importance for cost-minimising market participants are hourly price forecasts, which are an essential input for bidding algorithms. As a consequence, price forecasting has received a lot of attention in the literature in recent years. In more detail, excellent reviews of the plethora of models and techniques proposed in the literature can be found in [Aggarwal et al., 2009, Weron, 2014]. However, as discussed in Sections 2.3.3 and 2.4, not only hourly prices but also price impact is key for large market participants, such as—but not limited to—EV aggregators and coalitions thereof. We would like to note that price impact curves are intrinsically different from prices as they are functional quantities, as opposed to scalar.

There are several works in the literature investigating the issue of price impact curve forecasting [Calmarza and de la Fuente, 2002, Villar et al., 2001, Baillo et al., 2004, Ugedo et al., 2004, 2006, Aneiros et al., 2011, 2013]. In particular, ARIMA methods from classical time series statistics

are utilised in [Calmarza and de la Fuente, 2002]. In a different vein, clustering algorithms that model the considered forecast by looking for similar days in the past history are used in [Villar et al., 2001, Baillo et al., 2004]. Furthermore, decision trees, modelling the probability of finding different patterns in the residual curves, are considered in [Ugedo et al., 2004, 2006]. Finally, non-parametric functional and semi-functional prediction models are studied in [Aneiros et al., 2011, 2013]. Importantly, most of these works do not strictly focus on the forecasting issues. Instead, they present a brief residual curve forecasting technique and then focus on designing optimal strategies to participate in the considered electricity market. Among these, only [Aneiros et al., 2011, 2013] focus on forecasting and present a systematic study of their results. However, none of these works consider modern forecasting techniques which can yield a significant advantage. Moreover, none of these works exploit the correlations between the residual supply curves of different hours.

Focusing now on modern forecasting techniques, neural networks have received huge amounts of interest in recent years and have achieved spectacular results in many fields [LeCun et al., 2015]. Moreover, current software implementations offer great flexibility and computational efficiency. This translates into important advantages, such as the ability to efficiently process large amounts of data and the opportunity to use complex models. There is ample literature about the use of neural network predictive models in electricity market settings, such as price and demand forecasting. Excellent summaries of these efforts and their performance can be found in [Weron, 2014, 2006]. However, despite this vast literature on price and demand forecasting, to date there exist no studies about the performance of neural networks predicting residual supply curves.

2.13 Summary

This chapter presents the literature review necessary to put the research objectives and challenges into context. We start by providing an overview of the smart grid paradigm and the energy markets that rule the electricity operation in most countries. This is key to understanding the implications of the electrification of the transportation sector, which provides both great opportunities and substantial challenges. We then focus on day-ahead markets, which are the underlying framework of most of our contributions. More specifically, we describe the rules governing these markets and describe how to quantify the impact of larger orders, the so-called price impact.

Next, we focus on Research Challenge 1 and present a literature review on EV aggregation, with a focus on participation in day-ahead markets. Although this scenario has been widely studied in the literature in recent years, the vast majority of these works fail to account for the EV aggregator's price impact, an essential feature for large market participants. On the other hand, the works that do account for price impact present several shortcomings. Firstly, price impact is often modelled in a linear fashion, which results in a loss of accuracy. Secondly, some of the models explicitly solve the market clearing process, introducing computational complexity into the models. Thirdly, all the works are formulated in terms of complex optimisation algorithms with cumbersome constraints. This limits the scalability to large problem sizes and the applicability of the coordination mechanisms discussed below.

Then, we discuss Research Challenge 2, namely the possibility of inter-aggregator coordination for day-ahead bidding, a scenario that has not been previously addressed in the literature. Even though several multi-aggregator scenarios have studied in previous work, none consider the issue of coordinating self-interested aggregators in order to achieve optimal joint operation. Moreover, in order to address the issue of the self-interested nature of the aggregators, we consider two different sets of techniques. Firstly, we consider the field of mechanism design in order to design payment mechanisms that incentivise the EV aggregators to truthfully report their energy requirements to the coordinator. While mechanism design has been widely applied in smart grid scenarios, including EV-related works, it has not been applied for inter-aggregator coordination. Secondly, we consider the EV aggregator problem from the perspective of cooperative game theory, where arbitrary coalitions of EV aggregators can be formed. Although these techniques have applied to coordination mechanisms in several smart grid scenarios, they have not been applied in the multi-EV aggregator setting considered in this work.

Next, we turn our attention to decentralised algorithms (Research Challenge 3), which offer significant advantages with respect to their centralised counterparts, such as better privacy and transparency. Even though there exist works in the literature addressing these issues in related smart grid scenarios, the challenge of decentralised multi-aggregator cooperation has not been addressed. Moreover, we describe the potential applications of blockchain technology in our challenge. In more detail, although blockchains have received a huge amount of attention in the smart grid literature in the last two or three years, they have not been considered for the decentralised coordination of EV aggregators.

However, despite the benefits arising from the decentralised mechanisms described above, they also introduce new challenges. More specifically, the fact that not only private information, but also computation, are distributed among the participating aggregators, introduces novel manipulation opportunities. Although this is an important problem characteristic of distributed optimisation algorithms in general, and thus not limited to our considered scenario, it has received little attention in the literature. In more detail, none of the existing works study Research Challenge 4, namely the issue of strategic manipulation coming from a internal participating agent.

Next, we focus on the issue of online energy allocation (Research Challenge 5). This and related topics have been widely studied in a variety of fields, such as EV management and scheduling. However, none of the existing works consider the same setup we consider in our work. In more detail, an online scenario with dynamic arrivals and departures, soft completion times, and where each element of capacity is perishable. Moreover, rather than employing a mechanisms with money and valuations, we consider efficiency and fairness objectives instead, such as envy-freeness.

Finally, and in order to support participation in day-ahead electricity markets, we turn our attention to the area of forecasting. More specifically, we study the forecasting of hourly price impact curves, as captured in Research Challenge 6. Although there exists a large body of literature addressing price and demand forecasting issues, price impact has received less attention in the literature. Among the existing works, only a few present a systematic description of the technique employed for forecasting, and none employ modern machine learning techniques.

Moreover, none of the existing works exploit the correlations between the residual supply curves of different hours.

We are now ready to present our first contribution.

Chapter 3

A Novel Price-Maker Day-Ahead Bidding Algorithm

In this chapter we present a novel price-maker day-ahead bidding algorithm for EV aggregators, with the aim of addressing Research Challenge 1 (see Section 1.1). In more detail, an EV aggregator is modelled as a cost-minimiser who can make charging and scheduling decisions, and purchases electricity in the day-ahead market to meet its clients' charging needs. The proposed algorithm considers the price impact of the aggregator's orders, while being scalable and computationally tractable. Specifically, our algorithm extends the price-taker algorithm from [Bessa et al., 2012] by accounting for price impact by using residual supply curves [Herranz et al., 2012]. In order to achieve this, we consider two approaches. The first utilises raw historical supply and demand data, while the second formulation employs a quadratic convex approximation which results in an optimisation algorithm with useful theoretical properties. Lastly, we evaluate the performance of the proposed algorithm in an empirical case study which employs real market and driver data from the Iberian Peninsula.

The rest of the chapter is structured as follows. Firstly, we describe the proposed EV aggregator and the mathematical formalism to capture its operation. Secondly, we present our novel bidding algorithm. Lastly, we present the results from our empirical case study.

3.1 The EV Aggregator Model

As detailed in Section 2.4, an EV aggregator manages the charging of a fleet of EVs which uses its services, acting as an intermediary between the drivers and the day-ahead electricity market. The objective of the aggregator is to minimize costs for itself and its customers [Bessa et al., 2012, Gerding et al., 2016, Shafie-khah et al., 2016, Gonzalez Vaya and Andersson, 2015]. Examples of EV aggregators range from a private car park, to a residential micro-grid or a virtual cooperative with no physical proximity.

3.1.1 Modelling Energy Requirements

In our model, following [Bessa et al., 2012, Gonzalez Vaya and Andersson, 2015], EVs arrive and depart dynamically over time. When an EV i arrives to the charging point, it communicates the desired departure time, t_d^i , and desired state of charge at departure, SoC_d^i , to the aggregator. We assume that arrival time and state of charge, t_0^i and SoC_0^i can be automatically inferred by the aggregator. Each EV has a maximum charging speed, P_{\max}^i in kW, which depends on two factors: the available physical infrastructure, and the EV's battery. 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 3.3. The charging schedule of the EV is then left at the aggregator's discretion, which can choose when to perform the charging while guaranteeing the desired state of charge by departure time. This flexibility allows charging the battery in an informed way, rather than randomly, or at arrival, providing cheaper electricity costs and optimised operation.

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.

3.1.2 Forecasting of Energy Requirements

As day-ahead markets close around 12h prior to electricity delivery time (see Section 2.3), 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 [Bessa et al., 2012, Liu et al., 2017, Gonzalez Vaya and Andersson, 2015].

In more detail, the energy requirements of each EV driver 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 3.1. 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 3.2.

Then, in order to provide mathematical tractability, two global energy requirement vectors, \mathbf{R}^{\min} and \mathbf{R}^{\max} , can be obtained by summing the hourly requirements of all the N EVs associated to the particular aggregator, *i.e.* $R_t^{\min} = \sum_{i=1}^N r_t^{\min,i}$ and $R_t^{\max} = \sum_{i=1}^N r_t^{\max,i}$. Note that, although these aggregated constraints do not exactly capture the individual requirements of each EV, they have been widely employed in the literature [Bessa et al., 2012, Bessa and Matos, 2013a,b, Gonzalez Vaya and Andersson, 2015]. The reasons are the fact that considering constraints

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

TABLE 3.1: 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 3.2: Example of requirement vector $\mathbf{r}^{\max,i}$

for each individual EV renders the problem unfeasible even for moderate problem sizes and the intrinsic uncertainty of the forecasts used for bidding. However, in order to compare this approximate formulation with the exact one that considers each EV individually, we present an empirical evaluation in Section 3.3.4.3.

We will denote the quantities that need to be forecasted with a hat: hourly energy requirements, \hat{R}_t^{\min} and \hat{R}_t^{\max} , hourly number of available EVs, \hat{N}_t , and hourly price impact functions, \hat{P}_t . Given that considering advanced forecasting techniques is outside of the scope of this chapter, we will consider two simple methods commonly used in the literature [Bessa et al., 2012, Gonzalez Vaya and Andersson, 2015]. The first technique uses data from the day before as the forecast for the day after. The second considers a perfect forecast, in which exact future information is available. Note that, as will be shown in Section 3.3, the proposed bidding algorithm achieves good results with both forecasting techniques. However, accurate forecasting does present reduced energy costs. In order to address these forecasting issues, we turn our attention to day-ahead market forecasting techniques in Chapter 9.

We are now ready to formalise our proposed bidding algorithm.

3.2 A Novel Price-Maker Day-Ahead Bidding Algorithm

Now that the characteristics of the day-ahead market (see Section 2.3) and the role of the EV aggregator (see Section 3.1) have been defined, we proceed to detail our proposed price-maker day-ahead bidding algorithm. As mentioned above, the algorithm extends the price-taker algorithm from [Bessa et al., 2012] by accounting for price impact by using residual supply curves [Herranz et al., 2012].

The exact problem is as follows. Given an EV aggregator with capacity for N EVs, find the optimal distribution of energy quantities to bid across the 24 hourly slots in the day-ahead market, $\mathbf{E} = (E_0, \dots, E_{23})$, 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:

$$\min_{\mathbf{E}} \sum_t \hat{\mathcal{P}}_t(E_t) \cdot E_t \quad (3.1a)$$

$$\sum_{j=0}^t E_j \geq \sum_{j=0}^t \hat{R}_j^{\min}, \quad \forall t = 0, \dots, 23 \quad (3.1b)$$

$$\sum_{j=0}^t E_j \leq \sum_{j=0}^t \hat{R}_j^{\max}, \quad \forall t = 0, \dots, 23 \quad (3.1c)$$

$$E_t / 1 \text{ hour} \leq \hat{N}_t P_{\max}, \quad \forall t = 0, \dots, 23 \quad (3.1d)$$

$$E_t \geq 0, \quad \forall t = 0, \dots, 23 \quad (3.1e)$$

In more detail, the objective function (Eq. 3.1a) minimizes the total cost of the purchased energy. As explained in Sections 2.3.3 and 3.1.2, $\hat{\mathcal{P}}_t(E_t)$ 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 (Eq. 3.1b), that it is not purchased before the forecasted arrival of the EVs (Eq. 3.1c), 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 (Eq. 3.1d) and that energy quantities are non-negative (Eq. 3.3e). It is worth noting that the number of constraints is always 96, independent on the fleet size. However, in contrast with the original algorithm without price impact [Bessa et al., 2012], the novel optimization problem given by Eqs. 3.1a, 3.1b, 3.1c, 3.1d, 3.3e is no longer linear, making the problem more complex to solve.

Specifically, due to the arbitrary nature of the price impact functions obtained from real market historical data, the optimisation problem is non-convex and non-linear. As a consequence, it presents a complex optimisation landscape with multiple local minima, preventing solvers from obtaining the global minimum. We will henceforth refer to this minimisation problem as the *raw formulation*. In order to overcome these issues, we propose a convex approximation, in which any local minimum is actually globally optimal. The advantage of this approximation, the *convex formulation*, is that the global minimum is guaranteed to be found, which has attractive theoretical properties for the coordination mechanisms proposed in Chapters 4 and 5.

In more 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




FIGURE 3.1: 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.

terms takes the following form:

$$\min_{\mathbf{E}} \sum_t \hat{p}_t^{\text{convex}}(E_t) \cdot E_t \quad (3.2a)$$

$$\sum_{j=0}^t E_j \geq \sum_{j=0}^t \hat{R}_j^{\min}, \quad \forall t = 0, \dots, 23 \quad (3.2b)$$

$$\sum_{j=0}^t E_j \leq \sum_{j=0}^t \hat{R}_j^{\max}, \quad \forall t = 0, \dots, 23 \quad (3.2c)$$

$$E_t / 1 \text{ hour} \leq \hat{N}_t P_{\max}, \quad \forall t = 0, \dots, 23 \quad (3.2d)$$

$$E_t \geq 0, \quad \forall t = 0, \dots, 23 \quad (3.2e)$$

This resulting optimisation problem is then convex and can be readily solved optimally with little computational burden by using, for example, the SLSQP algorithm [Nocedal and Wright, 2006].

However, this convex approximation introduces a deviation from real data which will affect the accuracy of the solution. The degree of deviation depends on the price impact curves for the considered day, ranging from very good fit to larger discrepancies, as shown in Figure 3.1. The performance of the *raw* and *convex* approximations is evaluated and compared in Section 3.3.4.1.

Finally, in a slightly different vein, we now turn our attention to considering individual EV constraints, as discussed in Section 3.1.2. Specifically, the minimisation problem will now decide on an energy allocation for each individual EV, while also imposing individual constraints. Formally, let E_t^i be the energy allocation for EV i at hour t and $\mathbf{E}^i = (E_0^i, \dots, E_{23}^i)$. Moreover, let n_t^i be the boolean variable (0 or 1) representing whether EV i is available for charging at hour

t . The exact bidding algorithm is then given by:

$$\min_{\mathbf{E}^1, \dots, \mathbf{E}^N} \sum_t \left[\left(\sum_{i=1}^N E_t^i \right) \cdot \hat{\mathcal{P}}_t^{\text{convex}} \left(\sum_{i=1}^N E_t^i \right) \right] \quad (3.3a)$$

$$\sum_{j=0}^t E_j^i \geq \sum_{j=0}^t \hat{r}_j^{\text{min}, i}, \quad \forall t = 0, \dots, 23, \quad \forall i = 1, \dots, N \quad (3.3b)$$

$$\sum_{j=0}^t E_j^i \leq \sum_{j=0}^t \hat{r}_j^{\text{max}, i}, \quad \forall t = 0, \dots, 23, \quad \forall i = 1, \dots, N \quad (3.3c)$$

$$E_t^i / 1 \text{ hour} \leq \hat{n}_t^i P_{\text{max}}, \quad \forall t = 0, \dots, 23, \quad \forall i = 1, \dots, N \quad (3.3d)$$

$$E_t^i \geq 0, \quad \forall t = 0, \dots, 23, \quad \forall i = 1, \dots, N \quad (3.3e)$$

Note that, in this case, the number of decision variables and the number of constraints are, respectively, $24N$ and $96N$, where N is the size of the aggregator. Recall that, in contrast, the aggregated formulation presents 24 decision variables and 96 constraints, independent of the number of EVs. As a consequence, the scaling of the exact formulation is much worse than the approximation. We discuss these issues in more detail in Section 3.3.4.3, where we present an empirical evaluation comparing the exact and approximate formulations.

3.3 Case Study

This section describes an empirical case study, based on real data, with the aim of testing the performance of the bidding algorithms proposed in this Chapter. We will start by describing the empirical scenario, detailing the considered EV aggregators and market and real driver behaviour data. Then, the results concerning the performance of the proposed algorithms will be presented and described. All the data from this empirical case study is available in [Perez-Diaz et al., 2018f].

3.3.1 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, similarly to [Bessa et al., 2012]. With respect to the EV fleet, we consider medium-size electric vehicles with battery capacities of 24kWh [Heydarian-Forushani et al., 2016, Gerding et al., 2016]. Charging speed is considered to be the same for all EVs and set to $P_{\text{max}} = 3.7\text{kW}$ [Gerding et al., 2016, Bessa et al., 2012]. 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 [Yao et al., 2016, Su and Chow, 2012].

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 [ECCC, 2016].

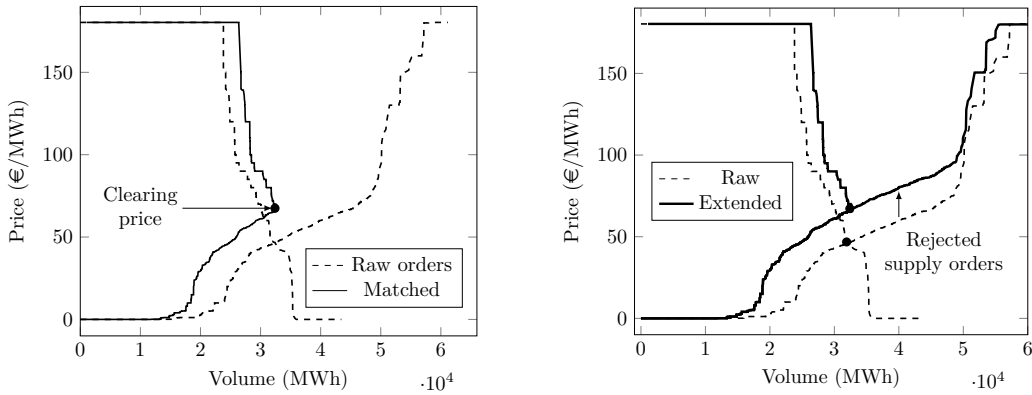


FIGURE 3.2: (LHS) Raw and matched market orders, together with the hourly clearing price. (RHS) Appending the rejected supply orders (dashed) to the accepted matched supply curve (solid). Source: OMIE, 02/11/2016, 11th hour.

3.3.2 Day-Ahead Market Data

Real market data from the Iberian day-ahead market, OMIE, is employed in the simulations [OMIE, 2017]. 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 [Herranz et al., 2012].

In more detail, all European markets employ the EUPHEMIA [EUPHEMIA, 2015] 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.

Specifically, 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 3.2. Given that modelling the exact clearing process is impossible, given that the available data does not contain the details of order types, approximations are employed in the literature [Vázquez et al., 2014].

The procedure works 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 3.2. 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.

Finally, residual supply curves are employed to build price impact functions as described in Section 2.3.3. In order to obtain the convex approximation of price impact functions described

in Section 3.2, 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 3.2.

3.3.3 Driver Behavioural Data

In order to model the behaviour of the EVs in our simulations, we use real data from the Spanish driver behaviour survey MOVILIA [Ropero Ortega, 2014]. 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 [Alonso et al., 2014], 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 3.3 and 3.4. 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 [Yao et al., 2016].

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

TABLE 3.3: 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 3.4: Possible departure times rounded to the nearest hour, with their respective probabilities.

The results from the simulations are described next, analysing the performance of the bidding algorithm proposed in Section 3.2.

3.3.4 Results: Performance of the Bidding Algorithm

Now that the empirical setting has been detailed, we are ready to present the results of the experiments. Firstly, we compare the two formulations of the novel bidding algorithm proposed in Section 3.1.2, namely, the *raw* and *convex* formulations. Secondly, we assess the performance of our algorithm by direct comparison with two other existing algorithms. Third and lastly, we compare the approximate and exact formulations.

3.3.4.1 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. 3.1a, 3.1b, 3.1c, 3.1d, 3.3e and the convex approximation

given by Eqs. 3.2a, 3.2b, 3.2c, 3.2d, 3.2e. Both formulations are solved using the SLSQP algorithm, specifically the implementation from [Jones et al., 2001]. Moreover, we present two separate results in the raw case. This is motivated by the fact that the problem is non-convex and the solution provided by the solver is not globally optimal in general, instead depending on the solver’s initial guess. A common technique to address this issue is basin-hopping [Wales, 2003]. This method sequentially solves the non-convex minimisation problem using an slightly randomised previous solution as initial guess. However, there is no guarantee of global optimality. In this empirical evaluation, we will solve the *raw* optimisation problem with and without basin-hopping. 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 [Bessa et al., 2012].

In more detail, we present a comparison of EV aggregators of different sizes employing both formulations and perfect forecasting, during four months of day-ahead trading, as specified in Section 3.3.2. The results are shown in Figure 3.3. 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. Given the larger computational complexity of basin-hopping (which needs to repeatedly solve the optimisation problem changing the initial guess), the very close performance achieved by both formulations, and the theoretical guarantees offered by convexity, the convex formulation will be employed in all the subsequent Chapters.

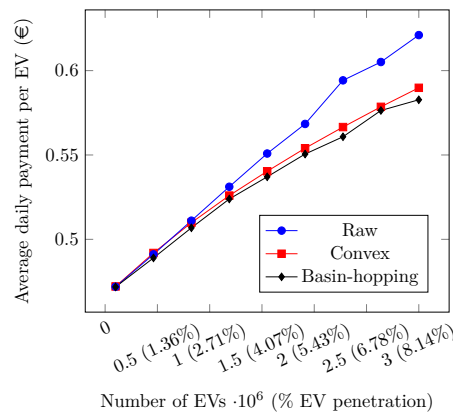


FIGURE 3.3: Average daily payments per EV, for both bidding algorithm formulations and for different EV aggregator sizes, employing perfect forecasts.

3.3.4.2 Comparison with Benchmark Algorithms

We now compare the proposed bidding algorithm (using the convex approximation) with two existing bidding algorithms. The first of these algorithms is a simple zero-intelligence 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 [Bessa et al., 2012, Alonso et al., 2014]. The second is the price-taker algorithm proposed by Bessa et al. [2012], which inspired our proposed price-maker algorithm. We will henceforth refer to these strategies

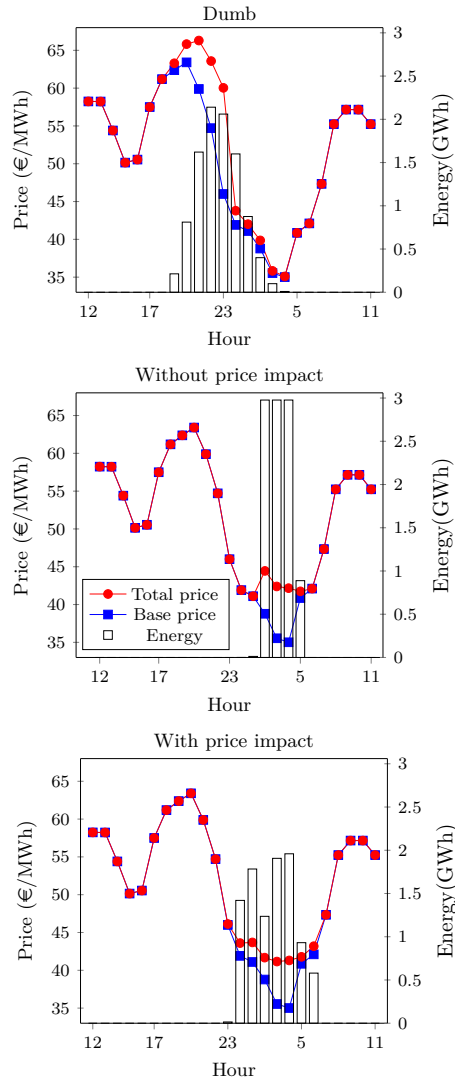


FIGURE 3.4: Comparison of the three bidding strategies for a single day: dumb, without price impact and with price impact for the same scenario where the aggregators manage one million EVs. Curves represent hourly prices, both with and without the EV aggregator bids, and correspond to the LHS axis. Bars represent the amount of energy purchased at each hourly slot by the EV aggregator and correspond to the RHS axis. Market data corresponds to OMIE, 09-10/11/2016.

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 3.1.2), 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 3.4. 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 3.4 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




FIGURE 3.5: Average daily payments per EV, for each different bidding strategy and for different EV aggregator sizes. Both with naive and perfect forecasts.




FIGURE 3.6: Percentage payment reduction when using the novel algorithm with price impact, compared to its counterpart without price impact.

EVs stay idle during the night and only need to be charged by the morning. By forecasting the EV requirements as explained in Section 3.1.2, 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 reduced energy 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 3.5. 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 3.6. 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 3.1.2). 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*.

3.3.4.3 Comparison Between the Exact And Approximate Formulations




FIGURE 3.7: (LHS) Runtimes for the approximate algorithm. (Center) Runtimes for the exact algorithm. (RHS) Normalised error between the solutions of the exact and approximate algorithm. Results are averaged across the first 15 days of November 2015 and error bars represent standard error.

In previous experiments, we have studied the performance of our proposed bidding algorithm in comparison with previous algorithms from the literature, and the effect of the convex approximation. Now, we turn our attention to the effects of aggregating together individual EV constraints, as described at the end of Section 3.2. To this end, we run both the approximated algorithm (Eqs. 3.2a, 3.2b, 3.2c, 3.2d) and the exact algorithm (Eqs. 3.3a, 3.3b, 3.3c, 3.3d) on the first 15 trading days of November 2016, in the same night-time scenario considered in the previous two experiments.

Given the computational complexity of the exact algorithm, it is unfeasible to use it even in moderate problem sizes, so we will focus the analysis on small scenarios varying the number of EVs. For each of this scenarios, we will compute the exact and the approximated solutions, $\mathbf{E}^{\text{exact}}$ and $\mathbf{E}^{\text{approx}}$, together with runtimes for each algorithm. In order to assess the approximation error, we compute the normalised mean square errors for each instance, given by:

$$\text{error}(\mathbf{E}^{\text{exact}}, \mathbf{E}^{\text{approx}}) = \frac{1}{24} \sum_{t=0}^{23} \left(\tilde{E}_t^{\text{exact}} - \tilde{E}_t^{\text{approx}} \right)^2$$

where $\tilde{\mathbf{E}}^{\text{approx}} = \mathbf{E}^{\text{approx}} / \|\mathbf{E}^{\text{approx}}\|_2$ and $\tilde{\mathbf{E}}^{\text{exact}} = \mathbf{E}^{\text{exact}} / \|\mathbf{E}^{\text{exact}}\|_2$ are the normalised solution vectors. This normalisation is introduced to remove dependence on the number of EVs and thus allow fair comparison across simulations. Results are presented in Figure 3.7. In more detail, runtimes for both algorithms are shown in the left-hand side and centre plots, and the dependence of the normalised error with the number of EVs in the right-hand side. With respect to runtimes, the approximated algorithm can be run very quickly (~ 1 sec) and does not show a dependence on the number of EVs. On the contrary, the exact algorithm scales exponentially with the number of EVs taking up to 20 hours to compute the solution in scenarios with 100

EVs. With respect to the approximation error, and in this particular experiment, we do not see any particular dependence on the number of EVs. Of course, one cannot extrapolate conclusions for the larger problem sizes considered in this Chapter and in the rest of this thesis. However, the poor scaling and hence limited applicability of the exact algorithm motivate the need for the proposed approximated algorithm, as detailed in Section 3.2.

3.4 Summary

This chapter addresses Research Challenge 1 with respect to day-ahead markets, by presenting a novel price-maker day-ahead bidding algorithm for EV aggregators. By taking into account the price impact of the aggregator's bids, our algorithm provides greatly reduced energy costs for large aggregators when compared to existing price-taker algorithms. With respect to existing price-maker algorithms, the proposed algorithm is formulated in terms of simple linear constraints, resulting in a very scalable optimisation problem. In more detail, two different formulations of the optimisation problem are proposed: the raw formulation and a convex approximation. While the raw formulation achieves slightly better results, the convex approximation is computationally faster and has the theoretical guarantee of finding the global minimum, which will prove essential for some of the results presented in later chapters.

Moreover, we evaluate the performance of the proposed algorithm in an empirical case study with real market and driver data from the Iberian Peninsula. More specifically, we benchmark the bidding algorithm in a night-time scenario where drivers arrive in the evening and depart in the morning. We use historical market data to simulate four months of day-ahead trading. Also, EV requirements are based on real driver behaviour data, including the distribution of arrival times in the evening and departure times in the morning. We compare our proposed bidding algorithm with two existing models, which serve as a benchmark. We find that, while for small aggregators the results are comparable, as the aggregator grows in size, significant cost reductions are achieved by our proposed algorithm. More specifically, we achieve over 10% savings for an EV penetration rate of 8% in the Iberian Peninsula.

This will be key when we consider joint multi-aggregator bidding, as detailed in Chapters 4 and 5, where bidding volume is high and price impact very significant. Finally, we evaluate the possibility of using an analogous minimisation formulation which considers each EV separately and thus produces an exact solution, at the expense of losing computational tractability. Results indicate that this exact formulation presents very poor scaling even with small problem sizes with hundreds aggregators. With respect to the approximation error incurred by the proposed approximate formulation, our experiments do not show any dependence on the number of aggregators.

Chapter 4

Inter-Aggregator Coordination: a Mechanism Design Approach

The previous chapter explored the day-ahead bidding of an EV aggregator in order to obtain the energy needed to charge its clients' vehicles. By employing a bidding algorithm, the aggregator can obtain the day-ahead bidding schedule that minimises energy costs while being able to satisfy its clients' needs. This way, the aggregator optimises the charging of the vehicles it manages, achieving lower energy costs and less grid impact. In addition, given the decentralised structure of electricity markets in most countries, any entity can become a market participant. This fact, together with the fast growth of the EV industry and the ambitious worldwide penetration targets for the near future, motivates studying a multi-aggregator scenario.

