



TECHNICAL UNIVERSITY OF CRETE
School of Electrical and Computer Engineering

MASTER'S THESIS

A Novel Smart Grid Flexibility Aggregation Framework

Stavros Orfanoudakis

Thesis Committee

Associate Professor Georgios Chalkiadakis

Associate Professor Fotios Kanellos

Professor Michail G. Lagoudakis

Chania, December 2022

Abstract

The increasing number of Distributed Energy Resources (DERs) in the emerging Smart Grid, has created an imminent need for intelligent multiagent frameworks able to utilize these assets efficiently. Additionally, the constant depletion of fossil fuels and carbon dioxide emissions have rendered the transition to Smart Grid a definite necessity. Electric vehicles, Battery Energy Storage Systems (BESS), interruptible load users, and renewable energy generators, such as solar panels and wind turbines, are only a few of the most common examples of DERs that will help in the transition to Smart Grid. In particular, DERs are essential for the smooth operation of the future Smart Grid since DERs enable flexible loads to be utilized, hence improving the stability of the Grid and enabling consumer-side demand management.

To address the significant aforementioned problems, in this MSc thesis we propose a novel DER aggregation framework, encompassing a multiagent architecture and various types of mechanisms for the effective management and efficient integration of DERs in the Grid. One critical component of our architecture is the Local Flexibility Estimators (LFEs) agents, which are key for offloading the Aggregator from serious or resource-intensive responsibilities—such as addressing privacy concerns and predicting the accuracy of DER statements regarding their offered demand response services.

The proposed aggregation framework allows the formation of efficient and effective LFE cooperatives. To this end, we developed and deployed a variety of cooperative member selection mechanisms, including (a) scoring rules, and (b) (deep) reinforcement learning. We use data from the well-known PowerTAC simulator to systematically evaluate our framework in various scenarios based on Smart Grid settings, so the efficiency of the framework can be properly assessed. Our experiments verify its effectiveness for incorporating heterogeneous DERs into the Grid in an efficient manner—showing that the use of appropriate mechanisms results in higher payments for competent LFEs managed by the Aggregator.



ΠΟΛΥΤΕΧΝΕΙΟ ΚΡΗΤΗΣ

Σχολή Ηλεκτρολόγων Μηχανικών και Μηχανικών
Υπολογιστών

Μεταπτυχιακή Εργασία

**Ένα Καινοτόμο Πλαίσιο
Συνάθροισης της Ευελιξίας Έξυπνου
Ηλεκτρικού Δικτύου**

Σταύρος Ορφανουδάκης

Εξεταστική Επιτροπή

Αναπληρωτής Καθηγητής Γεώργιος Χαλκιαδάκης

Αναπληρωτής Καθηγητής Φώτιος Κανέλλος

Καθηγητής Μιχαήλ Λαγουδάκης

Χανιά, Δεκέμβριος 2022

Περίληψη

Ο αυξανόμενος αριθμός εισαγωγής ολοένα και περισσότερων Κατανεμημένων Ενεργειακών Πόρων (ΚΕΠ) στο μελλοντικό Έξυπνο Δίκτυο Ηλεκτροδότησης (ΕΔΗ) έχει δημιουργήσει την ανάγκη για την δημιουργία ευφών πολυπρακτορικών συστημάτων ικανά να αξιοποιήσουν αποτελεσματικά αυτούς τους ενεργειακούς πόρους. Επιπλέον, η επικείμενη εξάντληση των ορυκτών καυσίμων και οι αυξανόμενες εκπομπές διοξειδίου του άνθρακα έχουν καταστήσει τη μετάβαση στο Έξυπνο Δίκτυο μια πρώτιστος ερευνητική, αλλά και πρακτική, διαδικασία υψίστης σημασίας. Μερικά παραδείγματα ΚΕΠ που θα βοηθήσουν στην μετάβαση στο ΕΔΗ είναι: τα ηλεκτρικά οχήματα, τα συστήματα αποθήκευσης ενέργειας μπαταριών, οι καταναλωτές διακοπτόμενου φορτίου και οι γεννήτριες ανανεώσιμων πηγών ενέργειας, όπως τα φωτοβολταϊκά πάνελ και οι ανεμογεννήτριες. Ειδικότερα, τα ΚΕΠ είναι απαραίτητα για την ομαλή λειτουργία του μελλοντικού ΕΔΗ, καθώς τα ΚΕΠ επιτρέπουν τη χρήση ευέλικτων φορτίων, βελτιώνοντας έτσι τη σταθερότητα του Ηλεκτρικού Δικτύου και επιτρέποντας τη διαχείριση της ζήτησης από την πλευρά των καταναλωτών.

Για την αντιμετώπιση των προαναφερθέντων προβλημάτων, σε αυτή τη μεταπτυχιακή διατριβή προτείνουμε ένα νέο πλαίσιο συνάθροισης ενεργειακής ευελιξίας προερχόμενη από ΚΕΠ, που περιλαμβάνει μια πολυπρακτορική αρχιτεκτονική και διάφορους τύπους μηχανισμών για την αποτελεσματική διαχείριση και την ομαλή ενσωμάτωση των ΚΕΠ στο υπάρχον Δίκτυο Ηλεκτροδότησης. Ένα σημαντικό συστατικό της αρχιτεκτονικής μας είναι οι πράκτορες Εκτίμησης Τοπικής Ευελιξίας (ΕΤΕ), οι οποίοι βοηθούν στην μείωση των ευθυνών και των λειτουργιών του Συναθροιστή ευελιξίας (Aggregator), όπως η αντιμετώπιση προβλημάτων απορρήτου σχετικά με τα προσωπικά στοιχεία κατανάλωσης και συναλλαγών των ΚΕΠ, καθώς και η υπολογιστικά κοστοβόρα εκτίμηση της ευελιξίας του κάθε ΕΤΕ.

Το προτεινόμενο πλαίσιο συνάθροισης επιτρέπει τη δημιουργία συνασπισμών ΕΤΕ με στόχο την αποτελεσματική χρήση τους για τις ανάγκες του ΕΔΗ, αλλά και για την αύξηση των χρηματικών απολαβών τους. Για το σκοπό αυτό, αναπτύξαμε και εφαρμόσαμε μια ποικιλία μηχανισμών επιλογής μελών συνασπισμού, συμπεριλαμβανομένων μεταξύ άλλων κανόνων βαθμολόγησης και αξιολόγησης μέσω βαθιάς ενισχυτικής μάθησης. Για να επιτύχουμε μια συστηματική και ενδεδειγμένη πειραματική διαδικασία αξιολόγησης του προτεινόμενου καινοτόμου πλαισίου συνάθροισης ευελιξίας ΚΕΠ, σχεδιάσαμε και αξιολογήσαμε την υλοποίηση μας σε διάφορα σενάρια που βασίζονται σε περιβάλλοντα ΕΔΗ, έτσι ώστε να προκύπτει μια σφαιρική εικόνα σχετικά με την αποτελεσματικότητα του πλαισίου μας. Η πειραματική μας διαδικασία αξιοποιεί δεδομένα από τον γνωστό προσομοιωτή για το ΕΔΗ, PowerTAC.

Εν τέλει, τα πειράματά μας επαληθεύουν την αποτελεσματικότητά του καινοτόμου πλαισίου συνάθροισης ευελιξίας που προτείνουμε για την ενσωμάτωση ετερογενών ΚΕΠ στο Δίκτυο Ηλεκτροδότησης, δεικνύοντας ότι η χρήση των κατάλληλων μηχανισμών έχει ως αποτέλεσμα υψηλότερες πληρωμές για τα αρμόδια ΕΤΕ που διαχειρίζεται ο προτεινόμενος Συναθροιστής ενεργειακής ευελιξίας.

Acknowledgements

I would like to convey my sincere gratitude and thanks to my supervisor Prof. Georgios Chalkiadakis for his consistent support and guidance throughout this Master's Program.

Furthermore, I would like to thank all members of the “inteLLigence” laboratory, especially Antonis and Andreas, for our excellent collaboration, from which I learned a lot over the past two years at the Technical University of Crete. Finally, I would like to thank Dina and Prof. Michail Zervakis for trusting me and allowing me to participate in exciting research projects.

Contents

Abstract	ii
Acknowledgements	iii
List of Figures	vi
List of Tables	ix
List of Abbreviations	x
1 Introduction	1
1.1 Thesis Contributions	2
1.2 Thesis Outline	3
2 Background & Related Work	4
2.1 The Smart Grid	4
2.1.1 Differences with the traditional Grid	4
2.1.2 Benefits of the Smart Grid	5
2.1.3 Distributed Energy Resources (DERs)	6
2.2 Flexibility	6
2.3 Flexibility Traders in the Grid	7
2.3.1 The Role of the Aggregator	7
2.3.2 Challenges faced by the Aggregators	8
2.3.3 Existing Flexibility Market Models in Europe	9
2.4 Scoring Rules	9
2.4.1 Continuous Ranked Probability Score (CRPS)	10
2.5 Reinforcement Learning in the Smart Grid	10
2.5.1 Q-learning	11
2.5.2 Deep Reinforcement Learning (Deep-RL)	12
2.5.3 Deep Q-Network (DQN)	12
2.5.4 RL applications for the Smart Grid	13
2.6 PowerTAC: The Smart Grid simulator	14
2.6.1 PowerTAC Overview	14
2.6.2 Simulation Time	15
2.6.3 Brokers	15
2.6.4 Customer Models	17
2.6.5 Weather Reports	17
2.6.6 Day-Ahead Wholesale Market	18
2.6.7 Balancing Market and Balancing Fees	18

2.6.8	Distribution utility and Transmission Capacity Fees	18
2.7	Related Work	19
2.7.1	Aggregator Business models	19
2.7.2	Demand-Response Aggregators	20
2.7.3	Privacy concerns of Aggregators	21
2.7.4	Electric Vehicle Aggregators	22
2.7.5	Aggregators viewed in the context of Multiagent Systems and Cooperative Game Theory	23
3	A Novel DER Aggregator Framework	25
3.1	Local Flexibility Estimators (LFEs)	25
3.2	The Aggregator agent	27
3.3	Flexibility Estimation	27
3.4	Selection mechanisms	28
3.4.1	Simple Selection mechanism	29
3.4.2	Continuous Ranked Probability Score (CRPS)	29
3.4.3	Reinforcement Learning: DQN	30
3.5	Pricing Mechanisms	32
3.5.1	Prediction Accuracy mechanism	32
3.5.2	CRPS based mechanism	33
3.5.3	Simple Grid-to-Aggregator Payment	35
4	Experimental Setup	36
4.1	Experimental Scenarios	36
4.1.1	Simple Selection	38
4.1.2	Using CRPS only	38
4.1.3	DQN	39
4.1.4	Singleton LFEs	41
4.1.5	All LFEs	42
4.2	Technical properties of the LFEs	42
4.2.1	PowerTAC Simulator	43
4.2.2	LFEs Specifications	44
4.3	Static and Dynamic LFE accuracy	47
4.3.1	Static LFE accuracy	48
4.3.2	Dynamic LFE accuracy	49
5	Experimental Evaluation	51
5.1	Results: CRPS-based Grid Payments	52
5.1.1	Static LFE accuracy	53
5.1.2	Dynamic LFE accuracy	58
5.2	Results: Simple Grid-to-Aggregator Payment	65
5.2.1	Static LFE accuracy	66
5.2.2	Dynamic LFE accuracy	72
5.3	Summarizing the Results	78
6	Conclusions and Future Work	79
6.1	Future Work	80

List of Figures

2.1	The main differences between the classical Grid and the Smart Grid can be distinguished in five main levels [27]	5
2.2	A classic Reinforcement Learning environment [48]	11
2.3	Deep Reinforcement Learning [53]	12
2.4	Main PowerTAC components [62]	15
2.5	Available actions of a broker during a timeslot [62]	16
3.1	The component diagram of the proposed aggregator architecture consists of three levels; the Distributed Energy Resources, the Aggregator Level, and the Energy Markets	26
3.2	The architecture of the Reinforcement Learning model used for the selection of the best LFE cooperative.	31
3.3	The "Simple Scoring" accuracy factor graph for errors ranging in $[-1,1]$	33
3.4	The CRPS accuracy factor graph for errors ranging in $[-1,1]$	34
4.1	An activity diagram presenting the five different experimental scenarios developed for the evaluation of the proposed Aggregator framework	37
4.2	The activity diagram of Experimental Scenario: "Simple Selection method"	39
4.3	The activity diagram of Experimental Scenario: "Using CRPS only"	40
4.4	The activity diagram of Experimental Scenario: "Using DQN"	41
4.5	The activity diagram of Experimental Scenario: "Singleton LFEs"	42
4.6	The activity diagram of Experimental Scenario: "All LFEs"	43
4.7	The weather simulation data of the PowerTAC decomposed to four parameters temperature, wind speed, wind direction, and cloud cover	44
4.8	The flexibility of an LFE decomposed to EV, battery charge, interruptible loads, and production	46
4.9	An example of the total flexibility of all 12 LFEs of the Aggregator	47
4.10	The error Probability Density Function (PDF) of the prediction probability of every LFE	48
4.11	Standard Deviation of the LFEs' flexibility estimations in the "Static LFE accuracy" use case	49
4.12	The fluctuating Standard Deviation of the LFEs' flexibility estimations in the "Dynamic LFE accuracy" use case	50

4.13	An example of how the fluctuating standard deviation impacts the CRPS of the LFEs throughout the simulation	50
5.1	“Static LFE accuracy”: Total profits (€) of the LFEs in every experimental scenario	54
5.2	“Static LFE accuracy”: Comparison of the average payment (€) per MWh sold of every LFE in all the experimental scenarios	55
5.3	“Static LFE accuracy”: Comparison of the amount of money (€) LFEs got via the Aggregator and the Grid	56
5.4	“Static LFE accuracy”: Comparison of the Total LFE Energy sold(GWh) for every experimental scenario	57
5.5	Static LFE accuracy: Comparing the LFE participation percentages (in the Aggregator cooperative) for every experimental scenario	57
5.6	“Static LFE accuracy”: Comparing the flexibility selling channel (directly to Grid or via the Aggregator) and the total flexibility (GWh) sold	58
5.7	“Dynamic LFE accuracy”: Comparison of the amount of money (€) LFEs got via the Aggregator and the Grid	60
5.8	“Dynamic LFE accuracy”: Total profits (€) of the LFEs in every experimental scenario	61
5.9	“Dynamic LFE accuracy”: Comparison of the average payment(€) per MWh sold of every LFE	62
5.10	“Dynamic LFE accuracy”: Comparison of the Total LFE Flexibility sold(GWh) for every experimental scenario	63
5.11	“Dynamic LFE accuracy”: Comparing the LFE participation percentages (in the Aggregator cooperative) for every experimental scenario	63
5.12	“Dynamic LFE accuracy”: Comparing the flexibility selling channel (directly to Grid or via the Aggregator) and the total flexibility (GWh) sold	64
5.13	“Static LFE accuracy”: Total profits (€) of the LFEs in every experimental scenario	67
5.14	“Static LFE accuracy”: Comparison of the average payment (€) per MWh sold of every LFE in all the experimental scenarios	68
5.15	“Static LFE accuracy”: Comparison of the amount of money (€) LFEs got via the Aggregator and the Grid	69
5.16	“Static LFE accuracy”: Comparison of the Total LFE Energy sold(GWh) for every experimental scenario	70
5.17	Static LFE accuracy: Comparing the LFE participation percentages (in the Aggregator cooperative) for every experimental scenario	70
5.18	“Static LFE accuracy”: Comparing the flexibility selling channel (directly to Grid or via the Aggregator) and the total flexibility (GWh) sold	71
5.19	“Dynamic LFE accuracy”: Comparison of the amount of money (€) LFEs got via the Aggregator and the Grid	73

5.20	“Dynamic LFE accuracy”: Total profits (€) of the LFEs in every experimental scenario	74
5.21	“Dynamic LFE accuracy”: Comparison of the average payment(€) per MWh sold of every LFE	75
5.22	“Dynamic LFE accuracy”: Comparison of the Total LFE Flexibility sold(GWh) for every experimental scenario	75
5.23	“Dynamic LFE accuracy”: Comparing the LFE participation percentages (in the Aggregator cooperative) for every experimental scenario	76
5.24	“Dynamic LFE accuracy”: Comparing the flexibility selling channel (directly to Grid or via the Aggregator) and the total flexibility (GWh) sold	77

List of Tables

4.1	The technical characteristics of the 12 LFEs formulated for the experiments	45
5.1	The total profits of LFEs via the Aggregator and the Grid at the end of the simulations for every experimental scenario	53
5.2	The total profits of LFEs via the Aggregator and the Grid at the end of the simulations for every experimental scenario	59
5.3	The total profits of LFEs via the Aggregator and the Grid at the end of the simulations for every experimental scenario	66
5.4	The total profits of LFEs via the Aggregator and the Grid at the end of the simulations for every experimental scenario	72

List of Abbreviations

DER	Distributed Energy Resource
LFE	Local Flexibility Estimator
BESS	Battery Energy Storage System
DR	Demand Response
EV	Electric Vehicle
CRPS	Continuous Ranked Probability Score
PowerTAC	Power Trading Agent Competition
DSOs	Distribution System Operators
TSOs	Transmission System Operators
BRPs	Balance Responsible Parties
MAE	Mean Absolute Error
RL	Reinforcement Learning
MDP	Markov Decision Process

Chapter 1

Introduction

The depletion of fossil fuels has created an imminent need to deploy even more renewable energy power generators [1]. However, the state of the current energy grid makes it hard to efficiently optimize the performance of intermittent assets [2], such as solar panels and wind turbines; thus, the need for a “smarter” electricity grid has been created [3], [4]. The Smart Grid [5], [6], with its bidirectional electricity and information flow, is envisaged to deliver electrical power in very resourceful ways, and successfully exploit all the Distributed Energy Resources (DERs) that are continuously emerging. DERs are the various electricity supply or demand assets that are spread across the Grid; and which, however small, when combined, can enhance the Grid’s ability to seamlessly provide power, even if it largely originates from intermittent renewable energy sources.

Furthermore, recent developments regarding environmental policies and the emergence of a multitude of (distributed) energy markets have turned the attention of the electricity stakeholders to the research on DERs’ flexibility [7]. In the literature, *flexibility* is defined as the elasticity property of the DERs that can provide ancillary services to support the stability of the Grid [8]. Essentially, flexibility corresponds to the DERs’ ability to either offer produced/-stored energy or consumption reduction services to the Grid. Thus, many works are studying various ways to estimate the flexibility of DER assets [9], [10]; while the use of *aggregators* [11]–[13] is one of the most important mechanisms that utilize the flexibility of the DERs in the Smart Grid.