This will be the objective of this and subsequent chapters, which will address Research Challenge 2 (see Section 1.1). In more detail, we consider a number of independent and self-interested EV aggregators competing in the same day-ahead market. Similarly to the benefits of EV aggregation versus individual EV operation, we will study the potential benefits arising from inter-aggregator cooperation. This is not without challenges, given the self-interested nature of the EV aggregators. In order to overcome this issue, in this chapter we propose a novel coordination mechanism based on mechanism design techniques which incentivises cooperation rather than strategic manipulation. In more detail, we propose the use of a centralised third-party coordinator, which incentivises the participating aggregators to truthfully report their energy needs, and bids on their behalf in the day-ahead market. More specifically, it proceeds in three stages. Firstly, the individual requirements reported by the EV aggregators are combined together, 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. This coordination mechanism is visually depicted in Figure 4.1.

This proposed coordination mechanism detailed in this chapter is generic and independent of the particular bidding algorithm employed. However, the simple constraint formulation and little computational burden characteristic of the novel bidding algorithm introduced in Chapter

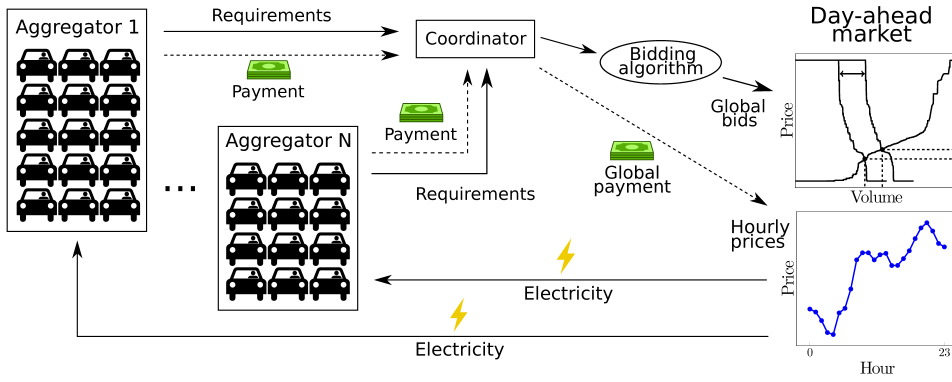


FIGURE 4.1: Graphical depiction of the proposed coordination mechanism.

3 guarantee good scaling with the 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.

The rest of the chapter is structured as follows. In Sections 4.1, 4.2 and 4.3 we detail the three stages of the proposed coordination mechanism. Then, in Section 4.4, we briefly discuss the details of a potential real implementation. Finally, in Section 4.5, we evaluate the performance of the novel coordination mechanism in an empirical case study which uses real market and driver data. Throughout this chapter, we consider an EV aggregator as proposed in Chapter 3. Specifically, energy requirements are modelled as in Section 3.1.1. Despite the proposed coordination mechanism being independent of the underlying bidding algorithm, we will focus on the bidding algorithm proposed in Section 3.2.

4.1 Stage 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 3.1.1. The combined requirements of all the aggregators are then:

$$\hat{R}_t^{\min} = \sum_{i=1}^n \hat{R}_t^{\min,i} \quad (4.1)$$

$$\hat{R}_t^{\max} = \sum_{i=1}^n \hat{R}_t^{\max,i} \quad (4.2)$$

$$\hat{N}_t = \sum_{i=1}^n \hat{N}_t^i \quad (4.3)$$

In order to find the optimal global energy bids, we can use the bidding optimisation algorithm proposed in the previous Chapter, given by Eqs. 3.2a, 3.2b, 3.2c, 3.2d, but now using

the combined requirements given by Eqs. 4.1, 4.2, 4.3. This will result in obtaining a global day-ahead energy volume E_t^{global} for each hour t .

4.2 Stage 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, focusing on the bidding algorithm proposed in Chapter 3, 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 t , find E_t^i for $t = 0, \dots, 23$ and $i = 1, \dots, n$ satisfying the following constraints:

$$\sum_{j=0}^t E_j^i \geq \sum_{j=0}^t \hat{R}_j^{\min,i}, \quad \forall t = 0, \dots, 23; \forall i = 1, \dots, n \quad (4.4a)$$

$$\sum_{j=0}^t E_j^i \leq \sum_{j=0}^t \hat{R}_j^{\max,i}, \quad \forall t = 0, \dots, 23; \forall i = 1, \dots, n \quad (4.4b)$$

$$E_t^i / \Delta t \leq \hat{N}_t^i P_{\max}, \quad \forall t = 0, \dots, 23; \forall i = 1, \dots, n \quad (4.4c)$$

$$E_t^i \geq 0, \quad \forall t = 0, \dots, 23; \forall i = 1, \dots, n \quad (4.4d)$$

$$\sum_{i=1}^n E_t^i = E_t^{\text{global}}, \quad \forall t = 0, \dots, 23 \quad (4.4e)$$

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 [Tsang, 1996, Gent et al., 2006]. In more detail, Eqs. 4.4a, 4.4b, 4.4c, 4.4d ensure that each EV aggregator has enough energy to satisfy its requirements for each hour. Eq. 4.4e makes sure the sums of the allocated hourly energies add up to the available global energy.

Existence of a solution is guaranteed by definition of the optimization problem, Eqs. 3.2a, 3.2b, 3.2c, 3.2d, 3.2e. However, uniqueness is not guaranteed. For example, consider the case with two identical EV aggregators who report the same preferences to the coordinator. 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. 4.4a, 4.4b, 4.4c, 4.4d, 4.4e grows linearly with the number of participating EV aggregators, n , being the number of constraints $96n + 24$. With this particular choice of model constraints, given by the bidding algorithm proposed in Chapter 3, 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.

4.3 Stage 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 [Nisan et al., 2007] 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 4.5) we find that the typical VCG mechanism, known as VCG with Clarke pivot payments [Nisan et al., 2007] and henceforth referred to here as *pure VCG*, results in too high payments in this setting. This is a common issue, given that VCG is not strongly budget-balanced Guo and Conitzer [2009]. 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.

4.3.1 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 [Vlassis, 2007, Nisan et al., 2007]. 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 optimality¹ 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 $E(\hat{\mathbf{R}}^{\min}, \hat{\mathbf{R}}^{\max}, \hat{\mathbf{N}}) = (E_0^{\text{global}}, \dots, E_{23}^{\text{global}}) = \mathbf{E}^{\text{global}}$ be the energy schedule provided by the bidding algorithm (Eq. 3.2a) 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 $\mathcal{P}(\mathbf{E}) = (P_0, \dots, P_{23})$ be the real hourly clearing prices after market closure.

Then, the total aggregated forecasted cost incurred by the coordinator is given by:

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

¹More precisely, the allocation needs to be *optimal in range* [Nisan and Ronen, 2007].

and the real total aggregated cost incurred by the coordinator (when the market is cleared) 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 intrinsic uncertainty of the forecasts, and 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:

$$p_i = \text{Cost}^{\text{convex}}(\hat{\mathbf{R}}^{\min}, \hat{\mathbf{R}}^{\max}, \hat{\mathbf{N}}) - \text{Cost}^{\text{convex}}(\hat{\mathbf{R}}_{-i}^{\min}, \hat{\mathbf{R}}_{-i}^{\max}, \hat{\mathbf{N}}_{-i}) \quad (4.5)$$

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). Consequently, in order for this coordination mechanism to be scalable and applicable, the bidding algorithm needs to be computationally quick and have good scaling properties. Both requirements are satisfied by the bidding algorithm presented in Chapter 3 and used throughout this chapter. As a result, we can 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 the fairness of the payments. Moreover, this Pure VCG payment mechanism is *truthful*, meaning that it is in the best interest of every participating aggregator to truthfully report their energy requirements [Nisan et al., 2007].

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 forecasted price curve. As a result, the coordinator can, depending on the accuracy of the prediction, incur a loss. However, as shown in Section 4.5, 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* [Nisan et al., 2007], also known as the participation constraint (meaning EV aggregators have no incentive to participate if given the choice). Hence, to address this problem, 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.

4.3.2 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 [Cavallo, 2006, Guo and Conitzer, 2008, 2010]. Formally, following the results by Cavallo [2006], the redistributed amount to participant i is given by:

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

where $\hat{\mathbf{R}}_{-i-j}^{\min}$, $\hat{\mathbf{R}}_{-i-j}^{\max}$ and $\hat{\mathbf{N}}_{-i-j}$ are the vectors of requirements without the contributions from the i -th and j -th EV aggregators. Finally, given the payments defined by Eq. 4.5, the total payment of each EV aggregator i is given by:

$$p_i = \left[\text{Cost}^{\text{convex}}(\hat{\mathbf{R}}^{\min}, \hat{\mathbf{R}}^{\max}, \hat{\mathbf{N}}) - \text{Cost}^{\text{convex}}(\hat{\mathbf{R}}_{-i}^{\min}, \hat{\mathbf{R}}_{-i}^{\max}, \hat{\mathbf{N}}_{-i}) \right] - \frac{1}{n} \sum_{\substack{j=1 \\ j \neq i}}^n \left[\text{Cost}^{\text{convex}}(\hat{\mathbf{R}}_{-i}^{\min}, \hat{\mathbf{R}}_{-i}^{\max}, \hat{\mathbf{N}}_{-i}) - \text{Cost}^{\text{convex}}(\hat{\mathbf{R}}_{-i-j}^{\min}, \hat{\mathbf{R}}_{-i-j}^{\max}, \hat{\mathbf{N}}_{-i-j}) \right] \quad (4.6)$$

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 [Guo and Conitzer, 2008]. However, in our experiments (see Section 4.5), 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.

4.3.3 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:

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

$$p_i^2 = p_i^1 - \frac{N_i}{\sum_{j=1}^n N_j} \left[\sum_{j=1}^n p_j^1 - \text{Cost}(\hat{\mathbf{R}}^{\min}, \hat{\mathbf{R}}^{\max}, \hat{\mathbf{N}}) \right] \quad (4.7)$$

It is worth noting that, in contrast to the redistribution approach discussed in Section 4.3.2, 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 practice, 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 4.5, 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.

4.4 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 3.2. 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 4.2 can be applied, using the same standard computer. Energy delivery to the aggregators is then managed by the distribution system operator (DSO) [Bessa et al., 2012] in the same way as in the individual bidding case.
3. Payment distribution: payments are computed by the coordinator (see Section 4.3), 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.

4.5 Case Study

This section describes a case study based on real data, in order to test the performance of the coordination mechanism proposed in this chapter. The considered scenario is the same as in the case study from Chapter 3: a night-time scenario where EVs arrive in the evening and depart in the morning. In more detail, the employed market and driver data are described in Sections

3.3.2 and 3.3.3, respectively. Moreover, all the data from this empirical case study is available in [Perez-Diaz et al., 2018f].

Differently from Chapter 3, we now consider the setting where a number of self-interested and independent EV aggregators participate in the same day-ahead market. In order to assess the performance of the proposed coordination mechanism, we compare the setting with *uncoordinated* bidding, where each EV aggregator does independent bidding based on past historical data, and

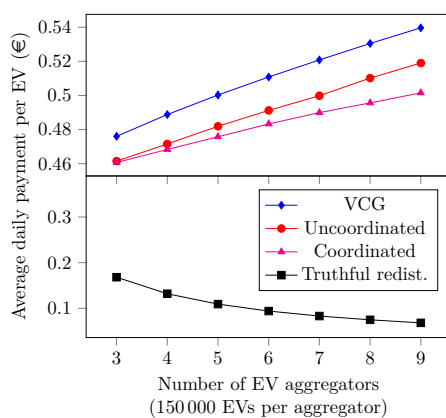


FIGURE 4.2: Average daily payments per EV when using a perfect forecast.

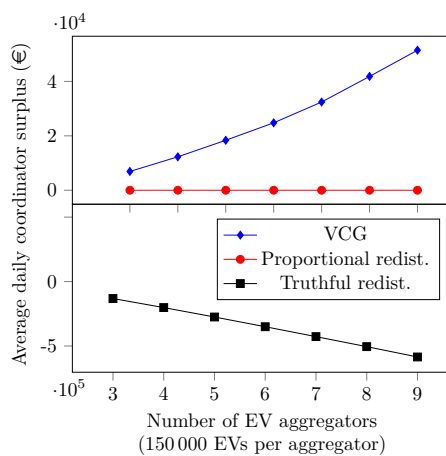


FIGURE 4.3: Average daily monetary surplus for the coordinator when using a perfect forecast.

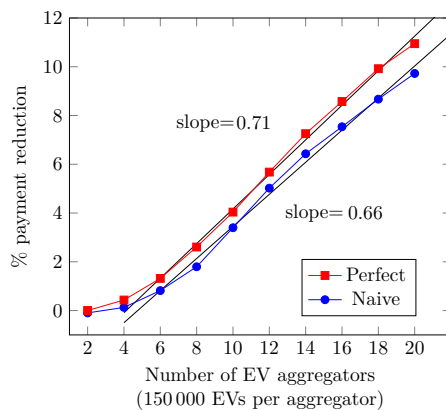


FIGURE 4.4: Percentage payment reduction when using the proportional redistribution mechanism w.r.t. uncoordinated bidding, using both naive and perfect forecasts.

coordinated bidding, where all the EV aggregators participate in the coordination mechanism presented in this chapter. Furthermore, for the coordinated setting, we compare the three different payment mechanisms, *i.e.*, the VCG mechanism (Eq. 4.5), the truthful redistribution (Eq. 4.6) and proportional redistribution (Eq. 4.7).

Specifically, Figures 4.2 and 4.3 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 4.4 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 4.4), we already see that the coordination mechanism achieves around a 10% payment reduction.

4.6 Summary

In this chapter we have proposed the first multi-aggregator coordination mechanism for joint bidding in day-ahead markets, targeting a generalised scenario where a number of EV aggregators of arbitrary nature compete in the same day-ahead market. In more detail, each aggregator is assumed to be a profit maximiser entity and hence self-interested. In order to extend the benefits of coordinated bidding to this multi-aggregator setting, we employ tools from the field of mechanism design to incentivise the participating EV aggregators to cooperate.

The proposed coordination mechanism consists of three stages. Firstly, the EV aggregators report their energy requirements to a third-party coordinator. The coordinator is then able to compute aggregated requirements and bid in the day-ahead market. Given its low computational complexity, and its accounting for price impact, the novel bidding algorithm proposed in Chapter 3 is well suited for this task. Secondly, the coordinator solves a constraint satisfaction problem in order to distribute the globally purchased energy among the EV aggregators, where each gets an energy schedule satisfying their energy requirements. Thirdly and lastly, the coordinator distributes the joint energy costs among the aggregators in a fashion such each aggregator pays for the impact of their energy requirements on the overall bids. Mechanism design is employed in this last stage in order to incentivise inter-aggregator cooperation and prevent strategic manipulation.

The performance of the proposed coordination mechanism is finally evaluated in an empirical case study which uses real market and driver data from the Iberian Peninsula. Results indicate that

the coordination mechanism (employing the proposed bidding algorithm in Chapter 3) 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.

This work partially addresses Research Challenge 2, from a centralised mechanism design perspective. However, the proposed coordinator is not necessarily unique, and different coalitions of aggregators may form. This different viewpoint will be studied in detail in the next chapter.

Chapter 5

Inter-Aggregator Coalition Formation

As discussed before, any entity with sufficient resources can become an electricity market participant. This is also true in the electric vehicle industry, where many different companies already manage fleets of EVs, such as car sharing companies, car parks, large electricity retailers, *etc.* In order to study this multi-aggregator scenario, we have so far proposed an inter-aggregator coordination mechanism in Chapter 4, and showed that it is able to reduce electricity costs by spreading energy consumption in time. However, we have assumed that the third-party coordinator is a unique entity that coordinates all the EV aggregators present in a given market. This is not necessarily true, as the proposed coordination mechanism has few requirements and is computationally inexpensive. This makes it easy for any trusted entity to assume the role of the coordinator, and this raises questions about the scenario where different EV aggregators are aggregated under different coordinators, which we will address in this chapter in order to further explore Research Challenge 2.

In more detail, we study the coordination mechanism proposed in Chapter 4 in a scenario where multiple coordinators may be present, and the aggregators can choose to join smaller cooperative groups. In order to address this issue, we employ techniques from the field of cooperative game theory (see Section 2.7). Specifically, we focus on finding payoff allocations (*i.e.* payment mechanisms) which result in fair and stable coalitions.

Furthermore, and similarly to the contributions of the previous chapter, it is worth noting that the multi-aggregator coalitional analysis proposed in this chapter is generic and can be applied employing different underlying bidding algorithms. However, we will focus on the novel bidding algorithm presented in Chapter 3 due to its good scalability and computational efficiency.

The rest of the chapter is structured as follows. We start by presenting an introduction to cooperative game theory in Section 5.1. Then, we define the multi-aggregator coalitional game in Section 5.2. We discuss the presence of coalitional externalities in Section 5.3. We then turn our attention to analysing and proving several theoretical properties of the proposed coalitional game in Section 5.4. The issue of choosing appropriate payment allocations is then discussed

in Section 5.5. Finally, we present an empirical case study with real market and driver data in Section 5.6.

5.1 A Brief Introduction to Cooperative Game Theory

We start by presenting the basic concepts of cooperative game theory, following the exposition in [Chalkiadakis et al., 2012]. Throughout the rest of this chapter, consider a set of players $N = \{1, \dots, n\}$, which in our setting correspond to different aggregators.

5.1.1 Characteristic Function Games

We start by defining a coalitional game.

Definition 5.1 (Coalition). A *coalition* is any subset of players $C \subseteq N$. The number of players in the coalition C is given by its cardinality $|C|$. All possible coalitions are denoted by the power set of N , 2^N . The *grand coalition* is the set of all players, N .

Definition 5.2 (Coalition structure). A *coalition structure* over N is a collection of non-empty subsets $CS = \{C_1, \dots, C_k\}$ such that $\cup_{j=1}^k C_j = N$ and $C_i \cap C_j = \emptyset \forall i \neq j$.

Definition 5.3 (Characteristic function game). A *characteristic function game* G is given by a pair (N, v) , where N is a finite and non-empty set of players, and $v : 2^N \rightarrow \mathbb{R}$ is a *characteristic function*. The value $v(C)$ is usually referred to as the *value* of the coalition C .

Note that the characteristic function assigns a value to the whole coalition, not to its individual members. Games in which a coalition value, $v(C)$, can be divided in any way among its members are called *transferable utility* (TU) games.

Definition 5.4 (Superadditive game). A coalitional game (N, v) is *superadditive* if for every pair of disjoint coalitions $C_1, C_2 \subset N$ such that $C_1 \cap C_2 = \emptyset$, we have $v(C_1) + v(C_2) \leq v(C_1 \cup C_2)$.

In other words, in a superadditive game, the grand coalition has the incentive to form, as the agents can earn at least as much profit by working together. Now that the basics have been defined, we will focus on how the value of a given coalition is distributed among its members.

5.1.2 Payoff Allocations

The distribution of the value of a given coalition among its members can be done in an arbitrary way. However, certain allocations have special properties, such as efficiency and individual rationality, which make them desirable, as defined below.

Definition 5.5 (Payoff allocation). A vector $\mathbf{x} = (x_1, \dots, x_n)$ is a payoff allocation vector for a coalition structure $CS = \{C_1, \dots, C_k\}$ over N if $x_i \geq 0$ for all $i \in N$, and $\sum_{i \in C_j} x_i \leq v(C_j)$ for any $j = 1 \dots, k$.

1. (*Efficiency*) An allocation \mathbf{x} is *efficient* if $\sum_{i \in C_j} x_i = v(C_j)$ for any $j = 1 \dots, k$.

2. (*Individually rational*) An allocation is *individually rational* if $x_i \geq v(\{i\})$ for all $i \in N$.

Definition 5.6 (Imputation). A payoff allocation for the grand coalition N is said to be an *imputation* if it is both efficient and individually rational.

Definition 5.7 (The Core). Given a TU characteristic function game (N, v) , the *core*, \mathcal{C} , is defined as the set of imputations such that no sub-coalition can obtain a payoff which is better than the sum of the members current payoffs. Explicitly:

$$\mathcal{C} := \left\{ \mathbf{x} \in \mathbb{R}^N \mid \sum_{i \in N} x_i = v(N), \sum_{i \in S} x_i \geq v(S), \forall S \subseteq N \right\}$$

In other words, payoff allocations lying in the core provide stability, as no sub-coalition has an incentive to deviate from the grand coalition in order to increase its profit. However, the core of coalitional game can be empty [Chalkiadakis et al., 2012]. A classical result guarantees the non-emptiness of the core for certain games, the so-called *balanced* games.

Definition 5.8 (Balanced function). A function $\alpha : 2^N \rightarrow \mathbb{R}$ is said to be *balanced* if for all $i \in N$, we have $\sum_{C \in 2^N} \alpha(C) \mathbf{1}\{i \in C\} = 1$, where $\mathbf{1}$ is the indicator function.

Definition 5.9 (Balanced game). A coalitional game (N, v) is balanced if for any balanced function α , we have $\sum_{C \in 2^N} \alpha(C) v(C) \leq v(N)$.

Theorem 5.10 (Bondareva – Shapley Theorem, [Shapley, 1967]). *A coalitional game has a non-empty core if and only if it is balanced.*

We will now focus on one of the most commonly used payoff allocations, the Shapley value.

5.1.3 The Shapley Value

One of the most widely used allocations is the *Shapley value*, defined as:

Definition 5.11 (Shapley value [Shapley, 1971]). Given a coalitional game (N, v) , the *Shapley value* assigns a payoff SV_i to each player $i \in N$ given by:

$$SV_i(v) = \sum_{C \subseteq N \setminus \{i\}} \frac{|S|!(N - |S| - 1)!}{N!} [v(C \cup \{i\}) - v(C)]$$

The Shapley value is the only payment allocation satisfying the *efficiency*, *symmetry*, *dummy action* and *additive* axioms [Chalkiadakis et al., 2012]. Hence, it is traditionally considered to present a fair payoff distribution. Moreover, there exist the following positive results which guarantee that the Shapley value lies in the core of a particular type of coalitional games, the so-called *convex* games.

Definition 5.12 (Convex game). A coalitional game (N, v) is *convex* if it has a supermodular value function:

$$v(C_1) + v(C_2) \leq v(C_1 \cup C_2) + v(C_1 \cap C_2), \forall C_1, C_2 \subseteq N$$

Theorem 5.13 ([Shapley, 1971]). *If (N, v) is a convex game, then the Shapley value is the barycentre of its core.*

5.1.4 The Nucleolus

Another widely used imputation is the *nucleolus*. It employs a different approach than the Shapley value, trying to minimise agent dissatisfaction, which is defined next.

Definition 5.14 (Excess). Given a coalitional game (N, v) and a payment allocation $\mathbf{x} \in \mathbb{R}^N$, the dissatisfaction of coalition C is measured by the *excess* defined as: $e(\mathbf{x}, C) = v(C) - \sum_{i \in C} x_i$.

Any payoff vector \mathbf{x} generates an excess vector, $\mathbf{e}(\mathbf{x}) = (e(\mathbf{x}, C_1), \dots, e(\mathbf{x}, C_{2^N})) \in \mathbb{R}^{2^N}$, where C_1, \dots, C_{2^N} is the list of subsets of N ordered in non-increasing order by their excess under \mathbf{x} . Then, two deficit vectors can be compared lexicographically. Given two payoff vectors \mathbf{x}, \mathbf{y} , we have $\mathbf{e}(\mathbf{x}) \leq_{\text{lex}} \mathbf{e}(\mathbf{y})$ if there exists $k \in \mathbb{R}$ such that for all $i < k$, $e_i(\mathbf{x}) = e_i(\mathbf{y})$, and $e_k(\mathbf{x}) \leq e_k(\mathbf{y})$.

Definition 5.15 (Nucleolus). Given a coalitional game (N, v) , its *nucleolus* is given by the lexicographically minimal imputation.

The nucleolus is in the core of any coalitional game with non-empty core, as the core is the set of imputations with negative excess [Solymosi, 2014]. However, computing the nucleolus for a game with n players requires solving 2^n linear programs [Sankaran, 1991], which is prohibitive for all but the smallest coalitional games. In order to maintain computational tractability and develop a scalable system, one can use an approximation to the nucleolus which lies in the core and presents better scaling properties, namely, the *least-core* [Maschler et al., 1979]. This imputation only minimises the *worst-case excess* for all coalitions, instead of finding the imputation that lexicographically minimises the vector of excesses. As a result, the least-core is much less computationally expensive than computing the nucleolus. Moreover, if the core is non-empty, the least-core belongs to it [Solymosi, 2014]. Formally, following the exposition in [Baeyens et al., 2013], we can define:

Definition 5.16 (Least-core). Given a coalitional game (N, v) , its *least-core* is the solution, \mathbf{x} , to the following linear program with $n + 1$ variables and $2^n + 1$ constraints:

$$\min_{\mathbf{x} \in \mathbb{R}^n, e \in \mathbb{R}} e, \text{ s.t. } \begin{cases} v(C) - \sum_{i \in C} x_i - e \leq 0, \forall C \subset N \\ v(N) - \sum_{i \in N} x_i = 0 \end{cases} \quad (5.1)$$

We now have the theoretical foundations to study our multi-EV aggregator scenario from a coalitional perspective.

5.2 Defining the Multi-EV Aggregator Coalitional Game

Focusing on our considered scenario, let $N = \{1, \dots, n\}$ be a set of EV aggregators and $C \subseteq N$ a coalition. Following the notation from Section 4.1, let $\hat{R}_t^{\min, i}$ and $\hat{R}_t^{\max, i}$ be aggregator i 's forecasted energy requirements for hour t , and \hat{N}_t^i the number of available EVs from aggregator i at hour t . The combined requirements of all the aggregators in the coalition C are then given by:

$$\hat{R}_t^{\min} = \sum_{i \in C} \hat{R}_t^{\min,i} \quad (5.2) \quad \hat{R}_t^{\max} = \sum_{i \in C} \hat{R}_t^{\max,i} \quad (5.3) \quad \hat{N}_t = \sum_{i \in C} \hat{N}_t^i \quad (5.4)$$

Once the requirements of all the aggregators in the coalition C have been aggregated, we can apply the bidding algorithm given by Eqs. 3.2a, 3.2b, 3.2c, 3.2d, 3.2e in Section 3.2 with constraints given by Eqs. 5.2, 5.3 and 5.4. This will provide a global day-ahead energy volume E_t^{global} for each hour t , which can be then distributed among the aggregators in C , as detailed in Section 4.2. This will be the basis of the proposed coalitional game.

Moreover, let's consider a realisation of the day-ahead market with hourly prices $\mathbf{p} = (p_0, \dots, p_{23})$. Then, the aggregators purchasing an energy schedule given by $\mathbf{E} = (E_0, \dots, E_{23})$ will incur a total electricity cost given by:

$$\text{cost}(\mathbf{p}, \mathbf{E}) = \sum_{t=0}^{23} p_t \cdot E_t \quad (5.5)$$

This provides a natural way to define the value function of our coalitional game. In more detail, for a coalition C , $v(C)$ must represent the electricity costs paid by the members of C when they perform coordinated bidding. However, the price impact present in our market model introduces an extra layer of complexity, as any market participant affects the resulting prices with their bids. More specifically, the cost paid by a coalition C depends not only on the members of the coalition itself, but on all the other aggregators as well. This situation is treated in detail in the next section.

5.3 Value Function with Externalities

The first thing to note is that our setting deviates from traditional characteristic function games. This is due to the presence of *externalities* [Chalkiadakis et al., 2012]. Specifically, in classical game theory, the value of a coalition C , $v(C)$, only depends on the coalition itself. However, in our market structure with price impact (see Section 2.3.3), a given coalition C is also affected by the aggregators not in the coalition. In more detail, any market participant will affect the resulting market prices, hence affecting every other participant's costs. Formally, the resulting prices depend on the whole coalition structure, $\mathbf{p} = \mathbf{p}(CS)$, and thus so does the value function of our game: $v(C, CS)$. Games with such value functions are called *partition function games* [Thrall and Lucas, 1963].

A coalitional game with externalities can be studied in partition function form. However, the resulting game has poor theoretical properties and does not yield useful results, as we show in Section 5.3.1. Another usual procedure when dealing with a coalitional game with externalities is to introduce a *conjecture* on the behaviour of the outsider agents [Aumann, 1961]. In more detail, when considering a coalition C , the behaviour of the outsider agents, $N \setminus C$, is assumed to be deterministic, and to follow the chosen conjecture, hence recovering the classical theory where the value of the coalition only depends on the coalition itself. The earliest proposed conjecture is the so-called α -*conjecture* [Aumann, 1961], which assumes that the outsider players act as to minimise the payoff of the deviated coalition. However, this conjecture is not appropriate in our setting, as an aggregator trying to minimise a coalition's payoff through price impact would automatically harm itself as well. More recent conjectures proposed in the literature include the

γ -conjecture and the *outsider coalition conjecture*. Both are reasonable in our setting and are further explored in the next two subsections.

5.3.1 The outsider coalition (oc) conjecture

Introduced in [Funaki and Yamato, 1999], it assumes that, when a coalition C deviates from the grand coalition, all the outsiders join together and form a counter coalition $N \setminus C$. Hence, the resulting coalition structure is $C^{oc} = \{C, N \setminus C\}$. Formally, the resulting prices depend on the coalition structure, and thus we can write $\mathbf{p} = \mathbf{p}(C^{oc})$. Similarly, the amount of energy purchased by the members of coalition C depends on C itself and on the coalition structure C^{oc} , $\mathbf{E} = \mathbf{E}_C(C^{oc})$. Therefore, the value function of this game can be defined as:

$$v_{oc}(C) := -\text{cost}(\mathbf{p}(C^{oc}), \mathbf{E}_C(C^{oc}))$$

Even though this conjecture seems reasonable in our scenario, the resulting coalitional game (N, v_{oc}) has poor stability properties given that it is not superadditive, as proven below.