An aggregator is a mediator between DERs and the energy markets [11], with the mission to trade the flexibility obtained from the DERs by participating in the markets on behalf of the DERs’ owners [13]. Generally, aggregators offer stability guarantees to the Grid by delivering flexible loads. Currently, the existing legal frameworks of many countries, especially in the European Union (EU) and United States of America (USA), have been updated to allow the existence of such aggregator mechanisms [14], [15].

Nonetheless, there are still many open research topics regarding the functionality and mechanisms of the aggregator. For example, there are privacy concerns about the metering information constantly transmitted between the

DERs and the aggregator [16], [17]. Also, there is an interest in designing efficient Demand-Response (DR) aggregator mechanisms to improve the Smart Grid’s technical, economic, and market aspects [18], [19].

1.1 Thesis Contributions

In this thesis, we employ mechanism design and cooperative game theory ideas to propose a novel aggregator framework for the efficient integration of DERs in the Grid. Our framework provides an aggregation architecture along with mechanisms for its effective and efficient operation and aims to (and, as our experiments show, succeeds in) increase the flexibility offered by the aggregator to the Grid and the profits of the participating agents.

In our multiagent architecture, we introduce the so-called *Local Flexibility Estimators (LFEs)* that allow us to address some severe aggregator issues, such as privacy concerns and evaluation of the DERs’ flexibility accuracy. LFEs essentially serve as DER coalition managers, coordinating their members’ market activities. Given this, our work’s focus is the creation of efficient LFE cooperatives intending to increase the profits of every stakeholder. To achieve this, we have populated our framework with various selection mechanisms—some of which are *scoring rules* [20], and some are (deep) *reinforcement learning (RL)* [21] techniques. To the best of our knowledge, using RL for this purpose is entirely novel. An Aggregator agent can then use these selection mechanisms to decide which LFEs to include in its (flexibility) offers to the day-ahead markets. We provide a systematic experimental evaluation using data from the PowerTAC [22] simulator in various experimental scenarios that we formulated to test the different aspects of our aggregator framework. Last but not least, our work in this paper extends the PowerTAC simulator with aggregator-enabling functionality, which we designed and coded.

Arguably, our aggregator framework contributes to the smooth DERs’ integration into the Grid since (a) it allows smaller DERs to participate in the Smart Grid markets; (b) it selects which LFEs to participate in the energy transactions, increasing the expected accuracy of the promised offers, thus indirectly aiding the Grid’s stability; and (c) as verified via our experiments, the use of certain designed selection and pricing mechanisms leads to higher payments for (at least the competent in terms of prediction accuracy, but also sometimes the not-so-accurate) LFEs that the aggregator manages. Specific findings of our systematic experimentation are:

(1) The use of the truthfulness-incentivizing CRPS [20], [23] mechanism rewards effectively LFEs that have reliable flexibility estimates and results to the highest aggregator-to-LFEs payments for those LFEs, compared to those achieved with other selection mechanisms; or compared to assets’ earnings in “baseline settings” when they either participate in a “traditional” aggregator that manages all available DERs or when they trade only with the Grid.

(2) A *Simple Selection* mechanism we put forward ranks as a close second to CRPS. However, this mechanism is easier for non-specialists to understand. This result implies a trade-off between using a highly efficient yet complex scoring rule vs. a slightly less efficient yet easy-to-understand selection mechanism since using the latter can motivate the participation of small DERs (e.g., corresponding to small & medium-sized enterprises or private homes).

(3) The RL selection mechanisms were better than the aforementioned baseline settings only for certain settings in which DER accuracy does not fluctuate dynamically over time.

(4) Low-accuracy LFEs prefer to participate in larger LFE cooperatives so their errors can be balanced out by the team.

(5) The proposed DER aggregation framework rewards better the most accurate LFEs, and, most of the time, the other less-accurate LFEs both in “CRPS-based” Grid-to-Aggregator and in “Simple” Grid-to-Aggregator payments (defined in Section 3.5.3). However, when using “Simple” payments, the second best-rewarding method, after the CRPS selection method, is when every LFE is alone.

(6) CRPS is extremely cost-efficient in every scenario with “Simple” Grid-to-Aggregator payments. Furthermore, the CRPS selection method, along with the CRPS-based payments it deploys, achieves the best results in environments with highly fluctuating settings too.

1.2 Thesis Outline

The rest of this thesis is structured as follows. In Chapter 2, we present all the necessary background for the methods and mechanisms presented in this thesis. In detail, we provide the required theoretical background on Smart Grid components directly affecting the aggregator mechanisms. At the end of Chapter 2, we also present a review of the existing works that study various aspects of the aggregators and their mechanisms. In Chapter 3, we present the technical details about the proposed LFE agents, and also discuss the configuration of the various aspects of our proposed aggregator framework. Additionally, we provide a detailed explanation regarding the selection mechanisms, the pricing mechanisms we used, and how we calculated the total flexibility. In Chapter 4, we provide a detailed description of the experimental setup we used, along with the specifications of the PowerTAC simulator and of the experimental scenarios we designed. In Chapter 5, we present the experimental results for various experimental scenarios using many different methods. Finally, in Chapter 6, we discuss our contributions and outline future work that can extend some aspects of our proposed DER aggregation framework.

Chapter 2

Background & Related Work

In this Chapter, we provide the required theoretical background on Smart Grid and its components that directly affect the aggregator mechanisms we have developed and deployed. Moreover, we explain the role of flexibility traders in the Smart Grid and provide a thorough background on the PowerTAC simulator. Finally, we also present an extensive review of the existing works that study various aspects of the aggregators.

2.1 The Smart Grid

The Smart Grid is a modern and actively-researched electricity network allowing a two-way flow of electricity and data, with digital communications technology enabling it to detect, react, and pro-act to customer load changes to maintain the Grid's stability. The Smart Grid has, in principle, a self-healing capability and gives the opportunity to electricity customers to become active participants that provide ancillary services to the Grid while also making a profit [24].

2.1.1 Differences with the traditional Grid

There are many differences between the traditional Grid and the Smart Grid both on the operational and implementation levels, as seen in Figure 2.1. The most apparent modification is the replacement of the large power plants with many more distributed, smaller renewable power producers, hence every participant can contribute to the power generation problem, even with small-size solar panels located on rooftops of residential buildings. Furthermore, centralized energy markets, usually in the scope of a nation or a coalition of nations, will be superseded by decentralized markets that ignore boundaries. Therefore, electricity transmission systems based on large power lines and pipelines will be rendered obsolete since the Smart Grid will contain many small-scale transmission systems that locally satisfy the energy supply [25]. The previous entails that electricity and smart metering data will flow both from the producers to the consumers, similar to how the current Grid operates, and from consumers back to the Grid in the form of flexibility, hence contributing to the application of smaller independent Smart Grids.

Finally, the once ordinary consumers would be able to actively participate in the energy system by providing flexibility to the Grid [26].

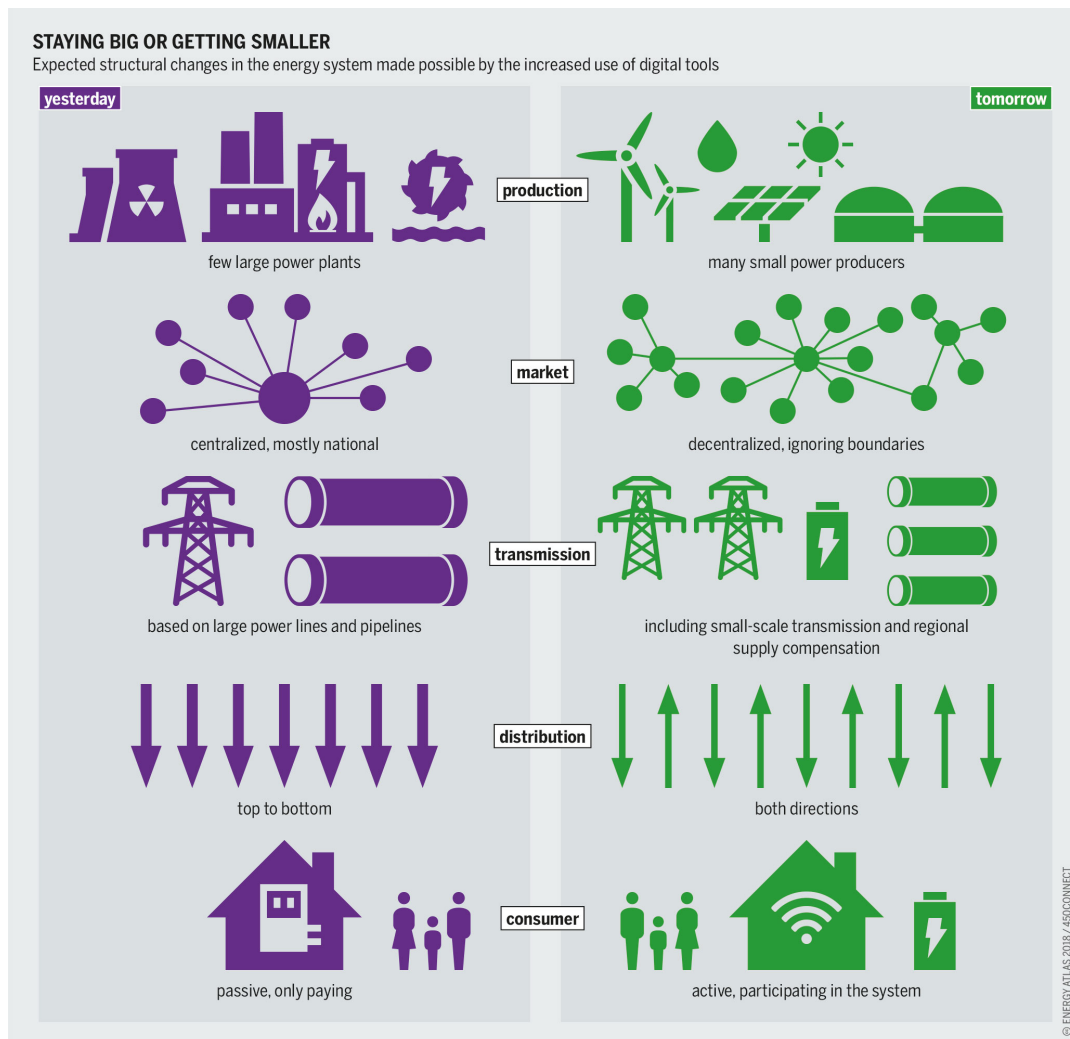


FIGURE 2.1: The main differences between the classical Grid and the Smart Grid can be distinguished in five main levels [27]

2.1.2 Benefits of the Smart Grid

Replacing the traditional Grid with a Smart Grid entails many benefits for the Grid itself and its participants. One of the most important benefits of the Smart Grid is the improved efficiency and reliability of the electricity supply since there will be many distributed electricity providers at a closer range that can provide ancillary services to the Grid quickly. On the one hand, Smart Grid's metering infrastructure will allow for new solutions for customers to be developed, so their electricity consumption will be optimized [28]. On the other hand, decentralized energy generation practically indicates that more and more energy gets generated (and stored) in various forms that are closer to the consumer that needs the power. Hence, if energy consumers generate their own energy more often, less money will be spent on larger transmission and distribution installations [29].

Smart Grid and its technologies will significantly reduce carbon emissions and help save energy. In particular, Smart Grid will help efficiently integrate more renewable energy systems into the existing network, thus playing an essential role in the deprecation of traditional, wasteful, and hazardous fossil fuel power plants. Additionally, carbon emissions, directly affecting the greenhouse effect and the rise in global temperature, will be reduced by the large-scale support of electric vehicles by the Smart Grid. Therefore, it could be argued that the development of the Smart Grid should be considered a priority because of its evident affiliation with the environment [30].

2.1.3 Distributed Energy Resources (DERs)

In general, DERs, are electricity supply or demand resources that are located across the Smart Grid. One of the most common DERs that can be found in the Smart Grid is the Battery Energy Storage System (BESS), these can be either in the form of dedicated batteries or embedded batteries within electric vehicles [31]. BESS are very important because they can store energy with the intention to shift the demand at other hours of need [32]. Furthermore, there are many interruptible load users; usually, these are households that have the ability to willingly limit their electricity usage at some specified points in the day in order to contribute to the demand-shifting process [33]. Also, there are many small-scale renewable energy generators, the most common are wind turbines and solar panels that can be located on the rooftops of both apartments and electric vehicles [34]. Last but not least, in the future Smart Grid, there will be many electric vehicles (EVs) that will have huge energy charging demands, but as mentioned earlier, they would be able to help in the demand-shifting process, making them very helpful for aggregator mechanisms.

2.2 Flexibility

In the past, flexibility has been defined as the ability of a power system to adapt its operation in response to variability or uncertainty by modifying electricity demand on the consumer side or power generation of the renewable energy systems [35]. Traditionally, flexibility can be obtained using the following four means: dispatchable power plants, demand response, energy storage, and interconnection between different parts of the Grid, probably located in different countries [36].

In particular, flexibility obtained from *dispatchable power plants* is admitted by traditional generation sources, such as diesel electricity generators, where production can be ramped up and down easily and in minimum time. However, usually dispatchable power plants are not environmentally friendly.

Flexibility acquired by *Demand Response* actions is one of the most studied methods in the literature of the past decade since it requires a significant effort with respect to optimization and planning. It is defined as the shifts in

electric usage by end-use customers from their typical consumption habits in response to changes in the price of electricity over time or incentive payments designed to induce lower electricity use at times of high wholesale market prices or when system reliability is endangered [37]. In detail, Smart Grid will support consumer tariffs or similar functionalities that allow utility operators to curtail the electricity usage of the agreeing customers to address extraordinary demand peaks.

Energy Storing methods refer to the process of converting energy from one form (mainly electrical energy) to a storable form and reserving it in various mediums; then, the stored energy can be converted back into electrical energy when needed [38]. This allows storing electricity to use it when Grid needs it to maintain its stability. In this way, electricity consumption and/or generation can be shifted in time to provide flexibility. Currently, there are four major groups of energy storing methods, these are: mechanical (pumped hydroelectric storage, compressed air energy storage, and flywheels), electrical (capacitors, supercapacitors), thermal (low temperature, high temperature), chemical (batteries, flow batteries, fuel cells) [13]. These technologies possess diverse characteristics in terms of energy density, power density, efficiency, energy capacity, volume, etc., making some more suitable depending on the application.

Finally, flexibility can be “imported” to a local Grid by using cross-border interconnections and networks. Usually, countries with significant energy surpluses can sell their electricity to other distant places in case it is cost-efficient for the recipient; however, proper facilities to transport huge amounts of energy loads are required. So, electricity can be transported from where it is produced at the lowest cost to where it is needed.

2.3 Flexibility Traders in the Grid

The newly established energy laws have made it easier for Flexibility traders to participate in the energy markets and make a profit by helping in the stability of the Grid. In general, flexibility traders are considered the Grid entities that buy or/and supply electric loads, usually in the form of flexibility Aggregators.

2.3.1 The Role of the Aggregator

In general, the main functionality of an aggregator is to perform demand-response and load-balancing to maximize the customers’ total profit while supporting the Grid’s stability. This is possible since an aggregator can be considered to be a reliable participant in the energy markets because of the huge flexibility of loads it can control [13].

The current regulations and the state of the energy markets allow existing actors, such as energy suppliers [39] and balance responsible parties (BRPs) [40],

to take the role of the aggregator and trade flexibility. An energy supplier has the ability to purchase and sell electricity for consumers by trading in electricity markets, while a BRP is responsible for submitting energy programs that indicate the net energy that is scheduled to be taken from/fed into the Grid for the next day [41]. Any variation between the energy designed to be taken from/fed into the Grid, and actual energy taken from/fed into the grid, is called the individual imbalance of the BRP, where the BRP needs to pay imbalance costs for their individual imbalances.

In addition to suppliers and BRPs, an independent actor, usually being organizations, that are not associated with a supplier or BRP and can also become an Aggregator. In the future, Distribution System Operators (DSOs) will be eligible, too, to form and support aggregator operations. However, at the moment, DSOs are heavily regulated, and they are not able to trade flexibility in the energy markets [42].

Summing up, suppliers, BRPs, and independent actors can take up the flexibility trader function to become an Aggregator. This is possible because, in practice, the aggregator's portfolio consists of DER assets usually owned by other stakeholders and utilized by the aggregator to implement its business model. However, there are also independent aggregators [43] which are not affiliated with any other entity like suppliers or balancing utilities. The framework we propose can support all three types of existing aggregators.

2.3.2 Challenges faced by the Aggregators

There are many open research challenges relating to the business model and the smooth operation of Aggregators. Initially, Aggregators are constrained by the number of contracts they can offer. For instance, as the number of contracts between the aggregator and the other actors increases, information exchange between them also becomes a serious issue. Actors may need information from the aggregator in order to enable accurate forecasting or calculating consumers' electricity bills. However, some of this information may contain commercial interests. Therefore, it is essential that the actors agree on what information will be disclosed [13]. The DER aggregation framework we propose in the next Chapter considers privacy a priority, hence it takes measures to address these privacy concerns.

The assessment of an Aggregator's business model can also be a serious challenge. Indeed, an Aggregator's main function is to make a profit when implementing a business model, so they need to ensure that the business model that will be incorporated is economically feasible. In order to evaluate whether a business model is profitable, it is important to consider all the financial relations the Aggregator has with the other actors. These financial relations may impact the economic feasibility of the business model. Not including these financial relations makes the assessment of economic feasibility incomplete, thus leading to wrong conclusions.

2.3.3 Existing Flexibility Market Models in Europe

The market for flexibility services was estimated, in 2019, to be worth more than £2.2 billion every year, with traditional procurement models typically used to contract large-scale, centrally-managed assets to dispatch their services on-demand to help balance the grid [44].

The market for local flexibility services is currently growing every day, but it is already recognized that assets that can adjust their energy generation and/or patterns of consumption in response to external signals will become part of the drive for a more decentralized energy system. In Europe, the digitalization of networks and smart metering implementations allow consumers and Distribution System Operators (DSOs) to know, almost in real-time, the load and generation patterns. In response to this situation, new digital platforms that implement new market models are arising. Under these market models, and by using these platforms, consumers and aggregators exploiting flexible distributed resources can provide services to DSOs and Transmission System Operators (TSOs) or trade energy between them [14]. In general, these platforms may differ widely between them in terms of the services they provide, the functions they perform, the required coordination between system operators (TSOs and DSOs), their ownership, or the interrelations with existing market factors.