Theorem 5.17. *The coalitional game (N, v_{oc}) is not superadditive.*

Proof. This is found empirically by using the case study detailed in Section 5.6. However, for clarity, we present a counter-example which employs a simplified market structure with three hours and synthetic prices. In more detail, consider hourly prices given by: $\hat{P}_1(E_1) = 10 + E_1$, $\hat{P}_2(E_2) = 5 + E_2/2$, $\hat{P}_3(E_3) = 10 + E_3$. Consider nine identical EV aggregators, $N = \{1, \dots, 9\}$, with the following individual energy requirements: $R^{\max} = (1, 0, 0)$, $R^{\min} = (0, 0, 1)$ and a maximum charging speed $P_{\max} = 1$. Considering the following pair of coalitions, $C_1 = \{1, 2\}$ and $C_2 = \{3, 4\}$, it holds: $v_{oc}(C_1 \cup C_2) < v_{oc}(C_1) + v_{oc}(C_2)$. Hence the coalitional game (N, v_{oc}) is not superadditive. \square

Thus, the grand coalition does not necessarily form, in which case full coordination is not achieved. This is a common issue in games considering a shared pool resource, as described in [Funaki and Yamato, 1999]. Furthermore, note that this counter-example also applies to the partition function game described in Section 5.3, which is not superadditive either.

5.3.2 The γ -conjecture

Another common choice is the so-called γ -conjecture [Chander and Tulkens, 1997], in which the outsider agents select their individual best strategies. Hence, the resulting coalition structure is $C^\gamma = \{C\} \cup \{\{i\} \mid i \in N, i \notin C\}$. Formally, we can write $\mathbf{p} = \mathbf{p}(C^\gamma)$ and $\mathbf{E} = \mathbf{E}_C(C^\gamma)$. Then, the value function can be defined as:

$$v_\gamma(C) := -\text{cost}(\mathbf{p}(C^\gamma), \mathbf{E}_C(C^\gamma)) \tag{5.6}$$

This conjecture is also reasonable in our setting and will be adopted throughout the rest of the chapter. For convenience, we will drop the subscript γ and write v instead of v_γ henceforth.

5.4 Properties of the Coalitional Game

The coalitional game using the γ -conjecture proposed in the previous section, (N, v) , has several desirable properties. Specifically, we will show that it is superadditive and balanced, hence it has a non-empty core. Thus, all the EV aggregators are incentivised to cooperate together (the grand coalition forms), and a payment mechanism can be implemented which results in a stable game where no sub-coalition having an incentive to deviate. We will formally prove these claims next. For convenience, and extending the notation presented in the previous subsection, let $\text{cost}(\mathbf{p}(C_2^\gamma), \mathbf{E}_{C_1}(C_2^\gamma))$ be the total electricity cost paid by the members of $C_1 \subseteq N$ when coalition $C_2 \subseteq N$ performs coordinated bidding, and all other participants perform individual bidding.

We are now ready to show that the game is superadditive.

Lemma 5.18. *For all coalitions $C_1, C_2 \subseteq N$ such that $C_1 \subset C_2$, it holds that:*

$$\text{cost}(\mathbf{p}((C_1 \cup C_2)^\gamma), \mathbf{E}_{C_1}((C_1 \cup C_2)^\gamma)) \leq \text{cost}(\mathbf{p}(C_1^\gamma), \mathbf{E}_{C_1}(C_1^\gamma))$$

Proof. This lemma trivially follows from the fact that coordinated bidding with more participants can only decrease the total costs. Hence the price paid by members of coalition C_1 when coordination happens inside $C_1 \cup C_2$ can only be lower, or equal, than when coordination happens only inside C_1 . The equality case happens only when the members of C_1 and C_2 have non-overlapping energy requirements, or when the price impact of their combined bids is not high enough. \square

Theorem 5.19. *The proposed coalitional game (N, v) is superadditive.*

Proof. Consider any two disjoint coalitions, $C_1, C_2 \subseteq N$. Then,

$$\begin{aligned} v(C_1 \cup C_2) &\geq v(C_1) + v(C_2) \Leftrightarrow \text{cost}(\mathbf{p}((C_1 \cup C_2)^\gamma), \mathbf{E}_{C_1 \cup C_2}((C_1 \cup C_2)^\gamma)) \\ &\leq \text{cost}(\mathbf{p}(C_1^\gamma), \mathbf{E}_{C_1}(C_1^\gamma)) + \text{cost}(\mathbf{p}(C_2^\gamma), \mathbf{E}_{C_2}(C_2^\gamma)) \end{aligned}$$

Given the following identity:

$$\begin{aligned} &\text{cost}(\mathbf{p}((C_1 \cup C_2)^\gamma), \mathbf{E}_{C_1 \cup C_2}((C_1 \cup C_2)^\gamma)) \\ &= \text{cost}(\mathbf{p}((C_1 \cup C_2)^\gamma), \mathbf{E}_{C_1}((C_1 \cup C_2)^\gamma)) \\ &\quad + \text{cost}(\mathbf{p}((C_1 \cup C_2)^\gamma), \mathbf{E}_{C_2}((C_1 \cup C_2)^\gamma)) \end{aligned}$$

the expression above reads:

$$\begin{aligned} &\text{cost}(\mathbf{p}((C_1 \cup C_2)^\gamma), \mathbf{E}_{C_1}((C_1 \cup C_2)^\gamma)) \\ &\quad + \text{cost}(\mathbf{p}((C_1 \cup C_2)^\gamma), \mathbf{E}_{C_2}((C_1 \cup C_2)^\gamma)) \\ &\leq \text{cost}(\mathbf{p}(C_1^\gamma), \mathbf{E}_{C_1}(C_1^\gamma)) + \text{cost}(\mathbf{p}(C_2^\gamma), \mathbf{E}_{C_2}(C_2^\gamma)) \end{aligned}$$

which is always true, applying Lemma 5.18. \square

This result shows that overall costs are minimised when the grand coalition forms. The main issue is now how to distribute the value of the grand coalition, *i.e.* the resulting costs, among

its members. To this end, we will next prove that stable coalitions are possible given that the game has a non-empty core.

Theorem 5.20. *The proposed coalitional game (N, v) is balanced. As a result, it has a non-empty core.*

Proof. Let $\alpha : 2^N \rightarrow \mathbb{R}$ be an arbitrary balanced function. Balancedness of the coalitional game follows from Lemma 5.18:

$$\begin{aligned} \sum_{C \in 2^N} \alpha(C) v(C) &= - \sum_{C \in 2^N} \alpha(C) \text{cost}(\mathbf{p}(C^\gamma), \mathbf{E}_C(C^\gamma)) \\ &\leq - \sum_{C \in 2^N} \alpha(C) \text{cost}(\mathbf{p}(N^\gamma), \mathbf{E}_C(N^\gamma)) \\ &= - \sum_{C \in 2^N} \sum_{i \in N} \alpha(C) \mathbf{1}_{\{i \in C\}} \text{cost}(\mathbf{p}(N^\gamma), \mathbf{E}_{\{i\}}(N^\gamma)) \\ &= - \sum_{i \in N} \text{cost}(\mathbf{p}(N^\gamma), \mathbf{E}_{\{i\}}(N^\gamma)) = -\text{cost}(\mathbf{p}(N^\gamma), \mathbf{E}_N(N^\gamma)) = v(N) \end{aligned}$$

Thus the coalitional game is balanced and, by the Bondareva-Shapley Theorem (Th. 5.10), it has a non-empty core. \square

The non-emptiness of the core of our proposed coalitional game guarantees the existence of at least one payoff allocation which stabilises the grand coalition.

5.5 Payoff Allocations

After proving that the proposed game has a non-empty core, we now seek a payoff allocation lying in it. As described in Section 5.1.3, a common choice is the Shapley value. However, it turns out that our proposed coalitional game is not convex hence the Shapley value is not in the core, as proven by the counter-example in the following result.

Theorem 5.21. *The coalitional game (N, v) is not convex.*

Proof. This is found empirically by using the case study detailed in Section 5.6. However, for clarity, we present a counter-example which employs a simplified market structure with three hours and synthetic prices. In more detail, consider hourly prices given by: $\hat{\mathcal{P}}_1(E_1) = 10 + 20E_1$, $\hat{\mathcal{P}}_2(E_2) = 0.1 + 20E_2$, $\hat{\mathcal{P}}_3(E_3) = 10 + 20E_3$. Consider three identical EV aggregators, $N = \{1, 2, 3\}$, with the following individual energy requirements: $R^{\max} = (1, 0, 0)$, $R^{\min} = (0, 0, 1)$ and a maximum charging speed $P_{\max} = 1$. Considering the following pair of coalitions, $C_1 = \{1, 2\}$ and $C_2 = \{2, 3\}$, it holds: $v(C_1 \cup C_2) + v(C_1 \cap C_2) < v(C_1) + v(C_2)$. Hence the coalitional game is not convex. \square

As a result, the Shapley value is not guaranteed to be in the core of our proposed coalitional game (N, v) . Therefore, we seek a different payoff allocation which is guaranteed to be in the core. As described in Section 5.1.4, given that our game has a non-empty core, the nucleolus and the least-core lie in it. Given the computational cost of the former, the latter provides a viable and scalable option to stabilise the grand coalition.

Following this theoretical analysis of the considered game, we now turn our attention to an empirical case study where we will compare the performance of the two proposed payoff allocations, namely the Shapley value and the least-core.

5.6 Case Study

This section describes a case study based on real data with the aim of testing the performance of the coalitional game proposed in this chapter. The considered scenario is the same as in the case studies from Chapters 3 and 4: a night-time scenario where EVs arrive in the evening and depart in the morning. In more detail, the aggregator, market and driver characteristics are described in Sections 3.3.1, 3.3.2 and 3.3.3, respectively. Throughout this section, naive forecasts where the data from the day before is employed as a forecast for the day after are employed (see Section 3.1.2).

The main goal of this analysis is to compare the two different payment allocations proposed in the previous section: the Shapley value and the least-core. Recall that the Shapley value is considered as a fair payment allocation (see Section 5.1.3), but is not in the core. Furthermore, the least-core minimises agent dissatisfaction (see Section 5.1.4), and it lies in the core. Finally, in order to maintain computational tractability, an error-bounded approximation of the Shapley value is employed, as described in Section 5.6.1. All the data from this empirical case study is available in [Perez-Diaz et al., 2018e].

5.6.1 Approximating the Shapley Value

The Shapley value (see Section 5.1.3) is known to be computationally expensive. Specifically, its computational complexity is $2^n \cdot \mathcal{O}(v)$, where $\mathcal{O}(v)$ is the complexity of the value function. In order to improve its computational tractability, approximations are commonly used in the literature. In this chapter, we apply the state-of-the-art error-bound approximation proposed by [Maleki, 2015] for superadditive games.

In more detail, let $SV_i(v)$ denote the Shapley value for a given agent $i \in N$, and $\widetilde{SV}_i(v)$ the approximated Shapley value for the same agent. Instead of considering the $v(C \cup \{i\}) - v(C)$ contributions from *all* subsets of $N \setminus \{i\}$, the approximation randomly samples a number of them, $m_{\epsilon, \delta}$, which depends on the desired level of precision:

$$\mathbb{P}\left(\left|SV_i(v) - \widetilde{SV}_i(v)\right| \geq \epsilon\right) \leq \delta$$

The specific formula for $m_{\epsilon, \delta}$ is detailed in [Maleki, 2015, Th. 4.3]. Throughout our simulations, we employ $\epsilon, \delta = 5\%$, in order to balance precision and computational tractability.

5.6.2 Experimental Results

We present the results from three different simulations, all of them considering every weekday of November 2016. Firstly, we study the global cost reductions provided by coordination, when the

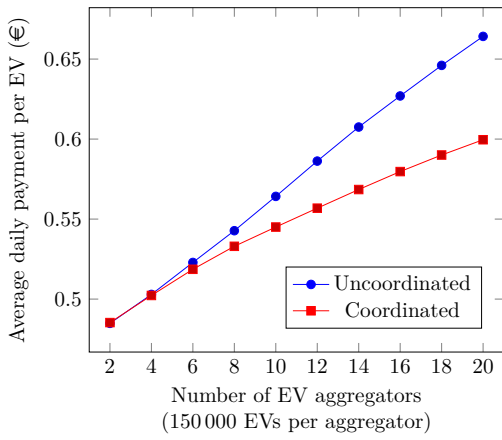


FIGURE 5.1: Comparison between uncoordinated and fully coordinated bidding when varying the number of aggregators.

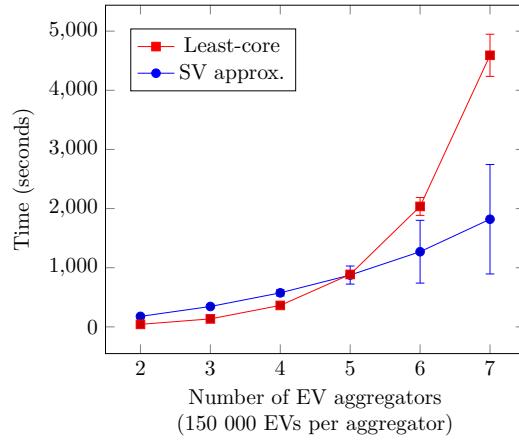


FIGURE 5.2: Runtimes for the approx. SV and least-core payment mechanisms. Error bars represent standard deviations.

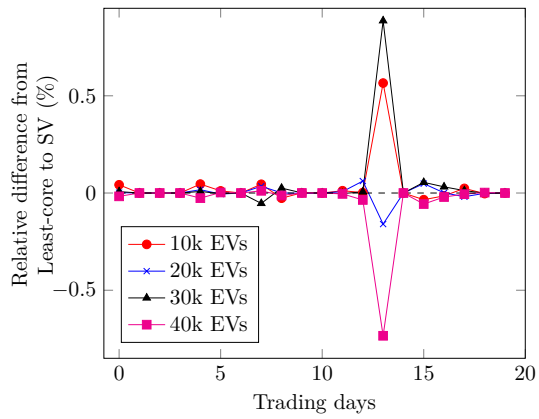



FIGURE 5.3: (LHS) Daily least-core payment allocations for four aggregators with different sizes. (RHS) Daily difference between the least-core and approximate Shapley value payments


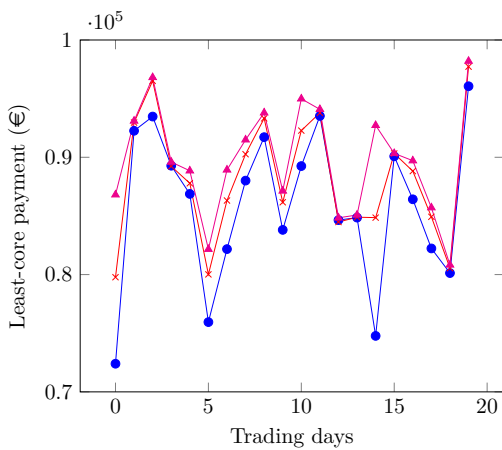


FIGURE 5.4: (LHS) Daily least-core payment allocations for three aggregators of the same size, but with different flexibilities. (RHS) Daily difference between the least-core and approximate Shapley value payments.

grand coalition performs joint bidding (see Section 5.2), in comparison to individual, uncoordinated, bidding. These results are the same as in Section 4.5, as the grand coalition performing joint bidding is analogous to all market participants performing joint bidding under a single coordinator. Results indicate that significant cost reductions can be achieved by coordination, as shown in Figure 5.1. Now that the efficacy of coordinated inter-aggregator bidding has been shown, we turn our attention to the two proposed payment mechanisms, the least-core and the approximated Shapley value.

In more detail, we present the results from two different simulations. Firstly, we consider a scenario where four different EV aggregators of different sizes, namely 50, 100, 150 and 200k EVs participate in the market. Apart from size, they share the same characteristics, as described in Section 3.3.3. The purpose of this simulation is to study how the proposed payment mechanisms capture the greater electricity costs of larger market participants. Results are shown in Figure 5.3. More specifically, the left-hand-side presents the daily prices assigned to each aggregator by the least-core payment mechanism. We can see that it correctly assigns larger payments to larger aggregators. Moreover, the right-hand-side presents the daily difference between least-core and approximated Shapley value payments for each aggregator. We can see that the differences are of very small magnitude, concluding that both payment mechanisms are both very close in all cases.

Lastly, we consider a similar scenario, where three different aggregators participate in the market. In this case, they all have the same size, 150k EVs each, but have different charging flexibilities. In more detail, instead of employing the real arrival and departure data provided in Table 3.4, we employ synthetic data in order to capture the effects of charging flexibility on allocated prices. Specifically, we consider an evening charging scenario where all EVs depart at midnight. The first aggregator receives EVs at 14h, the second at 16h and the third at 18h. Thus, the first aggregator has the highest flexibility and the third is the most constrained. This choice of arrival and departure times is motivated by the typical hourly prices found in the OMIE market, where electricity is cheaper around 15-17h, and more expensive afterwards. Results are presented in Figure 5.4. We can see that first aggregator, the most flexible, is able to obtain energy at cheaper hours, hence being allocated cheaper payments. Conversely, the third aggregator which is forced to operate at expensive hours, incurs larger payments. The second aggregator lies in between. Similarly to the previous scenario, we can see that least-core and approximated Shapley value payments are very close together for all aggregators and every trading days.

As just described, the least-core and Shapley value payments are very close in all considered scenarios. This suggests that both payment mechanisms present good stability and fairness properties. Moreover, given their computational complexity (see Figure 5.2), the approximated Shapley value presents much better scaling properties, and is more suitable for the application of this framework to large scenarios.

We note that the presented results do not depend on the particular trading days shown here. Longer simulations have been run utilising different trading months, and the results obtained follow the same trend.

5.7 Summary

In this chapter, we present the first study of inter-EV aggregator cooperation from the perspective of cooperative game theory. More specifically, we define a cooperative game in which self-interested EV aggregators can cooperate with each other in order to effectively bid in the day-ahead market. Given the presence of externalities, we propose employing a γ -conjecture in order to obtain a game in characteristic function form. Next, we show that the resulting game is superadditive, hence the larger the formed coalition, the larger electricity cost reductions for the EV aggregators. Moreover, we prove that our game is balanced and thus has a non-empty core. The payment mechanism given by the least-core is proposed in order to distribute payments among the grand coalition of aggregators. As the least-core belongs to the core of our game, this payment mechanism stabilises the grand coalition. Lastly, we present numerical simulations which employ real market and driver behaviour data, in order to show the efficacy of coordinated bidding. The least-core payments are compared with a more computationally tractable approximation to the Shapley value. The good agreement between the two suggests that they both present good stability and fairness properties.

In conclusion, Chapters 4 and 5 provide a detailed study of the coordination opportunities available in multi-EV aggregator participation in day-ahead markets, addressing Research Challenge 2. However, the coordination mechanisms proposed so far are centralised, which introduces some problems such as private information sharing, lack of transparency and the need for trust in the coordinator. In order to address these issues, we will next turn our attention to Research Challenge 3 and design a decentralised coordination approach.

Chapter 6

Decentralised Inter-Aggregator Coordination

In the last two chapters, we have addressed Research Challenge 2 from a centralised perspective, proposing EV aggregator coordination mechanisms based on central coordinators. However, these centralised approaches present certain drawbacks, such as the need for private information sharing, a lack of transparency and the need for trust in the coordinator. These issues are captured in Research Challenge 3, and can be addressed by employing a decentralised coordination mechanism, which is the object of this chapter.

In more detail, the idea is to reformulate the centralised coordination mechanisms proposed in earlier chapters as a distributed optimisation algorithm. More specifically, each aggregator will take an active computational role, rather than just reporting energy requirements, which will remove the need of reporting private information. Moreover, the use of blockchain technology, for which decentralised algorithms are well suited, can provide enhanced transparency and anti-tampering guarantees, and remove the need for trust on a centralised coordinator.

The rest of this chapter is structured as follows. We start by reformulating the centralised coordination mechanism in Section 6.1. Then, we turn our attention to the application of blockchain technologies in Section 6.2. Finally, we present an empirical evaluation of the proposed decentralised coordination algorithm in Section 6.3, using real market and driver data.

6.1 Decentralised Optimisation Algorithm

Our goal is to reformulate the optimisation problems given by Eqs. 3.2a, 3.2b, 3.2c, 3.2d, 3.2e, 4.4a, 4.4b, 4.4c, 4.4d, 4.4e as an iterative decentralised optimisation algorithm. In comparison with the centralised approaches proposed in Chapters 4 and 5, where each aggregator reports their energy requirements to the centre who then optimises joint bidding, in the decentralised case the aggregators do not report their private requirements but solve local optimisation problems instead. By iteratively solving these local problems and updating some shared variables via a global *consensus* step, the globally optimal solution can be obtained.

In more detail, we propose using the Alternating Direction Method of Multipliers (ADMM) [Boyd et al., 2010]. This type of algorithm is appropriate in our setting for several reasons: (i) given that our optimisation problem is convex, it is guaranteed to converge to the global optimum [Boyd et al., 2010]; (ii) it enables coordination without the aggregators revealing their energy requirements, *i.e.* $\mathbf{R}^{\min,i}$ and $\mathbf{R}^{\max,i}$; (iii) it is particularly well suited for blockchain implementation, providing transparency and anti-tampering guarantees [Musing et al., 2017].

Following the notation introduced in earlier chapters, let $\mathbf{E}^i = (E_0^i, \dots, E_{23}^i)$ be the energy schedule for aggregator i . Moreover, let $\mathbf{E} = (\mathbf{E}^1, \dots, \mathbf{E}^n)$ be the joint vector encapsulating each individual energy schedule. We can now rewrite Eq. 3.2a as:

$$\min_{\mathbf{E}} \sum_{t=0}^{23} \left[\hat{\mathcal{P}}_t^{\text{convex}} \left(\sum_{i=1}^n E_t^i \right) \cdot \sum_{i=1}^n E_t^i \right] = \min_{\mathbf{E}} \sum_{i=1}^n \left[\sum_{t=0}^{23} \left(E_t^i \cdot \hat{\mathcal{P}}_t^{\text{convex}} \left(\sum_{j=1}^n E_t^j \right) \right) \right] \quad (6.1)$$

This way the objective function is expressed as a sum of n terms, as required by the ADMM formulation. Note that, given that the price impact of each aggregator affects everybody else, we cannot separate Eq. 6.1 in the variable i , *i.e.* the equation is coupled and the sum's terms cannot be independently distributed among the aggregators. This type of problem is suited to be formulated as a *global variable consensus problem* [Boyd et al., 2010], which works as follows. Consider a minimisation problem in the following form:

$$\min_{\mathbf{x}} \sum_{i=1}^n f_i(\mathbf{x})$$

where the goal is that each term in the sum can be handled independently. In the cases where the variable \mathbf{x} is not separable in i , we can introduce *local* variables \mathbf{x}^i and a *global* variable \mathbf{z} and rewrite the problem as:

$$\begin{aligned} & \min_{\{\mathbf{x}^i\}} \sum_{i=1}^n f_i(\mathbf{x}^i) \\ & \text{subject to: } \mathbf{x}^i - \mathbf{z} = 0, \forall i = 1, \dots, n \end{aligned}$$

As mentioned above, the problem constraints require all local variables to agree with each other and with the global variable. This way, global consensus on the solution is achieved. Also, note that any individual constraints can be embedded into each f_i .

In a similar vein and focusing on our scenario, let \mathbf{E} and $\mathbf{E}^{(i)}$ denote the global and local variables respectively, each of which comprises a vector with dimension $24n$ *i.e.* $\mathbf{E}^{(i)} = (\mathbf{E}^{(i),1}, \dots, \mathbf{E}^{(i),n})$ and $\mathbf{E}^{(i),j} = (E_0^{(i),j}, \dots, E_{23}^{(i),j})$. Following Eq. 6.1, in our case the functions f_i are given by:

$$f_i(\mathbf{E}^{(i)}) = \begin{cases} \sum_{t=0}^{23} \left[E_t^{(i),i} \cdot \hat{\mathcal{P}}_t^{\text{convex}} \left(\sum_{j=1}^n E_t^{(i),j} \right) \right], & \text{if constraints (3.2b),} \\ & \text{(3.2c), (3.2d), (3.2e) are} \\ & \text{met by } \mathbf{E}^{(i),i} \\ \infty & \text{, otherwise} \end{cases}$$

The resulting ADMM algorithm is then given by the following iterative equations:

$$\mathbf{E}_{[k+1]}^{(i)} = \arg \min_{\mathbf{E}'} \left(f_i(\mathbf{E}') + \boldsymbol{\xi}_{[k]}^{(i)T} \cdot (\mathbf{E}' - \mathbf{E}_{[k]}) + \frac{\rho}{2} \|\mathbf{E}' - \mathbf{E}_{[k]}\|_2^2 \right) \quad (6.2a)$$

$$\mathbf{E}_{[k+1]} = \frac{1}{n} \sum_{i=1}^n \left(\mathbf{E}_{[k+1]}^{(i)} + \frac{1}{\rho} \boldsymbol{\xi}_{[k]}^{(i)} \right) \quad (6.2b)$$

$$\boldsymbol{\xi}_{[k+1]}^{(i)} = \boldsymbol{\xi}_{[k]}^{(i)} + \rho \left(\mathbf{E}_{[k+1]}^{(i)} - \mathbf{E}_{[k+1]} \right) \quad (6.2c)$$

where the subscript $[k]$ denotes iteration number, and $\boldsymbol{\xi}$ and ρ are the dual variable and the augmented Lagrangian parameter, respectively [Boyd et al., 2010]. Intuitively, ρ controls the trade-off between each aggregator solving its own local problem and achieving global consensus. In more detail, if ρ is set too high, the algorithm forces consensus *too much*, resulting in very slow convergence. Conversely, if ρ is set too small, each aggregator solves its local problem and consensus is not reached. Examples of this effect are presented in Section 6.3.

Given this, the iterative algorithm works as follows: first, each EV aggregator solves their local problem, Eq. 6.2a, and update their local copy of the energy schedule, $\mathbf{E}^{(i)}$. Then, an aggregation step, Eq. 6.2b, collects all the local solutions proposed by each aggregator and updates the global energy schedule, \mathbf{E} , reporting this vector back to all the aggregators. Lastly, each aggregator updates their local copy of the dual variable, $\boldsymbol{\xi}^{(i)}$, as per Eq. 6.2c and proceeds to the new iteration.

This iterative process is stopped when the primal and dual residuals reach some user-specified tolerances, ϵ_{pri} and ϵ_{dual} [Boyd et al., 2010, Munsing et al., 2017]. Specifically, the primal residual is denoted by $\mathbf{r}_{[k]} = (\mathbf{r}_{[k]}^1, \dots, \mathbf{r}_{[k]}^n)$, where $\mathbf{r}_{[k]}^i = \mathbf{E}_{[k]}^{(i)} - \mathbf{E}_{[k]}$. Similarly, the dual residual is given by $\mathbf{s}_{[k]} = \mathbf{E}_{[k]} - \mathbf{E}_{[k-1]}$. The stopping criterion then takes the following form:

$$\|\mathbf{r}_{[k]}\|_2^2 \leq \epsilon_{\text{pri}} \quad (6.3a)$$

$$\|\mathbf{s}_{[k]}\|_2^2 \leq \epsilon_{\text{dual}} \quad (6.3b)$$

and the algorithm stops when both conditions have been met.

We are now ready to discuss how the proposed decentralised coordination algorithm can be implemented in a blockchain, hence allowing execution in a completely trustless environment.

6.2 Proposed Blockchain Implementation

So far, we have motivated and formulated a decentralised algorithm addressing the coordination of multi-EV aggregator bidding in day-ahead markets. Specifically, the proposed algorithm addresses the first part of Research Challenge 3, namely the need for disclosing private information. However, the second aspect of this challenge, *i.e.* trust-less and transparent operation, has not been addressed so far. In order to do so, we propose employing a blockchain to implement the consensus step of the proposed decentralised algorithm, similarly to [Munsing et al., 2017]. In more detail, this consensus step requires communication between each aggregator and the coordinator, hence is susceptible to tampering. Also, if the coordinator remains a black-box, *i.e.* an entity that receives messages from the aggregators and responds with some other message

without its internal workings being transparent, the aggregators cannot be sure that they are billed and allocated a fair electricity schedule.

Recall that a blockchain is a decentralised ledger and computation environment, protected by cryptographic techniques, which allows the participants to agree on the state of the system at all times. Moreover, apart from monetary transactions, blockchains can be seen as a general computing system, by using *smart contracts*. In more detail, a smart contract is simply a piece of code hosted publicly and immutably in the blockchain, which can receive messages from other blockchain users and send its own following its internal logic. Thus, smart contracts are a suitable architecture for implementing the consensus step of the proposed decentralised algorithm. Moreover, if desired, the proposed coordination mechanism can be fully implemented in the blockchain by using cryptocurrencies. This way, no centralised entity is needed, and all the bidding and coordination process can run transparently on the blockchain. Finally, it is worth noting that, when running on a blockchain, the operation of the coalition of EV aggregators is transparent and can easily be externally audited, perhaps by the government or the electricity market operator.

Formally, the proposed blockchain implementation is described in Algorithm 1, closely following [Munsing et al., 2017]. In more detail, the local optimisation problems \mathbf{P}_i are executed locally by each aggregator, sending the results to the smart contract \mathbf{S} , who performs the consensus step. This process is iterated until the algorithm has converged.

```

Data: initialise  $\vec{E}_{[0]}$ ,  $\vec{E}_{[0]}^{(i)}$ ,  $\vec{\xi}_{[0]}^{(i)}$ ,  $\epsilon_{pri}$ ,  $\epsilon_{dual}$ 
while  $\|\vec{r}_{[k]}\|_2^2 \leq \epsilon_{pri}$  and  $\|\vec{s}_{[k]}\|_2^2 \leq \epsilon_{dual}$  do
  begin  $\mathbf{P}_i$ : private optimisation problem, compute locally
    Update  $\vec{\xi}_{[k]}^{(i)}$ , Eq. 6.2c
    Update  $\vec{E}_{[k]}^{(i)}$ , Eq. 6.2a, and send to smart contract  $\mathbf{S}$ 
  end
  begin  $\mathbf{S}$ : ADMM consensus, smart contract
    Update  $\vec{E}_{[k]}$ , Eq. 6.2b
    if  $\|\vec{r}_{[k]}\|_2^2 \leq \epsilon_{pri}$  and  $\|\vec{s}_{[k]}\|_2^2 \leq \epsilon_{dual}$  then
      ADMM algorithm finished
      Report final allocations
    end
  end
end

```

Algorithm 1: Decentralised optimal multi-EV aggregator day-ahead bidding algorithm implemented in a blockchain.