An extensive review [14] conducted in early 2021 identified eighteen European initiatives that allow flexible resources connected to distribution networks to act as flexibility traders/providers. Furthermore, in early 2021, there were four more initiatives classified as aggregator platforms. In particular, two of them were represented as aggregator platforms specially devoted to clustering small flexible resources that were connected to the Aggregator so that this flexibility can be offered to TSO markets of ancillary services. Moreover, the other two developed platforms aimed to create new supplier business opportunities in the retail market by promoting peer-to-peer (P2P) transactions and taking advantage of solar panels installations located at prosumer¹ premises.

2.4 Scoring Rules

One of the goals of statistical analysis is to make forecasts for the future and provide suitable measures of the uncertainty associated with them. Therefore, forecasts should be probabilistic in nature, taking the form of probability distributions over future quantities or events [45]. Scoring rules provide summary measures for evaluating probabilistic forecasts by assigning a numerical score based on the predictive distribution and the event or value that materializes. In terms of elicitation, the role of scoring rules is to encourage the assessor to make careful assessments and to be honest [46]. In terms of

¹A prosumer is an individual who both consumes and produces electrical energy.

evaluation, scoring rules measure the quality of the probabilistic forecasts, reward probability assessors for forecasting jobs, and rank competing forecast procedures [20].

2.4.1 Continuous Ranked Probability Score (CRPS)

Probabilistic forecasts assign a probability to every possible future. Yet, all probabilistic forecasts are not equally accurate, and metrics are needed to assess the respective accuracy of specific probabilistic forecasts. Simple accuracy metrics such as Mean Absolute Error (MAE) are not directly applicable to probabilistic forecasts [20]. On the other hand, the Continuous Ranked Probability Score (CRPS) generalizes the MAE to the case of probabilistic forecasts. Along with the cross entropy, the CRPS is one of the most widely used accuracy metrics where probabilistic forecasts are involved [47]. We provide more details on the functioning and our use of CRPS in Sections 2.7.5, 3.4.2, and 3.5.2 of this thesis.

2.5 Reinforcement Learning in the Smart Grid

Reinforcement learning (RL) refers to a sub-field of machine learning that enables AI-based systems to take actions in a dynamic environment through trial and error to maximize the collective rewards based on the feedback generated for individual activities. In the RL context, feedback refers to a positive or negative notion reflected through rewards or punishments. In the most interesting and challenging cases, actions may affect not only the immediate reward but also the following states and, indirectly, all subsequent rewards. These two characteristics—trial-and-error search and delayed reward—are the two most important distinguishing features of reinforcement learning [21].

Interestingly, Reinforcement learning is not defined by characterizing learning methods but by characterizing a learning problem [21]. Any method well suited to solving that problem is considered a reinforcement learning method. Usually, RL problems are formulated using Markov decision processes where a learning agent interacts with its environment, trying to accomplish an objective. Indeed, such an agent must be able to sense the *state* of the environment to some extent and must be able to take *actions* that affect the state. The agent also must have a *goal*—and reward function—relating to the state of the environment. Therefore, an RL problem can be solved as a Markov Decision Process (MDP), as depicted in Figure 2.2.

One of the main challenges that arise in reinforcement learning and not in other kinds of learning is the trade-off between exploration and exploitation [21]. To obtain larger rewards, a reinforcement learning agent must select actions that it has tried in the past and found to be effective in producing rewards. But to discover such actions, it has to try actions that it has not selected before. The agent has to exploit what it already knows to obtain




FIGURE 2.2: A classic Reinforcement Learning environment [48]

the reward, but it also has to explore to make better action selections in the future. Therefore, the dilemma is that neither exploration nor exploitation can be pursued exclusively without failing at the task [49]. The agent must try a variety of actions and progressively favor those that appear to be best. On a stochastic task, each action must be tried many times to gain a reliable estimate of its expected reward.

2.5.1 Q-learning

Q-learning is a model-free RL method, which can learn based on experience in an unknown environment and explore the optimal strategy [50]. So, it can be applied to discrete MDP to find the optimal action strategy. The main parameters of the Q-learning model include the action a_t , state s_t and reward r_t , and an agent that participates in the Q-learning process. Q-learning is an off-policy algorithm that learns from random actions (greedy policy). Q in Q-learning refers to the quality of activities that maximize the rewards generated through the algorithmic process.

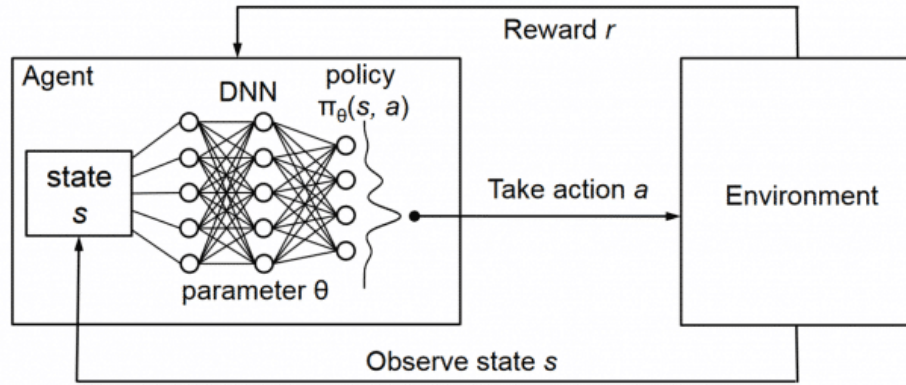
The Q-learning algorithm uses the Q-table to store the earned rewards. These values are updated using methods such as policy iteration and value iteration. Policy iteration refers to policy improvement or refinement through actions that amplify the value function. In detail, in a value iteration, the values of the value function are updated as shown in Equation 2.1:

$$Q^{new}(s_t, a_t) = Q(s_t, a_t) + \alpha \cdot (r_t + \gamma \cdot \max(Q(s_{t+1}, a_t)) - Q(s_t, a_t)) \quad (2.1)$$

where α is the learning rate, and γ is the discount factor. The training phase finishes after many iterations when the most action-state pairs (s, a) are explored, and the Q values have converged.

2.5.2 Deep Reinforcement Learning (Deep-RL)

Deep Learning is one of the best tools to handle unstructured environments with either continuous or highly dimensional action and/or state spaces. For instance, the state space can consist of continuous variables or even whole images, and the action space can be a continuous variable, such as the force required for an autonomous vehicle to steer. Traditional RL methods would not be able to assist in solving the aforementioned problems hence methods combining deep neural networks and classic RL algorithms have enabled the application of these methods in complex problems [51], [52].



Initially, the training process is split into M episodes, each with T steps. Secondly, the action space is explored in a greedy way using the probability ϵ , which can change with the course of the simulation. Then, after the action a_t is selected, the environment simulates the action so the next state s_{t+1} and the reward r_t are generated. Finally, in every step, with sufficient transitions stored in the buffer \mathcal{D} , a gradient descent step is performed to train the Q network.

Algorithm 1: Deep Q-Network (DQN) with Experience Replay Buffer [52]

Initialize replay memory \mathcal{D} to capacity N

Initialize action-value function Q with random weights

for $episode = 1$ to M **do**

Initialize sequences $s_1 = \{x_1\}$ and preprocessed sequenced $\phi_1 = \phi(s)$

for $t = 1$ to T **do**

With probability ϵ select a random action a_t

otherwise select $a_t = \max_a Q^*(\phi(s_t), a; \theta)$

Execute action a_t , and observe reward r_t and state s_{t+1}

Set s_{t+1} and preprocess $\phi_{t+1} = \phi(s_{t+1})$

Store transition (s_t, a_t, r_t, s_{t+1}) in \mathcal{D}

Sample random minibatch of transitions $(\phi_j, a_j, r_j, \phi_{j+1})$ from \mathcal{D}

$$\text{Set } y_j = \begin{cases} r_j & \text{for terminal } \phi_{j+1} \\ r_j + \gamma \cdot \max_{a'} Q(\phi_{j+1}, a'; \theta) & \text{for non terminal } \phi_{j+1} \end{cases}$$

Perform a gradient descent step on $(y_j - Q(\phi_j, a_j; \theta))^2$

2.5.4 RL applications for the Smart Grid

Smart Grid has many components that require complex decision-making processes to efficiently solve, such as energy trading, dynamic pricing, tariff design, EV charging, etc. Therefore, the application of Deep-RL algorithms in these complex environments was inevitable.

Zhang et al. [55] proposed a deep reinforcement learning-based energy double auction trading strategy. By utilizing the deep reinforcement learning algorithm, buyers and sellers can gradually learn the environment by treating the three elements: total supply, total demand, and their own supply and demand as states, in addition, regarding both bidding price and quantity as bidding strategy. Furthermore, Zhang et al. [56] also applied reinforcement learning (RL) to model the supply-demand relationship between power providers and consumers in a Smart Grid. The dynamic pricing problem of the Smart Grid was modeled as a discrete Markov decision process (MDP), and the decision process was solved by Q-learning.

In another recent work, Lu et al. [57] proposed a reinforcement learning-based decision system for assisting in the selection of electricity pricing plans,

which can minimize the electricity payment and consumption dissatisfaction for the individual Smart Grid end user. Furthermore, Qian et al. [58] proposed a deep reinforcement learning (DRL)-based EV charging navigation, aiming at minimizing the total travel time and the charging cost at EVCS. At first, they utilized the deterministic shortest charging route model (DSCRM) to extract feature states out of collected stochastic data, and then they formulated EV charging navigation as a Markov Decision Process (MDP) with an unknown transition probability.

However, other works also focus on the Smart Grid's retail part. Nastaran et al. [59] developed and designed an autonomous retailer in which a Sequence-to-Sequence (Seq2Seq) algorithm was employed to predict the customers' net demand. Furthermore, using Reinforcement Learning (RL), the proposed retailer designs tariff mechanisms based on other retailers' behavior and customers' load profiles. Finally, the proposed design of the retailer was evaluated on a retail market simulation platform called Power TAC, in which autonomous retailers compete in retail, wholesale, and balancing markets to maximize their profits.

2.6 PowerTAC: The Smart Grid simulator

The Power Trading Agent Competition (PowerTAC) [22] is an international trading agents competition conducted on a rich competitive economic simulation platform of future energy markets featuring several Smart Grid components (e.g., DERs, retail and wholesale energy markets, etc.). With the help of this simulator, researchers can better understand the behavior of future customer models as well as experiment with retail and wholesale market decision-making by creating competitive agents and benchmarking their strategies against each other. In this way, a host of useful information is extracted, which can be used by policymakers and industries in order to prepare for the upcoming market changes.

2.6.1 PowerTAC Overview

PowerTAC was developed in 2011, and since then, the developers have been updating various elements of the simulator, such as the customer models, to increase the accuracy of the simulation as well as to catch up with the latest scientific research. The PowerTAC vision consists of competitive agents that will harness the energy supply and demand of the simulation environment to make a profit [60]. Specifically, the "broker"-agents buy and sell energy through consumption tariffs with individual retail customers, such as households, retail stores, or even bigger enterprises, as well as electric vehicles. At the same time, the agents interact and trade within the wholesale market, which is a real-world replica of the European and North American wholesale energy markets [61].

Moreover, this simulation is designed to model the energy trading environment mainly from an economic and not a technical viewpoint. Through the years, the *main* elements of PowerTAC have not changed and can be seen in Figure 2.4 below. More details about these components will be further discussed in the rest of this section.

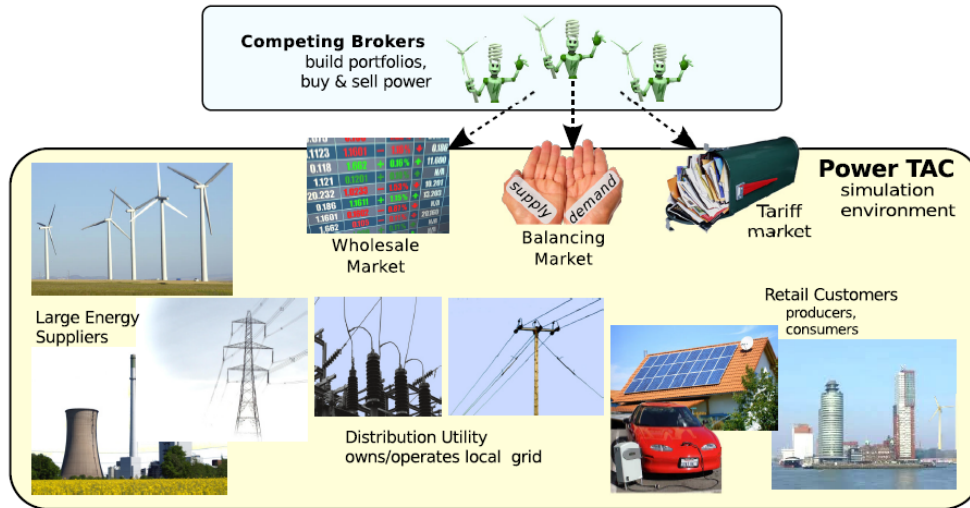


FIGURE 2.4: Main PowerTAC components [62]

2.6.2 Simulation Time

The simulation steps need to be discrete for the simulation to work, thus some discrete time blocks are created. These are called “timeslots”, and each one represents an hour of simulated time while each takes almost 5 seconds of real-time. Each game consists of at least 1440 timeslots (two months of simulated time and at least 2 hours of real-time per game). This means that in each PowerTAC game at any time, there is an active timeslot and a set of future timeslots for which the brokers can reserve and trade energy. The main objective of the brokers is to try to balance the demand and supply in each of the future timeslots to avoid getting monetary penalties.

2.6.3 Brokers

Broker agents are the real-life analogy to energy retailers and can take similar actions. In each timeslot, every agent can decide and perform any of the actions seen in Figure 2.5 below.

Specifically, these actions are:

- **Publish** new tariffs in the retail market.
- **Modify** existing tariff terms by revoking the old tariff and publishing a new one in its place.
- **Adjust** tariff prices, if tariff terms allow it.




FIGURE 2.5: Available actions of a broker during a timeslot [62]

- **Trade** in the Wholesale Market by placing *bids* or *asks* to sell or procure energy for future timeslots.
- **Curtail** demand by issuing economic control orders which apply to customers with incorruptible consumption tariffs.
- **Submit** balancing orders in the balancing market, which consists of offering controllable capacities for actual time balancing.

However, there is also a lot of information available to the brokers in order to help them make the correct decisions. In a few words, agents know about:

- **Most of the game parameters**, such as multipliers, are cost-specific parameters that are stable through the course of a game and are essential to calculate some values in advance.
- **Broker identities**, only the names of the competitor agents.

- **Bootstrap data**, consists of the net demand and the electricity usage of each customer as well as the bid and ask prices that are cleared in the Double Auction of the Wholesale Market through the duration of the bootstrap period. The bootstrap period is a time in which only the default broker is active and is used to initialize a game. The Bootstrap data is usually used by the agents to generate accurate electricity usage and customer models.
- **Weather reports/forecasts.**
- **Active Tariffs of each broker.**
- **Wholesale market clearing data.**
- **Wholesale market order-books.**
- **Total aggregate energy consumption/production.**
- **Every Transaction is private data for each agent**, such as cash positions, market positions, portfolio supply-demand, balancing, and distribution transactions.

2.6.4 Customer Models

One of the most important features of the PowerTAC simulator is that of the customer models. Customer models interact with the retail market to find the best tariffs and generate the values for the energy consumption or generation during a timeslot taking into account many different variables, such as weather, market state, etc. Each customer model is defined by its name, population, power type, controllable capacity, and multi-contracting ability. The customers' models themselves are of high complexity, so these could be the sole focus of other scientific research, therefore no more technical details will be presented in this work. More insights can be acquired by studying the corresponding chapter of PowerTAC definition [62].

2.6.5 Weather Reports

Another significant feature of PowerTAC is the weather reports. Customer models depend on weather reports to produce realistic production and consumption data. In each timeslot, a weather report for the current timeslot is sent to the brokers along with a weather forecast about the next 24 timeslots. This information can be used by the agents to predict the consumption and production curves of each DER asset since the customer models are directly affected by the weather. More specifically, weather data are drawn from real weather reports of the past, which means that during the competition, agents are not aware of the game location for obvious reasons.

2.6.6 Day-Ahead Wholesale Market

The wholesale market is where brokers primarily buy and sell energy for every available future timeslot (usually 24 future timeslots are enabled). Specifically, the PowerTAC wholesale market is a “day ahead” periodic double auction, which clears in each timeslot and tries to imitate the current core wholesale trading foundations like FERC, EEX, and NordPool (wholesale trading utilities in Europe and North America).

In real life, energy wholesale markets serve larger regions consisting of many energy providers and millions of customers. On the other hand, the PowerTAC can currently simulate only one city from such an area. Therefore, in order to make the *wholesale market* more realistic, three more entities were created to trade in that market along with the brokers.

The first entity is called “Grid Genco” and represents the wide population of generating facilities that can supply the simulated city. This entity is very important because it provides the market with a realistic supply curve deriving from statistics observed in North America’s MISO and PJM LMP markets. The second entity is the “Grid Buyer” which simulates the regional demand also based on real-life metrics drawn from a time series trained on two years of MISO North actual demand. So, this entity’s responsibilities are to buy enough energy so it can satisfy its energy demands. The third and last entity is created to provide the market with liquidity, using a stochastic behavior that mimics a population of buyers and speculators who are only engaging in the Wholesale Market of PowerTAC and not in the Retail Market.

2.6.7 Balancing Market and Balancing Fees

The balancing market is the real-life equivalent of an Independent System Operator (ISO) in the U.S. and a Transmission Systems Operator (TSO) in Europe. Its main responsibilities are to monitor the electricity grid and maintain balance by keeping voltage and frequency within some bounds. However, in PowerTAC, technical aspects like voltage and power factor are not taken into consideration, so the balancing market’s only responsibility is to balance the supply and demand in each timeslot by exercising capacity controls on behalf of the agents. In this way, when a broker fails to procure the required energy in time, balancing utility comes and finally charges the broker for the missed energy at a much higher price, acting as a penalty. That fee is called *Balancing Fee* and can vary from smaller charges to very high destructive penalties depending on the wholesale market prices.

2.6.8 Distribution utility and Transmission Capacity Fees

The Distribution Utility (DU) is primarily responsible for three different operations. The first one, as its name suggests, is to distribute energy to each customer while it charges each broker for using the Grid by implying distribution fees that are relevant to the energy transmitted through it. Second,

DU is responsible for issuing the Transmission Capacity Fees (TSF). These fees represent the amount of money a broker should pay for its customers' contribution to demand peaks. This means that when there is a demand peak, each broker will have to pay for a portion of the exceeding energy. TSF charges the *three* highest demand peaks at the end of a 168 timeslot period (1 week of simulated time).

In the current PowerTAC competition, these fees are the main challenge an agent faces when it tries to dominate the retail market [63]. The last function of the Distribution utility is to publish some default tariffs for the times when there are no published tariffs by any other broker.

2.7 Related Work

Here, we provide useful information about existing studies on various aspects of aggregators. Initially, general aggregator architectures and business models will be discussed. Right after, influential Demand-Response aggregator schemes will be mentioned, followed by works on privacy preservation for Aggregators. Next, EV Aggregators techniques will be thoroughly presented for the significant smooth integration of EVs in the Smart Grid since the number of EVs is constantly rising. Finally, a few works that combine the multiagent systems theory with Aggregators will be discussed.

2.7.1 Aggregator Business models

Aggregators and their mechanisms have become a very active research topic because of the increase of flexible loads in the Grid. Several studies about the business models of aggregators and the way they should operate exist.

In [12], the authors review the existing aggregator models' operational and economic aspects. In particular, in their analysis, they review the existing aggregator business models and split their economic analysis into two parts. The first is about the financial aspects and studies how much profit the aggregator makes while also considering how much money the consumers (DERs) save. The second is about the operational aspects of the business model and corresponds to how efficient the management of the consumers' assets by the aggregator is, thus resulting in higher financial gains. Finally, based on that analysis, they propose various strategies to be adopted by the existing aggregators.

Another work [7] proposes a hierarchical control framework that enables the provision of flexible services in power systems through aggregation entities. That work focuses on both the wholesale energy markets, such as day-ahead markets, and ancillary services, such as demand-response. More specifically, they put forward a versatile and scalable control framework that can support all kinds of operations, from the lower level, i.e., the user, to the distribution

and transmission levels. It is also notable that they establish some procedures to be taken in case of competing objectives when procuring ancillary services. They use historical data from residential buildings in the Netherlands, where they are usually equipped with solar panels and battery energy storage systems.

In order for a residential aggregator demand-response scheme to be efficient, optimized day-ahead customer incentive pricing (CIP) and load-shifting protocols should be established. Zheng et al. [64] have developed an aggregator-based resource allocation system using an artificial neural network and other optimization techniques. In detail, the day-ahead CIP is generated by the artificial neural network, which is trained on historical data, and the load-scheduling is addressed as a day-ahead optimization problem solved using a blocked sliding window technique and parallel computing.

In the Ph.D. dissertation of Carducci [65], tools and methods to support the research community, as well as industrial entities and policymakers, are introduced while also elaborating on the role that flexibility assets and aggregators will play in the future Smart Grid. Interestingly, [65] proposes, among other things, a flexibility aggregation architecture that uses so-called “minimum flexibility units” that operate at a local level and represent single flexibility assets, industrial microgrids, or multiple end users. These units, however, always correspond to a single meter and do not manage heterogeneous DERs nor evaluate their accuracy.

2.7.2 Demand-Response Aggregators

One of the main responsibilities of an aggregator is to provide flexibility in demand-response operations by aggregating the flexible loads of the individual DERs it controls. So, there are also many works about the demand response operations of Aggregators.

Demand Response Programs (DRP) are modern strategies for improving the control of the generation, distribution, and consumption of electrical energy and affecting the demand on the consumer side. Currently, multiple DRP strategies are proposed in the literature, with each trying to address a subset of the challenges; however, it is significant to perform a holistic review showing the implications of the objectives on each other. Ibrahim et al. [18] propose a consolidated model with assigned weights where intensive “what-if” scenarios will be applied to reach the optimal model settings. The authors decided to focus on the technical, economic, and marketing aspects with the consideration of architectures and business intelligence. Therefore, their mechanism assists in planning adequately and supporting the decision-making process of DR Aggregators.

Another work studying DR Aggregators is that of Gkatzikis et al. [11], focusing on efficient DR mechanisms for the residential sector, which can be an extremely hard task since each home user has negligible impact on the

market. They introduce a hierarchical market model that can be used in the Smart Grid, where a set of competing Aggregators are trading flexibility between the home users and the Grid. In that study, the designed Aggregators compete to sell cost-efficient DR services to the Grid while also trying to compensate the home users in a way that would incentivize them to alter their consumption patterns. In the end, they manage to capture the diverse objectives of the involved entities and, compared to flat pricing, guarantee notable profits for each. An interesting observation was that users who are willing to modify their consumption patterns do not derive maximum profits.

SEMIAH [66] is an aggregator framework designed to support European demand response programs. It uses a component-based architecture and focuses on the functionality of the virtual power plant. In particular, they focus on the impact of deploying an automated residential DR program on the quality and stability of a low-voltage grid. A similar study ([19]) proposes a two-stage interactive, responsive load scheduling model developed between a demand response aggregator and the distribution system operator. There, the distribution system operator has two objectives: to minimize the network losses and maximize the revenue while ensuring that the Aggregator gets payments relevant to the profits of the first stage. More specifically, they formulated the first stage as a linear programming optimization problem, while the second is considered to be a second-order cone programming problem.

Additionally, there are distributed algorithms developed for large-scale demand response aggregation. Initially, [67] propose a fully distributed optimal aggregation algorithm to provide peak-hours DR to the Grid through peer-to-peer cooperation of a large fleet of Distributed Energy Storage Devices. The main features of this algorithm include the ability to support a large number of assets, the prolongation of the battery life of the assets, private transactions between the assets and the Aggregator entity without disclosing each device's local private information, and elimination of single-point-of-failure in aggregation systems because of the distributed aspect of the algorithm. Another study focuses on the design of a fast-distributed algorithm for aggregating a large number of residential buildings with a mixture of discrete and continuous energy levels. One of the novelties is that the non-convex DR problem is decomposed in terms of households as opposed to devices, which allows incorporating more intricate couplings between energy storage devices, appliances, and distributed energy resources. Therefore, the method they proposed is a fast-distributed algorithm that can be applied to the double-smoothed dual function of the adopted DR model.

2.7.3 Privacy concerns of Aggregators

The smart meters deployed on the assets of the Smart Grid collect the users' power usage data regularly and upload it to the control centers so the data can be processed by the responsible entities. The control center can evaluate the supply and demand of the power grid through aggregated data from users and then dynamically adjust the power supply and price, etc. However,

since the Grid data collected from users may disclose the user's electricity usage habits and daily activities, privacy concerns have become a critical issue in Smart Grid data aggregation [68], [69].

To address the aforementioned issues, Gai et al. [17] propose an aggregation scheme with local differential privacy that can efficiently and practically estimate power supply and demand statistics while preserving any individual participant's privacy. In their analysis, they found that their scheme is computationally efficient and has minimum communication overhead.

Usually, privacy-preserving protocols and techniques either have a high computational overhead or are dependent on a single dedicated aggregator, making them prone to single points of failure. Wagh et al. [16] address the problem of cybersecurity by developing a distributed privacy preserving framework that aggregators can adapt. That framework is designed to prevent dedicated Aggregators and the electrical utility from linking the aggregated reading to a specific smart meter.

2.7.4 Electric Vehicle Aggregators

The gradual replacement of traditional vehicles with electric ones has created many new opportunities for Smart Grid research; however, there are also plenty of challenges concerning the efficient charging schedules of EVs. Therefore, EV aggregation schemes have been proposed to help smooth EVs' integration in the Smart Grid.

For example, a key element in achieving sustainable energy systems is the efficient utilization of EVs in parking lots with respect to the energy and reserve markets. [70] developed a model to derive optimal strategies of parking lots, as responsive demands, in both price-based and incentive-based demand response programs (DRP). The proposed model reflects the impacts of different DRPs on the operational behavior of parking lots and optimizes the participation level of parking lots in each DRP. There the uncertainties of EVs and the electricity market are also considered by using a stochastic programming approach. In their experiments, they found out that EV parking lots can benefit from selective participation in DRPs.

The problem of scheduling the charging of EVs in coordination with the dynamic electricity tariffs and load changes throughout the day is important. Goyal et al. [71] in order to satisfy the economy of both the EV owners and the aggregator, they mathematically formulated this scheduling problem and solved it by different optimization techniques to arrive at the optimal cost and revenue for the customer and the aggregator, respectively. In detail, they used heuristic techniques based on evolutionary optimization algorithms to solve the optimal charge scheduling problem. Another study ([72]) proposed a framework for real-time (RT) charging management of an electric vehicle aggregator that participates in electric energy and regulation markets. The developed models, which assign charging set points to the electric vehicles

(EVs) based on evolving EV charging priorities, are formulated as linear programs that can be solved very fast.

Another crucial aspect of EV Aggregators' research is about payment mechanisms that are employed. Perez-Diaz et al. [73] study the participation of large fleets of EVs in electricity day-ahead markets. Specifically, they consider a scenario where several independent and self-interested EV aggregators participate in the day-ahead market to purchase energy to satisfy their clients' driving needs. To address this issue, they employed techniques from mechanism design to develop a coordination mechanism that incentivizes self-interested EV aggregators to report their energy requirements truthfully to a third-party coordinator. This coordinator is then able to employ a day-ahead bidding algorithm to optimize the global bids on their behalf, extending the benefits of smart bidding to groups of competing EV aggregators. To ensure scalability and computational tractability, a novel price-maker day-ahead bidding algorithm is proposed, which is formulated in terms of simple energy requirement constraints.

Furthermore, the previous work is extended in [74], where they study the possibility of inter-aggregator cooperation trying to reduce electricity costs and lower the impact on electricity markets. In detail, they modeled the system as a coalitional game and proved that the resulting game is superadditive and balanced, hence having a non-empty core. However, the Shapley value is not guaranteed to lie in the cooperative game theoretic stability concept of the *core* [75] because the game is not convex. As an alternative, they proposed employing the payment mechanism provided by the least core.

2.7.5 Aggregators viewed in the context of Multiagent Systems and Cooperative Game Theory

Last but not least, several works in the multiagent systems literature [75], [76] approach the aggregation problem as a coalition formation one, proposing game theoretic and mechanism design solutions to form DER cooperatives so that these are able to participate in demand-response tasks.

To begin, Chalkiadakis et al. [77] proposed a game-theoretic solution to achieve cost-efficient integration of the many distributed energy resources (DERs) that are starting to emerge in Smart Grid. They designed cooperatives of rational, autonomous DER agents representing small-to-medium-size renewable electricity producers, which coalesce to sell their energy to the electricity grid profitably. By so doing, we help to counter the fact that individual DERs are often excluded from the wholesale energy market due to their perceived inefficiency and unreliability. Also, they proposed a pricing mechanism with certain desirable properties. Specifically, their mechanism guarantees that Cooperative Virtual Power Plants (CVPPs) have the incentive to truthfully report to the grid accurate estimates of their electricity production and that larger rather than smaller CVPPs form, thus improving CVPP efficiency and reliability. Most importantly, they provided a scheme to allocate payments

within the cooperative and show that, given this scheme and the pricing mechanism, the allocation is in the core, and, as such, no subset of members has a financial incentive to break away from the CVPP.

Extending the previous work, Robu et al. [23] combine the findings of [77] with the Continuous Ranked Probability Score (CRPS), which is a strictly proper rule, to incentivize the provision of accurate predictions from the CVPPs—and in turn, the member DERs—which aids in the planning of the supply schedule at the Grid. They showed that the proposed mechanism incentivizes real DERs to form CVPPs, while also outperforming the, at that time, state-of-the-art payment mechanism developed for this problem.

Moreover, multiagent systems theory has contributed to the solution of the coordinated consumption shifting for electricity prosumers. In particular, Akasiadis and Chalkiadakis [78] have shown that individual optimization concerning electricity prices does not always lead to minimized costs, thus necessitating a cooperative approach. The authors proposed a prosumer cooperative that employs an internal cryptocurrency mechanism to coordinate members' decisions and distribute the collectively generated profits. The mechanism generates crypto coins in a distributed fashion and awards them to participants according to the contribution impact and accuracy between stated and final shifting actions. In particular, it was shown that when a scoring rules-based distribution method is employed, participants are incentivized to be accurate.

Chapter 3

A Novel DER Aggregator Framework

Privacy concerns, impartial and fair scoring of the LFEs, and efficient management of all stakeholders' assets are the main problems our novel DER aggregation approach addresses. In particular, our framework provides an *aggregation architecture* along with *mechanisms* for its effective and efficient operation, increasing the actual flexibility offered by the aggregator to the Grid and the profits of the participating agents. Additionally, the proposed architecture can be distinguished into three levels, as depicted in Figure 3.1. The first is the *Distributed Energy Resources* level consisting of many Smart Grid assets. Specifically, our framework supports smaller and bigger scale DERs utilized by the Smart Grid. The second level of this architecture is termed as the *Aggregator Level*, and it consists of Local Flexibility Estimators (LFEs) organizing the DERs into coalitions, and a central Aggregator agent which selects which LFEs to include in a cooperative to participate in flexibility trading in the Smart Grid. The third level of this architecture represents the Smart Grid itself—specifically, the energy markets to which the Aggregator-coordinated LFEs cooperative participates. We now proceed to describe the key parts of our overall framework in detail.

3.1 Local Flexibility Estimators (LFEs)

An LFE agent acts as a coordinator of a coalition consisting of a varying number of heterogeneous DERs, effectively offering them visibility to the Grid. This means that all DER assets, regardless of size, can now participate indirectly in a flexibility aggregation process: though originally potentially too small to bid in the energy markets directly [79], DERs can now form LFE-coordinated coalitions; which in turn can be selected by the Aggregator to participate in LFE cooperatives to trade the accumulated flexibility.

The rationale and particular method an LFE uses to select its participating DERs are not absolute, as it depends on a multitude of constraints, priorities, or other factors. One factor could be locality limitations. For instance,

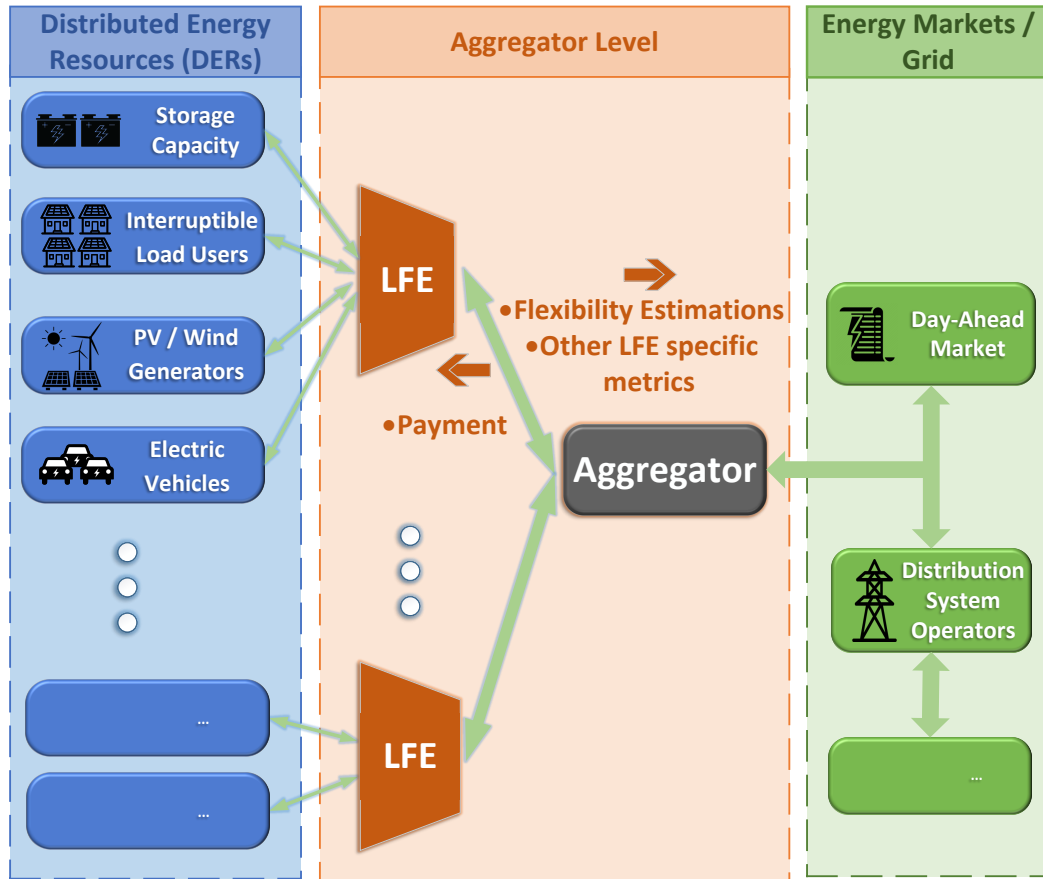


FIGURE 3.1: The component diagram of the proposed aggregator architecture consists of three levels; the Distributed Energy Resources, the Aggregator Level, and the Energy Markets

it might be easier to pick all DER assets lying close to each other in a physical neighborhood because of how the smart meters are placed [80]. Alternatively, an LFE might be formed for the convenience of or to serve the privacy requirements of a single stakeholder or a group of stakeholders with the same goals, like Local Energy Communities [81] or small companies. LFEs can then employ any cooperative formation method of their choosing that respects such requirements; for instance, they can use the very same member selection mechanisms that an Aggregator may employ, which we describe in detail in Section 3.4.

One of the primary responsibilities of an LFE is to monitor the consumption and production of the DERs it controls. Importantly, it can then use the historical consumption/production data to generate accurate estimations of the total flexibility for all participating DERs. In addition to the flexibility predictions, an LFE also provides the aggregator with other necessary metrics, such as confidence in the predictions (see Section 3.4.2). As a result, each LFE handles a portion of the total aggregator flexibility estimation problem; LFEs provide the aggregator with all information necessary for it to participate in the market, thus reducing the computational complexity of the aggregator

optimization process [82].

Most importantly, all the DER-related information is constrained and protected under the LFEs, so only the corresponding LFE can access the private data of each DER. Additionally, the proposed distributed aggregator scheme enables the usage of Smart Grid blockchain technologies [83], [84], which have been used for Smart Grid applications to secure the transactions made, while also protecting the private information required for the transactions.

3.2 The Aggregator agent

The aggregator in our framework possesses all the properties of a usual Smart Grid aggregator [11]; however, it has some additional novel properties, too, such as the ability to form LFE cooperatives. The framework we propose is comprised of a single aggregator that directly controls the assigned LFEs and indirectly manages all the DERs that the LFEs closely monitor. Additionally, the aggregator is responsible for gathering all the smart-metering data, e.g., flexibility estimations already pre-processed by the LFEs, to efficiently manage the assets and trade their excess flexibility. To address the problem of increasing the profits of each stakeholder, our Aggregator deploys a variety of scoring rules so it can fairly rank the LFEs according to their historic flexibility estimations' accuracy.