We are now ready to empirically evaluate the performance of the proposed algorithm.

6.3 Empirical Evaluation

The experiment setup described in this section closely follows the case studies presented in Chapters 3, 4 and 5. Recall that this consists of a night-time residential scenario in which EVs arrive in the evening and need to be charged by the next morning. In more detail, the aggregator, market and driver characteristics are described in Sections 3.3.1, 3.3.2 and 3.3.3, respectively.

The goal of this empirical evaluation is to study the convergence of the proposed decentralised algorithm and hence its practical applicability. As discussed in Section 6.1, key determinant of convergence is the augmented Lagrangian parameter ρ (see Eqs. 6.2a, 6.2b, 6.2c). Intuitively, it controls the *weight* that the similarity of local and global solutions has in the local minimisation algorithms. If it is set too large or too small, the algorithm will not converge. For every problem there exists a range of values providing convergence but, as already mentioned, it can be very slow in some cases. Also, the number of participating aggregators affects the convergence of the algorithm: the higher the number of participants, the more fragmented the optimisation problem is, so more iterations may be required. Thus, a suitable value for ρ needs to be found in order to make the algorithm converge fast, a key point for its practical applicability.




FIGURE 6.1: Convergence of the ADMM decentralised algorithm to the optimal centralised solution, for different values of ρ . (Top) Simulations with two aggregators, each with 150 000 EVs. (Bottom) Simulations with ten aggregators, each with 150 000 EVs. Market data from 1/11/2016.

Figure 6.1 shows the convergence for different values of ρ , and for two scenarios, with two and ten EV aggregators respectively. Each plot shows the convergence of the decentralised solution to the optimal solution. We can see similar convergence behaviour for the two scenarios, although the case with two EV aggregators is faster and more uniform. These results show evidence of good computational scaling with the number of EV aggregators, something key for tackling larger problem sizes. Moreover, for both scenarios, convergence starts slow for a value of $\rho = 10^{-3}$, becoming fastest for a value $\rho \sim 10^{-5}$, and diverging for larger values. This suggests that a value about $\rho = 10^{-5}$ presents the best convergence for these scenarios, although this may vary for larger problem sizes. Also, we would like to note that these results are consistent across different trading days. Lastly, note that there exist recent extensions of the ADMM algorithm which include an adaptive parameter ρ and can provide faster convergence and rule out the need for parameter tweaking [Xu et al., 2017]. We discuss this possibility as future work in Section 10.

6.4 Summary

In this chapter, we have presented the first decentralised coordination mechanism for multi-EV aggregator bidding in the day-ahead market. In more detail, this algorithm uses the Alternating Direction Method of Multipliers algorithm in order to decouple the centralised mechanisms proposed in Chapters 4 and 5. This new decentralised formulation removes the need for the

aggregators to communicate private requirement information to the coordinator, as each aggregator solves its own local private optimisation problem with their own requirements. Moreover, this algorithm is well-suited for implementation in a blockchain, providing a trust-less execution environment with greatly increased transparency and anti-tampering guarantees. Finally, in order to study the appropriateness of the proposed algorithm, we present an empirical evaluation using real market and driver data. Results show the convergence of the decentralised method to the optimal solution for two scenarios, with two and ten cooperating EV aggregators respectively. Convergence can be achieved in around 50 iterations in the first case and around 80 in the second case. Therefore, although problem complexity increases with the number of participants, these numbers suggest the applicability of the algorithm in large settings.

This chapter has contributed towards Research Challenges 2 and 3, but also introduces an important challenge of its own. More specifically, by decentralising computation among self-interested aggregators, new strategic manipulation opportunities arise, as described in Research Challenge 4, an issue which will be addressed in the next chapter.

Chapter 7

Strategic Manipulation of Decentralised Optimisation Algorithms

In the previous chapter we proposed a decentralised coordination mechanism for EV aggregators, with the aim of addressing Research Challenges 2 and 3. Indeed, by decentralising computation, the need for sharing private information and trusting a central coordinator are removed. However, as hinted in the previous chapter, a new risk for potential strategic manipulation arises. More specifically, in centralised approaches, such as the ones proposed in Chapters 4 and 5, self-interested aggregators can only manipulate the coordination algorithm by misreporting their energy requirements. In order to avoid this, payment mechanisms can be carefully designed to incentivise the aggregators to cooperate and report truthful requirements (see Sections 4.3 and 5.5). However, when considering decentralised mechanisms, the participating aggregators' roles are not limited to reporting preferences, but also play an active part in computation. As a consequence, the opportunities for strategic manipulation are greatly increased, as detailed in Research Challenge 4. Again, we would like to note that, although we will focus on our multi-aggregator scenario, these issues exist in any decentralised optimisation algorithm.

In order to address these issues, we will describe and analyse how the decentralised coordination mechanism from Chapter 6 can be manipulated, and propose a manipulation detection algorithm which detects deviating aggregators. The rest of this chapter is structured as follows. We start by describing several attack vectors which can be used by a deviating aggregator in order to increase its personal utility in Section 7.1. Then, in Section 7.2 we turn our attention to the issue of detecting strategic manipulation, and present frameworks for measuring and detecting deviating behaviour. Finally, we present a case study using real driver and market data in order to assess the effects of the proposed attacks in the coordination algorithm and the performance of the detection algorithm in Section 7.3.

7.1 Strategic Manipulation of the ADMM Algorithm

The decentralised ADMM-based algorithm proposed in Chapter 6, given by Eqs. 6.2a, 6.2b, 6.2c, has nice convergence properties and asymptotically reaches the global optimum for suitable values of ρ [Boyd et al., 2010]. However, this requires every participating agent to run the algorithm faithfully. In our case, where agents are assumed to be self-interested, an aggregator could deviate from their assigned local algorithm and/or misreport their local solutions with the aim of improving their allocation. More specifically, we assume that a potential attacker aims to reducing its energy costs (*i.e.* increase its utility) and we will show how this misbehaving aggregator can significantly affect the algorithm's outcome for its own benefit. Note that we do not look at all possible manipulation vectors, as this is not feasible, but instead focus on several intuitive and specific types of manipulation that are beneficial for the attacker in our setting.

Formally, considering a set of n aggregators and following the notation from Chapter 6, the electricity costs incurred by aggregator i when a global allocation $\mathbf{E} = (\mathbf{E}^1, \dots, \mathbf{E}^n)$ is reached are given by:

$$\text{cost}_i = \sum_{t=0}^{23} \left[E_t^i \cdot \hat{P}_t \left(\sum_{j=1}^n E_t^j \right) \right] \quad (7.1)$$

In order to reduce these costs, the attacker aims to minimise the price impact on their desired hours, which in turn can be achieved by moving other aggregators' overlapping allocations to different hours. To this end, we consider different attack vectors, namely *Shift*, *Proportional*, *Freeze*, *FreezeShift*, *FreezeProp* and *Adversarial* attacks, which are described next. These capture different ways that an attacker can try to increase its own utility and present very different outcomes and efficacy, as detailed in Section 7.3.1. Note that, throughout the rest of this chapter, we assume that an attacker performs a given attack vector with a given intensity in every round. The study of more sophisticated attacks is outlined as future work (Section 7.4). For quick reference, a summary of the considered attack vectors can be found in Table 7.1. Finally, note that an empirical evaluation of all the proposed attack vectors can be found in Section 7.3.1.

Attack name	Short description
Shift(All)	Shift the proposed allocation for one attacked aggregator (or all aggregators) to more expensive hours
Proportional(All)	Scale down the proposed allocation for the attacked aggregator (or all aggregators)
Freeze	Propose its individually optimal allocation for itself (without considering the competitor aggregators)
FreezeShift(All)	Freeze + Shift(All)
FreezeProp(All)	Freeze + Proportional(All)
Adversarial	Attempt to incriminate a benign aggregator as deviator by artificially favouring their allocations

TABLE 7.1: Summary of the proposed attack vectors.

7.1.1 Shift Attack

In this type of attack, the deviating aggregator i runs its local optimisation problem (*i.e.* Eq. 6.2b), but artificially modifies its local schedule allocation for another aggregator j , $\mathbf{E}^{(i),j}$, by




FIGURE 7.1: (a) Truthful allocation from aggregator i to aggregator j following Eq. 6.2b, $\mathbf{E}_{[k+1]}^{(i),j}$. (b) Attacked $\hat{\mathbf{E}}_{[k+1]}^{(i),j}$ employing a shift attack with $\mu = 2$ as given by Eq. 7.2.




FIGURE 7.2: Real example of the effects of *ShiftAll* on the resulting energy allocations. The scenario consists of three aggregators of the same size with the third aggregator being the attacker and employing $\mu = 1$. Results for both vanilla and attacked scenarios are presented.

shifting it outside of aggregator i 's preferred hours. In this work and without loss of generality, we will focus on a particular case where the deviating aggregator splits the energy schedule of the attacked aggregator j by its mid-hour, and then shifts the first half outwards by a number of hours (the *strength* of the attack) $\mu = 1, 2, \dots$, as depicted in Figure 7.1. This is motivated by the fact that, normally, the cheapest prices lie somewhat in the middle hours of the day. Note that the analogous attack where both halves are shifted outwards was also considered and its empirical results were found to be very similar.

In more detail, let t^* be the median hour with non-negative energy allocation for agent j , $\mathbf{E}_{[k+1]}^{(i),j}$. Then, given $\mathbf{E}_{[k+1]}^{(i)}$ from Eq. 6.2b, the allocation of aggregator j is modified as follows:

$$\hat{E}_{[k+1], t}^{(i),j} = \begin{cases} E_{[k+1], t+\mu}^{(i),j}, & \text{if } t \leq \lfloor t^* \rfloor - \mu \\ 0, & \text{if } t \in (\lfloor t^* \rfloor - \mu, \lfloor t^* \rfloor] \\ E_{[k+1], t}^{(i),j}, & \text{if } t > t^* \end{cases} \quad (7.2)$$

Note that, in the mathematical formulation presented in Eq. 7.2, the allocation can be pushed beyond the 24h day interval for large values of μ , but this does not happen for the range of values employed in the empirical evaluation presented in Section 7.3.

Finally, this attack vector can be extended to target all the competing aggregators, rather than aggregator j only, and will then be referred to as *ShiftAll* attack.

For illustrative purposes, the effects of *ShiftAll* in a scenario with three aggregators are shown in Figure 7.2. The third aggregator deviates from the vanilla algorithm and manages to shift the allocation of the first two aggregators out of one of its preferred hours (3 am). As a result, the attacker is able to obtain more energy at 3 am and reduce its allocation during the more expensive hours 5 and 6 am.




FIGURE 7.3: Real example of the effects of *ProportionalAll* on the resulting energy allocations. The scenario consists of three aggregators of the same size with the third aggregator being the attacker and employing $\lambda = 0.66$. Results for both vanilla and attacked scenarios are presented.

7.1.2 Proportional Attack

In this type of manipulation, the deviating aggregator i runs its local optimisation problem (*i.e.* Eq. 6.2b), but scales down its allocation for another aggregator j , $\mathbf{E}^{(i),j}$, by a factor $\lambda \in [0, 1]$, which indicates the *strength* of the attack. Formally, $\mathbf{E}_{[k+1]}^{(i)}$ is obtained from Eq. 6.2b, and then modified as:

$$\hat{\mathbf{E}}_{[k+1]}^{(i),j} = \mathbf{E}_{[k+1]}^{(i),j} \cdot (1 - \lambda) \quad (7.3)$$

The effect of this attack is to flatten the the allocation of aggregator j , resulting in less overlap with aggregator i 's desired schedule, in a similar way to the *Shift* attack. Note that the attacked aggregator still enforces its own constraints, Eqs. 3.2b, 3.2c, 3.2d, 3.2e, so that the amount of energy it receives is enough to satisfy its requirements.

As an example, the effects of *ProportionalAll* in a scenario with three aggregators are shown in Figure 7.3. The third aggregator deviates from the vanilla algorithm and manages to flatten the allocation of the first two aggregators out of one of its preferred hours (3, 4, 5 am). As a result, the attacker is able to obtain more energy at the cheap hours and reduce its consumption during more expensive ones.

Analogously to the previous attack vector, *Proportional* can be targetted to all the competing aggregators. In such a case it will be denoted by *ProportionalAll*.

7.1.3 Freeze Attack

In this case, attacker i *freezes* its own allocation to the individually optimal one, *i.e.* the allocation that can be obtained by solving Eqs. 3.2a, 3.2b, 3.2c and 3.2d without taking into account the other aggregators. Formally, being \mathbf{E}_i^* the optimal individual allocation for aggregator i :

$$\hat{\mathbf{E}}_{[k+1]}^{(i),i} = \mathbf{E}_i^* \quad (7.4)$$

This way, the attacker hopes to get its individually optimal allocation and get the other competing aggregators to arrange their allocations around. Importantly, this attack vector can be combined with the others presented so far. Specifically, we call the attack vectors combining Eqs. 7.2 and 7.4 *FreezeShift* and *FreezeShiftAll*. Similarly, the attacks combining Eqs. 7.3 and 7.4 will be called *FreezeProp* and *FreezePropAll*. Intuitively, these combinations should yield

more benefit to the attacker since its own allocation will be less affected by the smoothing effect of the competing aggregators.

7.1.4 Adversarial Attack

This last type of attack we consider is different from the previously described ones, as the deviating aggregator i does not seek to directly manipulate another aggregator's allocation. Instead, it will try to incriminate a benign aggregator j to make it appear as a deviator, hoping it will be a false positive of the manipulation detection algorithm (discussed in Section 7.2) and penalised accordingly. Depending on the imposed penalty, this could consist on banning aggregator j 's participation on the current trading day, thus benefiting aggregator i as competition is reduced. Otherwise, this can be seen as a purely malignant adversarial attack. Note that this sort of attacks have attracted a lot of recent interest in the field of machine learning [Huang et al., 2011, Kurakin et al., 2016] and could present a serious drawback for incentivising agents to participate in cooperative schemes such as the one proposed in this work.

Formally, this attack can be performed by proposing a schedule for aggregator j , $\mathbf{E}_{[k+1],t}^{(i),j}$, artificially close to the allocation of aggregator j to itself in the previous round, $\mathbf{E}_{[k]}^{(j),j}$. Hence aggregator j appears to deviate from the algorithm as it breaks the balance in aggregators i - j interaction. This will become clear in Section 7.2.1 where we describe how to quantify manipulation. This attack can be parametrised by a parameter $\lambda \in [0, 1]$ which determines a linear combination between the schedule proposed by aggregator j to itself, and the schedule allocated by aggregator i to j as a result of Eq. 6.2a. Formally, given $\mathbf{E}_{[k+1]}^{(i),j}$ from Eq. 6.2b, modify the allocation to aggregator j as follows:

$$\hat{\mathbf{E}}_{[k+1]}^{(i),j} = \mathbf{E}_{[k+1]}^{(i),j} \cdot (1 - \lambda) + \mathbf{E}_{[k]}^{(j),j} \cdot \lambda$$

In more detail, an attack with parameter $\lambda = 1$ proposes an allocation to aggregator j equal to what j proposed for itself in the previous round. This is likely to be beneficial for aggregator j 's schedule, as it will contribute towards maintaining the more favourable schedules characteristic of early rounds before convergence. However, as will be detailed in Section 7.2.1, this will make benign aggregator j seem a deviator, with the subsequent penalty. Conversely, as λ tends to zero, we recover the vanilla ADMM algorithm.

7.2 Detecting Manipulation

In this section, we detail a mathematical framework for quantifying the influence of a given ADMM participant, *i.e.* an aggregator, onto the rest of participants. The aim is to be able to detect outliers that are symptom of strategic manipulation in the system. Note that this framework is general, and can be applied to any ADMM (or variant) scenario, although we focus on our particular case for ease of exposition.

7.2.1 Quantifying Manipulation

The basic idea is that any group of aggregators with overlapping energy requirements should influence each other's schedules with *similar* intensity. If a particular aggregator i is self-interested and wants to improve its allocation by deviating from the ADMM algorithm, it will exert a heavier influence onto its competitors' allocations. Conversely, as happens in the adversarial attack detailed in Section 7.1.4, an aggregator that tries to wrongly flag another benign aggregator as deviator would exert too little influence.

A key point is that each aggregator i produces a (local) proposed schedule for all the n participating aggregators. Formally, following the notation from Section 6.1:

$$\mathbf{E}_{[k+1]}^{(i)} = \left(\mathbf{E}_{[k+1]}^{(i),1}, \dots, \mathbf{E}_{[k+1]}^{(i),n} \right)$$

Hence, this local solution proposed by aggregator i at iteration k contains its own schedule, $\mathbf{E}_{[k+1]}^{(i),i}$, and all the schedules for all the other participants, $\mathbf{E}_{[k+1]}^{(i),j}$ for $j \neq i$. We assume that each aggregator, benign or deviator, is truthful about their own allocations in their proposed local solutions. The reason for this is that every aggregator wants the best energy schedule given their requirements and hence would report the optimal schedule arising from their minimisation problem. Moreover, note that, given that an aggregator is unlikely to frequently change size and needs electricity everyday, past behaviour can also be used to roughly infer some of the aggregator's characteristics. The study of more general manipulation settings is out of the scope of this chapter and is outlined as future work (Section 7.4). Also, without loss of generality, we assume that the deviating behaviour starts from the second ADMM round, when every aggregator has seen the proposals from each aggregator. This allows us to focus on the first two iterations ($k = 0, 1$) for ease of exposition.

Formally, let d be a square matrix of dimension n , the *difference matrix*, storing how much each aggregator affects its competitors' self-proposed allocations. In more detail, every i, j entry quantifies how much aggregator i modifies the self-assigned schedule of agent j , and is given by:

$$d^{i,j} = \|E_{[1]}^{(i),j} - E_{[0]}^{(j),j}\|$$

As mentioned above, we expect benign aggregators to affect each other's schedules in a similar way, but aggregator size significantly affects this. More precisely, there are natural magnitude deviations in $d^{i,j}$ and $d^{j,i}$ when the sizes of the benign aggregators i and j differ. An example of this effect is shown in the top row of Figure 7.4 in dark grey.

These natural differences are an obstacle for detection algorithms. In order to overcome this issue, we normalise the matrix d employing the total amount of energy allocated by each aggregator to itself, as a proxy to potentially unknown aggregator size. Note that other proxies can be used instead, such as actual size and total amount of electricity purchased in previous trading days. Formally, we can write:

$$\text{size}_i = \sum_{t=0}^{23} E_{[0]}^{(i),i}$$




FIGURE 7.4: *Difference matrix* (top) and *normalised difference matrix* (bottom), d and \bar{d} respectively, for a scenario with two aggregators of size 50 000 EVs (1 and 4) and two aggregators of size 150 000 EVs (2 and 3). One of the small aggregators (Aggregator 4) performs a *Shift* attack with $\mu = 1$ against the other small aggregator (Aggregator 1), displayed as light grey.

and the proportion of the size of aggregator i among the whole group of aggregators is given by:

$$p_i = \frac{\text{size}_i}{\sum_j \text{size}_j}$$

Then, the *normalised difference matrix*, \bar{d} , is given by:

$$\bar{d}^{i,j} = \|E_{[1]}^{(i),j} - E_{[0]}^{(j),j}\| \cdot \frac{\sqrt{p_i}}{\text{size}_i + \text{size}_j} \quad (7.5)$$

This scaling function was chosen as it empirically *flattens* the entries of the matrix \bar{d} corresponding to benign aggregators, eliminating most of the dependence on aggregator size. In more detail, extensive simulations were performed, studying a variety of scenarios with different number of aggregators of different sizes, and different attack vectors and attack strengths. The selected normalisation approach provides the best results. An example of the normalisation effect is pictured in the bottom row of Figure 7.4. In this plot, the effect of the manipulating aggregator is shown in light grey, whereas the rest of the dark grey bars correspond to benign behaviour. In the top plots, corresponding to the *difference matrix* d , we can see large differences between different entries, arising from the large size differences between the aggregators. Importantly, these natural differences are larger than the effect of the manipulating aggregator. On the contrary, the normalised bottom plots manage to nearly flatten all the natural differences, and the effect of the deviator aggregator clearly stands out.

Lastly, for the rest of the chapter, we assume that $n - 1$ aggregators are benign and only one of them can potentially be a deviator. This is motivated by the fact that, with a perfect detection algorithm, there exists a Nash equilibrium in which no-one wants to deviate. Note that the proposed detection algorithm, which we are now ready to introduce, could be extended to deal with the more general case of having any number of deviators.

7.2.2 Detecting Manipulation

The overall idea is to be able to detect deviating aggregators in order to penalise and discourage manipulation. As it is usually the case in complex stochastic environments, the aim here is to reduce false positives and false negatives, while keeping true positives and true negatives as high

as possible. In this work we consider a *positive* to be an aggregator detected as deviator, and a *negative* an aggregator classified as benign.

```

Input :  $\bar{d}, \alpha$ 
Output: list with the detected manipulating aggregator, if any

/* off-diagonal */
consider the off-diagonal elements: offDiag;
compute the median:  $\mu_{1/2} = \text{median}(\text{offDiag})$ ;
compute distances from each element in offDiag to  $\mu_{1/2}$ ;
find max distance  $\rightarrow \text{maxOffDiag}$ ;

/* on-diagonal */
consider the on-diagonal elements: onDiag;
compute the median:  $\mu_{1/2} = \text{median}(\text{onDiag})$ ;
compute distances from each element in onDiag to  $\mu_{1/2}$ ;
find max distance  $\rightarrow \text{maxOnDiag}$ ;

/* threshold-based detection */
max(maxOffDiag, maxOnDiag)  $\rightarrow \text{max}$ ;
aggregator index: index(max)  $\rightarrow i$ ;
if max >  $\alpha$  then
| deviator  $\rightarrow [i]$ ;
else
| deviator  $\rightarrow []$ ;
end

return deviator

```

Algorithm 2: Threshold-based strategic manipulation detection algorithm for a scenario with at most one deviator.

As explained in previous sections, the idea is that manipulating behaviour will stand out, as it exerts a larger or smaller influence in other aggregators allocations, compared to the scenario's average. Formally, one can use the normalised difference matrix \bar{d} defined in the previous section in order to quantify this mathematically: manipulating behaviour from aggregator i towards aggregator j is translated into a too large or too small entry $\bar{d}^{i,j}$. We propose applying a threshold-based algorithm, with threshold parameter α , as described in Algorithm 2. In more detail, the algorithm looks at the difference matrix \bar{d} , computes the medians of the matrix entries, and then finds the entry that deviates the most from the median. This is done separately for off- and on-diagonal elements (as there are intrinsic magnitude differences between $\bar{d}^{i,i}$ and $\bar{d}^{i,j}$ even when all aggregators are benign) and only the highest deviation of the two is taken as final candidate. Lastly, this candidate is classified as deviator if its deviation from the median is greater than the user-defined threshold α .

The choice of threshold α is critical and we empirically study the performance of different thresholds in Section 7.3.2. Also, although the presented algorithm is designed to work in scenarios with at most one manipulating agent, by selecting the aggregator that deviates the most, it can be easily adapted to a general scenario. The most straightforward way would be to simply classify as deviator any aggregator i with $|\mu_{1/2} - \bar{d}^{i,j}| > \alpha$ for some j . This extended algorithm is conceptually the same as Algorithm 2 and will be studied in future work.

We are now ready to present the empirical evaluation.

7.3 Empirical Evaluation

In this section we present an analysis of the performance of the different attack vectors proposed in Section 7.1 and the manipulation detection algorithm presented in Section 7.2.2. This empirical evaluation closely follows the scenarios considered in Chapters 3, 4, 5, 6. In more detail, we consider a night-time residential scenario in which EVs arrive in the evening and need to be charged by the next morning. In more detail, the aggregator, market and driver characteristics are described in Sections 3.3.1, 3.3.2 and 3.3.3, respectively.

7.3.1 Attack Vectors: Utility and Convergence

Similarly to the analysis presented in the the previous chapter, Section 6.3, where we studied the convergence of the proposed ADMM algorithm, we will now turn our attention to analysing the behaviour of the algorithm under the effect of the manipulative attacks presented in Section 7.1. To this end, we identify two key quantities to analyse in order to assess the efficacy of a given attack vector: the effectiveness of the considered attack, in terms of utility increase (*i.e.* energy cost reduction) for the attacker, and the convergence of the algorithm under attack. This is motivated by the fact a given manipulative aggregator would try to improve its own allocation by manipulating the ADMM algorithm in a subtle way that goes unnoticed.

In order to evaluate these two quantities, we run simulations for each of the proposed attack vectors over the ten first days of November 2016 using $\rho = 10^{-5}$, and present the results in Figure 7.5. Here, we can see that both *Proportional* and *ProportionalAll* manage to achieve reduced costs for the attacker with good convergence rates in the scenarios with three aggregators. However, for settings with four aggregators, both these attacks destabilise the algorithm which outputs non-optimal allocations with increased costs for all the aggregators. Interestingly, both *FreezeProp* and *FreezePropAll* present very good results, consistently providing reduced costs for the attacker and close to 100% convergence rates for all attack strengths. On the other hand, *Shift*, *ShiftAll*, *FreezeShift* and *FreezeShiftAll* present very large cost reductions but fail to converge in the vast majority of cases. We would like to note that even a 1% cost reduction in the scenarios considered in this work represents savings in the order of hundreds of thousands of Euros per year. Finally, in a slightly different vein, *Adversarial* presents 100% convergence rates for all attack strengths and very small cost alterations. Recall that the aim of this attack is not to directly increase the attacker's utility, but to incriminate a benign aggregator as deviator (see Section 7.1.4). As a consequence, the efficacy of *Adversarial* is more appropriately measured by looking at the number of false positives (benign aggregators incorrectly classified as deviators) that the attack is able to generate. Results from this analysis are detailed in Section 7.3.3, as they require parts of the explanation about the performance of the proposed detection algorithm, which is detailed in the next section.

Summarising, we have studied the effects of each of the attack vectors on the convergence and outcome of the proposed ADMM algorithm. While some attacks (*Shift* and variants) prevent the algorithm from converging thus limiting their practical applicability, others, particularly *FreezeProp* and *FreezePropAll* are successful in decreasing the attacker's electricity costs while barely affecting the algorithms convergence. We would like to note that these interesting results all arise from myopic attack vectors that are simply repeated iteration after iteration. Given

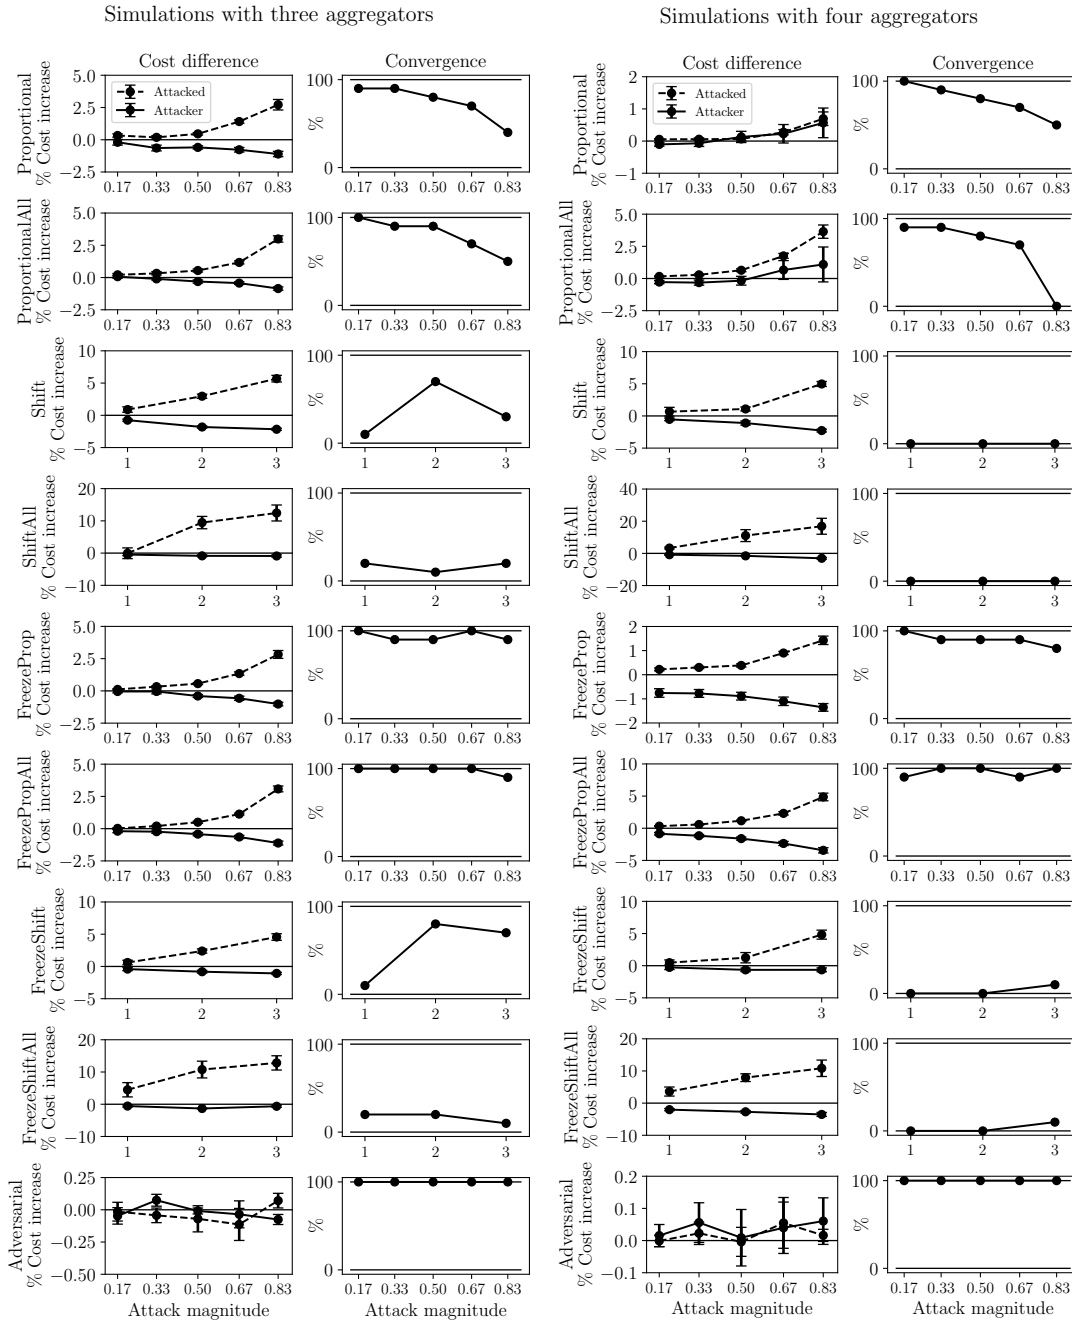


FIGURE 7.5: Cost and convergence analysis for each of the attack vectors described in Section 7.1. Scenarios with three aggregators (LHS) and four aggregators (RHS) and averaged over the ten first days of November 2016. All aggregators have capacity for 150 000 EVs.

that an attacker is potentially able to monitor convergence and other metrics during successive iterations of the algorithm, it is likely that greater cost savings and convergence rates could be obtained by more sophisticated algorithms.