Ranking and selecting LFEs are key since LFEs with unreliable (low-accuracy predictions regarding their) flexibility estimations should not participate in the Aggregator's flexibility offers since they can damage both the Aggregator's profits and overall reputation, while other LFEs might be more accurate individually. This does not mean that low-accuracy LFEs will be excluded from the markets. Instead, they can trade with the Grid directly.

The Aggregator is thus able to calculate the total energy flexibility it can offer by selecting which LFEs will participate in the upcoming flexibility trades, using scoring and ranking mechanisms such as the ones we propose below. Moreover, the Aggregator is responsible for splitting the profits back to the LFEs, based on their contribution and appropriate scoring mechanisms that may also take into account the accuracy of LFE flexibility predictions.

3.3 Flexibility Estimation

In a Smart Grid context, total aggregator flexibility is its capability of shifting electrical loads either from itself to the Grid or in the opposite direction [7], [9]. At a given timeslot t , the flexibility provided by each DER in our work is calculated as follows.

To begin, BESS's available flexibility is proportional to its current energy level and specific charge/ discharge speed in KWh (Equation 3.1).

$$flex_{BESS}(t) = Charge(or\ Discharge)\ Speed(t) \quad (3.1)$$

The same principle applies to EVs (Equation 3.2); with the difference that to calculate EVs' flexibility $flex_{EV}(t)$, one has to account for a minimum battery level they should maintain in order to continue traveling.

$$flex_{EV}(t) = Charge(or\ Discharge)\ Speed(t) \quad (3.2)$$

The flexibility that interruptible load users can provide was set, following the bibliography [9], to 10% of the load, they are currently using. Hence, upon request of the aggregator through the LFEs, interruptible load users can alter their energy consumption by 10%:

$$flex_{InterruptibleLoadUsers}(t) = 10\% Load(t) \quad (3.3)$$

Most *renewable energy-producing* DERs do not have the ability to halt their production, so the flexibility of such DER assets is represented by the amount of energy they produce (Equation 3.4):

$$flex_{EnergyGenerators}(t) = Energy\ Generated(t) \quad (3.4)$$

Then, the total flexibility of an LFE is the aggregate flexibility of all DER assets it controls. Hence, the flexibility of LFE_i that controls a set K of DERs at timeslot t is defined as follows (Equation 3.5):

$$flex_{LFE_i}(t) = \sum_{\forall k \in K} flex_{DER_k}(t) \quad (3.5)$$

where $flex_{DER_k}(t)$ is calculated given the DER's type.

Finally, the total flexibility of the Aggregator is calculated as in Equation 3.6 below, where S is the set of LFEs selected by the Aggregator to contribute at timeslot t .

$$flex_{Agg}(t) = \sum_{\forall s \in S} flex_{LFE_s}(t) \quad (3.6)$$

3.4 Selection mechanisms

A key problem our framework addresses is the formation of efficient LFEs cooperatives to participate in the flexibility trading. To facilitate this in an impartial manner, the Aggregator first ranks the LFEs using certain scoring

functions. All our proposed scoring methods calculate the $Score_{LFE_i}$ for each LFE_i of the Aggregator, regardless of their prior participation in the latest flexibility tradings of the Aggregator. The Aggregator needs to score the LFEs regularly to have accurate information regarding the performance of every LFE it controls. Additionally, it is possible that during some specific periods of the year, the accuracy of the flexibility predictions can fluctuate [85], so this can also be an essential factor in the aggregator's decision process.

In our experiments, the Aggregator selects the LFEs with the LFE_i scores that exceed an Aggregator-specified threshold, in order to calculate its $flex_{Agg}(t)$ flexibility at time t . We now present our scoring mechanisms in detail.

3.4.1 Simple Selection mechanism

The first selection mechanism we developed uses the Mean Absolute Error (MAE) of the flexibility prediction \widetilde{flex}_{LFE_i} . Specifically, for an LFE_i that estimates its flexibility for the next k hours, we calculate the MAE as in Equation 3.7. Similarly, we define the average flexibility of an LFE as in Equation 3.8.

$$MAE_{LFE_i}(t) = \frac{\sum_{j=1}^k |\widetilde{flex}_{LFE_i}(j) - flex_{LFE_i}(j)|}{k} \quad (3.7)$$

$$AvgFlex_{LFE_i}(t) = \frac{\sum_{j=1}^k |flex_{LFE_i}(j)|}{k} \quad (3.8)$$

Then, we calculate the $Score_{LFE_i}(t)$ for each LFE_i of the Aggregator by using Equation 3.9. The first step is to divide the $MAE_{LFE_i}(t)$ by the average actual flexibility, then we subtract that value from 1 to have a straightforward confidence percentage. Finally, to avoid negative scores, we bound the score value, $Score_{LFE_i}(t) \in [0, 1]$.

$$Score_{LFE_i}(t) = \max \left(1 - \frac{MAE_{LFE_i}(t)}{AvgFlex_{LFE_i}(t)}, 0 \right) \quad (3.9)$$

As we can see, the score is at maximum when the predictor is perfect (LFE_i has an MAE of 0). By contrast, the score is 0 when the predictions' mean absolute error is worse than that of having a "dummy" predictor that always outputs 0 values, i.e., when the prediction MAE is greater than the average flexibility. The central concept of this metric is to get a simple and easy-to-compute confidence percentage so that we can use it as a "naive" selection method to compare with other, more sophisticated solutions.

3.4.2 Continuous Ranked Probability Score (CRPS)

The second selection mechanism we deployed uses the Continuous Ranked Probability Score (CRPS) [20], which assesses the accuracy of a probabilistic

prediction over the actual occurrence. CRPS has been previously used for virtual power plant formation [23], and we use it in our aggregator framework to score the LFEs predictions aiming to optimize aggregator profits. CRPS is a strictly proper scoring rule, meaning that the expected score is maximized only if predictors accurately report their expectation over the prediction error they can potentially make [20]. CRPS has been shown to incentivize energy suppliers to be truthful and accurate [23], and we use it here to incentivize truthful and accurate LFE flexibility predictions.

With the Simple Selection mechanism, the LFEs only had to send their flexibility predictions to the aggregator. When using the CRPS selection mechanism, LFEs also need to provide the uncertainty over the prediction error, and are rewarded accordingly: estimates that are both accurate and highly confident will be the ones achieving higher CRPS scores. The $CRPS_{LFEi}(t)$ is defined as follows:

$$CRPS_i(t) = \sigma(t) \left[\frac{1}{\sqrt{\pi}} - 2\phi\left(\frac{e_i(t)}{\sigma(t)}\right) - \frac{e_i(t)}{\sigma(t)}(2\Phi\left(\frac{e_i(t)}{\sigma(t)}\right) - 1) \right] \quad (3.10)$$

where ϕ and Φ denote the probability density and the cumulative distribution function of a *standard Normal* variable. In our case, e_i is the relative prediction error, as shown in Equation 3.11.

$$e_i(t) = \frac{flex_{LFEi}(t) - \widetilde{flex}_{LFEi}(t)}{\widetilde{flex}_{LFEi}(t)} \quad (3.11)$$

(which we ensure to take values in $[-1,1]$). Note that in Equation 3.11 $flex_{LFEi}$ is the random variable to be predicted (unknown in advance), while \widetilde{flex}_{LFEi} is the prediction for this variable. Additionally, along with every prediction, the LFEs also send a Normal distribution $\mathcal{N}(0, \sigma(t)^2)$ to describe their uncertainty over their error, where $\sigma(t)$ changes over time according to the LFEs uncertainty on the predictions. (We assume, as in [23], that random errors, over a long enough period, will be normally distributed around a mean of 0.)

3.4.3 Reinforcement Learning: DQN

The final selection mechanism we developed deploys the celebrated Q-Networks (DQN) [52] RL algorithm. We formulate the aggregator's decision-making problem as a decision process, aiming to find the action with the highest Q-value—corresponding to the long-term utility of taking a (selection) action, i.e., selecting a set of LFEs at this time step. We assume a continuous state-space defined via providing as inputs the $CRPS_{LFEi}(t)$ and the $\widetilde{flex}_{LFEi}(t)$ for every LFE of the aggregator; and define a discrete action-space containing all the different possible selections of LFEs, using one-hot encoding. For example, for n LFEs, there will be n bits of information, with 1 meaning the LFE was selected, and 0 the opposite, so when the aggregator selects among n LFEs, the action-space will have 2^n possible actions.

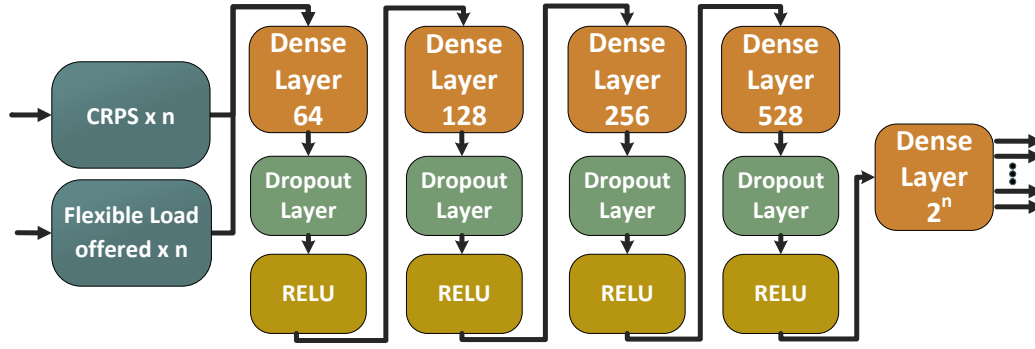


FIGURE 3.2: The architecture of the Reinforcement Learning model used for the selection of the best LFE cooperative.

DQN is a value-based method that does not store any explicit policy but only a value function in the form of a deep neural network, hence the policy is implicit and can be derived directly from the neural network as the action with the highest value. In Figure 3.2, we can see our neural network architecture. We used five fully-connected layers with increasing nodes, so we could better handle this complex optimization problem. Also, we put three "Drop-Out" layers followed by rectified linear activation units (ReLU) to address the randomness of the samples during training.

Two reward functions were developed for the training of the DQN model. The goal of the first reward function is to maximize the profits of the aggregator during the evaluation time (Equation 3.12). Z is a normalization constant that keeps the reward values closer to 0, but without an upper bound. The term $V_{Grid,Agg}(t)$ refers to the Grid payments towards the Aggregator for the flexibility trading it has completed in timeslot t .

$$Reward_1(t) = \frac{V_{Grid,Agg}(t)}{Z} \quad (3.12)$$

The second reward function we tested tries to balance the profits of both the Aggregator and the LFEs (Equation 3.13). This function takes values in the $[0, 1]$ interval. Here too, J is the set of LFEs selected to participate in the aggregator's latest flexibility trades, and S is a set comprised of every LFE in the aggregator. We define $V_{Agg,LFE_i}(t)$ to refer to Aggregator payments to each LFE_i for participating in its flexibility trading. Also, we define $V_{Grid,LFE_i}(t)$ to refer to the payments received by LFE_i at t while trading flexibility directly with the Grid (when not selected by the Aggregator in timeslot t).

$$Reward_2(t) = \frac{V_{Grid,Agg}(t)}{V_{Grid,Agg}(t) + \sum_{i \in \{S-J\}} V_{Grid,LFE_i}(t)} \quad (3.13)$$

3.5 Pricing Mechanisms

Another problem our framework addresses is the distribution of the aggregators' profits to each individual LFE. We have deployed two different pricing mechanisms inspired by previous studies. At first, using these mechanisms, we calculate the payment from the energy markets to the aggregator for trading its flexibility and then payments by the aggregator to its LFEs. The (per KWh) price of the traded energy in the markets at timeslot t is denoted as $p(t)$.

3.5.1 Prediction Accuracy mechanism

Our first pricing mechanism calculates the payments based only on "point estimates" of the accuracy of the flexibility predictions (i.e., no uncertainty-related distribution is reported), as in [77]. The more accurate the flexibility estimators, the higher the payments it awards them.

To begin, Equation 3.14 shows how the Aggregator is rewarded for trading its flexibility in the energy markets.

$$V_{Grid,Agg}(t) = \frac{\log|flex_{Agg}(t)| \cdot flex_{Agg}(t)}{1 + \alpha \cdot e_{Agg}(t)^\beta} \cdot p(t) \quad (3.14)$$

The logarithmic term increases with the provided flexibility. This incentivizes the Aggregator to include a large number of LFEs in its offers. At the same time, the Aggregator has to proceed in its LFEs selection with caution, as its flexibility prediction error, denoted by e_{Agg} and calculated as in Equation 3.11, plays a role in its final reward: notice that the parameters of the denominator resemble a bell-shaped function, so the value is maximized when the prediction error is zero. Parameters α and β determine the exact shape of the curve [77]; in our experiments, we set $\alpha = 1.6$ and $\beta = 4$. In particular, Figure 3.3 illustrates a bell-shaped graph depicting the accuracy factor $\frac{1}{1 + \alpha \cdot e_{Agg}(t)^\beta}$ of this payment method. We can observe that for errors in $[-0.25, 0.25]$, there is not any significant penalty resulting from this accuracy factor, however, as the absolute error increases, the accuracy factor decreases too in a linear fashion until it reaches values close to zero. Overall, this payment mechanism can be considered forgiving for predictions with smaller errors, but it is also fair because it implies penalties for predictions with greater errors.

The LFEs not selected by the Aggregator trade directly with the Grid, and are also rewarded according to Equation 3.14 by swapping the terms $e_{Agg}(t)$ with $e_{LFEi}(t)$ and $flex_{Agg}(t)$ with $flex_{LFEi}(t)$.

After the Aggregator is paid, it distributes the profits to the contributing LFEs. Equation 3.15 displays the pricing function used to distribute $V_{Grid,Agg}$ to the LFEs based on their prediction error e_i .




FIGURE 3.3: The “Simple Scoring” accuracy factor graph for errors ranging in $[-1,1]$

$$V_{Agg, LFE_i}(t) = \frac{Z}{1 + \alpha \cdot e_i(t)^\beta} \cdot \frac{flex_i(t)}{flex_{Agg}(t)} \cdot V_{Grid, Agg}(t) \quad (3.15)$$

The term $flex_i(t)/flex_{Agg}(t)$ represents the flexibility contribution percentage of the LFE_i to the Aggregator. Also, Z is a normalization parameter calculated to split the payment to every LFE completely. One can adjust Z to allow for some portion of the total payment to be withheld by the Aggregator.

3.5.2 CRPS based mechanism

The second pricing mechanism [23] uses the CRPS score. This mechanism encourages the formation of larger LFE cooperatives while giving incentives for accurate and truthful flexibility predictions. A key difference with the Prediction Accuracy-only pricing mechanism is that the CRPS-based mechanism can be more forgiving to low-accuracy predictors. This is because LFEs also provide their flexibility predictions’ uncertainty distributions, accounted for in their CRPS scores (see Sec. 3.4.2).

In particular, Figure 3.4 depicts the CRPS accuracy factor for various uncertainty values σ_i . We can clearly see that the uncertainty parameter of CRPS (σ_i) creates various curves with different characteristics each. For instance, when the provided uncertainty is close to zero, hence the predictor is stating that it is very confident in the accuracy of its predictions, then the maximum

possible accuracy factor value is close to the maximum only when the actually observed error is tiny too. Meanwhile, it suffers from a significant decrease in case the error is actually more remarkable, for example, when the stated uncertainty is $\sigma_i = 0.1$, and the actual error is maximum, it results in accuracy values close to zero. On the contrary, for predictors that state that their uncertainty is very high, the CRPS is more forgiving, leading to higher overall values. For instance, when a predictor reports that its uncertainty is at maximum, hence $\sigma_i = 1$, the maximum accuracy value this predictor can get might be the lowest at around only 0.7 out of 1. In contrast, if the error is at maximum, this predictor will get the highest CRPS score in that case, with the actual accuracy value being over 0.4 out of 1.




FIGURE 3.4: The CRPS accuracy factor graph for errors ranging in $[-1,1]$

The "Grid-to-Aggregator" payment shown in Equation 3.16 is calculated similarly to the previous mechanism. However, the accuracy factor is the aggregator's CRPS value (normalized to $[0,1]$ similarly to what is done in [23], [86]) instead of a bell-shaped function.

$$V_{Grid,Agg}(t) = CRPS_{Agg}(t) \cdot \log|flex_{Agg}(t)| \cdot flex_{Agg}(t) \cdot p(t) \quad (3.16)$$

Here too, the LFEs that were not selected by the aggregator use Equation 3.16 but with the relevant LFE_i terms to calculate their "Grid-to-LFE" payment. After the aggregator is paid, it distributes the profits to each selected LFE member (set J of LFEs) using Equation 3.17.

$$V_{Agg,LFE_i}(t) = \frac{CRPS_{Agg}(t) \cdot flex_{LFE_i}(t) \cdot V_{Grid,Agg}(t)}{\sum_{j \in J} (CRPS_{LFE_j}(t) \cdot flex_{LFE_j}(t))} \quad (3.17)$$

This pricing mechanism ensures that each participant is awarded a weighted portion of the total payment based on their contribution and their individual CRPS score. Additionally, the CRPS is helpful because it shows how beneficial the LFE estimates were for the total flexibility trading of the aggregator.

3.5.3 Simple Grid-to-Aggregator Payment

As mentioned, CRPS has been proposed for Smart Grid applications and was proven to reward more generously the truthful and accurate predictors, regardless if they were small DERs or large aggregators [23], [86]. However, nowadays, this complex but strictly proper payment mechanism is not used for the calculation of payments in energy trading. Instead, energy markets use simple mechanisms that are mainly based on the flexibility contribution and not on the accuracy of the predictions. For example, a large flexibility platform in the United Kingdom named “Flexible Power” [87] uses “simple payments” [88] to reward the participating DERs. Specifically, this “simple payment” mechanism has two steps. The first step is for each energy markets participant, either a single DER or a DER aggregator, to promise to deliver a specific amount of flexibility \widetilde{flex}_i at a future timeslot t at a predetermined price $p(t)$. The second step occurs after the timeslot t has passed and the actual flexibility $flex_i(t)$ has been successfully traded.

$$flex_{diff}(t) = flex_i(t) - \widetilde{flex}_i(t) \quad (3.18)$$

Initially, we calculate the difference $flex_{diff}(t)$ between the promised flexibility and the actually delivered flexibility as seen in Equation 3.18. Then, if the $flex_{diff}(t)$ is lower than 0, meaning that the flexibility delivered is less than the promised one, the Grid simply pays each KWh delivered using the predetermined price $p(t)$. By contrast, when the $flex_{diff}(t)$ is positive, the Grid pays only the pre-agreed amount of flexibility $\widetilde{flex}_i(t)$ using the price $p(t)$. The rest of the flexibility delivered is paid using the current electricity prices $p_{standard}(t)$, which is usually much lower than the predetermined energy. Remember that, in our case, we have made the assumption that DERs and the aggregator always sell their generated energy to the Grid so that no energy can be wasted. The final payment $V_{Grid,i}$ is calculated using Equation 3.19

$$V_{Grid,i} = \begin{cases} flex_i(t) \cdot p(t) & flex_{diff}(t) \leq 0 \\ \widetilde{flex}_i(t) \cdot p(t) + flex_{diff}(t) \cdot p_{standard}(t) & flex_{diff}(t) > 0 \end{cases} \quad (3.19)$$

Chapter 4

Experimental Setup

Here, we describe in detail the experimental setup of this study. The main goal was to evaluate the performance of the LFEs and the aggregator in a variety of experimental scenarios that we especially formulated to test the different aspects of our aggregator framework while maintaining the realistic aspects of the current electricity Grid. Initially, we will showcase each experimental scenario in great detail, so it will be much easier to distinguish the differences between them. Furthermore, there is a discussion about the technical characteristics of the simulation and the assets comprising them. In the end, we will also present the additional use cases we have formulated to evaluate the proposed framework under realistic conditions, for instance in case a natural disaster occurs or when weather is unpredictable.

4.1 Experimental Scenarios

It is very important to provide realistic simulations, so we can accurately assess the performance of the proposed Aggregator framework before implementing it in the actual Smart Grid. For this reason, we have created five unique experimental scenarios in order to observe how each stakeholder reacts, while also evaluating the performance of the mechanisms we used.

We provide a detailed activity diagram (Figure 4.1) depicting the actions taken by the various stakeholders in a period of time t . As we can see, the main activity flow remains the same for all experimental scenarios, but there are some important variations that will be discussed explicitly later.

As we can see, at all times, the LFEs monitor the electricity production and consumption of all their DER members with the intention to process the available data and train their flexibility estimators. At the start of every trading cycle, the LFEs deliver their flexibility predictions for the next 24 hours to the Aggregator with the intention to trade in a day-ahead wholesale market before the markets close.

Then the Aggregator using a scenario-specific selection method, chooses which

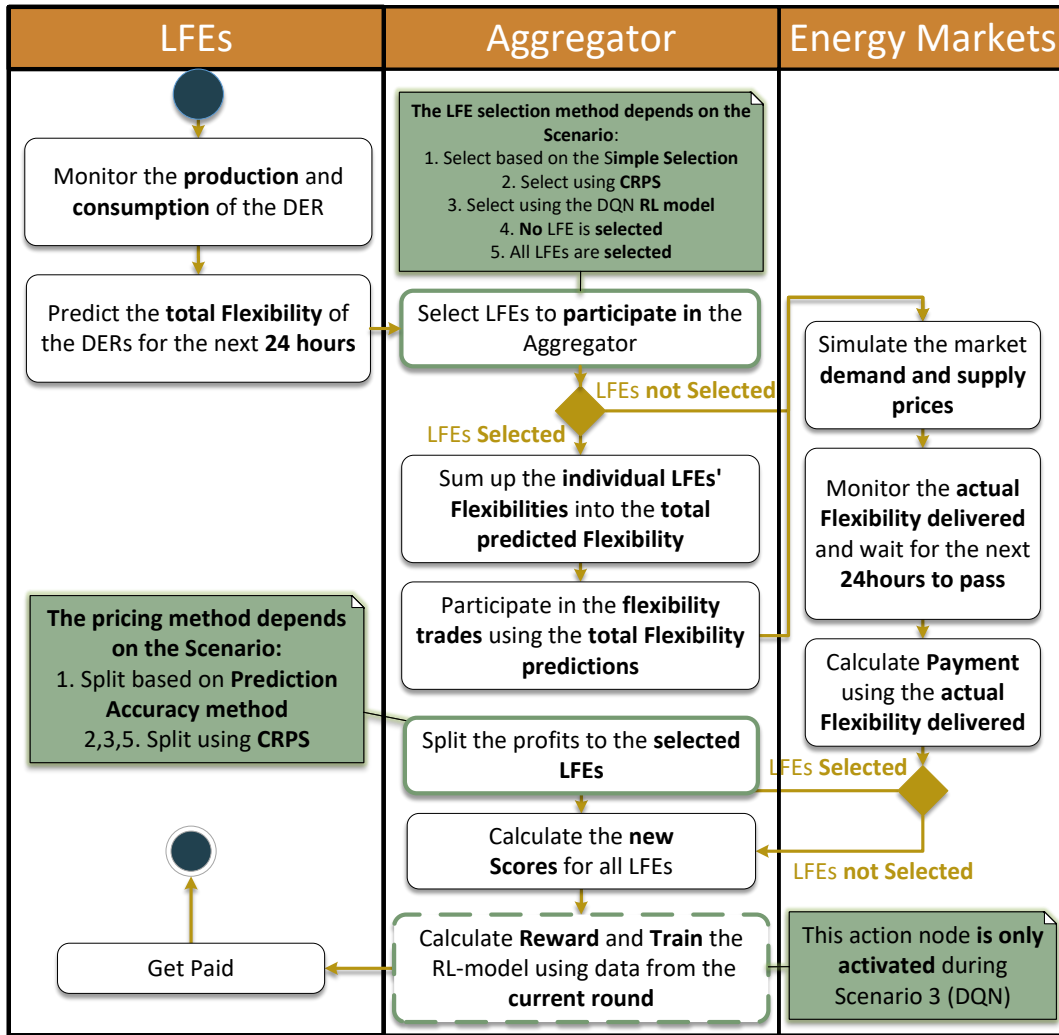


FIGURE 4.1: An activity diagram presenting the five different experimental scenarios developed for the evaluation of the proposed Aggregator framework

LFEs to participate in the upcoming flexibility trades. Right after, the Aggregator calculates its total flexibility estimation and participates in the flexibility tradings. At the same time, PowerTAC simulates the demand and supply prices so that the energy market participants can buy or sell energy. After that, the day-ahead market closes and continues to monitor the actual flexibility delivered until the market opens again. After 24 hours pass and the market opens again it calculates the payment of the Aggregator based on the actually delivered flexibility and the prediction error, as seen in the previous section about the pricing mechanisms.

Now that the Aggregator has gathered its payment, it distributes the profits to the LFEs using scenario-specific pricing mechanisms. In the meanwhile, the LFEs that were not selected by the Aggregator have also traded their flexibility directly with the energy markets. The final step for the Aggregator is to assess the prediction accuracy of every LFE, both the ones that were

selected and those that were not, by re-calculating the Simple Scoring and the CRPS metrics. It is important to re-evaluate after every trading cycle because some LFEs might have better or worse prediction performance so the Aggregator will know which LFEs to select for the next tradings. For the third experimental scenario, the Aggregator also has to calculate the RL reward and train the RL selection model using the latest information from the current trading cycle.

As mentioned earlier, we have formulated five different experimental scenarios with the intention of evaluating different aspects of the proposed framework. In practice, we are testing five different LFE selection methods as seen below.

4.1.1 Simple Selection

The first experimental scenario, “Simple Selection”, aims to assess the performance of the proposed DER aggregation framework when it combines a straightforward selection method with a “smart” Aggregator. Additionally, the pricing mechanism used for this scenario is based only on the flexibility contribution and the historic prediction accuracy, hence is easier to understand for the nonspecialists.

$$AvgMAE_{LFE_i}(t) = \frac{\sum_{l=t-w}^t MAE_{LFE_i}(l)}{w} \quad (4.1)$$

The first experimental scenario uses the Simple Selection method to decide which LFEs to participate in the Aggregator. In detail, the Aggregator calculates the average Simple Score of LFE_i over a time period w for several trading cycles of the past, as seen in Equation 4.1. Then the Aggregator checks if the average score is over the desired threshold τ . We have set $\tau = 0.7$ in our experiments. As we can see in Figure 4.2, if $Score_{LFE_i} > \tau$ then the LFE_i is not selected to participate in the upcoming flexibility tradings of the Aggregator, however the rejected LFEs trade directly with the Grid.

The other significant variation of this scenario is the pricing mechanism deployed. Specifically, we use the prediction-accuracy-only pricing mechanism described in the previous Chapter. In detail, the Aggregator uses this pricing mechanism to split the profits after the Aggregator is paid by the Grid for its flexibility trades, using only the LFEs’ flexibility accuracy. For all the other experimental scenarios, we employed the CRPS pricing mechanism since it was proven to give incentives for truthful and reliable LFE predictions, which makes it a perfect fit for our Aggregator framework.

4.1.2 Using CRPS only

The second experimental scenario uses the more sophisticated CRPS selection method and the CRPS-based pricing mechanism. Here too, we calculate

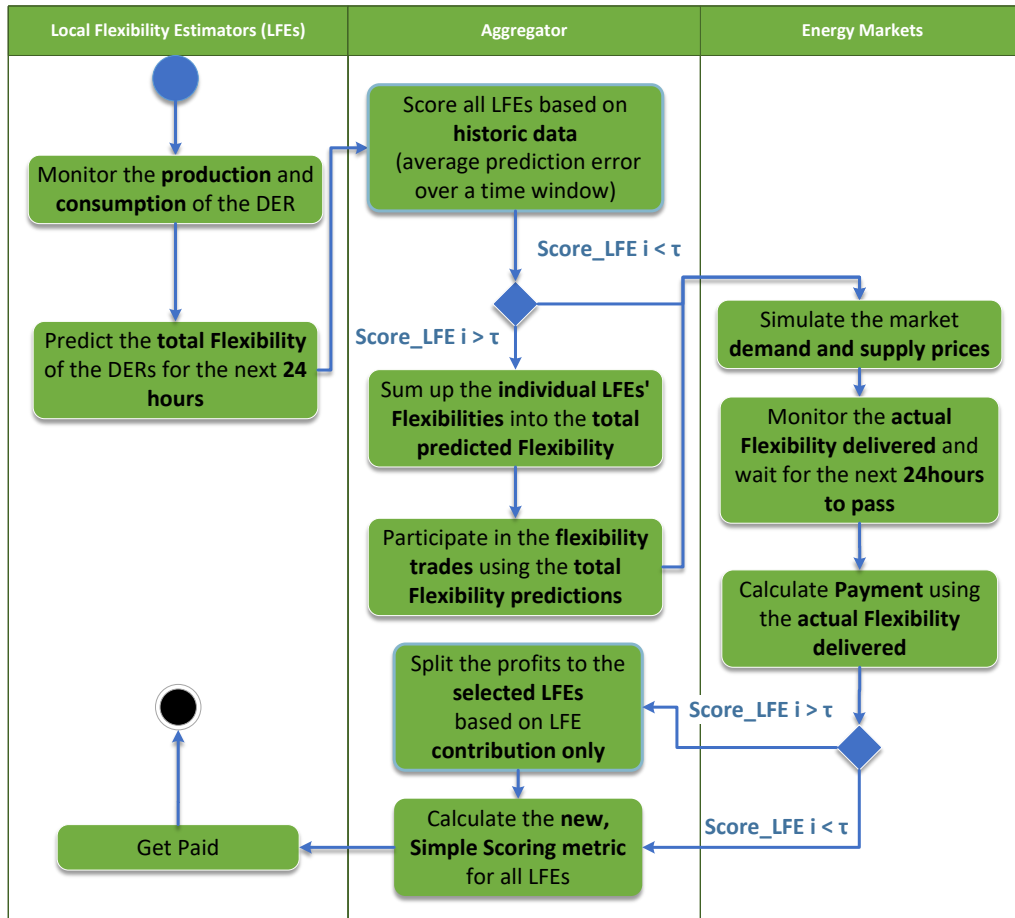


FIGURE 4.2: The activity diagram of Experimental Scenario: "Simple Selection method"

the average CRPS score of a time period w and check if the final CRPS score is less than a specific threshold as we can observe in Figure 4.3.

In particular, we have set the threshold $\tau = 0.77$ to mimic the selection rates of the previous Simple selection method, so we could better compare them together. We considered it very important to set the thresholds τ of the first two experimental scenarios in a way that the results would be comparable; thus, we experimented and came up with the aforementioned values. Additionally, having similar LFEs selection behaviors can help us to concentrate on the comparison of the profits and not the selection methods themselves, as we will see later.

4.1.3 DQN

For the third experimental scenario, we deployed the DQN RL model intending to see if there are any more dynamic selection patterns resulting from the training of distinguished A.I. techniques that could probably give us better overall results. We have dedicated two experimental scenarios for this variation because we also wanted to compare the two proposed reward methods.

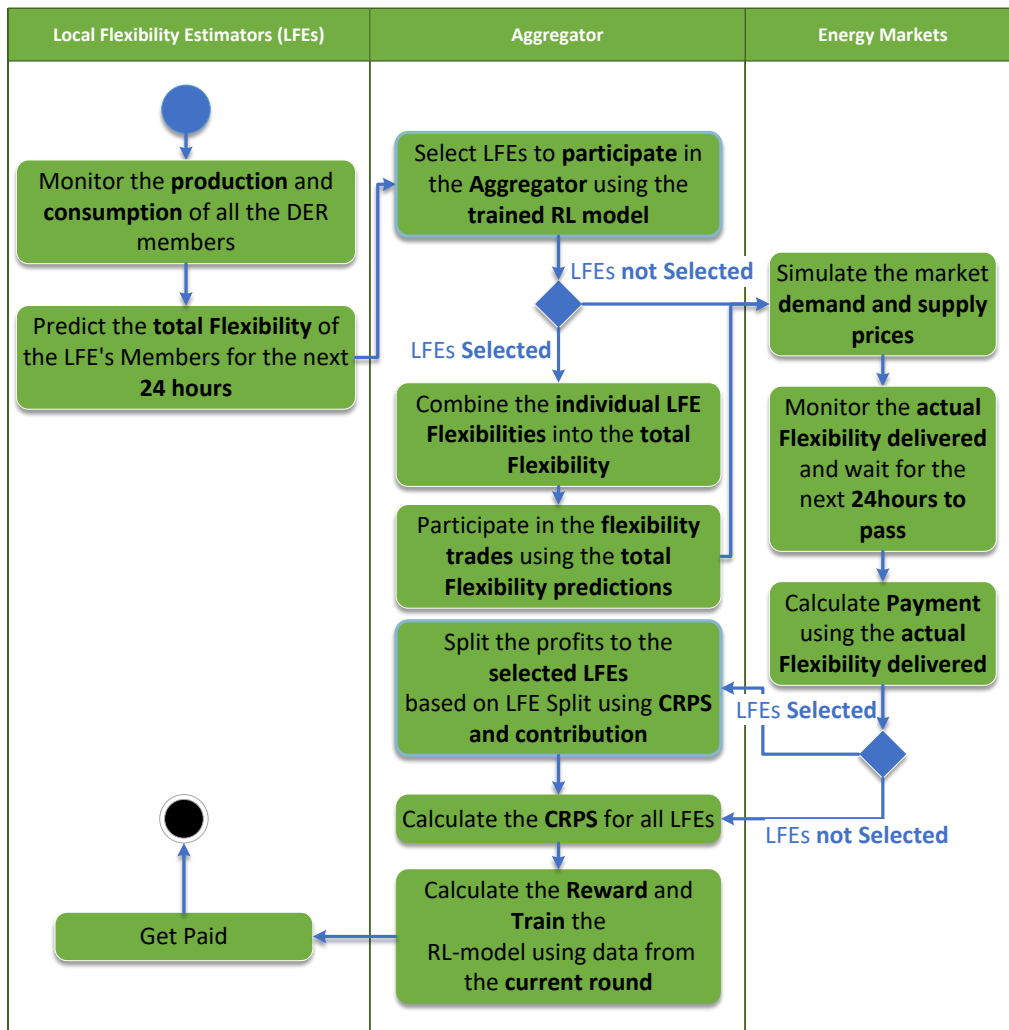


FIGURE 4.4: The activity diagram of Experimental Scenario: “Using DQN”

The second phase is the evaluation of the DER aggregation framework in the same game seeds as all the other experimental scenarios so that we can have a stable and trustworthy comparison.

4.1.4 Singleton LFEs

The two following experimental scenarios were designed to showcase the two opposite edge cases of the DER aggregation problem; hence the results of these two can act as proper baselines with which we can compare our results.

In the fourth scenario, as we can see in Figure 4.5, the Aggregator does not exist; hence every LFE interacts directly with the energy markets using the designated pricing mechanisms. This use-case acts as the first baseline method

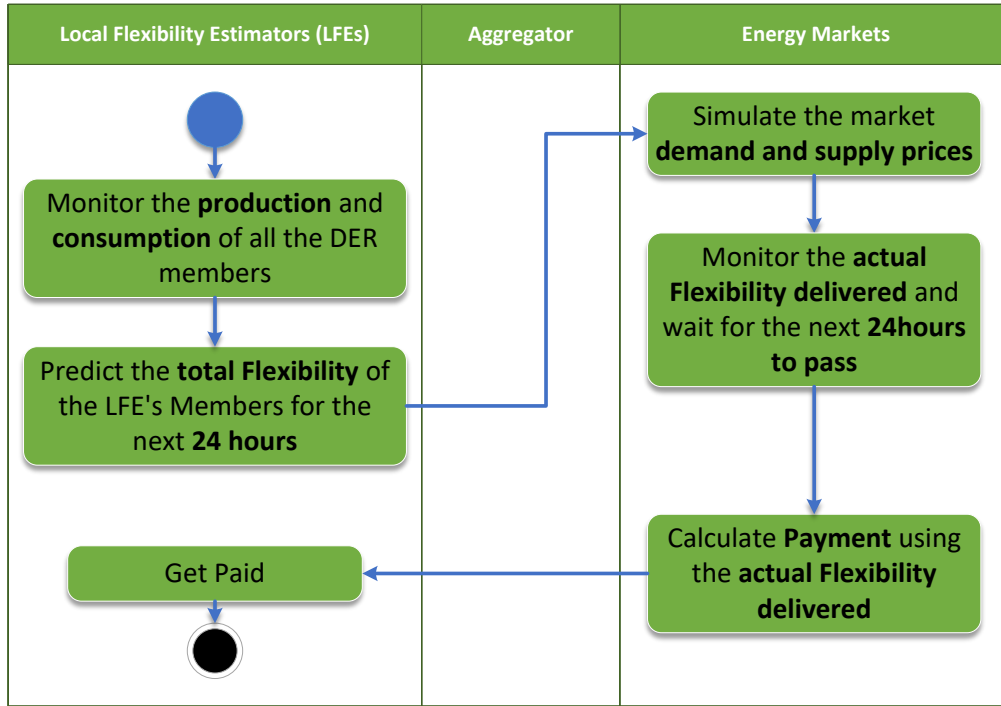


FIGURE 4.5: The activity diagram of Experimental Scenario: “Singleton LFEs”

depicting how the LFEs would have performed if they had never participated in our Aggregator framework. Therefore, we expect our Aggregator framework to always give better results than this scenario, at least in the future Smart Grid settings.