7.3.2 Threshold-Based Detection Results

So far, we have shown the efficacy of the different proposed attack vectors. Next, we present the empirical performance of the proposed detection algorithm. As mentioned in Section 7.2.2, the

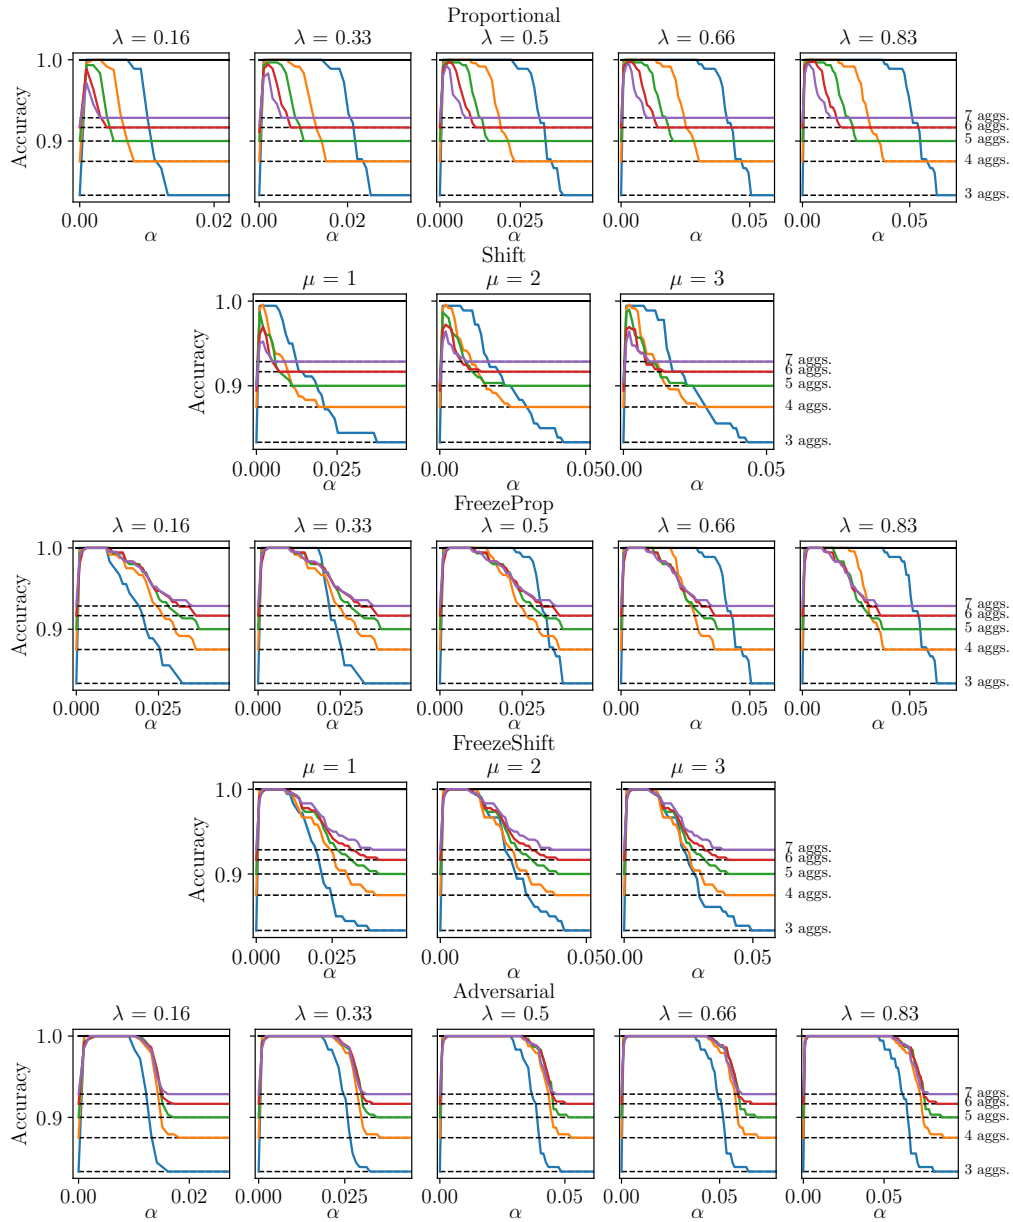


FIGURE 7.6: Accuracy analyses for each of the proposed attack vectors, attack strengths and number of aggregators. Results averaged over every day of November 2016. All aggregators have capacity for 150 000 EVs. Dashed lines represent the naive benchmark for each scenario, which considers every aggregator to be benign.

choice of the threshold parameter α is critical for good performance. Recall that the ultimate aim is to maximise correct classifications, *i.e.* true positives and true negatives. If α is set too high, we will fail to detect deviating behaviour. Conversely, if α is set too low, the detection algorithm would be too sensitive and misclassify benign aggregators as manipulating.

In order to quantify the performance of the detector we employ the *accuracy* metric [Metz, 1978], given by:

$$Accuracy = \frac{True\ Positives + True\ Negatives}{Total\ population}$$

Specifically, accuracies 0 and 1 correspond to detectors which are always wrong and always right, respectively. Moreover, we will compare the accuracy of our proposed algorithm with that of a

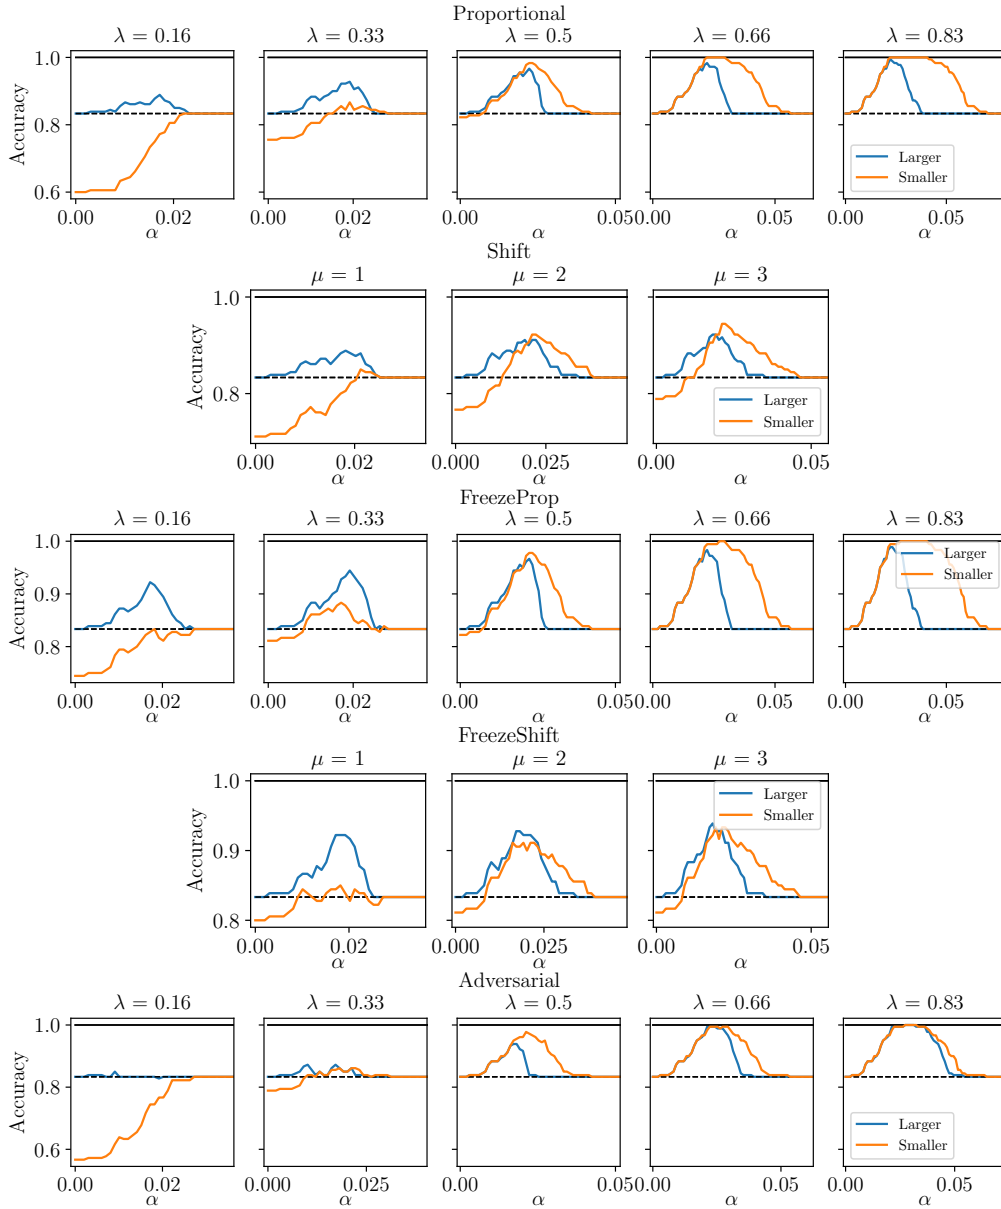


FIGURE 7.7: Accuracy analyses for each of the proposed attack vectors, attack strengths and aggregator size. Results averaged over every day of November 2016. Dashed lines represent the naive benchmark for each scenario, which considers every aggregator to be benign.

naive detector that classifies all the aggregators as benign. This is motivated by the fact that our dataset is unbalanced given that we consider at most one deviator per simulation (see Section 7.2). A well performing algorithm should present increased accuracy from this naive benchmark.

In more detail, we present the results from two different experiments. First, we consider scenarios with varying numbers of aggregators of the same size. This smoothens the natural discrepancies in the difference matrix, d , and allows us to focus on each of the different attacks and attack strengths. Second, we present an experiment considering aggregators of different sizes. In this case, the natural discrepancies arising from the size difference play an important role (see Section 7.2.1), and it is more difficult for a detection algorithm to distinguish between size effects and manipulation. For each scenario, we simulate every day of November 2016 and use $\rho = 10^{-5}$. Recall that we assume that the deviating aggregator performs uses the same attack vector and




FIGURE 7.8: Analysis of the efficacy of *Adversarial* for two attack strengths in the *Smaller* scenario, as described in Section 7.3.2. Dashed lines correspond to the percentage of times the attacked aggregator is incorrectly classified as deviator, in the case with no manipulation. Solid lines represent the same quantity when the aggregator is the target of *Adversarial*. Results averaged over every day of November 2016.

attack strength in all rounds, and focus our detection algorithm in the first and second rounds. The results for each of the two experiments are shown in Figs. 7.6 and 7.7, respectively.

The results from the first experiment indicate that our proposed algorithm is able to detect the majority of attacks with an accuracy of, or very close to, 1. This shows that our algorithm significantly outperforms the naive benchmark. Moreover, these results are consistent across different numbers of aggregators ranging from 3 to 7. Focusing on each different attack, detection is easier as the strength increases. This is intuitive as stronger attacks have a more pronounced effect on the algorithm and stand out. Interestingly, the attack vector which is most difficult to detect is *Shift*, which is the only one where the algorithm is not able to reach perfect accuracy for any number of aggregators. Finally, the presence of greater number of aggregators makes detection slightly more difficult in most cases. Note that, as explained above, these simulations consider aggregators with the same size. As we will see next, size effects make detection more challenging. Finally, note that for large values of α , the detection algorithm classifies everyone as benign thus converging to the naive benchmarks.

In the second experiment, we consider two different settings with three aggregators. Firstly, the attacker has capacity for 450 000 EVs, whereas the other two aggregators are smaller and have capacity for 150 000 EVs. Secondly, the attacker and one of the benign aggregators have capacity for 150 000 EVs, while the attacked aggregator has capacity for 450 000 EVs. These two scenarios are called *Larger* and *Smaller* in Figure 7.7, respectively. We can see that size effects make detection considerably more challenging. Our proposed algorithm still outperforms the naive detector in the vast majority of cases, and perfect accuracy is still achieved for high attack strength in some cases, such as *Adversarial*, *Proportional* and *FreezeProp*. However, smaller attacks are difficult to detect and in some cases the performance of our algorithm is comparable to the naive detector (see for example *Adversarial* with $\lambda = 0.16$ in the *Smaller* scenario. This suggests that the discrepancies in the difference matrix arising naturally due to size differences are not smoothed enough by the proposed normalisation scheme. However, despite this extra challenge, the proposed detection algorithm significantly outperforms the naive benchmark and presents very good detection results across a variety of settings.

7.3.3 Performance of *Adversarial*

As mentioned in Section 7.3.1, the efficacy of the *Adversarial* attack vector is appropriately measured in terms of its success rate in incriminating a benign aggregator as deviator. We perform this analysis using the same datasets as in the previous section: namely the scenarios with varying number of aggregators of the same size and the ones including varying size aggregators. Results indicate that this attack vector is unsuccessful in most cases, not being able to incriminate the attacked aggregator.

The most interesting behaviour is obtained in scenarios containing aggregators of different sizes and where the attacked aggregator has a larger size than the rest, *i.e.* the scenarios labelled as *Smaller* in the previous section (the attacker is *smaller* than the attacked aggregator). In these, the natural discrepancies in the difference matrix cause the detection algorithm to incorrectly classify the attacked aggregator as deviator for small α values. This happens as, for small α 's, the algorithm will pick up the aggregator that stands out the most in the difference matrix, even if only by a very small margin. Interestingly, if the attacker performs an *Adversarial* attack on the large benign aggregator, the detection algorithm is able to detect it in some cases, hence eliminating the false positive that would be obtained without manipulation. This situation is depicted in Figure 7.8. In more detail, we can see that weak *Adversarial* attacks ($\lambda = 0.16$) achieve up to 80% success rates. However, the same scenario without manipulation presents higher false positive rates due to the reasons explained above. Moreover, increasing the values of α reduces the false positives to zero. Similar results can be observed for stronger attacks ($\lambda = 0.33$). For the rest of scenarios, the attacker is not able to incriminate the attacked aggregator at all, and is in turn detected as deviator.

Importantly, we would like to remark the fact that this chapter focuses on scenarios where there is, at most, one deviating aggregator (see 7.2.1). In more general settings, which will be subject of study in future work, it is likely that this attack vector will present much higher success rates. In more detail, the proposed detection algorithm will choose at most one aggregator as deviator, so the natural deviations arising from size differences compete with the manipulation effects. If size difference stands out more, the detection algorithm will incur a false positive without any manipulation. Conversely, if the effects of an *Adversarial* overcome the natural discrepancies, the algorithm will correctly detect the attacker. In the more general case where there are any number of deviating aggregators, and with an algorithm (extending the one proposed in this chapter or otherwise) that can detect any number of aggregators as deviators, both size effects and malignant incrimination from an *Adversarial* attack could lead to multiple detections.

7.4 Summary

In this chapter, we present the first study about strategic manipulation of ADMM algorithms by self-interested internal agents. Even though ADMM and related decentralised optimisation algorithms are widely applied in many disciplines, little work has focused on studying how these algorithms can be disrupted by internal attackers. In order to address this issue, we study how a deviating agent can alter their local algorithm in order to increase their own utility at the expense of their competitors. Focusing on the decentralised multi-EV aggregator setting

and algorithm considered in Chapter 6, we introduce several attack vectors that a self-interested aggregator can employ in order to alter the outcome of the ADMM algorithm. Moreover, in order to prevent strategic manipulation, and working towards resilient decentralised optimisation, we study how deviating behaviour can be detected. In more detail, we propose a mathematical framework which measures the effects that different agents exert onto each other when employing the ADMM algorithm. Furthermore, we propose a threshold-based algorithm which employs this formalism in order to classify the participating aggregators as benign or deviators. Although we focus on an energy setting, the proposed detection framework is general and can be applied in general.

In order to study the proposed decentralised algorithm and detection mechanism, we present an empirical evaluation using real market and vehicle usage data from Spain. We first focus on the performance of the proposed attack vectors and analyse their impact on attacker utility (reduced energy costs) and on algorithm convergence. Specifically, we show that an attacker is able to effectively alter the outcome of the algorithm for their own benefit. Secondly, we turn our attention to the proposed detection framework, and present an accuracy study to assess its performance. Results show that our algorithm achieves very high accuracy, close and up to 1 in all cases when the considered aggregators have the same size, significantly outperforming the naive benchmark. However, considering aggregators with different sizes makes detection more difficult. Although our proposed algorithm outperforms the naive benchmark in most cases, some small attacks are very challenging to detect.

By addressing Research Challenge 4 in particular, this chapter concludes our investigation on the coordination opportunities of EV aggregators, namely Research Challenges 2 and 3. In the next chapter, we will turn our attention to Research Challenge 5, that is, the issue of energy units allocation.

Chapter 8

Online Allocation of Energy Units

Once a given EV aggregator has purchased an energy allocation from the day-ahead market, it faces the problem of how to schedule this energy among the available EVs, as captured by Research Challenge 5. This issue is not trivial, given that the aggregator does not know about future EV arrivals, and decisions need to be taken on the go. In order to address this challenge, we will consider an online model with dynamic arrivals and departures, and propose different online scheduling algorithms. The overall aim is to maximise resource utilisation, customer satisfaction (*i.e.* number of EVs with the desired state of charge by departure time) and fairness.

The rest of the chapter is structured as follows. We formally define the online model in Section 8.1. Then, we describe the considered efficiency and fairness objectives in Section 8.2. Next, in Section 8.3, we study our scenario from an offline perspective, that is, assuming perfect knowledge of future arrivals. The main contribution of this chapter is presented in Section 8.4, where we study the online scenario and present existing and novel scheduling algorithms. Finally, we evaluate the proposed algorithms in a realistic case study in Section 8.5.

8.1 Model

We consider a setting with n agents $A = \{1, \dots, n\}$ (e.g. EVs) who arrive in an online manner with arrival time $a_i \in T$ and departure time $d_i \in T$ ($d_i \geq a_i$) and who require up to a certain quantity of resource q_i on departure. The value q_i could be seen as the total energy needed to fully charge the battery of an electric vehicle. We assume that both time and quantity are discrete, *i.e.* $T = \{1, \dots, \tau\}$ and $q_i \in \mathbb{N}^+$. Let $x_{i,t} \in \mathbb{N}^+$ denote the amount of resource allocated to agent i at time t . Moreover, we say an agent is *satisfied* if its allocation is equal to its required quantity, *i.e.* if $\sum_{t=a_i}^{d_i} x_{i,t} = q_i$. We assume that agents can consume resources at a maximum charging rate of $\bar{r}_i \in \mathbb{N}^+$. However, since no algorithm will allocate more resources than an agent needs, the *effective* maximum charging rate is given by $\bar{r}_{i,t} = \min(\bar{r}_i, q_i - \sum_{t'=a_i}^{t-1} x_{i,t'})$. Finally, at each point in time $t \in T$ a supply $s_t \in \mathbb{N}^+$ of the resource becomes available, which can be allocated only to agents who are in the market, *i.e.* $A_t = \{i \in A : a_i \leq t \leq d_i\}$. We assume that supply is *perishable* meaning that supply needs to be allocated immediately and any unallocated supply is lost. Consequently, we have that $\forall t : \sum_i x_{i,t} \leq s_t$.

This results in the following *problem constraints* for any allocation we consider:

$$\forall i \in A, \forall t < a_i : x_{i,t} = 0 \quad (8.1)$$

$$\forall i \in A, \forall t > d_i : x_{i,t} = 0 \quad (8.2)$$

$$\forall i, a_i \leq t \leq d_i : x_{i,t} \leq \bar{r}_i \quad (8.3)$$

$$\forall i \in A : \sum_{t=a_i}^{d_i} x_{i,t} \leq q_i \quad (8.4)$$

$$\forall t : \sum_{i \in A} x_{i,t} \leq s_t \quad (8.5)$$

Example 8.1. Consider two EVs arriving at a charging station with parameters $a_1 = 1$, $d_1 = 3$ and $a_2 = 2$, $d_2 = 4$ respectively. See Figure 8.1. The first vehicle has a maximum charging rate $\bar{r}_1 = 2$ and needs three units of energy, $q_1 = 3$. The second vehicle has a maximum charging rate $\bar{r}_2 = 1$ and needs two units of energy, $q_2 = 2$. The available supply to the charging station is one unit at time steps 1, 2, 4, and two units at time step 3. We give an example allocation that fully satisfies both EVs by allocating $x_{1,1} = x_{1,2} = x_{1,3} = 1$ and $x_{2,3} = x_{2,4} = 1$, and zero otherwise.




FIGURE 8.1: Example allocation. For each EV and time step, each empty square represents one potential unit of energy, while a grey square represents an allocated unit.

8.2 Objectives

This work is motivated by the fact that agent preferences are often not easily obtainable in practice, e.g. because of the additional burden on the user and/or the difficult for users to express their willingness to pay. Hence, although it is reasonable to assume the agent constraints (a_i, d_i, q_i and \bar{r}_i) can be elicited with relatively little effort, we would like to avoid assumptions about the utility functions, and hence objectives such as social welfare (sum of utilities) maximisation are no longer applicable. Instead, we consider other objectives a system designer may reasonably want to achieve in our setting. In so doing, our goal is to study their computational complexity, as well as the theoretical and empirical trade offs between the various objectives. Specifically, we consider the following set of objectives:

MaxDelivered: Maximise the total resource allocated: $\sum_{i \in A} \sum_{t=a_i}^{d_i} x_{i,t}$.

MaxSatisfied: Maximise the number of agents satisfied: $\sum_{i \in A} \{ \sum_{t=a_i}^{d_i} x_{i,t} = q_i \}$.

Envy-Freeness: We say agent i is envious of agent j 's allocation if $\min \left(q_i, \sum_{t=a_i}^{d_i} \min(\bar{r}_i, x_{j,t}) \right) > \sum_{t=a_i}^{d_i} x_{i,t}$. In words, the allocation to agent j while agent i is in the market is greater than the allocation i receives, subject to i 's maximum charging rate and total demand. We say that an allocation is *envy-free* if no agent is envious of another agent's allocation.




FIGURE 8.2: Example showing that MaxDelivered does not imply MaxSatisfied.




FIGURE 8.3: Example showing that MaxSatisfied does not imply MaxDelivered.

The first objective is arguably the closest to the common social welfare maximisation objective and, indeed, is equivalent if agents have the same preferences and these are linear in the number of units received. The second objective is reasonable in cases where agents have some complementary preferences, e.g. in the EV setting a driver will need a minimum charge to complete a day's journeys and it is also reasonable to expect that users would want their vehicle fully charged by the time they need it. Similarly, in a cloud computing setting, partial completion of a computational task may not be feasible. The third objective is about allocation fairness, where we build on the established notion of envy-freeness, which requires that no agent should prefer another agent's allocation. Our definition is the first generalisation to online settings where agents may be in the markets at different times (and so it is not considered unfair when an agent is receiving more resources while another agent is not present). We note that, although our definition is less restrictive than the original offline one, it would be interesting to study even less restrictive extensions, and we briefly return to this in our future work discussion.

Note that, in general, not all objectives can be satisfied simultaneously, and none are equivalent. We explore this in more detail in the remainder of this chapter and present now an illustrative example.

Example 8.2. *In this example, we show that MaxDelivered and MaxSatisfied are not equivalent and neither implies the other. First, consider the setup depicted in Figure 8.2. The proposed allocation satisfies MaxDelivered, as all the available supply is allocated, but fails to satisfy any of the two EVs. Alternatively, if the unit of energy allocated to EV 2 at $t = 3$ was instead allocated to EV 1, we would have both MaxDelivered and MaxSatisfied.*

Secondly, consider the setup depicted in Figure 8.3. In this case, the proposed allocation verifies MaxSatisfied, given that it satisfies EV 1 and no other allocation could possibly satisfy the two EVs given the available supply. However, this allocation is not MaxDelivered, as there is an unused unit of supply at time $t = 3$.

8.3 Offline Algorithms

We start by considering the complexity of solving the offline problem (i.e. where we have complete foresight of the agents arriving in the future). Specifically, for the MaxSatisfied objective, this is generally hard, but it can be computed in polynomial time when the supply is a single unit per timestep:

Theorem 8.1. *A MaxSatisfied allocation can be computed in polynomial time when $s_t = 1$ at every time step, and we have $\forall i : \bar{r}_i = 1$.*

Proof. The problem corresponds to the classic scheduling problem $1|pmtn, r_i|\sum U_i$; the case where we have a single machine, tasks can be preempted, there are release times and deadlines for every task, and we want to minimize the *unit penalty* which is $U_i := 0$ if i is completed before d_i and 1 otherwise. Here, r_i denotes the *release time* of job i , following standard notation in the literature. This problem can be solved in polynomial time in the number of agents using an algorithm from [Lawler, 1990]. \square

Theorem 8.2. *Computing a MaxSatisfied allocation is NP-hard, even when $s_t = 2$ at every time step, and we have $\forall i : \bar{r}_i = 1$.*

Proof. We reduce from the scheduling problem $P2|pmtn, r_i|\sum U_i$, which is NP-hard [Du et al., 1992]. This is the scheduling problem with two identical machines, tasks that can be preempted, where each task has its own release time and deadline, and the objective is to minimize the number of tasks for which the deadline is not met. This corresponds to our setting where we have one agent for each task whose arrival time equals the task's release time, whose departure time equals the task's deadline and whose charging requirement q_i equals the task's processing time. Moreover, r_i and U_i are defined analogously to the previous Theorem. Finally, we also have that $s_t = 2$ at every time step t , and the charging rates are $r_i \in \{0, 1\}$ for every agent i . \square

Interestingly, the MaxDelivered objective is easier computationally:

Theorem 8.3. *A MaxDelivered allocation can be computed in polynomial time.*

Proof. We present a polynomial-time reduction to that of computing a feasible flow of a flow network with upper capacity constraints [Goldberg and Tarjan, 1988].

We build a network flow with nodes as follows:

$$\{s, t\} \cup A \cup \{v_t \mid t \in T\} \cup \{a_i^t \mid t \in T, i \in A\}.$$

The arcs between the nodes are as follows:

- Source s points to the n agents.
- Each agent node $a_i \in A$ points to each node a_i^t for all time slots $t \in T$.
- A node a_i^t only points to node v_t corresponding to time slot t if agent a_i is active in time slot t .

- Each node v_t points to sink t .

The only capacities of the nodes are as follows:

- Each node a_i has upper capacity q_i .
- Each node v_t has upper capacity s_t .

We then compute a maximum flow of the network. Moreover, an actual schedule can be found by looking at the corresponding feasible flow: if an agent node a_i send some flow f through a vertex v_t , then a_i gets f units of electricity in time slot t . \square

We now consider the tradeoffs between MaxDelivered and MaxSatisfied. Although a MaxDelivered allocation process need not optimize MaxSatisfied, as shown in Figure 8.2, we now prove that it is always possible to modify a MaxSatisfied schedule to also be MaxDelivered, and so achieve both objectives simultaneously.

Proposition 8.4. *There always exists a solution which maximizes both MaxDelivered and MaxSatisfied.*

Proof. To find a schedule that is both MaxSatisfied and MaxDelivered we first find a MaxSatisfied schedule (e.g. using our linear programming formulation). We can then take this schedule, which may not allocate all units of charge s_t at each time step, and augment it by assigning any unused charge to any available agent not covered in the returned MaxSatisfied schedule. If no such agent exists, then we are allocating as much charge in each time step as possible and the allocation is MaxDelivered. \square

We note that this result critically depends on the foresight given by the offline setting. In an online setting, it will be impossible for any algorithm to always maximise MaxDelivered or MaxSatisfied, as we prove in the next section.

Finally, in terms of the envy-freeness objective, a solution can be trivially obtained in polynomial time (both in online and offline settings) by an *equal contention* algorithm, i.e. an algorithm which, subject to charging constraints, allocates everyone the same amount at each time step. Formally:

Definition 8.5 (Equal Contention (EC)). A mechanism satisfies equal contention if, at every time step, the charging rate of any agent in the market which is below their maximum effective rate, is at least as fast as that of any other agent in the market. Formally, let $\bar{r}_{i,t}$ be the effective maximum charging rate (as defined in Section 8.1). Then, $\forall t, i \in A_t$, whenever $x_{i,t} < \bar{r}_{i,t}$ we have that $\max_{i' \in A_t} x_{i',t} = x_{i,t}$.

8.4 Online Algorithms

We now consider online algorithms, where, at every time step, the algorithm only has knowledge of the agents arrived so far and has no information about future arrivals. We will assume

that future supply is known. We start by adapting the following well known online scheduling algorithms to our setting: earliest deadline first, least laxity first and value density. In these algorithms, at each time t , the agents are prioritised based on different criteria, and the agent $i \in A_t$ with the highest priority is allocated resources at the maximum possible rate, i.e. $x_{i,t} = \min(s_t, \bar{r}_{i,t})$. If there is still supply remaining (i.e. if $s_t - x_{i,t} > 0$), the next agent in the priority list is allocated resources in the same manner until no more supply or agents are available. Ties can be broken in any way desirable. Where unspecified, in our empirical evaluation (Section 8.5) we use a random tie breaking rule.

- Earliest Deadline First (EDF): Prioritise agents based on their deadline, with earlier deadline having higher priority. In our implementation any ties are resolved using arrival time (earlier is better) followed by random.
- Least Laxity First (LLF): Let τ_i^{\max} denote the latest possible time that agent i can be scheduled any units in order to fully satisfy its demand, given its current the effective maximum charging rate and the available supply over time. Allocate the agents with the lowest τ_i^{\max} first. If there are agents who cannot be fully satisfied they get top priority and least satisfied agents are prioritised first. Ties are broken based on arrival time.
- Value Density (VD): Prioritise agents by *density* defined as: $\frac{q_i - \sum_{t' < t} x_{i,t'}}{(d_i - t + 1) \cdot \bar{r}_i}$. Ties are broken according to EDF.