4.1.5 All LFEs

The last experimental scenario acts as the second baseline with the intention to compare the revenue of the LFEs and the Aggregator when every LFE is participating in the Aggregator without any selection criteria.

After all, this scenario is very straightforward because it does not deploy any selection methods, as we can observe by looking at Figure 4.6. With this scenario, we intend to directly compare our novel Aggregator architecture and the traditional aggregators proposed in the relevant literature.

4.2 Technical properties of the LFEs

In this section, we are going to discuss the technical properties of the Power-TAC simulation by focusing on the LFEs we studied.

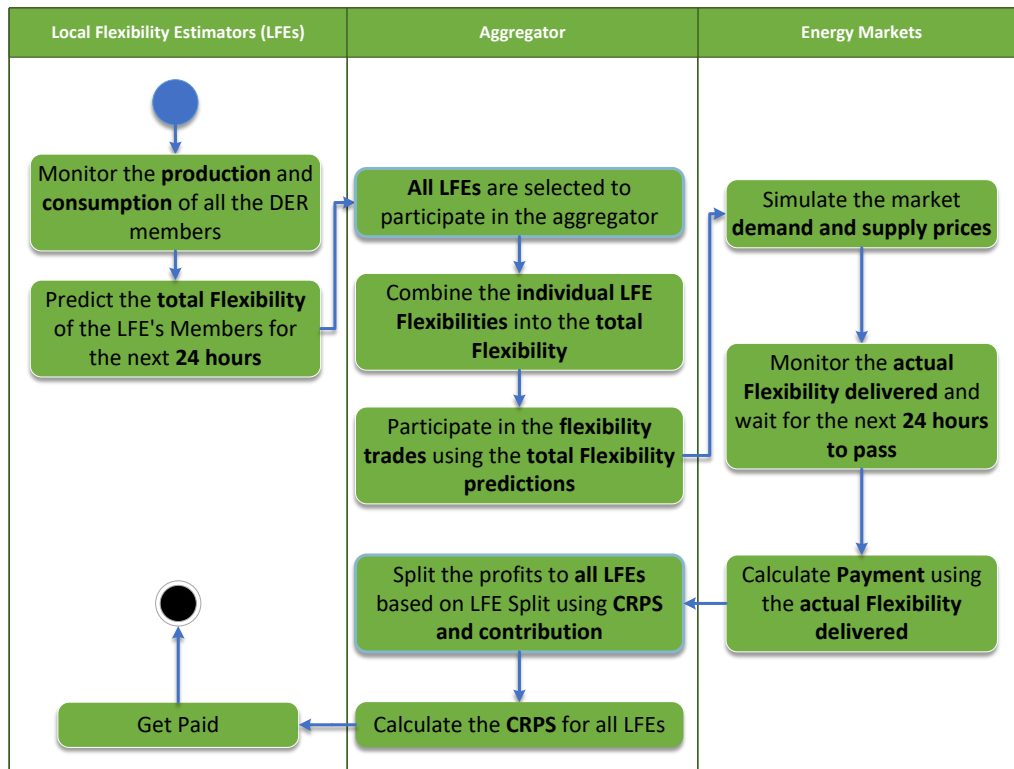


FIGURE 4.6: The activity diagram of Experimental Scenario: "All LFEs"

4.2.1 PowerTAC Simulator

The PowerTAC [22] platform, where the PowerTAC competition is conducted, is a rich competitive economic simulation of future energy markets, featuring several Smart Grid components (e.g., DERs, retail and wholesale energy markets, etc.). With the help of this simulator, researchers can better understand the behavior of future customer models and experiment with retail and wholesale market business models or strategies so helpful information can be extracted.

In our case, we have modified the PowerTAC platform so it could support the addition of our Aggregator-specific mechanisms but without altering the behavior of the realistic DER and market models. Another reason we selected this simulation platform is that it employs a realistic day-ahead wholesale market simulating the supply and the demand of future energy markets in a pretty accurate way.

Another important feature of PowerTAC is that it uses real weather data to determine the behavior of the Smart Grid models; hence this increases the accuracy of the simulations. Figure 4.7 depicts an example of the four weather variables of PowerTAC along the hour of the day. The first weather variable is the temperature (in Celsius). As we can see, it can fluctuate a lot during a simulation depending on the geographical location where the weather

data were gathered from. Additionally, there are two variables regarding the wind. The first is the wind speed, and the second is the wind direction that has a significant impact on wind generation. Finally, another very important variable is cloud cover which has a direct impact on solar energy production. However, the PowerTAC models are not only affected by these four weather variables, they also take into account the hour of the day, the day of the week, and the month.




FIGURE 4.7: The weather simulation data of the PowerTAC decomposed to four parameters temperature, wind speed, wind direction, and cloud cover

4.2.2 LFEs Specifications

In general, PowerTAC has a plethora of heterogeneous DER assets modeled differently, trying to resemble the various aspects of the future Smart Grid. Overall, the entirety of the PowerTAC assets are designed to model the production and consumption of a small city; hence this means that there are many non-DER assets, such as households and small enterprises, that might not have either any available solar panels, for example, if they live in apartments or any flexible loads such as heat pumps. However, aggregators, by default, require flexible loads to operate, so we preferred to exclude “traditional Grid” assets from participating in any LFE and indirectly in an Aggregator.

	Storage Capacity (kWh)		Controllable load (kWh)		Maximum Prod. (kWh)	Maximum Cons. (kWh)
	BESS	EVs	BESS	EVs	Generators	Interruptible
LFE 1	270	437	120	30	1000	- 1000
LFE 2	270	535	120	35	5000	-1000
LFE 3	180	369	80	23	1000	-10000
LFE 4	180	361	80	20	1000	-250
LFE 5	180	485	80	29	1250	-250
LFE 6	180	429	80	26	500	-200
LFE 7	180	404	80	23	500	-200
LFE 8	180	418	80	28	500	-200
LFE 9	270	449	120	26	1000	-250
LFE 10	270	449	120	26	500	-200
LFE 11	270	362	120	22	500	-200
LFE 12	270	450	120	129	500	-15000

TABLE 4.1: The technical characteristics of the 12 LFEs formulated for the experiments

To be precise, it was very important for us to include every DER in the Aggregator regardless of its scale. In particular, our proposed Aggregator framework supports both small and big-scale DER assets. For our current experiments, we have randomly divided all the DER assets of PowerTAC into 12 LFE teams with varying characteristics each. However, we could have used the cooperative formation mechanisms we proposed for our Aggregator to determine the DER each LFE can select. In particular, Table 4.1 illustrates the technical specifications of every LFE team that remained the same for every experimental scenario.

We can compare the main characteristics of each LFE by dividing them into four categories. The first column describes the storage capacity of the BESS and the EVs of every LFE. The controllable load depicts the rate that the batteries of BESS and EVs can charge or discharge electrical energy per hour (kWh). The next characteristic, “Maximum Production”, shows the maximum amount of energy that renewable energy DERs can generate per hour, indirectly showing the size and the generation potential of every LFE. Finally, the last column shows the maximum energy consumption per hour for every LFE.

We can notice that there are LFEs with higher storage capacities like LFE_2 with a total capacity of 805 KW (270 KW from BESS and 535 KW from EVs) and others that have much lower like LFE_4 with 541 KW storage capacity in total. However, LFEs also differ in the number of controllable loads they can charge or discharge. For example, if we assume that every battery has enough energy to discharge for an hour continuously, then LFE_{12} would be able to deliver 249 kWh. At the same time, LFE_4 would only be able to deliver 100 kWh. Furthermore, there are LFEs with potential renewable energy

generation from 500 kWh up to 5000 kWh, such as LFE_2 . Additionally, LFEs also consume load in a different manner. For instance, most LFEs consume around 100 to 250 kWh while LFE_3 , which hosts a few electricity-intensive facilities, can consume up to 10000 kWh in peak hours. This heterogeneity among the LFEs is a significant factor in making the experiments more realistic.

These technical specifications also impact the amount of flexibility the LFEs can offer. Figure 4.8 provides us with a detailed example of how the total LFE flexibility is calculated. We can distinguish the four DER categories, Electric vehicles, Battery Energy Storages(BESS), Interruptible load users, and renewable energy generators, along with the two variations of flexible loads.




FIGURE 4.8: The flexibility of an LFE decomposed to EV, battery charge, interruptible loads, and production

It is visible that in this example, the majority of the flexibility derives from renewable energy generators, while a smaller portion originates from interruptible load users, batteries, and EVs. However, it is noteworthy that the flexibility provided by the generators is not stable and has major changes depending on the weather and time of day. For example, Wind generators can only produce energy when the wind speed is over a specific threshold; similarly, solar panels provide negligible energy when the sun is covered or when the sun is covered. Consequently, the rest of the DERs may contribute to the total flexibility with a much smaller flexibility portion (around 1/5 of the total), however, there is always some amount of energy ready to help in the balancing of the Grid. To sum up, every DER is important in its own way because even the smallest amount of flexibility offered can help.

Furthermore, after reviewing the individual flexibility of a good LFE, we can now look at Figure 4.9 that illustrates the total flexibility offered by the 12 LFEs over some period of time. We can see that various flexibility patterns greatly depend on the scale of the DER assets in every LFE. For example,

most LFEs, such as LFE_2 display many flexibility spikes in their waveforms. These spikes are usually the outcome of solar panels that can only produce renewable energy only when the sun is in the sky. On the other hand, other LFEs such as LFE_1 and LFE_4 , have flexibility curves that illustrate spikes but are generally smoother. In these cases, we could infer that the wind turbines are responsible for these curves since they could produce energy for days if there is a sufficient wind stream blowing.

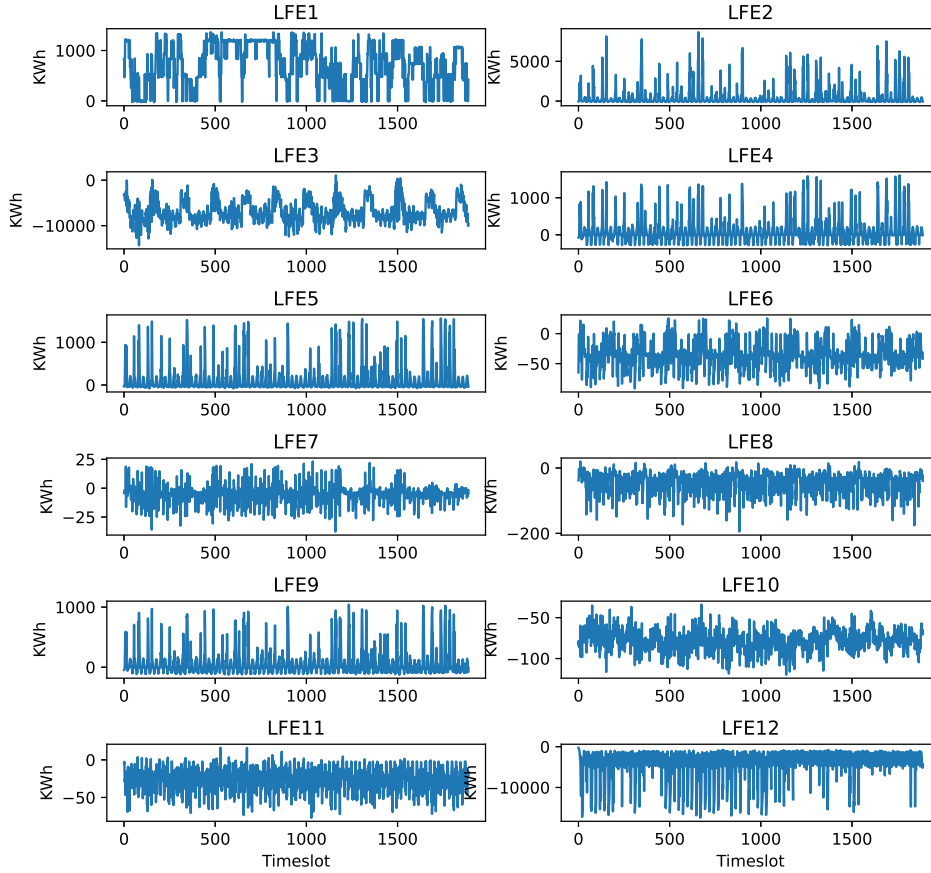


FIGURE 4.9: An example of the total flexibility of all 12 LFEs of the Aggregator

4.3 Static and Dynamic LFE accuracy

As their name suggests, local Flexibility Estimators (LFEs) are responsible for estimating the available flexibility of the DER assets they control and then providing the Aggregator with accurate flexibility predictions for the next hours. Additionally, when CRPS is used, LFEs also have to provide the Aggregator with their uncertainty distributions so the strictly proper CRPS can be computed by the Aggregator.

The flexibility predictions of the LFEs are the outcome of a Gaussian random process, $\mathcal{G}(\mu, \sigma \cdot \mu)$, with μ being the actual flexibility of the LFEs and the




FIGURE 4.10: The error Probability Density Function (PDF) of the prediction probability of every LFE

variance σ being an LFE-specific variance with $\sigma \in [0, 1]$. Additionally, we have modeled the LFEs to have fluctuating prediction accuracy through the course of the simulations so they can be more realistic. Figure 4.10 shows the probability density function of the prediction error of every LFE.

In particular, this figure shows the probability an LFE makes an inaccurate prediction. We can see that LFEs with higher variances have more chances to make inaccurate predictions, while LFEs with lower σ usually have smaller prediction errors. In addition to these heterogeneous LFEs of our simulation, we also wanted to evaluate the performance of the proposed DER aggregation framework in situations where the variance of the LFEs can change dynamically due to unforeseen reasons, trying to resemble the occurrence of natural phenomena. Specifically, we tested two different LFE prediction accuracy use cases.

4.3.1 Static LFE accuracy

The first use-case we are testing with is the “Static LFE accuracy”, there, the LFEs have a static standard deviation in their flexible load predictions. More specifically, LFE_1 resembles the best flexibility predictor with $\sigma \approx 0$ and LFE_{12} the worse with $\sigma \approx 1$. All the rest LFEs after LFE_1 have gradually worse accuracy, for example with $c = 1/11$ we have $\sigma_{LFE_2} = 1 \cdot c$, $\sigma_{LFE_3} = 2 \cdot c$, $\sigma_{LFE_4} = 3 \cdot c$ and so on, as depicted at Figure 4.11.

Basically, we have formulated the random Gaussian distributions $\mathcal{G}_{LFEi}(Flex_{LFEi}, \sigma \cdot Flex_{LFEi})$ of the LFE in this specific way so that we could associate σ with the prediction error percentage e_i . Specifically, when $\sigma = 1$, it means that the average prediction error is equal to 1 after infinite predictions. On the other hand, when the standard deviation of an LFE is $\sigma = 0$ its




FIGURE 4.11: Standard Deviation of the LFEs' flexibility estimations in the "Static LFE accuracy" use case

average error will be 0 after many simulations, resembling a highly-accurate flexibility predictor.

4.3.2 Dynamic LFE accuracy

The second use case we evaluate is the exact opposite of the first. In particular, here, the prediction accuracy of the LFEs gradually changes throughout the simulations, hence the name "Dynamic LFE accuracy". In this use case, we tried to design experiments where the best LFE predictors would eventually get worse. The LFE predictors that were initially inaccurate could become very accurate after some time. Figure 4.12 gives us an example of how the standard deviation of the LFEs can fluctuate during the games. We have put some boundaries so the value of the standard deviation will not be over one or negative.

Specifically, we wanted to test how our novel Aggregator framework copes with a dynamic environment where everything can change without notice. This use case also provides another real-life scenario where the flexibility predictions' accuracy may vary instantly, trying to resemble the occurrence of natural phenomena and equipment faults. Figure 4.13 provides us with an example of how the fluctuating LFE prediction deviation can impact the CRPS value of the LFEs during the games. For this example, we have set the CRPS selection threshold $\tau = 0.68$, which means that LFEs with an average CRPS score over the line will not participate in flexibility trades of the Aggregator; instead, these LFEs will trade directly with the Grid. At first, we can notice that a low-accuracy LFE, such as LFE_{11} , initially started with CRPS values around 1, however, as the game progressed, it managed to improve its CRPS score. As a result, at the end of the simulation, it was selected to participate with the Aggregator. Another example is that of LFE_0 , which




FIGURE 4.12: The fluctuating Standard Deviation of the LFEs' flexibility estimations in the "Dynamic LFE accuracy" use case

initially started with CRPS values that were always under the threshold, and over time the CRPS was getting higher.




FIGURE 4.13: An example of how the fluctuating standard deviation impacts the CRPS of the LFEs throughout the simulation

Chapter 5 below presents our interesting results regarding the efficiency of the selection mechanisms we developed in this real-life experimental scenario.

Chapter 5

Experimental Evaluation

In this chapter, we present the experimental evaluation of the proposed Aggregator framework. The results are averages of 30 simulations with different properties and 2000 timeslots each, so that we can get statistically significant results. We divided PowerTAC's customer assets into 12 heterogeneous LFEs comprising various DER assets, such as solar panels, batteries, electric vehicles, households, and interruptible consumption users. This makes the simulations much more realistic, since every LFE has different attributes. Some of these attributes refer to the electricity consumption/production profiles and the total storage capacity.

We use PowerTAC's day-ahead wholesale market to trade flexibility in our experiments. Specifically, the simulation is divided into timeslots, with each day having 24 timeslots resembling the hours of the day. Like in many real energy markets, in our experimental scenarios too, all offers should be submitted before 12:00 pm, so there is a need for each LFE to provide the Aggregator with at least 24-hour-ahead flexibility predictions. Furthermore, we make the very important assumption that when an LFE generates renewable energy, it will always sell it to the Grid, with the price depending on the experimental scenario, so no electricity is wasted.