These algorithms are simple as they are very short-sighted. Moreover, they do not directly optimise our objectives of MaxDelivered and MaxSatisfied. Therefore, we now introduce novel algorithms which are based on their offline counterparts, but without considering future arrivals. That is, an optimal schedule is computed based on the agents currently in the market, assuming no new agents arrive. The schedule is then recomputed at each timestep, considering any new agents who have arrived (noting that past allocations cannot be revoked). A challenge is that, typically, many such schedules exist and a random tie breaking rule turns out to perform poorly empirically. Therefore, our algorithm prioritises schedules where more resources are allocated to the agents earlier on in the schedule. Formally, for MaxDelivered the algorithm is as follows:

Definition 8.6 (OnlineMaxDelivered (OMDel)). Let t be a current time step and A_t the agents currently in the market. The algorithm proceeds in two steps. First, find the maximum number of units, \bar{x}_t , we can allocate at the current time, t , i.e. find $\bar{x}_t = \max \sum_{i \in A_t} x_{i,t}$ subject to the problem constraints. Then, an online MaxDelivered solution is found by maximising $\sum_{i \in A_t} \sum_{t'=a_i}^{d_i} x_{i,t'}$ subject to the problem constraints and $\sum_{i \in A_t} x_{i,t} = \bar{x}_t$.

Note that, first, the additional constraint of allocating \bar{x}_t units at time t does not reduce the optimality of the solution. That is, OnlineMaxDelivered will allocate the same number of units regardless of the constraint (we leave out a formal proof due to space). Second, imposing similar constraints at later timesteps is superfluous since any future constraints do not influence the choice at time t . Finally, it is easy to see that the additional constraint does not affect the computational time:

Theorem 8.7. *OnlineMaxDelivered runs in polynomial time.*

Proof. The argument is similar as that of Theorem 8.3 except that, since we are in an online setting and the allocation in previous time steps is irrevocable, we enforce the prior history of the allocation by having lower capacity constraints for the a_i^t nodes representing the agents for a particular time step. \square

We can similarly define an algorithm for MaxSatisfied. Due to Proposition 8.4 we can always achieve an allocation which maximises both MaxDelivered and MaxSatisfied, and this can be achieved by combining both criteria as follows:

Definition 8.8 (OnlineMaxSatisfied (OMSat)). Find \bar{x}_t as in Definition 8.6. Then, an online MaxSatisfied solution is found by maximising both MaxDelivered as well as MaxSatisfied, i.e. $\sum_{i \in A_t} \left(\sum_{t'=a_i}^{d_i} x_{i,t'} + \{ \sum_{t=a_i}^{d_i} x_{i,t} = q_i \} \right)$ subject to the problem constraints and $\sum_{i \in A_t} x_{i,t} = \bar{x}_t$.

Although this algorithm satisfies both objectives, it can be trivially shown that its computational hardness is the same as that of solving MaxSatisfied offline (see Section 8.3).

Finally, we show that, unfortunately, even though these algorithms look ahead, none of the algorithms are always optimal in online settings. This is surprising, especially for MaxDelivered.

Theorem 8.9. *No online mechanism exists which always optimises MaxDelivered.*

Proof. The proof is by counterexample. Consider a setting with initially 2 agents entering the market at time $t = 1$, where $\bar{r}_1 = 1$, $q_1 = 2$ and $\bar{r}_2 = 2$, $q_2 = 2$. Agent 1 is in the market for 4 timesteps, and agent 2 for only 2 timesteps ($d_1 = 4$, $d_2 = 2$). Supply is 2 units per time step.

We consider two options. In option 1, the algorithm allocates 0 units to agent 1 and 2 units to agent 2. In option 2, the algorithm allocates 1 unit to both. It is easy to see that any other option (i.e. which does not allocate the full supply in the first round) is suboptimal.

If option 2 is chosen, suppose that a third agent enters the market at $t = 2$ with $\bar{r}_3 = 2$, $q_3 = 2$, $d_3 = 2$ (i.e. the agent is in the market for 1 timestep and requires 2 units). The maximum total units allocated is 5. However, if option 1 was chosen, 6 units could have been allocated in total, showing that option 2 is not optimal.

If option 1 is chosen, suppose that 2 agents enter the market at times $t = 3$ and $t = 4$ with $\bar{r}_3 = \bar{r}_4 = 2$, $q_3 = q_4 = 2$, $d_3 = 3$, $d_4 = 4$ (they stay in the market for 1 timestep each). At timestep 2 agent 2 requires no more resources, so the only option is to allocate 1 unit to agent 1. Then, at timesteps 3 and 4 at most 4 units can be allocated, making the total delivery 7. However, if option 2 is chosen, at timestep 2 we give 1 unit to each. The total delivery after 4 timesteps is then 8.

Hence, neither option is always optimal, showing that no online algorithm can always optimise MaxDelivered. \square

Theorem 8.10. *No online mechanism exists which always optimises MaxSatisfied.*

Proof. The proof is by counterexample. Consider the setting detailed in the proof of Theorem 8.9. If option 2 is chosen and, following the previous proof, a new agent arrives at $t = 2$ with

$\bar{r}_3 = 2$, $q_3 = 2$, $d_3 = 2$, then only two agents would be satisfied. It is easy to see that choosing option 1 would satisfy all three agents. Similarly, if option 1 is chosen and two agents enter the market at times $t = 3$ and $t = 4$ with $\bar{r}_3 = \bar{r}_4 = 2$, $q_3 = q_4 = 2$, $d_3 = 3$, $d_4 = 4$, only three agents would be satisfied. Again, choosing option 2 would satisfy all four agents. Hence, neither option is always optimal, showing that no online algorithm can always optimise MaxSatisfied. \square

8.5 Empirical Evaluation

Despite the negative theoretical results for online settings, the algorithms may perform well in practice. Hence, in this section, we compare the performance of the proposed algorithms in a realistic simulation. We will start by detailing the setup considered, and then proceed to describe the results.

We consider a 24-hour period, *i.e.* $T = \{0, 1, \dots, 23\}$, each time unit representing one hour. Specifically, we consider a night-time charging scenario, in which EVs arrive to the charging station in the evening and depart in the morning. For every considered day, random EVs and supply conditions are generated, as detailed in the next sections. Also, we run 50 instances of each given simulation and present averaged results. Note that all the data from this empirical evaluation is publicly available [Perez-Diaz et al., 2019b].

8.5.1 Supply Settings

Each supply unit represents 3 kWh [Binetti et al., 2015]. In terms of the available supply, high availability is characteristic of night hours. In more detail, [Robu et al., 2013] identify the average available energy in a small neighbourhood to be 615 kWh between 2:00 and 6:00, corresponding to 5-12 units/time step. On the contrary, reduced supply is available during the rest of the day, corresponding to 99 kWh, hence 0-4 units/time step. Thus, the supply available at time t , s_t , is drawn from a Gaussian distribution with parameters specified in Table 8.1, both for high and low supply intervals. Specifically, mean, standard deviation, and minimum and maximum limits are presented for each.

	min	average	std	max
high supply	1	10.5	2.25	15
low supply	1	5.25	6	6
arrival	0	17	2.3	23
online time	1	12.75	0.75	remaining hours
charging rate	1	2.5	1	5
quantity	1	5.5	1.5	10

TABLE 8.1: Limits, averages and standard deviations of the Gaussian distributions of each of the scenario's stochastic variables.

8.5.2 EV Settings

Similarly to the available supply, the characteristics of every EV are drawn independently from a Gaussian distribution with the parameters specified in Table 8.1.

Robu et al. [2013] reported that the majority of agents arrive in the time interval 16:00-19:00 (peak at 17:00). With regards to the standard deviation of this arrival time, values such as 1.5 and 3.13 have been reported [Benetti et al., 2015]. Similarly, Robu et al. [2013] find that EV departures peak between 8:00 and 10:00. Thus, most EVs are plugged-in for at most 15h every night. Average and standard deviations are set to 85% and 5% of the online time limit, respectively.

With respect to battery capacity, it is assumed to be approximately 25 kWh, which corresponds to 9 units of supply. In terms of required charge, it is assumed that the battery's health is taken into account by the EV owner. Specifically, Neaimah et al. [2015] report a minimum charge value of 39%, with a median of 53%, and a maximum value of 68%.

Finally, the maximum charging rate of an EV depends on the vehicle itself and on the charging infrastructure. Given that we consider a night-time charging scenario, presumably in residential areas, we consider that the EVs do not have access to industrial-grade charging speeds. Typical charging speeds in this type of scenario are around 3-6 kW and 11 kW [Heydarian-Forushani et al., 2016].

8.5.3 Results

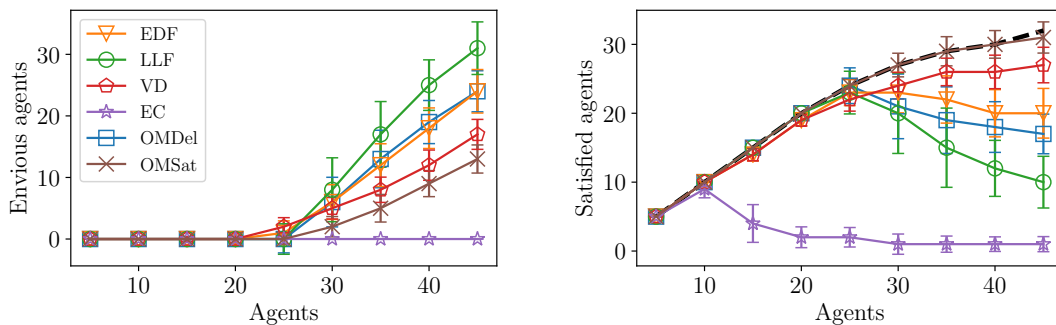


FIGURE 8.4: Number of envious agents (LHS) and number of satisfied agents (RHS) when varying the number of agents present in the simulation. The stochastic available supply units follow the Gaussian distribution specified in Table 8.1. The dashed line corresponds to the optimal offline MaxSatisfied allocation. Error bars indicate standard deviations.

We now compare how each of the considered algorithms performs with respect to the considered objectives. Figure 8.4 shows the number of envious agents (as per our envy-freeness definition, see Section 8.2) and the number of satisfied agents in a setting where we vary the total number of agents. We can see that, as the number of agents grows, it becomes more difficult to keep them envy-free and to satisfy their requirements. In fact, the number of satisfied agents drops as more agents are in the market for most algorithms, since more agents only get partially satisfied. Moreover, even though EqualContention is the only envy-free algorithm, it is extremely inefficient and scores very low on agent satisfaction. Interestingly, while LLF is the worst in terms of both measures, EDF and OnlineMaxDelivered are comparable, and Value Density is the best of

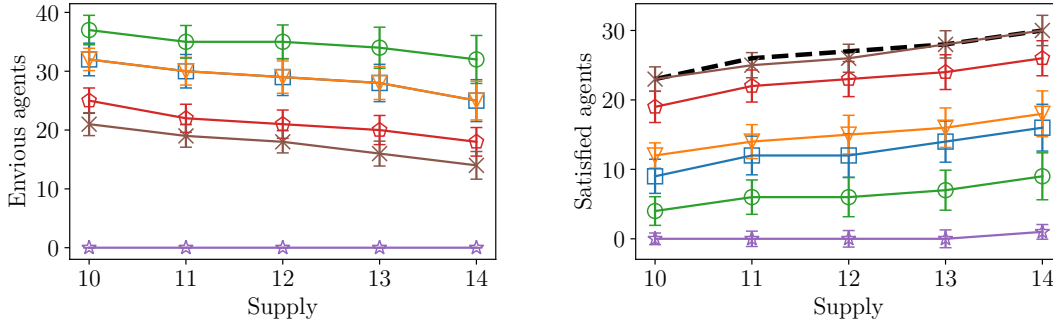


FIGURE 8.5: Number of envious agents (LHS) and number of satisfied agents (RHS) when varying the amount of supply available at each time step. Total number of agents is 45. The dashed line corresponds to the optimal offline MaxSatisfied allocation, averaged over 20 instances. Legend is shared with Figure 8.4. Error bars indicate standard deviations.




FIGURE 8.6: Runtimes for the proposed algorithms averaged over 50 instances.

the simple algorithms and performs surprisingly well. OnlineMaxSatisfied presents the highest efficiency in terms of the number of satisfied agents, achieving over 96% efficiency in all cases, and the best envy-freeness after EqualContention, but incurs more costly computation. Furthermore, all the algorithms perform similarly with respect to MaxDelivered, achieving between 95 and 100% of the offline optimal result, apart from EqualContention, which is only 13% optimal on average (graphs omitted).

These trends are similar when we vary the amount of available supply, as depicted in Figure 8.5. We can see that, as expected, overall results improve as the available supply grows, since there is less competition for the resources. Again, EqualContention is very inefficient, and OnlineMaxSatisfied presents, on the whole, the best results. Results with respect to MaxDelivered are also similar to before.

Finally, we present a runtime analysis of the proposed algorithms in Figure 8.6. We can see that the classical algorithms (EDF, LLF, VD) and EqualContention scale linearly with the number of agents and run in the order of milliseconds. In the contrary, our proposed algorithms, OnlineMaxSatisfied and OnlineMaxDelivered, present super-linear scaling and longer runtimes. In particular, as discussed above, OnlineMaxSatisfied presents the most expensive calculation. However, note that these optimisation algorithms can be parallelised in order to achieve faster runtimes, and thus allowing their application in larger settings.

8.6 Summary

In this chapter, we have addressed Research Challenge 5 by studying online allocations of energy units in a setting without money and where agents dynamically enter and leave the market. In doing so, we consider efficiency objectives (MaxDelivered and MaxSatisfied) as well as fairness (envy-freeness). Specifically, we provide computational complexity results as well as the trade-offs between the objectives. We also introduce novel online algorithms, called OnlineMaxDelivered and OnlineMaxSatisfied, and prove that their computational complexity is polynomial and NP-hard, respectively. Moreover, we prove that no algorithm exists which can always guarantee a solution satisfying either efficiency objective in online settings. Also, we extend the definition of envy-freeness to online settings, and show that an envy-free allocation can be found in polynomial time using an equal contention algorithm, but this algorithm performs poorly in terms of the other objectives. Finally, our empirical results show that the proposed OnlineMaxSatisfied algorithm achieves the best overall performance across the considered objectives, at the cost of harder computational complexity.

Finally, in the next chapter, we will focus on Research Challenge 6, namely the issue of forecasting price impact for participation in electricity markets.

Chapter 9

Forecasting Residual Supply Curves

In previous chapters, 3, 4, 5 and 6, we have explored the issues of optimising day-ahead market bidding and the design of coordination mechanisms for joint bidding. These efforts, addressing Research Challenges 1, 2 and 3, have been shown to reduce the energy costs associated with EV charging and, by reducing demand peaks, limit the stress on the electricity grid. However, due to the nature of electricity markets, such as day-ahead, bidding happens hours to days in advance, and relies heavily on forecasts. In the case of EV aggregators, these include future energy requirements (*i.e.* EV arrivals, departures and charging needs) and hourly prices. In particular, in this chapter, we will focus on the prediction of hourly price impact curves in day-ahead markets (see Section 2.3.3), as described in Research Challenge 6. In this setting, accurate predictions allow the optimisation and coordination mechanisms to perform at their best, and have been shown to reduce electricity costs by around 2% in the empirical evaluation presented in Section 3.3. Note that, although this percentage is small, given the size of the market players, it translates into millions of Euros per year.

In order to address these forecasting issues, we present a study that employs artificial neural networks as forecasting models, and compare their performance with existing models in the literature. More specifically, the rest of the chapter is structured as follows. First, we recall the notation employed to quantify price impact and detail the structure of the historical data used in this study in Section 9.1. Next, we describe the considered neural network models in Section 9.2. Section 9.3 describes the methodology used to pre-process the data and to train the neural networks. Finally, we present empirical results and compare our models to previous literature in Section 9.4.

9.1 The Data

Recall that the price impact curve for a given hour is a function of energy, E . As detailed in [Aneiros et al., 2013], this functional data can be converted to a multi-variate scalar problem by discretising the curve into a sequence of scalar points. In this study, we employ an equally-spaced

50-point discretisation of price impact curves between $E = 0$ and $E = 10$ GWh, $[E_1, \dots, E_{50}]$. This range is reasonable given the typical consumption volumes in the OMIE day-ahead market, and using 50 points provides a detailed discretisation of the curves. Note that the models presented in this chapter can be applied to arbitrary ranges and discretisations. The chosen parameters provide a balance between detail and computational complexity. Then, following the notation from Section 2.3.3, we can describe the price impact curve of a given day d and hour h by $[\mathcal{P}_h^d(E_1), \dots, \mathcal{P}_h^d(E_{50})]$.

9.2 Neural Network Models

This section describes the type of neural network models used in this chapter and details the form of the forecasting problem. We focus on multilayer perceptron (MLP) models, a type of feed-forward neural network that has achieved great results in many domains. More specifically, MLPs are a kind of neural network which consist of a sequence of layers of neurons with forward connections only. A detailed exposition can be found in, for example, [Bishop, 2006, Nunes et al., 2019].

9.2.1 Targets

The focus of this work is the prediction of the 24 price impact curves for the day-ahead. As explained in Section 9.1, these curves are discretised into 50 points, so for a given day there are 1200 points to forecast. There are several possibilities to achieve this. On the one hand, each of these points can be forecasted individually by using a single-task learning architecture. On the other hand, multiple points (either the 50 points of a given hour or the 1200 points of the whole day) can be forecasted together using multi-task learning. The latter benefits from potentially being able to capture the functional relation between different points, as they are forecasted using the same neural network. A more detailed discussion about these choices can be found in [Nunes et al., 2019]. In order to be able to capture the relations within the 50 points of every price impact curve, we choose to employ multi-task learning models with 50 outputs. This means that, for each hour of the day, one neural network is trained to forecast the whole price impact curve.

9.2.2 Features

A number of different features are utilised as input for our neural network models. Among these, the main feature is historical price impact data. Neural networks are trained to detect and exploit patterns in the time series of past price impact data, and these are then used to obtain a forecast for the day-ahead given previous data.

The other features considered are wind generation and total demand forecasts [Aneiros et al., 2011, 2013]. Official day-ahead forecasts of these two quantities can be obtained from the European Network of Transmission System Operators (ENTSOE) website ¹. These two exogenous features should have predictive power given that wind production accounts for a significant

¹<https://transparency.entsoe.eu>

amount of the total generation in the OMIE market, and that total electricity demand affects prices. Henceforth, we will refer to our considered models as *NN Simple*, *NN Wind*, *NN Demand* and *NN Both*, depending on whether they use price impact historical data alone, price impact historical data and wind forecasts, price impact historical data and demand forecasts, and price impact historical data, wind and demand forecasts, respectively.

Also, we assume that the sequence of price impact curves is Markovian, *i.e.* the current day only depends on the previous day. This means that, when forecasting the price impact curves for day d , price impact and, depending on the model, wind and demand forecasts, for the day $d - 1$ are used [Aneiros et al., 2011, 2013]. Note that additional information to help in the forecasting process is used in the models via the exogenous features mentioned above (wind generation and total demand forecasts). Moreover, additional features result from the use of inter-hour data (Section 6.2). Overall, this makes having more past time steps less important, but the Markovian assumption can be relaxed, and all the models used in this chapter can be seamlessly extended to employ any number of previous days as input for the forecast. However, this would increase the size of the neural networks and their computational complexity. Given the good results obtained by using one day, we will focus on this scenario and leave larger models for future work.

9.3 Methodology

In this section we discuss the methodology employed throughout the experiments and justify decisions and parameter choices.

9.3.1 Data Stationarity

A time series is said to be *stationary* if its statistical properties (*e.g.* mean, variance, etc.) do not vary over time. When dealing with a non-stationary time series, it is common practice to *difference* the data in order to remove trends [Aneiros et al., 2013]. Forecasting the differenced time series is then easier than forecasting the original data. Given the non-stationary patterns present in the considered data (see top of Figure 9.1), and following previous literature [Aneiros et al., 2011, 2013], we apply a one-day difference. Formally, given a price impact time series and following the notation introduced in Section 9.1 and [Aneiros et al., 2013], let $\chi_h^d(E_i) = \mathcal{P}_h^{d+1}(E_i) - \mathcal{P}_h^d(E_i)$ for all $h = 0, \dots, 23$ and $i = 1, \dots, 50$, be the differenced price impact time series. The effect of differencing and the original data are shown in Figure 9.1.

9.3.2 Inter-hour Correlations

As described in Section 2.12, previous literature forecasts price impact curves for each hour using historical data from that hour only (*i.e.* intra-hour). However, given that the day-ahead market comprises of all different hours and that market participants certainly participate in multiple hourly auctions, historical data from different hours should have valuable information for forecasting a given hour.




FIGURE 9.1: (Top) Time series $\{\mathcal{P}_{h=0}^d(E_1)\}$ for 2008 to 2018. (Bottom) Differenced time series $\{\chi_{h=0}^d(E_1)\}$ for 2008 to 2018

A pairwise correlation plot can be used to assess the relation between price impact data between different hours of the day, as shown in Figure 9.2. We can see high correlations for every pair of hours and, more specifically, two clusters with very high correlations of above 0.95%. This means that different hours are very interconnected and employing multi-hour data should yield improved forecasts. Finally, note that, although in the figure we show correlations for E_1 , similar correlations are obtained for every E_i .

In order to assess this empirically, we will explore single-hour models and compare their predictive performance with their multi-hour counterparts. Throughout the rest of the chapter, we will refer to single-hour models as *NN SingleHour*. Empirical results are detailed in Section 9.4.




FIGURE 9.2: Pairwise correlations between the time series $\{\mathcal{P}_h^d(E_1)\}$ for all d in 2016.

9.3.3 Data Normalisation

Neural networks perform best when all the input values are of similar magnitude [Bishop, 2006]. Hence it is an ubiquitous practice to normalise each input feature independently so that they

have zero mean and unit variance. In all the models discussed henceforth, input data will be normalised this way.

9.3.4 Sliding Window

Compared to traditional regression problems, time series forecasting has the particularity that new data is observed over time and added to the historical dataset. Specifically, in the day-ahead market setting, the resulting market data for a given day is made publicly available after market clearing. In practice, this means that the most recent market information can be included in the forecasting model to predict the next day. Given that most recent data is more relevant than old data, it is common practice to use a *sliding window*, in which the new daily data is added to the training dataset and the oldest data removed [Nunes et al., 2019]. This process is depicted in Figure 9.3 and used in all models described in this chapter. Note that the length of the training window can be varied and different values can yield different results.




FIGURE 9.3: Illustration of the sliding window procedure to assign a training window to two given test points, d and d' .

As a consequence, models are re-trained on the new training window on a daily basis. Note that this is computationally cheap and a standard practice in time series forecasting [Nunes et al., 2019]. We detail the results of our experiments in Section 9.4.

9.3.5 Validation and Test Sets

Our experiments use daily data for the years 2015, 2016 and 2017, and follow the standard training-validation-test split [Bishop, 2006]. In more detail, the test set for all our models is every day of 2017. This dataset includes all seasons of the year, thus accounting for seasonal factors such as temperatures, which affect the required energy and hence prices. Similarly, the validation set is the last two months of 2016. As is customary, this dataset is used to tune hyperparameters and to optimise the neural network models under consideration. Results concerning this process are detailed in Section 9.4.1.

Note that the training dataset is slightly different than most non-time series neural network use cases, as we consider the sliding window approach detailed in Section 9.3.4. This means that, for each data point in the validation or test datasets, the training data set the data window of the considered size immediately before the data point.

9.3.6 Optimising the Number of Epochs

The number of *epochs*, in a neural network context, refers to the number of times all training data is used to update the model. In general, too few epochs means not enough training and too

many epochs cause over-fitting [Bishop, 2006]. It is then essential to strike a balance in order to get good performance. In this work, given that each model is re-trained at each time step, and that different models are considered for each of the 24 different hours, an automated procedure to optimise the number of epochs was employed. In more detail, for a given experiment, a dataset immediately prior to the considered test dataset is considered. Then, the given neural network model is trained and evaluated on this dataset (following Section 9.3.4) keeping track of the loss function every training epoch. Finally, these curves are averaged among every day of the dataset and minimised to obtain an optimal number of epochs. This optimal value is then used in the final evaluation of the model. Details about the length of this dataset are presented with each experiment in Section 9.4.

9.3.7 Quantifying Prediction Errors

Given that the aim of this work is to accurately forecast price impact curves, measuring the *error* or *deviation* between a given forecast and the actual realisation is essential. Moreover, neural networks learn by iteratively minimising this error. We consider two common measures of error, the so-called *mean squared error* (MSE) and *mean absolute error* (MAE) [Nunes et al., 2019]. Specifically, MSE is used as a loss function for all the neural network models, and MAE as model performance metric, allowing standardised comparisons between different models. Formally, given forecasted and real price impact curves, $[\hat{\mathcal{P}}(E_1), \dots, \hat{\mathcal{P}}(E_{50})]$ and $[\mathcal{P}(E_1), \dots, \mathcal{P}(E_{50})]$ respectively, both errors are defined as:

$$\text{MSE} = \frac{1}{50} \sum_{i=1}^{50} \left(\hat{\mathcal{P}}(E_i) - \mathcal{P}(E_i) \right)^2$$

$$\text{MAE} = \frac{1}{50} \sum_{i=1}^{50} \left| \hat{\mathcal{P}}(E_i) - \mathcal{P}(E_i) \right|$$

Moreover, when evaluating a given model and presenting averaged results, both across different hours and different days, MSE and MAE errors for a given price impact curve will be averaged.

9.4 Results

In this section we present and discuss the results obtained with different neural network models and compare their performance to existing literature. First, we detail the experiments carried to choose suitable architectures and hyperparameters for the considered models: *NN Simple*, *NN Wind*, *NN Demand*, *NN Both* and *NN SingleHour*. Secondly, we discuss the performance of these optimised models on the test set. Thirdly, we compare the results from the intra-hour model with the rest of the inter-hour models. Finally, we compare our results to previous literature.

As mentioned above, a naive benchmark commonly used in the literature uses the price impact curves from the previous day as forecasts for the next day. Henceforth we will refer to this model as *DayBefore*.

Throughout all the experiments, the ADAM optimiser [Kingma and Ba, 2015] and ReLU activation functions were used. These are the most common choices in the modern neural networks

	Hidden neurons	Learning rate	Window size	Num. features
NN Simple	50	1e-4	200	1200
NN Wind	600	1e-4	200	1224
NN Demand	600	1e-4	200	1224
NN Both	600	1e-4	200	1248
NN SingleHour	25	1e-4	300	50

TABLE 9.1: Optimal hyperparameters found for each model. The number of features for *NN Simple* is equal to the number of points considered for the price impact functions (50) multiplied by the number of hours (24). For *NN SingleHour*, only one hour is considered as input, so the number of features is just 50. The models *NN Wind*, *NN Demand* and *NN Both* augment *NN Simple* by considering 24 extra data points (one for each hour) for wind and/or demand forecasts.

literature [LeCun et al., 2015]. Also, apart from *NN SingleHour*, all models use training data from all 24 hours (see Section 9.3.2). Finally, we used Tensorflow [Abadi et al., 2016] with the Keras API [Chollet and Al., 2015] in all the experiments.

9.4.1 Hyperparameters and Architecture Tuning

Following the procedure detailed in Section 9.3.5, we employ the validation set (last two months of 2016) to tune our neural network models with the aim of improving forecasting accuracy. In these experiments we vary the length of the training window size, the learning rate of the optimiser and the number and size of the hidden layers. Based on these experiments, we can select the best performing hyperparameters for each model, which are specified in Table 9.1. These selected models will be fully tested in the test set and the results detailed in Section 9.4.

9.4.2 Model Results

Using the optimal architecture for each model found in Section 9.4.1, we run each model on the full test set (every day of 2017), and present the obtained results in Figure 9.4. We can see that all our models, *NN Simple*, *NN Wind*, *NN Demand*, *NN Both* and *NN SingleHour*, perform significantly better than the naive *DayBefore* model. The performance of our five neural network models is close, although *NN Both* presents the best forecasts overall, followed by *NN Wind*, *NN Demand*, *NN Simple* and *NN SingleHour*. Exact forecasting accuracy of all models, aggregating every hour and averaged among all days of 2017 are presented in Table 9.2. These results are reasonable, given that *NN Both* has access to the largest set of features, and the models get smaller until *NN SingleHour*, which takes historical data for a single hour only.

Note that our models are able to extract the valuable information from the wind and total demand forecasts, improving the accuracy of their predictions. This is in contrast to previous literature where the considered models were not large enough to produce an improvement and results were mixed [Aneiros et al., 2011, 2013]. This is considered in more detail in the next section, where we present comparison of the accuracy of our models and existing models in the literature.

NN Both	4.051 (3.634)
NN Demand	4.072 (3.638)
NN Wind	4.073 (3.691)
NN Simple	4.113 (3.626)
NN SingleHour	4.377 (4.272)
DayBefore	9.653 (7.021)

TABLE 9.2: MAE errors (MAE stds) for each model, evaluated on the test set (every day of 2017). Units: EUR/MWh.




FIGURE 9.4: Results of the optimised architectures for each of the considered models.

9.4.3 Single-hour versus Inter-hour Models

As explained in Sections 2.12 and 9.3.2, while previous models in the literature use historical data from each separate hour in order to produce forecasts, we consider larger models where data from all 24 hours is taken into account. In this section we present results from experiments comparing the performance of single-hour and the larger intra-hour models.

As detailed in Table 9.2, results indicate that increased predictive performance, up to 6.028%, is obtained when using multi-hour models. Also, all the inter-hour models present lower spreads, *i.e.* reduced standard deviations. These improved results are to be expected given the high correlations between hours present in the data (see Section 9.3.2).

9.4.4 Comparison with Results in the Literature

Following the discussion from Section 2.12, we compare our models to the ones from existing literature [Aneiros et al., 2011, 2013]. More specifically, in those papers, the authors present three different models that are then combined into a fourth one, *SCOMB*, by using some heuristic rules. Their results indicate that *SCOMB* is the best performing and will be considered as literature benchmark henceforth.

Unfortunately, *SCOMB* (and the other models considered in these papers) are functional models, so direct performance comparison is not possible. In more detail, functional models work with functions rather than vectors, as our models do. Hence, *SCOMB* forecasts whole price impact curves. As a consequence, their error measures are also functional. Specifically, they employ the

NN Both	58.028%
NN Demand	57.815%
NN Wind	57.802%
NN Simple	57.392%
NN SingleHour	54.659%
SCOMB	7.805%

TABLE 9.3: Percentage improvements (using MAE) with respect to the *DayBefore* baseline. SCOMB is the best-performing model from existing literature [Aneiros et al., 2011, 2013].

L_1 error, which can be seen as the continuous version of the MAE we use in our work:

$$L_1 = \frac{1}{b-a} \int_a^b |\hat{\mathcal{P}}(E) - \mathcal{P}(E)| dE$$

where a and b are the considered E upper and lower limits, respectively.

However, the authors also present a comparison of *SCOMB* with the naive *DayBefore* benchmark. This allows us to use the percentage improvement of a given model from the naive *DayBefore* model as a uniform comparison metric. Finally, note that [Aneiros et al., 2011, 2013] use every day of 2009 as test set, while we use more recent 2017 data. This will introduce small discrepancies, given that the results of the *DayBefore* model will not be exactly the same. Nonetheless, given the very significant performance increase of our models, this effect is negligible.

Model comparison results are presented in Table 9.3. We can see that our five models largely improve on the results obtained by SCOMB by achieving up to 58.028% improvement from the *DayBefore* predictions, compared to the 7.805% obtained by SCOMB. As discussed before, the reasons for this improvement in performance include the use of modern machine learning techniques and larger datasets. As an example, in comparison with [Aneiros et al., 2011, 2013], we have considered window sizes up to five times larger than the experiments in their work. In terms of input data size, we have considered all 24 hours as features for our models, compared to the individual hours considered in the previous literature. Moreover, as reported by the authors, the models employed in [Aneiros et al., 2011, 2013] are very computationally expensive, which makes parameter tuning challenging. Comparatively, neural networks are quick to train on modern hardware which allows careful optimisation for peak performance.

9.5 Summary

In this chapter, we explore the forecasting of price impact curves in day-ahead electricity markets, as described in Research Challenge 6. To this end, we apply, for the first time in this setting, artificial neural network techniques. In particular, we focus on multilayer perceptron models. Moreover, we explore, for the first time, the effect of using historical data from different hours to improve the performance of our models. Following previous literature, we also consider external explanatory variables such as wind generation and total demand forecasts. Results using real data from the Spanish day-ahead market suggest that our models achieve very good performance, significantly outperforming previous results in the literature. In more detail, the largest considered model, which uses historical price impact data and wind and demand forecasts, presents the best predictive results improving the naive day-before benchmark by 58.028%. In

comparison, the best performing model from previous literature achieves only a 7.805% increase. Furthermore, we find that considering inter-hour models does improve forecasting performance by up to 6%.

This concludes the contributions presented in this thesis, and we are now ready to discuss the conclusions and outline future work directions.

Chapter 10

Conclusions and Future Work

This thesis is concerned with the challenge of integrating large fleets of EVs in the electricity market and power systems infrastructure. To this end, we have presented several studies that address distinct aspects of this broad challenge. Conclusions related to the findings of the considered studies and directions for future work are described in more detail in the rest of this chapter.

10.1 Conclusions

Following the research challenges exposed in Section 1.1, we have made the following contributions to the state-of-the-art. Firstly, we study the optimisation of the charging of a large fleet of EVs (Research Challenge 1). In order to address this issue, we propose a novel price-maker day-ahead bidding algorithm for EV aggregators. Importantly, this algorithm explicitly considers the price impact of the EV aggregators orders, an essential requirement for large market participants. More specifically, the algorithm takes the form of a non-linear minimisation problem which presents a complex landscape with multiple local minima. In order to guarantee global optimality, a desirable property for the coordination mechanisms introduced in subsequent chapters, we introduce an approximated convex minimisation formulation. In more detail, by approximating the hourly price impact functions with quadratic convex curves, we obtain a convex non-linear optimisation problem in which any local minimum is actually a global minimum. Moreover, the proposed algorithm is computationally inexpensive and is formulated in terms of linear constraints, presenting very good scalability with problem size and allowing the application of the coordination mechanisms detailed next. Finally, an empirical case study with real market and driver data is presented, in order to assess the performance of the proposed algorithm. Results show that the proposed algorithm presents very significant cost reductions with respect to existing price-taker algorithms, achieving over 10% savings for an EV penetration rate of 8% in the Iberian Peninsula.

Secondly, we turn our attention to Research Challenge 2, namely the coordination opportunities available for EV aggregators. To this end, we propose the first coordination mechanism that allows extending the benefits of coordinated EV charging to groups of independent private

EV aggregators, with respect to day-ahead market participation. In more detail, by employing techniques from mechanism design, we design a centralised third-party coordinator which bids on behalf of a group of EV aggregators, distributing the purchased energy and payments among them. This way, more informed bidding can take place, as the requirements of all the participating EV aggregators are known to the bidding algorithm. In order to ensure fairness and prevent strategic manipulation, in which an EV aggregator can try to cheat the system for its own benefit, we propose employing a VCG payment mechanism. This payment mechanism is truthful, *i.e.* it is in every participants interest to report true requirements to the coordinator. However, due to an undesirable characteristic of the VCG mechanism, it results in too high payments for it to be useful in this setup. This motivates the introduction of two further VCG-based payment mechanisms. Firstly, a VCG-based truthful mechanism which maintains truthfulness and minimises the undesired payment surplus. However, the payments assigned by this system are too little, and the coordinator incurs large monetary losses. Secondly and lastly, we propose VCG-based proportional payment mechanism, which achieves good performance at the expense of sacrificing full theoretical truthfulness. The performance of the proposed coordination mechanism is evaluated in an empirical case study which employs real market and driver data. Results show greatly reduced energy costs when compared to uncoordinated bidding.

Thirdly, addressing a slightly different aspect of Research Challenge 2, we present the first coalitional study of a competitive multi-aggregator day-ahead bidding scenario. Similarly to the previous contribution, we consider an arbitrary number of self-interested and independent EV aggregators who participate in the same day-ahead market trying to minimise energy expenditure. While the second contribution presented above proposes the use of a coordinator in order to reduce energy costs for all market participants, this coordinator does not necessarily need to be unique. Specifically, each EV aggregator is free to choose a number of other aggregators to coordinate with, possibly forming smaller coalitions if decreased energy costs are perceived. To study this issue, we turn our attention to the field of cooperative game theory, and model the consider scenario as a coalitional game. We are then able to prove the theoretical results of superadditivity and balancedness, which ensure that full cooperation (where the grand coalition forms) provides the lowest energy costs and that there is a payment mechanism which stabilises the grand coalition (no participating aggregator has the incentive to deviate). We propose a payment mechanism lying in the core, namely the least-core, and compare to the commonly used Shapley value in an empirical case study which uses real market and driver data. Results show that both payments are usually very close, suggesting a payment mechanism that is both fair and stable.

Fourthly, we turn our attention to Research Challenge 3, which is concerned with decentralising computation in multi-aggregator coordination mechanisms. In more detail, and in contrast with the two centralised coordination mechanisms proposed in the previous two contributions, decentralised cooperation provides important advantages such as limiting the need for sharing sensitive information and improving transparency and security. In order to address these issues, we reformulate the centralised coordination mechanisms proposed earlier as a distributed optimisation problem by using the alternating direction method of multipliers algorithm. In this new decentralised problem, each aggregator takes an active computational role, rather than just reporting energy requirements, which removes the need of reporting private information. Moreover, we propose the use of blockchain technology to implement the algorithm, which provides enhanced

transparency and anti-tampering guarantees, and removes the need for trust on a centralised coordinator. Finally, in a similar vein to previous contributions, we test the novel decentralised coordination algorithm in a case study using real market and driver data. Results show good convergence properties, suggesting the applicability of the algorithm in large settings.

Fifthly, we consider Research Challenge 4. Intimately related to the previous contribution, this challenge is concerned with the risk for potential strategic manipulation in distributed optimisation algorithms. In more detail, when considering centralised approaches (such as the ones presented in Chapters 4 and 5), self-interested aggregators can only manipulate the coordination algorithm by misreporting their energy requirements. However, when considering decentralised mechanisms such as the one presented in Chapter 6, the participating aggregators roles are not limited to reporting preferences, but also play an active part in computation. As a consequence, the opportunities for strategic manipulation are greatly increased. In order to address these issues, we present a study on the computational manipulation opportunities present in the decentralised coordination mechanism proposed in Chapter 6. More specifically, we propose several attack vectors which can be used by a deviating aggregator in order to modify its computation with the aim of increasing its personal utility. Furthermore, we present a manipulation detection mechanism which monitors the behaviour of each aggregator and identifies deviators. Finally, the performance of the proposed attack vectors and detection algorithm are evaluated in a case study using real market and driver data. Our findings show that a deviating aggregator is able to successfully affect the outcome of the coordination mechanism and artificially reduce its energy costs in detriment of the other participants. Moreover, the novel detection algorithm significantly outperforms a naive algorithm and presents very good detection accuracies. To conclude, we would like to note that, although we have focused on our multi-aggregator scenario, these issues exist in any decentralised optimisation algorithm, and the proposed algorithms can be easily generalised.

Sixthly, we address Research Challenge 5. Specifically, in order to support the real-time scheduling of energy resources, we study the online allocation of energy units in a setting without money and where agents dynamically arrive and depart over time. In more detail, we consider both efficiency and fairness objectives, such as maximising the number of allocated energy units or the number of fully satisfied customers, and envy-freeness. With respect to the first two objectives, we study their computational complexity both in offline and online settings and propose novel online algorithms (OnlineMaxDelivered and OnlineMaxSatisfied) which attempt to maximise them. However, we theoretically prove that no online algorithm exists that can always guarantee either efficiency objective. With respect to envy-freeness, we present the first extension of the classical definition to online settings, and prove that it can be obtained in polynomial time using an equal contention algorithm. Finally, we present a realistic case study where we evaluate existing and the novel online algorithms. Our empirical results show that the proposed OnlineMaxSatisfied algorithm achieves the best overall performance, at the cost of harder computational complexity.

Finally, in our seventh study we address Research Challenge 6. More specifically, we have focused on the prediction of hourly price impact curves in day-ahead markets, with the aim of supporting day-ahead bidding (Research Challenges 1, 2 and 3). In more detail, better forecasts translate into cheaper energy costs and, in order to address this issue, we present a study that employs artificial neural networks as forecasting models, and compare their performance with existing models in the literature. Moreover, we explore, for the first time, the effect of using

historical data from different hours to improve the performance of our models. Following previous literature, we also consider external explanatory variables such as wind generation and total demand forecasts. Results using real data from the Spanish day-ahead market suggest that our models achieve very good performance, significantly outperforming previous results in the literature. In more detail, the largest considered model, which uses historical price impact data and wind and demand forecasts, presents the best predictive results improving the naive day-before benchmark by 58.028%. In comparison, the best performing model from previous literature achieves only a 7.805% increase. Furthermore, we find that considering inter-hour models does improve forecasting performance by up to 6%.

10.2 Future Work

Despite these contributions, there is ample room for future work in order to address different important aspects of the research challenges considered in this thesis. For example, we have only considered participation on day-ahead markets so far. Despite the very similar structure between day-ahead and shorter-term markets, such as intra-day or reserve, an interesting challenge arises from sequential participation. More specifically, an EV aggregator can participate in both the day-ahead and several intra-day markets (see Section 2.2) in order to adjust for deviations between their forecasted and actual energy requirements. While this issue has been considered in the literature [Ugedo et al., 2006], it has not been explored for large EV penetration. Moreover, the issue of inter-aggregator cooperation has not been considered in this setting either. To this end, similar coordination mechanisms to the ones proposed in this thesis could be employed, using techniques from mechanism design or cooperative game theory. However, the interplay between the different markets adds complexity and new manipulation opportunities, which need to be addressed.

On a slightly different vein, but still addressing the issue of charging large fleets of EVs, are the issues of physical energy delivery. In more detail, after the electricity market is closed, the distribution system operator (DSO) needs to arrange electricity delivery between generators and consumers. In this setting, issues such as line capacity and transformer loading prevent the delivery of arbitrarily large energy quantities to a given geographical location. As a consequence, the delivery of energy to different self-interested aggregators can also be seen as a competitive game, and coordination mechanisms could be employed in order to prevent overloading the grid, a scenario which has not been addressed in the literature. In particular, the decentralised approach presented in Chapter 6 is particularly well suited for addressing this issue. More specifically, one can imagine an iterative process where a number of self-interested aggregators optimise their joint bidding (as per Chapter 3) and then this allocations are checked for network violations. If any, extra constraints can be applied to the aggregators as needed, and joint bidding is repeated. After a number of iterations, one would obtain optimal feasible energy schedules for all the aggregators.

Turning now our attention to the issue of online energy scheduling, we identify several directions for future work. First, studying the apparent incompatibility of the efficiency objectives, such as maximising the number of scheduled energy units and satisfied drivers, and the fairness

objective given by envy-freeness. More specifically, even though our definition of online envy-freeness (see Section 8.2) is quite relaxed (when compared to the classical definition of offline envy-freeness), it turns out to be quite restrictive. In particular, although it can be guaranteed by equal contention algorithms (see Section 8.3), they are extremely inefficient and waste the majority of the available resources. However, it is possible that other more efficient alternatives exist. Moreover, it would be interesting to consider approximately envy-free solutions. Second, there is the issue of departure time and charge requirements elicitation from drivers. Similarly to previous works, techniques from mechanism design can be employed to encourage truthful reporting. One of the challenges here is to study and devise monotonic algorithms (with respect to the tuple of driver requirements) [Gerding et al., 2011].

Finally, we consider the issue of forecasting, in support EV operation. To this end, given the enormous recent advances in machine learning, more sophisticated models can be utilised. For example, recurrent neural networks such as long short-term memory (LSTM) architectures have vastly outperformed classical architectures such as the multi-layer perceptrons considered in this thesis. Given that we consider a time-series problem, where recurrent architectures excel, it is to be expected that more accurate forecasts can be obtained from this and other new models. Also, the forecasting of EV arrivals and departures is equally important for optimised bidding, and can also be addressed in a similar way.

Bibliography

- M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mane, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viegas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems. *Arxiv*, 2016. URL <http://arxiv.org/abs/1603.04467>.
- S. K. Aggarwal, L. M. Saini, and A. Kumar. Electricity price forecasting in deregulated markets: A review and evaluation. *International Journal of Electrical Power and Energy Systems*, 31(1):13–22, 2009. ISSN 01420615. doi: 10.1016/j.ijepes.2008.09.003. URL <http://dx.doi.org/10.1016/j.ijepes.2008.09.003>.
- C. Akasiadis and G. Chalkiadakis. Decentralized Large-Scale Electricity Consumption Shifting by Prosumer Cooperatives. In *Proceedings of the European Conference on Artificial Intelligence (ECAI)*, 2016.
- M. Alam, S. D. Ramchurn, and A. Rogers. Cooperative Energy Exchange for the Efficient Use of Energy and Resources in Remote Communities. *Proceedings of the 12th International Conference on Autonomous Agents in Multiagent Systems (AAMAS)*, pages 731–738, 2013. URL <http://eprints.soton.ac.uk/346637/>.
- S. Albers. Energy-efficient algorithms. *Communications of the ACM*, 53(5):86–96, 2010.
- M. Aleksandrov and T. Walsh. Pure Nash equilibria in online fair division. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 42–48, 2017. ISBN 9780999241103. doi: 10.24963/ijcai.2017/7.
- M. Aleksandrov, H. Aziz, S. Gaspers, and T. Walsh. Online Fair Division: Analysing a Food Bank problem. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2540–2546, 2015.
- M. Alonso, H. Amaris, J. G. Germain, and J. M. Galan. Optimal charging scheduling of electric vehicles in smart grids by heuristic algorithms. *Energies*, 7(4):2449–2475, 2014. ISSN 19961073. doi: 10.3390/en7042449.
- E. J. Anderson and A. B. Philpott. Optimal Offer Construction in Electricity Markets. *Mathematics of Operations Research*, 27:82–100, 2002.

- G. Aneiros, R. Cao, J. M. Vilar-Fernandez, and A. Munoz-San-Roque. Functional Prediction for the Residual Demand in Electricity Spot Markets. In *Recent Advances in Functional Data Analysis and Related Topics*, pages 9–15. 2011. ISBN 978-3-7908-2735-4. doi: 10.1007/978-3-0348-8374-0. URL <http://www.springerlink.com/index/10.1007/978-3-7908-2736-1>.
- G. Aneiros, J. M. Vilar, R. Cao, and A. Munõz San Roque. Functional prediction for the residual demand in electricity spot markets. *IEEE Transactions on Power Systems*, 28(4):4201–4208, 2013. ISSN 08858950. doi: 10.1109/TPWRS.2013.2258690. URL <http://www.jstor.org/stable/4132848>.
- O. Ardakanian, S. Keshav, and C. Rosenberg. Real-Time Distributed Control for Smart Electric Vehicle Chargers: From a Static to a Dynamic Study. *IEEE Transactions on Smart Grid*, 5(5):2295–2305, 2014. ISSN 1949-3053. doi: 10.1109/TSG.2014.2327203. URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6871422>.
- R. J. Aumann. The Core of a Cooperative Game Without Side Payments. *Transactions of the American Mathematical Society*, 98(3):539–552, 1961.
- E. Baeyens, E. Y. Bitar, P. P. Khargonekar, and K. Poolla. Cooperative Aggregation of Wind Power. *IEEE Transactions on Power Systems*, 28(4):1–11, 2013.
- A. Baillo, M. Ventosa, M. Rivier, and A. Ramos. Optimal offering strategies for generation companies operating in electricity spot markets. *IEEE Transactions on Power Systems*, 19(2):745–753, 2004. ISSN 08858950. doi: 10.1109/TPWRS.2003.821429.
- A. Baillo, S. Cerisola, J. Fernandez-Lopez, and R. Bellido. Strategic bidding in electricity spot markets under uncertainty: a roadmap. *IEEE PES General Meeting General Meeting*, pages 1–8, 2006. doi: 10.1109/PES.2006.1708895.
- T. Baroche, P. Pinson, R. L. G. Latimier., and H. B. Ahmed. Exogenous Approach to Grid Cost Allocation in Peer-to-Peer Electricity Markets. pages 1–8, 2018. URL <http://arxiv.org/abs/1803.02159>.
- G. Benetti, M. Delfanti, T. Facchinetti, D. Falabretti, and M. Merlo. Real-Time Modeling and Control of Electric Vehicles Charging Processes. *IEEE Transactions on Smart Grid*, 6(3): 1375–1385, may 2015. ISSN 1949-3053. doi: 10.1109/TSG.2014.2376573.
- D. Berzal, J. I. De La Fuente, and T. Gómez. Building generation supply curves under uncertainty in residual demand curves for the day-ahead electricity market. In *2001 IEEE Porto Power Tech Proceedings*, volume 1, pages 350–355, 2001. ISBN 0780371399. doi: 10.1109/PTC.2001.964623.
- R. J. Bessa and M. Matos. The role of an aggregator agent for EV in the electricity market. *7th Mediterranean Conference and Exhibition on Power Generation, Transmission, Distribution and Energy Conversion (MedPower 2010)*, (November):126–126, 2010. ISSN 1949-3053. doi: 10.1049/cp.2010.0866. URL <http://digital-library.theiet.org/content/conferences/10.1049/cp.2010.0866>.
- R. J. Bessa and M. A. Matos. Economic and technical management of an aggregation agent for electric vehicles: A literature survey. *European Transactions on Electrical Power*, 22(3): 334–350, 2012. ISSN 1430144X. doi: 10.1002/etep.565.

- R. J. Bessa and M. A. Matos. Global against divided optimization for the participation of an EV aggregator in the day-ahead electricity market. Part I: Theory. *Electric Power Systems Research*, 95:309–318, 2013a. ISSN 03787796. doi: 10.1016/j.epsr.2012.08.013.
- R. J. Bessa and M. A. Matos. Global against divided optimization for the participation of an EV aggregator in the day-ahead electricity market. Part II: Numerical analysis. *Electric Power Systems Research*, 95:309–318, 2013b. ISSN 03787796. doi: 10.1016/j.epsr.2012.08.013. URL <http://dx.doi.org/10.1016/j.epsr.2012.08.013>.
- R. J. Bessa and M. A. Matos. Optimization models for EV aggregator participation in a manual reserve market. *IEEE Transactions on Power Systems*, 28(3):3085–3095, 2013c. ISSN 08858950. doi: 10.1109/TPWRS.2012.2233222.
- R. J. Bessa, M. A. Matos, F. J. Soares, and J. A. P. Lopes. Optimized bidding of a EV aggregation agent in the electricity market. *IEEE Transactions on Smart Grid*, 3(1):443–452, 2012. ISSN 19493053. doi: 10.1109/TSG.2011.2159632.
- A. Billh, K. Naik, and R. El-Shatshat. A Novel Online Charging Algorithm for Electric Vehicles Under Stochastic Net-Load. *IEEE Transactions on Smart Grid*, 9(3):1787–1799, 2018. ISSN 19493053. doi: 10.1109/TSG.2016.2599819.
- G. Binetti, A. Davoudi, D. Naso, B. Turchiano, and F. L. Lewis. Scalable Real-Time Electric Vehicles Charging With Discrete Charging Rates. *IEEE Transactions on Smart Grid*, 6(5): 2211–2220, sep 2015. ISSN 1949-3053. doi: 10.1109/TSG.2015.2396772.
- C. M. Bishop. *Pattern Recognition and Machine Learning*, volume 4. 2006. ISBN 9780387310732. doi: 10.1117/1.2819119. URL <http://www.library.wisc.edu/selectedtocs/bg0137.pdf>.
- J.-P. Bouchaud. Price Impact. In R. Cont, editor, *Encyclopedia of Quantitative Finance*, pages 1–11. John Wiley & Sons, 2010. URL <http://arxiv.org/abs/0903.2428>.
- S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2010. ISSN 10495258. doi: 10.1561/22000000016. URL <http://arxiv.org/abs/1408.2927>.
- J. Bremer and M. Sonnenschein. Estimating shapley values for fair profit distribution in power planning smart grid coalitions. *Lecture Notes in Computer Science*, 8076 LNAI:208–221, 2013. ISSN 03029743. doi: 10.1007/978-3-642-40776-5-19.
- A. M. Calmarza and J. de la Fuente. New Forecasting Method for the Residual Demand Curves using Time Series (ARIMA) Models. In *Proceedings of the 7th International Conference on Probabilistic Methods Applied to Power Systems*, pages 1–7, 2002.
- R. Cavallo. Optimal Decision-Making With Minimal Waste: Strategyproof Redistribution of VCG Payments. In *Proceedings of the 5th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 603–607, 2006. ISBN 1595933034. doi: 10.1145/1160633.1160790. URL cavallo-aamas-06.pdf.
- G. Chalkiadakis, V. Robu, R. Kota, A. Rogers, and N. Jennings. Cooperatives of Distributed Energy Resources for Efficient Virtual Power Plants. In *Proceedings of the 10th International*

- Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 787–794, 2011. URL <http://eprints.soton.ac.uk/271950/>.
- G. Chalkiadakis, E. Elkind, and M. Wooldridge. *Computational Aspects of Cooperative Game Theory*. 2012. ISBN 9781608456529. doi: 10.2200/S00355ED1V01Y201107AIM016. URL <http://www.morganclaypool.com/doi/abs/10.2200/S00355ED1V01Y201107AIM016>.
- P. Chander and H. Tulkens. The Core of an Economy with Multilateral Environmental Externalities. *International Journal of Game Theory*, 26(3):379–401, 1997. ISSN 0020-7276. doi: 10.1007/BF01263279. URL <http://ideas.repec.org/a/spr/jogath/v26y1997i3p379-401.html>
<http://rd.springer.com/content/pdf/10.1007/2F01263279>.
- A. C. Chapman, S. Mhanna, and G. Verbič. Cooperative Game Theory for Non-linear Pricing of Load-side Distribution Network Support. In *The 3rd IJCAI Algorithmic Game Theory Workshop*, 2017.
- Y. Chen, S. Kar, and J. M. Moura. Resilient Distributed Estimation Through Adversary Detection. *IEEE Transactions on Signal Processing*, 66(9):2455–2469, 2018. ISSN 1053587X. doi: 10.1109/TSP.2018.2813330.
- F. Chollet and E. Al. Keras, 2015. URL <https://keras.io>.
- K. Christidis and M. Devetsikiotis. Blockchains and Smart Contracts for the Internet of Things. *IEEE Access*, 4:2292–2303, 2016. ISSN 21693536. doi: 10.1109/ACCESS.2016.2566339.
- A. J. Conejo, J. Contreras, J. M. Arroyo, and S. De La Torre. Optimal response of an oligopolistic generating company to a competitive pool-based electric power market. *IEEE Transactions on Power Systems*, 17(2):424–430, 2002. ISSN 08858950. doi: 10.1109/TPWRS.2002.1007913.
- R. W. Conway, W. L. Maxwell, and L. W. Miller. *Theory of scheduling*. Courier Corporation, 2003.
- R. K. Dash, N. R. Jennings, and D. C. Parkes. Mechanism Design : A Call to Arms. *IEEE Intelligent Systems*, pages 40–47, 2003.
- S. De La Torre, J. M. Arroyo, A. J. Conejo, and J. Contreras. Price maker self-scheduling in a pool-based electricity market: A mixed-integer LP approach. *IEEE Transactions on Power Systems*, 17(4):1037–1042, 2002. ISSN 08858950. doi: 10.1109/TPWRS.2002.804945.
- M. De Weerd, M. Albert, V. Conitzer, and K. Van Der Linden. Complexity of scheduling charging in the smart grid. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 4736–4742, 2018. ISBN 9780999241127.
- Department of Energy and Climate Change. Smarter Grids: The Opportunity. (December): 30, 2009. URL http://www.decc.gov.uk/en/content/cms/meeting/_energy/network/strategy/strategy.aspx.
- J. Du, J.-T. Leung, and C. S. Wong. Minimizing the number of late jobs with release time constraint. *Journal of Combinatorial Mathematics and Combinatorial Computing*, 11:97–107, 1992.
- ECCC. 2020 Renewable Heat and Transport Targets. Technical report, 2016. URL <https://www.publications.parliament.uk/pa/cm201617/cmselect/cmenergy/173/173.pdf>.

- EIA. Monthly Energy Review. Technical report, 2016. URL <http://www.eia.gov/totalenergy/data/monthly/pdf/sec11{ }5.pdf>.
- Electricity Statistics. Energy Trends: Electricity. pages 1–10, 2017. URL <https://www.gov.uk/government/uploads/system/uploads/attachment{ }data/file/604090/Electricity.pdf>.
- EUPHEMIA. EUPHEMIA Public Description. Technical report, 2015. URL <https://www.nordpoolspot.com/globalassets/download-center/pcr/euphemia-public-documentation.pdf>.
- European Commission. Energy Roadmap 2050. Technical report, 2011. URL <https://ec.europa.eu/energy/en/topics/energy-strategy-and-energy-union/2050-energy-strategy>.
- J. Feigenbaum and S. Shenker. Distributed algorithmic mechanism design. In *Proceedings of the 6th international workshop on Discrete algorithms and methods for mobile computing and communications (DIALM)*, page 1, New York, New York, USA, 2002. ACM Press. ISBN 1581135874. doi: 10.1145/570810.570812.
- S. E. Fleten and E. Pettersen. Constructing bidding curves for a price-taking retailer in the Norwegian electricity market. *IEEE Transactions on Power Systems*, 20(2):701–708, 2005. ISSN 08858950. doi: 10.1109/TPWRS.2005.846082.
- Y. Funaki and T. Yamato. The core of an economy with a common pool resource: A partition function form approach. *International Journal of Game Theory*, 28(2):157–171, 1999. ISSN 0020-7276. doi: 10.1007/s001820050010.
- M. D. Galus, M. Zima, and G. Andersson. On integration of plug-in hybrid electric vehicles into existing power system structures. *Energy Policy*, 38(11):6736–6745, 2010. ISSN 03014215. doi: 10.1016/j.enpol.2010.06.043.
- L. Gan, U. Topcu, and S. H. Low. Optimal Decentralized Protocols for Electric Vehicle Charging. *IEEE Transactions on Power Systems*, 28(2):940–951, 2013. ISSN 0885-8950. doi: 10.1109/CDC.2011.6161220.
- J. Garcia-Gonzalez, J. Roman, J. Barquin, and A. Gonzalez. Strategic Bidding in Deregulated Power Systems. In *Proceedings of the 13th Power Systems Computation Conference*, pages 258–264, 1999.
- I. P. Gent, C. Jefferson, and I. Miguel. Minion: A Fast Scalable Constraint Solver. *Proceedings of the 17th European Conference on Artificial Intelligence (ECAI)*, 141:98–102, 2006. ISSN 09226389. doi: <http://dl.acm.org/citation.cfm?id=1567016.1567043>.
- E. H. Gerding, V. Robu, S. Stein, and D. Parkes. Online mechanism design for electric vehicle charging. *Proceedings of the 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 811–818, 2011. URL <http://eprints.soton.ac.uk/271907/>.
- E. H. Gerding, S. Stein, V. Robu, D. Zhao, and N. R. Jennings. Two-Sided Online Markets for Electric Vehicle Charging. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 989–996, 2013.

- E. H. Gerding, S. Stein, S. Ceppi, and V. Robu. Online mechanism design for vehicle-to-grid car parks. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 286–293, 2016.
- E. H. Gerding, A. Perez-Diaz, H. Aziz, S. Gaspers, A. Marcu, N. Mattei, and T. Walsh. Fair Online Allocation of Perishable Goods and its Application to Electric Vehicle Charging. In *Proceedings of the 28th International Conference on Artificial Intelligence (IJCAI)*, 2019.
- A. V. Goldberg and R. E. Tarjan. A New Approach to the Maximum-Flow Problem. *Journal of the Association for Computing Machinery*, 35(4):921–940, 1988.
- M. Gonzalez Vaya and 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):2375–2385, sep 2015. ISSN 0885-8950. doi: 10.1109/TPWRS.2014.2363159. URL <http://ieeexplore.ieee.org/document/6935028/>.
- M. Guo and V. Conitzer. Undominated VCG Redistribution Mechanisms. *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1039–1046, 2008. ISSN 15582914.
- M. Guo and V. Conitzer. Worst-case optimal redistribution of VCG payments in multi-unit auctions. *Games and Economic Behavior*, 67(1):69–98, 2009. ISSN 08998256. doi: 10.1016/j.geb.2008.06.007.
- M. Guo and V. Conitzer. Optimal-in-expectation redistribution mechanisms. *Artificial Intelligence*, 174(5-6):363–381, 2010. ISSN 00043702. doi: 10.1016/j.artint.2009.12.003.
- K. Hayakawa, E. H. Gerding, S. Stein, and T. Shiga. Online mechanisms for charging electric vehicles in settings with varying marginal electricity costs. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2610–2616, 2015. ISBN 9781577357384.
- R. Herranz, A. Muñoz San Roque, J. Villar, and F. A. Campos. Optimal demand-side bidding strategies in electricity spot markets. *IEEE Transactions on Power Systems*, 27(3):1204–1213, 2012. ISSN 08858950. doi: 10.1109/TPWRS.2012.2185960.
- E. Heydarian-Forushani, M. E. H. Golshan, and M. Shafie-khah. Flexible interaction of plug-in electric vehicle parking lots for efficient wind integration. *Applied Energy*, 179:338–349, 2016. ISSN 03062619. doi: 10.1016/j.apenergy.2016.06.145. URL <http://dx.doi.org/10.1016/j.apenergy.2016.06.145>.
- M. Honarmand, A. Zakariazadeh, and S. Jadid. Integrated scheduling of renewable generation and electric vehicles parking lot in a smart microgrid. *Energy Conversion and Management*, 86:745–755, 2014. ISSN 01968904. doi: 10.1016/j.enconman.2014.06.044. URL <http://dx.doi.org/10.1016/j.enconman.2014.06.044>.
- J. Horta, D. Kofman, and D. Menga. Novel paradigms for advanced distribution grid energy management. 2017. URL <http://arxiv.org/abs/1712.05841>.
- J. Hu, H. Morais, T. Sousa, and M. Lind. Electric vehicle fleet management in smart grids: A review of services, optimization and control aspects. *Renewable and Sustainable Energy*

- Reviews*, 56:1207–1226, 2016. ISSN 18790690. doi: 10.1016/j.rser.2015.12.014. URL <http://dx.doi.org/10.1016/j.rser.2015.12.014>.
- L. Huang, A. D. Joseph, B. Neslson, B. I. P. Rubinstein, and J. D. Tygar. Adversarial machine learning. In *Proceedings of the 4th ACM workshop on Security and artificial intelligence*, pages 43–58, 2011. ISBN 9781450310031. doi: 10.1109/MIC.2011.112.
- International Energy Agency. Global EV Outlook 2016: Beyond one million electric cars. Technical report, 2016. URL [http://www.eia.gov/forecasts/aeo/pdf/0383\(2016\).pdf](http://www.eia.gov/forecasts/aeo/pdf/0383(2016).pdf).
- E. Jones, E. Oliphant, and P. Peterson. SciPy: Open Source Scientific Tools for Python, 2001. URL <http://www.scipy.org/>.
- W. Kempton, J. Tomic, S. Letendre, A. Brooks, and 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):95, 2001. URL <http://escholarship.org/uc/item/0qp6s4mb.pdf>.
- D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- T. K. Kristoffersen, K. Capiion, and P. Meibom. Optimal charging of electric drive vehicles in a market environment. *Applied Energy*, 88(5):1940–1948, 2011. ISSN 03062619. doi: 10.1016/j.apenergy.2010.12.015. URL <http://dx.doi.org/10.1016/j.apenergy.2010.12.015>.
- A. Kurakin, I. Goodfellow, and S. Bengio. Adversarial Machine Learning at Scale. *Arxiv*, pages 1–17, 2016. URL <http://arxiv.org/abs/1611.01236>.
- E. L. Lawler. A dynamic programming algorithm for preemptive scheduling of a single machine to minimize the number of late jobs. *Annals of Operations Research*, 26(1):125–133, 1990.
- C. Le Floch, F. Belletti, S. Saxena, A. M. Bayen, and S. Moura. Distributed optimal charging of electric vehicles for demand response and load shaping. *Proceedings of the IEEE Conference on Decision and Control*, 54rd IEEE(Cdc):6570–6576, 2015. ISSN 07431546. doi: 10.1109/CDC.2015.7403254.
- C. Le Floch, F. Belletti, and S. Moura. Optimal Charging of Electric Vehicles for Load Shaping: A Dual-Splitting Framework With Explicit Convergence Bounds. *IEEE Transactions on Transportation Electrification*, 2(2):190–199, 2016. ISSN 2332-7782. doi: 10.1109/TTE.2016.2531025. URL <http://ieeexplore.ieee.org/document/7410071/>.
- Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, may 2015. ISSN 0028-0836. doi: 10.1038/nature14539. URL <http://www.nature.com/articles/nature14539>.
- M. Liao and A. Chakraborty. A Round-Robin ADMM algorithm for identifying data-manipulators in power system estimation. *Proceedings of the American Control Conference*, 2016-July:3539–3544, 2016. ISSN 07431619. doi: 10.1109/ACC.2016.7525462.
- M. Liao and A. Chakraborty. Identifying data-manipulators in power system mode estimation loops with noisy measurements. *Proceedings of the American Control Conference*, (2):2773–2778, 2017. ISSN 07431619. doi: 10.23919/ACC.2017.7963371.

- M. Liao and A. Chakraborty. Optimization Algorithms for Catching Data Manipulators in Power System Estimation Loops. *IEEE Transactions on Control Systems Technology*, pages 1–16, 2018. ISSN 10636536. doi: 10.1109/TCST.2018.2805294.
- Z. Liu, Q. Wu, S. Huang, L. Wang, M. Shahidepour, and Y. Xue. Optimal Day-ahead Charging Scheduling of Electric Vehicles through an Aggregative Game Model. *IEEE Transactions on Smart Grid*, 3053(c):1–1, 2017. ISSN 1949-3053. doi: 10.1109/TSG.2017.2682340. URL <http://ieeexplore.ieee.org/document/7879192/>.
- Z. Liu, Q. Wu, S. Huang, L. Wang, M. Shahidepour, and Y. Xue. Optimal Day-ahead Charging Scheduling of Electric Vehicles through an Aggregative Game Model. *IEEE Transactions on Smart Grid*, 9(5):5173–5184, 2018. ISSN 1949-3053. doi: 10.1109/TSG.2017.2682340. URL <http://ieeexplore.ieee.org/document/7879192/>.
- F. Lombardi, L. Aniello, S. De Angelis, A. Margheri, and V. Sassone. A Blockchain-based Infrastructure for Reliable and Cost-effective IoT-aided Smart Grids. *Living in the Internet of Things Conference: Cybersecurity of the IoT*, pages 1–6, 2018. URL https://eprints.soton.ac.uk/417098/1/BlockIT_{_}conf.pdf.
- J. A. P. Lopes, F. J. Soares, and P. M. R. Almeida. Integration of Electric Vehicles in the Electric Power System. *Proceedings of the IEEE*, 99(1), 2010. ISSN 00189219. doi: 10.1109/JPROC.2010.2066250.
- Z. J. Ma, D. S. Callaway, and I. A. Hiskens. Decentralized Charging Control of Large Populations of Plug-in Electric Vehicles. *IEEE Transactions on Control Systems Technology*, 21(1):67–78, 2013. ISSN 0743-1546. doi: Doi10.1109/Tcst.2011.2174059.
- S. Maleki. *Addressing The Computational Issues of the Shapley Value With Applications in The Smart Grid*. PhD thesis, University of Southampton, 2015. URL <https://eprints.soton.ac.uk/383963/>.
- M. Maschler, B. Peleg, and L. S. Shapley. Geometric Properties of the Kernel, Nucleolus, and Related Solution Concepts. *Mathematics of Operations Research*, 4(4):303–338, 1979.
- J. Mattila, T. Seppala, C. Naucler, R. Stahl, M. Tikkanen, A. Badenlid, and J. Seppala. Industrial Blockchain Platforms: An Exercise in Use Case Development in the Energy Industry. 2016. URL <http://pub.etla.fi/ETLA-Working-Papers-43.pdf>.
- R. Mehta, D. Srinivasan, A. M. Khambadkone, J. Yang, and A. Trivedi. Smart charging strategies for optimal integration of plug-in electric vehicles within existing distribution system infrastructure. *IEEE Transactions on Smart Grid*, 9(1):299–312, 2018. ISSN 19493053. doi: 10.1109/TSG.2016.2550559.
- C. E. Metz. Basic principles of ROC analysis. *Seminars in Nuclear Medicine*, 8(4):283–298, 1978. ISSN 0001-2998. doi: [http://dx.doi.org/10.1016/S0001-2998\(78\)80014-2](http://dx.doi.org/10.1016/S0001-2998(78)80014-2).
- J. C. Mukherjee and A. Gupta. Distributed Charge Scheduling of Plug-In Electric Vehicles Using Inter-Aggregator Collaboration. *IEEE Transactions on Smart Grid*, 8(1):331–341, 2017. ISSN 19493053. doi: 10.1109/TSG.2016.2515849.

- E. Munsing and S. Moura. Cybersecurity in Distributed and Fully-Decentralized Optimization: Distortions, Noise Injection, and ADMM. *Arxiv*, 2018. URL <http://arxiv.org/abs/1805.11194>.
- E. Munsing, J. Mather, and S. Moura. Blockchains for decentralized optimization of energy resources in microgrid networks. *2017 IEEE Conference on Control Technology and Applications (CCTA)*, pages 2164–2171, 2017. doi: 10.1109/CCTA.2017.8062773. URL <http://ieeexplore.ieee.org/document/8062773/>.
- B. Murray. *Power Markets and Economics*. John Wiley & Sons, 2009.
- M. Mylrea and S. N. G. Gourisetti. Blockchain for smart grid resilience: Exchanging distributed energy at speed, scale and security. In *Proceedings of the 2017 Resilience Week*, pages 18–23, 2017. ISBN 9781509060559. doi: 10.1109/RWEEK.2017.8088642.
- M. Neaimeh, R. Wardle, A. M. Jenkins, J. Yi, G. Hill, P. F. Lyons, Y. Hübner, P. T. Blythe, and P. C. Taylor. A probabilistic approach to combining smart meter and electric vehicle charging data to investigate distribution network impacts. *Applied Energy*, 157:688–698, 2015. ISSN 0306-2619. doi: <http://dx.doi.org/10.1016/j.apenergy.2015.01.144>. URL <http://www.sciencedirect.com/science/article/pii/S0306261915001944>.
- N. Nisan and A. Ronen. Computationally feasible VCG mechanisms. *Journal of Artificial Intelligence Research*, 29:19–47, 2007. ISSN 10769757. doi: 10.1613/jair.2046.
- N. Nisan, E. Tardos, T. Wexler, and V. Vazirani. *Algorithmic Game Theory*. 2007. ISBN 0-521-87282-0. doi: 10.1145/1785414.1785439.
- J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, 2006. ISBN 0387987932. doi: 10.5194/bgd-11-12733-2014.
- M. Nunes, E. Gerding, F. McGroarty, and M. Niranjan. A comparison of multitask and single task learning with artificial neural networks for yield curve forecasting. *Expert Systems with Applications*, 119:362–375, 2019. ISSN 09574174. doi: 10.1016/j.eswa.2018.11.012. URL <https://doi.org/10.1016/j.eswa.2018.11.012>.
- G. O’Brien, A. El Gamal, and R. Rajagopal. Shapley Value Estimation for Compensation of Participants in Demand Response Programs. *IEEE Transactions on Smart Grid*, 6(6):2837–2844, 2015. ISSN 1949-3053. doi: 10.1109/TSG.2015.2402194.
- Ofgem and Department of Energy and Climate Change. Smart Grid Vision and Routemap. *Report Number: URN 14D / 056*, (February), 2014. doi: URN14D/056.
- OMIE. Historical data from the Iberian day-ahead market, 2017. URL <http://www.omie.es/en/inicio>.
- ONS. UK Perspectives 2016: Energy and emissions in the UK. Technical report, 2016. URL <http://visual.ons.gov.uk/uk-perspectives-2016-energy-and-emissions-in-the-uk/>.
- G. Pasaoglu, D. Fiorello, A. Martino, G. Scarcella, A. Alemanno, A. Zubaryeva, and C. Thiel. *Driving and parking patterns of European car drivers - a mobility survey*. 2012. ISBN 9789279277382. doi: 10.2790/7028.

- G. Pasaoglu, D. Fiorello, L. Zani, A. Martino, A. Zubaryeva, and C. Thiel. *Projections for Electric Vehicle Load Profiles in Europe Based on Travel Survey Data Contact information*, volume 1. 2013. ISBN 9789279303876. doi: 10.2790/24108.
- B. Peleg and P. Sudhölter. *Introduction to the Theory of Cooperative Games*. 2007. ISBN 9783540729440.
- Q. Peng and S. H. Low. Distributed algorithm for optimal power flow on a radial network. In *Proceedings of the 53rd IEEE Conference on Decision and Control*, pages 167–172, 2014. ISBN 978-1-4673-6090-6. doi: 10.1109/CDC.2014.7039376. URL <http://ieeexplore.ieee.org/document/7039376/>.
- M. V. Pereira, S. Granville, M. H. C. Fampa, R. Dix, and L. A. Barroso. Strategic bidding under uncertainty: A binary expansion approach. *IEEE Transactions on Power Systems*, 20(1):180–188, 2005. ISSN 08858950. doi: 10.1109/TPWRS.2004.840397.
- A. Perez-Diaz. Coordination of Electric Vehicle Aggregator Participation in the Day-Ahead Market. *Proceedings of the 17th International Conference on Autonomous Agents and Multi Agent Systems (Doctoral Consortium)*, pages 1768–1769, 2018. URL <http://dl.acm.org/citation.cfm?id=3237383.3237969>.
- A. Perez-Diaz, E. Gerding, and F. McGroarty. Coordination and payment mechanisms for electric vehicle aggregators. *Applied Energy*, 212:185–195, 2018a. ISSN 03062619. doi: 10.1016/j.apenergy.2017.12.036. URL <http://linkinghub.elsevier.com/retrieve/pii/S0306261917317518>.
- A. Perez-Diaz, E. Gerding, and F. McGroarty. Coordination of Electric Vehicle Aggregators: A Coalitional Approach. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018)*, pages 676–684, 2018b. doi: 10.5258/SOTON/D0413. URL www.ifaamas.org.
- A. Perez-Diaz, E. Gerding, and F. McGroarty. Decentralised Coordination of Electric Vehicle Aggregators. In *International Workshop on Optimization in Multiagent Systems (OptMAS-18)*, pages 1–15, 2018c.
- A. Perez-Diaz, E. Gerding, and F. McGroarty. Cathing Cheats: Detecting Strategic Manipulation in Distributed Optimisation of Electric Vehicle Aggregators. pages 1–27, 2018d. URL <http://arxiv.org/abs/1810.07063>.
- A. Perez-Diaz, E. H. Gerding, and F. McGroarty. Dataset for "Coordination of Electric Vehicle Aggregators: A Coalitional Approach", 2018e. URL <https://eprints.soton.ac.uk/417940/>.
- A. Perez-Diaz, E. H. Gerding, and F. McGroarty. Dataset for "Coordination and payment mechanisms for electric vehicle aggregators" article, dec 2018f. URL <https://eprints.soton.ac.uk/416291/>.
- A. Perez-Diaz, A. Augustin, M. Nunes, E. H. Gerding, and F. McGroarty. Forecasting Residual Supply Curves in Electricity Markets with Neural Networks. (*under review*), 2019a.

- A. Perez-Diaz, A. Marcu, E. H. Gerding, H. Aziz, S. Gaspers, N. Mattei, and T. Walsh. Dataset for "Fair Online Allocation of Perishable Goods and its Application to Electric Vehicle Charging", 2019b. URL <https://eprints.soton.ac.uk/430960/>.
- A. B. Philpott and E. Pettersen. Optimizing demand-side bids in day-ahead electricity markets. *IEEE Transactions on Power Systems*, 21(2):488–498, 2006. ISSN 08858950. doi: 10.1109/TPWRS.2006.873119.
- L. Pieltain Fernández, T. Gómez San Román, R. Cossent, C. Mateo Domingo, and P. Frías. Assessment of the impact of plug-in electric vehicles on distribution networks. *IEEE Transactions on Power Systems*, 26(1):206–213, 2011. ISSN 08858950. doi: 10.1109/TPWRS.2010.2049133.
- M. Pinedo. *Scheduling: Theory, Algorithms, and Systems*. Springer, 2012.
- C. Pop, T. Cioara, M. Antal, I. Anghel, I. Salomie, and M. Bertocini. Blockchain based decentralized management of demand response programs in smart energy grids. *Sensors*, 18(162), 2018. ISSN 14248220. doi: 10.3390/s18010162.
- J. Portela González, A. Muñoz San Roque, E. F. Sanchez-Úbeda, J. García-González, and R. González Hombrados. Residual Demand Curves for Modeling the Effect of Complex Offering Conditions on Day-Ahead Electricity Markets. *IEEE Transactions on Power Systems*, 32(1):50–61, 2017. ISSN 08858950. doi: 10.1109/TPWRS.2016.2552240. URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7450198>.
- R. Porter. Mechanism design for online real-time scheduling. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 61–70, 2004.
- POST. Trends in Energy Production. Technical report, 2015. URL <http://researchbriefings.parliament.uk/ResearchBriefing/Summary/POST-PN-0503>.
- W. Qi, Z. Xu, Z. J. M. Shen, Z. Hu, and Y. Song. Hierarchical coordinated control of plug-in electric vehicles charging in multifamily dwellings. *IEEE Transactions on Smart Grid*, 5(3):1465–1474, 2014. ISSN 19493053. doi: 10.1109/TSG.2014.2308217.
- S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings. Putting the 'smarts' into the smart grid. *Communications of the ACM*, 55(4):86, 2012. ISSN 00010782. doi: 10.1145/2133806.2133825. URL <http://dl.acm.org/citation.cfm?doid=2133806.2133825>.
- E. S. Rigas, S. D. Ramchurn, and N. Bassiliades. Managing Electric Vehicles in the Smart Grid Using Artificial Intelligence: A Survey. *IEEE Transactions on Intelligent Transportation Systems*, 16(4):1619–1635, 2015.
- V. Robu, E. H. Gerding, S. Stein, D. C. Parkes, A. Rogers, and N. R. Jennings. An online mechanism for multi-unit demand and its application to plug-in hybrid electric vehicle charging. *Journal of Artificial Intelligence Research*, 48:175–230, 2013. ISSN 10769757. doi: 10.1613/jair.4064.
- J. M. Roperó Ortega. Encuesta Movilidad Cotidiana. *Ministerio de Fomento*, (1):1–5, 2014. doi: 10.1007/s13398-014-0173-7.2.
- J. Sankaran. On finding the nucleolus of a n-person cooperative game. *International Journal of Game Theory*, 19:329–338, 1991.

- M. R. Sarker, Y. Dvorkin, and M. A. Ortega-Vazquez. Optimal participation of an electric vehicle aggregator in day-ahead energy and reserve markets. *IEEE Transactions on Power Systems*, 31(5):3506–3515, 2016. ISSN 08858950. doi: 10.1109/TPWRS.2015.2496551.
- P. Scott and S. Thiébaux. Dynamic Optimal Power Flow in Microgrids using the Alternating Direction Method of Multipliers. *Arxiv*, pages 1–8, 2014. URL <http://arxiv.org/abs/1410.7868>.
- M. Shafie-khah, E. Heydarian-Forushani, M. E. H. Golshan, P. Siano, M. P. Moghaddam, M. K. Sheikh-El-Eslami, and J. P. S. Catalão. Optimal trading of plug-in electric vehicle aggregation agents in a market environment for sustainability. *Applied Energy*, 162:601–612, 2016. ISSN 03062619. doi: 10.1016/j.apenergy.2015.10.134. URL <http://dx.doi.org/10.1016/j.apenergy.2015.10.134>.
- A. D. Shah. *On Coordinating Electricity Markets: Smart Power Scheduling for Demand Side Management and Economic Dispatch*. PhD thesis, Harvard University, 2012.
- C. Shao, X. Wang, X. Wang, C. Du, and B. Wang. Hierarchical Charge Control of Large Populations of EVs. *IEEE Transactions on Smart Grid*, 7(2):1147–1155, 2016. ISSN 19493053. doi: 10.1109/TSG.2015.2396952.
- L. S. Shapley. On balanced sets and cores. *Naval Research Logistics Quarterly*, 14(4):453–460, 1967. ISSN 00281441. doi: 10.1002/nav.3800140404. URL <http://doi.wiley.com/10.1002/nav.3800140404>.
- L. S. Shapley. Cores of complex games. *International Journal of Game Theory*, 1(1):11–26, 1971.
- H. Shi, V. R. Prasad, E. Onur, and I. Niemegeers. Fairness in wireless networks: Issues, measures and challenges. *IEEE Communications Surveys*, 16(1):5–24, 2014.
- J. Shneidman and D. C. Parkes. Overcoming rational manipulation in mechanism implementations. 2003.
- SMMT. Electric & Alternatively-Fuelled Vehicle Registrations, 2017. URL <https://www.smm.co.uk/2017/05/april-2017-ev-registrations/>.
- T. Solymosi. The kernel is in the least core for permutation games. *Central European Journal of Operations Research*, pages 795–809, 2014. ISSN 1435-246X. doi: 10.1007/s10100-014-0342-y. URL <http://dx.doi.org/10.1007/s10100-014-0342-y>.
- J. A. M. Sousa, F. Teixeira, and S. Faias. Impact of a price-maker pumped storage hydro unit on the integration of wind energy in power systems. *Energy*, 69:3–11, 2014. ISSN 03605442. doi: 10.1016/j.energy.2014.03.039. URL <http://dx.doi.org/10.1016/j.energy.2014.03.039>.
- T. Sousa, H. Morais, J. Soares, and Z. Vale. Day-ahead resource scheduling in smart grids considering Vehicle-to-Grid and network constraints. *Applied Energy*, 96:183–193, 2012. ISSN 03062619. doi: 10.1016/j.apenergy.2012.01.053. URL <http://dx.doi.org/10.1016/j.apenergy.2012.01.053>.
- S. Stein, E. H. Gerding, V. Robu, and N. R. Jennings. A Model-Based Online Mechanism with Pre-Commitment and its Application to Electric Vehicle Charging. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 4–8, 2012.

- S. Stoft. *Power System Economics: Designing Markets for Electricity*. Wiley, 2002. ISBN 9780471150404.
- W. Su and M. Y. Chow. Computational intelligence-based energy management for a large-scale PHEV/PEV enabled municipal parking deck. *Applied Energy*, 96:171–182, 2012. ISSN 03062619. doi: 10.1016/j.apenergy.2011.11.088. URL <http://dx.doi.org/10.1016/j.apenergy.2011.11.088>.
- P. Sulc, S. Backhaus, and M. Chertkov. Optimal Distributed Control of Reactive Power Via the Alternating Direction Method of Multipliers. *IEEE Transactions on Energy Conversion*, 29(4): 968–977, 2014. ISSN 0885-8969. doi: 10.1109/TEC.2014.2363196. URL <http://ieeexplore.ieee.org/document/6963439/>.
- B. Sun, Z. Huang, X. Tan, and D. H. Tsang. Optimal Scheduling for Electric Vehicle Charging with Discrete Charging Levels in Distribution Grid. *IEEE Transactions on Smart Grid*, 9(2): 624–634, 2018. ISSN 1949-3053. doi: 10.1109/TSG.2016.2558585. URL <http://ieeexplore.ieee.org/document/7460201/>.
- O. Sundström and C. Binding. Flexible charging optimization for electric vehicles considering distribution grid constraints. *IEEE Transactions on Smart Grid*, 3(1):26–37, 2012. ISSN 19493053. doi: 10.1109/TSG.2011.2168431.
- R. M. Thrall and W. F. Lucas. n-Person Games in Partition Function Form. *Naval Research Logistics (NRL)*, 10:281–298, 1963.
- E. Tsang. *Foundations of constraint satisfaction: the classic text*. Books on Demand, 1996. ISBN 0127016104.
- A. Ugedo, E. Lobato, A. Franco, L. Rouco, J. Fernandez-Caro, J. De-Benito, J. Chofre, and J. De-la Hoz. Stochastic model of residual demand curves with decision trees. *Proceedings of the 2003 IEEE Power Engineering Society General Meeting*, 2:979–984, 2004. doi: 10.1109/pes.2003.1270443.
- A. Ugedo, E. Lobato, A. Franco, L. Rouco, J. Fernandez-Caro, and J. Chofre. Strategic Bidding in Sequential Electricity Markets. *IEEE Proceedings - Generation, Transmission and Distribution*, 153(4):431–442, 2006. ISSN 13502360. doi: 10.1049/ip-gtd. URL <http://digital-library.theiet.org/content/journals/10.1049/ip-gtd{ }20040098>.
- S. Vázquez, P. Rodilla, and C. Batlle. Residual demand models for strategic bidding in European power exchanges: Revisiting the methodology in the presence of a large penetration of renewables. *Electric Power Systems Research*, 108:178–184, 2014. ISSN 03787796. doi: 10.1016/j.epsr.2013.11.005. URL <http://dx.doi.org/10.1016/j.epsr.2013.11.005>.
- J. Villar, A. Muñoz, E. F. Sánchez-Úbeda, A. Mateo, M. Casado, A. Campos, J. Maté, E. Centeno, S. Rubio, J. J. Marcos, and R. González. SGO: Management information system for strategic bidding in electrical markets. *2001 IEEE Porto Power Tech Proceedings*, 1: 414–419, 2001. doi: 10.1109/PTC.2001.964634.
- N. Vlassis. *A Concise Introduction to Multiagent Systems and Distributed Artificial Intelligence*. 2007. ISBN 9781598295269. doi: <http://dx.doi.org/10.2200/S00091ED1V01Y200705AIM002>.
- D. J. Wales. *Energy Landscapes*. Cambridge University Press, Cambridge, UK., 2003.

- Y. Wang, L. Wu, and S. Wang. A Fully-Decentralized Consensus-Based ADMM Approach for DC-OPF with Demand Response. *IEEE Transactions on Smart Grid*, 8(6):2637–2647, 2017. ISSN 19493053. doi: 10.1109/TSG.2016.2532467.
- C.-k. Wen, J.-c. Chen, J.-h. Teng, and S. Member. Decentralized Plug-in Electric Vehicle Charging Selection Algorithm in Power Systems. *IEEE Intelligent Systems*, 3(4):1779–1789, 2012. ISSN 1949-3053. doi: 10.1109/TSG.2012.2217761.
- R. Weron. *Modeling and Forecasting Electricity Loads and Prices: A Statistical Approach*. John Wiley & Sons, nov 2006. ISBN 978-0-470-05753-7.
- R. Weron. Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting*, 30(4):1030–1081, 2014. ISSN 01692070. doi: 10.1016/j.ijforecast.2014.08.008. URL <http://dx.doi.org/10.1016/j.ijforecast.2014.08.008>.
- M. Wooldridge. *An introduction to multiagent systems*. John Wiley & Sons, 2009.
- World Energy Council. Average electricity consumption per electrified household, 2017. URL <https://www.wec-indicators.enerdata.eu/household-electricity-use.html>.
- H. Wu, M. Shahidepour, A. S. Alabdulwahab, and A. Abusorrah. A Game Theoretic Approach to Risk-Based Optimal Bidding Strategies for Electric Vehicle Aggregators in Electricity Markets With Variable Wind Energy Resources. *IEEE Transactions on Sustainable Energy*, 7 (January):118, 2016. doi: 10.1109/TSTE.2015.2498200.
- Z. Xu, M. A. T. Figueiredo, and T. Goldstein. Adaptive ADMM with Spectral Penalty Parameter Selection. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 54, pages 718–727, 2017. URL <http://arxiv.org/abs/1605.07246>.
- Ya’an Liu and Xiaohong Guan. Purchase allocation and demand bidding in electric power markets. *IEEE Transactions on Power Systems*, 18(1):106–112, 2003. ISSN 0885-8950. doi: 10.1109/tpwrs.2002.807063.
- Z. Yang, K. Li, and A. Foley. Computational scheduling methods for integrating plug-in electric vehicles with power systems: A review. *Renewable and Sustainable Energy Reviews*, 51:396–416, 2015. ISSN 18790690. doi: 10.1016/j.rser.2015.06.007. URL <http://dx.doi.org/10.1016/j.rser.2015.06.007>.
- F. Yao, A. Demers, and S. Shenker. A scheduling model for reduced CPU energy. In *Proceedings of IEEE 36th Annual Foundations of Computer Science*, pages 374–382. IEEE, 1995.
- L. Yao, W. Lim, and T. Tsai. A Real-Time Charging Scheme for Demand Response in Electric Vehicle Parking Station. *IEEE Transactions on Smart Grid*, 3053(c):1–1, 2016. ISSN 1949-3053. doi: 10.1109/TSG.2016.2582749. URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7496956>.
- J. Y. Yong, V. K. Ramachandaramurthy, K. M. Tan, and N. Mithulananthan. A review on the state-of-the-art technologies of electric vehicle, its impacts and prospects. *Renewable and Sustainable Energy Reviews*, 49:365–385, 2015. ISSN 18790690. doi: 10.1016/j.rser.2015.04.130. URL <http://dx.doi.org/10.1016/j.rser.2015.04.130>.

- S. You, J. Hu, and C. Ziras. An Overview of Modeling Approaches Applied to Aggregation-Based Fleet Management and Integration of Plug-in Electric Vehicles . *Energies*, 9(11):968, 2016. ISSN 1996-1073. doi: 10.3390/en9110968. URL <http://www.mdpi.com/1996-1073/9/11/968>.
- J. J. Q. Yu, J. Lin, A. Y. S. Lam, and V. O. K. Li. Maximizing Aggregator Profit through Energy Trading by Coordinated Electric Vehicle Charging. In *IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2016. URL <http://arxiv.org/abs/1508.00663>.
- G. Zizzo, E. Riva Sanseverino, M. G. Ippolito, M. L. Di Silvestre, and P. Gallo. A Technical Approach to P2P Energy Transactions in Microgrids. *IEEE Transactions on Industrial Informatics*, 3203(c), 2018. ISSN 15513203. doi: 10.1109/TII.2018.2806357.