There are various types of electricity markets currently operating; however, in this work we focus on the most common type, which is the day-ahead market. The day-ahead market is based on the predicted load or consumption of the DER assets. There, the base wholesale market prices typically reflect the electricity price when it can flow freely without transmission constraints across the Grid operators' territory. When that is impossible, Grid operators recalculate the costs to account for congestion on transmission lines by allowing prices to differ by location. As a result, areas with high demand and limited electric resources usually have higher prices than those with sufficient generation relative to the load. The PowerTAC simulator can simulate this behavior and provide realistic supply and demand models.

For this reason, we chose to evaluate our DER aggregation framework in two payment scenarios. First, we employ the CRPS pricing mechanism to compute the reward from the Grid to the Aggregator and the singleton LFEs,

because CRPS has been proven to be very resourceful in Smart Grid applications [23], [86]. For the second scenario, we are adopting a simple linear payment mechanism currently used by some European Aggregator platforms, such as United Kingdoms' "Flexible Power" [87]. The second payment function is based only on the pre-arranged KWh prices of the day-ahead market and the actual flexibility delivered to the Grid. Therefore, in order to have a better perspective, we evaluate the performance of our proposed DER aggregation framework both in future Smart Grid settings, using CRPS-based payments, and in the current Grid's settings too.

5.1 Results: CRPS-based Grid Payments

In the literature, we have already seen strictly proper scoring rules, such as "CRPS", to be successfully used in Smart Grid applications. Specifically, [23] have shown that when the Grid pays a DER asset using the CRPS metric, it incentivizes it to give honest load flexibility forecasts, thus improving the stability of the Grid in general. So, in this first batch of experimental scenarios, we will be using the following CRPS-based payments because we expect CRPS to be of value both to the Smart Grid in terms of stability and to the Aggregator in terms of monetary gains.

$$V_{Grid,Agg}(t) = CRPS_{Agg}(t) \cdot \log|flex_{Agg}(t)| \cdot flex_{Agg}(t) \cdot p(t) \quad (5.1)$$

The CRPS "Grid-to-Aggregator" payment function is shown in Equation 5.1, and the CRPS "Grid-to-LFE" is described in Equation 5.2. The payment for the Aggregator and the LFEs that trade directly with the Grid is calculated using the Aggregator's or LFE's CRPS value, along with a specific KWh price $p(t)$ and the actual flexibility contribution $flex(t)$.

$$V_{Grid,LFE_i}(t) = CRPS_{LFE_i}(t) \cdot \log|flex_{LFE_i}(t)| \cdot flex_{LFE_i}(t) \cdot p(t) \quad (5.2)$$

In the rest of this Section, we will present the experimental results when using the CRPS payment mechanisms to reward the Aggregator and the individual single LFEs that trade their flexibility directly with the Grid. Additionally, we have formulated two different use cases with respect to the accuracy of the LFEs' flexibility predictions. In the first use case, the Gaussian distribution that simulates the theoretical LFE predictions remains the same for the duration of the simulations. In the second use case, we try to take a more realistic approach where LFEs will not be either "high-accuracy" or "low-accuracy" predictors but a mix of both. This means that in the "Dynamic LFE-Accuracy" use case, the probability distribution of every LFE will gradually change during the simulations.

5.1.1 Static LFE accuracy

The first use case we experiment with is the “Static LFE accuracy”, where the LFEs have a constant standard deviation in their flexible load predictions. More specifically, LFE_1 resembles the best flexibility predictor with $\sigma \approx 0$ and LFE_{12} the worse with $\sigma \approx 1$.

Total Profits(€) of LFEs via the:	Static LFE accuracy	
	Aggregator	Grid
Experiment 1: Simple Selection	440 K	152 K
Experiment 2: Using CRPS only	426 K	161 K
Experiment 3: DQN Reward 1	563 K	19 K
Experiment 3: DQN Reward 2	525 K	57 K
Experiment 4: Singleton LFEs	0	519 K
Experiment 5: All LFEs participate	604 K	0

TABLE 5.1: The total profits of LFEs via the Aggregator and the Grid at the end of the simulations for every experimental scenario

As mentioned previously in Section 4.1.1, in the “Simple Selection” scenario, we evaluate the use of the Prediction Accuracy pricing mechanism 3.5.1 when the Aggregator pays the LFEs. In every other experimental scenario, the Aggregator utilizes the CRPS-based payment mechanism 3.5.2 to distribute its profits to the LFEs.

First, we start by comparing the total profits of the Aggregator in every experimental scenario. Table 5.1 shows the total profits of the LFEs that sold their energy via the Aggregator and the total profits for the energy sold directly via the Grid at the end of the simulations. We can see that the profits of the Aggregator are similar in the “Simple Selection” and “Using CRPS only” scenarios. This similarity in profits has to do with the selection mechanisms and the thresholds we used, as we will see in other figures later.

Additionally, when comparing the “Singleton LFEs” and the “All LFEs” we can see that the total profits of the LFEs via the Aggregator (604K) are higher than when they traded exclusively with the Grid (519 K), showing that our DER Aggregator guarantees greater payments in the static LFE-Accuracy use case. As expected, the Aggregator has the highest profits when every LFE




FIGURE 5.1: “Static LFE accuracy”: Total profits (€) of the LFEs in every experimental scenario

participates and trades its flexible loads through it; however, it is very important to note that this does not mean that because the Aggregator in total had the highest profits, every LFE was more profitable too. Specifically, Figure 5.1 shows the total amount of money each LFE accumulated in every experimental scenario. We can see that the best LFEs with the highest accuracy ($LFEs_{1-4}$) had higher profits in the first two experimental scenarios. By contrast, the remaining LFEs had the highest total profits either when every LFE was participating in the Aggregator, or for some LFEs, when we used the DQN selection method. Also, the total profits of the LFEs, especially for “good” predictors, were higher than their profits when being singleton LFEs, showcasing, once again, one of the benefits of participating in the Aggregator.

Figure 5.2 shows the average flexibility selling price for all the LFEs in different experimental scenarios where LFEs’ selection rules are different. At first, we can notice that the first 4 LFEs, which are better predictors than the rest, are selling at a higher average price than the other LFEs, this is a result of both the prediction accuracy and the amount of flexibility they traded. Additionally, the better performing scenarios for the “best predictors” are when using either the Simple Selection or the CRPS scoring rules, while the most profitable scenarios for the “worse predictors” are when every LFE participates in the Aggregator so their prediction error can be balanced out by the whole LFE team when the total flexibility of the Aggregator is computed.

Furthermore, Figure 5.3 depicts the selling channel and the total profits of each individual LFE. We can see that in this Static LFE-accuracy use case,




FIGURE 5.2: “Static LFE accuracy”: Comparison of the average payment (€) per MWh sold of every LFE in all the experimental scenarios

all LFEs traded their energy only through one channel, excluding LFE_5 in the first two experimental scenarios. In the next section, where we examine the dynamic LFE-accuracy case, we will have the chance to make a better comparison of the total profits of each LFE and their trading channels.

Another comparison that can be helpful can be derived from observing how the total LFE Energy was sold (Figure 5.4). In the first two scenarios, our Aggregator framework prefers to select only the best-predictor LFEs, hence guaranteeing a higher flexibility price, as seen in previous figures. Additionally, with the DQN selection method, the LFEs traded most of their energy via the Aggregator, leaving out only a small portion to be directly traded with the Grid. The latter explains the overall lower selling prices of the DQN selection method. Also, we can see that the LFEs produced the same energy in every case, which was intended because we wanted a fair comparison of every selection method.

However, for this analysis, we also need Figure 5.5 to understand better how often the LFEs were selected to participate in the Aggregator by the way they traded their flexibility. We can see in the Simple Selection, and the CRPS-only



FIGURE 5.3: “Static LFE accuracy”: Comparison of the amount of money (€) LFEs got via the Aggregator and the Grid

selection methods choose the best-performing LFE predictors to participate and trade flexibility with the Aggregator. At the same time, the rest are not selected because of their low accuracy. In scenarios 3 and 5, most LFEs are participating in the Aggregator, which is the main reason their average price is higher, as seen in the previous Figure 5.2.

Figure 5.6 compares the flexibility selling channel of every LFE, and it is directly related to the selection percentage depicted in Figure 5.5. Here again, we can observe that our aggregation framework only selects LFEs that it




FIGURE 5.4: “Static LFE accuracy”: Comparison of the Total LFE Energy sold(GWh) for every experimental scenario




FIGURE 5.5: Static LFE accuracy: Comparing the LFE participation percentages (in the Aggregator cooperative) for every experimental scenario

“trusts”, based on the scoring functions it used to trade flexibility together. On the contrary, it prompts “low-accuracy” LFEs to trade their flexibility directly with the Grid.

In summary, we observed that it is more profitable for the best-performing LFEs to use the Simple Selection and CRPS scoring rules because these scenarios value more truthful and reliable LFEs. On the contrary, the worse-performing LFEs prefer to participate in a cooperative so they can cover their inability to predict accurately. Probably the more efficient solution for this “Static LFE accuracy” use case would be to divide the LFEs into two teams based on their CRPS and create two different Aggregators, one for each team



FIGURE 5.6: “Static LFE accuracy”: Comparing the flexibility selling channel (directly to Grid or via the Aggregator) and the total flexibility (GWh) sold

using the most suitable payment mechanism for each. For instance, the most accurate predictor LFEs could be paid using the CRPS payments, shown in Section 3.5.2, and the less accurate using another.

5.1.2 Dynamic LFE accuracy

The second use case we are testing is the “Dynamic LFE accuracy”, where the LFE predictors with the lowest prediction accuracy gradually improve

while the best-performing ones slowly degrade. This use case provides another real-life scenario where the flexibility predictions' accuracy may change without any notice, trying to resemble the occurrence of natural phenomena and equipment faults. In this use case, we expect to observe different results when using different selection methods than when we used the "static LFE accuracy" scheme.

Total Profits(€) of LFEs via the:	Dynamic LFE accuracy	
	Aggregator	Grid
Experiment 1: Simple Selection	318 K	304 K
Experiment 2: Using CRPS only	302 K	335 K
Experiment 3: DQN Reward 1	482 K	73 K
Experiment 3: DQN Reward 2	516 K	52 K
Experiment 4: Singleton LFEs	0	579 K
Experiment 5: All LFEs participate	578 K	0

TABLE 5.2: The total profits of LFEs via the Aggregator and the Grid at the end of the simulations for every experimental scenario

The total profits of the Aggregator and the LFEs are important indicators showing us the essence of every experimental scenario. Table 5.2 depicts the total profits of all the LFEs of the Aggregator and the way they sold their flexibility. We can see that in this use case, too, the selection methods and the thresholds we have set have created similar profit margins for both Scenario 1 and Scenario 2. Specifically, the LFE teams have made around half of their profits (318K and 302K) through the Aggregator flexibility trades, while the other half (304K and 335K) was traded directly with the Grid. Furthermore, the two different DQN selection methods preferred to trade their overall flexibility via the Aggregator, thus resulting in 482K and 516K profits against the 73K and 52K profits via the direct Grid trading. In contrast with the "static LFE accuracy" use case, here, both the "Singleton LFEs" and "All LFEs participate" use cases have similar profits but, of course, from opposite sources: directly from the Grid the first, just from the Aggregator the second.

This similarity of the payments of the two exactly opposite baseline scenarios might seem counter-intuitive; however, it is a characteristic of the CRPS pricing mechanism, which is used for the calculation of the payments from the



FIGURE 5.7: "Dynamic LFE accuracy": Comparison of the amount of money (€) LFEs got via the Aggregator and the Grid

Grid to the singleton LFEs and the Aggregator. In detail, CRPS rewards better LFEs that are truthful and accurate, which means that even if an LFE has low-accuracy predictions, it informs the Grid/Aggregator about it with the provided uncertainty distribution. Therefore, that LFE can receive higher payments than a higher-accuracy LFE predictor that has miscalculated its uncertainty distribution. This CRPS attribute was demonstrated earlier when we discussed the CRPS pricing mechanism. In experimental Section 5.2, where we use simple usage fees to calculate the Grid payments, we expect to see many different results regarding the total profits of the LFEs.

Table 5.2 can only demonstrate one aspect of the results; however, with Figure 5.7, we can see the total profits of each LFE and the way they traded their flexibility. In contrast to the “static LFE accuracy” use case, here we can see that the Aggregator had a more dynamic selection strategy to adjust to the fluctuating LFE accuracy during the simulations. Here again, the first two Experimental scenarios have very similar distributions because of the threshold parameters τ we had designated.

The DQN selection method was less dynamic than we would have expected. Like the previous “static accuracy” use case, we can see that the DQN selection methods selected only specific LFEs to participate in the Aggregator. However, the selected LFEs are different in every case, showing the dynamic selection capability of the decision-making of the reinforcement learning algorithms. Additionally, this DQN selection method might have performed even better if we had experimented with a more detailed state definition. For example, adding other more suitable environmental variables into the state space of this DQN RL model might have helped to detect more efficiently the dynamic changes of the LFEs’ accuracy, resulting in optimal selection mechanisms that maximize the profits of the LFEs and the Aggregator.




FIGURE 5.8: “Dynamic LFE accuracy”: Total profits (€) of the LFEs in every experimental scenario

We continue breaking down the profit-inspection problem even further by observing the total profits of each LFE in every experimental scenario, clearly depicted in Figure 5.8. In this “dynamic-accuracy” use case, it is evident that the “CRPS only” selection method provides the higher total incomes for most LFEs, excluding $LFE_{6,7,8,9}$ where the difference is minimal and LFE_{12} where the “DQN reward-1” method resulted in almost double the profits. Also, the profits of the “Simple Selection” selection method were similar to the CRPS

method but not as good. This similarity in the profits shows that the “CRPS only” selection method is the most suitable for environments with highly unstable LFE-predictors. Additionally, both DQN methods achieved higher profits than the baseline scenarios for only specific LFEs, showing that the DQN method needs to be further developed, in order to be applied in this dynamic environment.




FIGURE 5.9: “Dynamic LFE accuracy”: Comparison of the average payment(€) per MWh sold of every LFE

As mentioned earlier, in this dynamic use case, the CRPS and Simple Selection Rules offer the best total flexibility profits for almost every LFE; however, these methods also provide the highest average payment per MWh for most LFEs. As we can see in Figure 5.9, only the $LFEs_{6,7,8}$ had another method nearly better than the first two scoring rules. Also, in this use case, the DQN selection methods work slightly better for some LFEs. Still, they result in low profits for some individual LFEs, which also indicates that this DQN method does not perform as well as expected when the LFEs’ prediction accuracy fluctuates as much. Additionally, we can see that neither baseline scenario (scenarios 4 and 5) can reward all the LFEs with a higher payment per MWh.

In the rest of this experimental results section, we will view the performance of our DER aggregation framework using each selection method from a Flexibility-sold point of view. Figure 5.10 displays the channel through which the LFEs traded their flexibility. We can observe that for the first two selection methods, which were the highest rewarding ones for most LFEs, only around 2 GWh out of the 5 GWh total, were sold via the Aggregator. This fact verifies that CRPS and the Simple Selection method thrive in dynamic environments by carefully selecting which LFEs to participate in the Aggregator cooperative, as shown in Figure 5.11. On the contrary, the DQN methods traded




FIGURE 5.10: “Dynamic LFE accuracy”: Comparison of the Total LFE Flexibility sold(GWh) for every experimental scenario




FIGURE 5.11: “Dynamic LFE accuracy”: Comparing the LFE participation percentages (in the Aggregator cooperative) for every experimental scenario

around 4 GWh of energy via the Aggregator and only 1 GWh directly with the Grid. This shows that trading with the Aggregator only sometimes guarantees the highest profits. Additionally, the results from the two baseline scenarios indicate that neither is the ultimate solution, which also verifies that DERs need “smarter” Aggregators schemes, such as the DER framework we propose, and not simplistic everyone-in or everyone-out DER aggregation strategies.

Moreover, Figure 5.12 depicts the Aggregator participation percentages for every LFE. We can see that the Simple Selection and the CRPS-only method

have quite a unique selection scheme. LFEs 1,7,8,9 and 12 were participating in the Aggregator tradings under 25% of the time. LFE_3 had around 30% participation percentage, and the rest were participating in the Aggregator flexibility trading around 50-75 % of the time. It is impressive that no LFE was selected at all available times. The other Selection methods had more static results; for example, in the “Experimental scenario 3.1: Reward 1”, LFEs 6,7,9 and 11 had a close to 0 % participation percentage, and all the rest LFEs had an overall 100% participation percentage.



FIGURE 5.12: “Dynamic LFE accuracy”: Comparing the flexibility selling channel (directly to Grid or via the Aggregator) and the total flexibility (GWh) sold

Figure 5.12 depicts how the LFEs traded their flexibility. We can observe that the reason the first two Selection methods performed better was due to the way they selected which LFEs to participate in. On the contrary, the DQN scoring rule did not adapt to this dynamic environment and, most of the time, selected the same LFEs to participate, resulting in lower selling prices we previously discussed.

Overall, in both use cases, we can see that it is inefficient for all the LFEs to trade directly with the Grid, judging by the results of the "Singleton LFEs" scenario. Additionally, we found out that the CRPS Scoring Rule performs the best when the accuracy of the LFE team is higher because it gives incentives for truthful predictions. Also, the Simple Selection method proved to be very useful since it can perform equally well or usually slightly worse than the CRPS. Finally, the "All LFEs" scenario, where every LFE participates in the Aggregator, is more profitable for LFEs with lower prediction accuracy because the weight of the prediction errors of some individuals is split between all the members of the cooperative. Thus, in that case, the average reward for better predictors gets lower, while the average reward for worse predictors gets higher.

5.2 Results: Simple Grid-to-Aggregator Payment

In the previous experimental-results section, we have seen how a strictly proper rule, such as CRPS, can contribute to the pricing mechanisms of the Smart Grid. In the past, CRPS has been used for applications in the Smart Grid and was proven efficient for truthful and accurate predictors, regardless of whether they were small DERs or large Aggregators. However, nowadays, this complex but strictly proper payment mechanism is not used for calculating payments in energy trading. Instead, energy markets use simple mechanisms based mainly on the flexibility contribution and not so much on the accuracy of the predictions. As we have seen in Section 3.5.3, the simple payment is calculated as shown in Equations 5.3 and 3.14.

$$flex_{diff}(t) = flex_i(t) - \widetilde{flex}_i(t) \quad (5.3)$$

Initially, the difference $flex_{diff}(t)$ between the estimated flexibility and the actually delivered flexibility is calculated. The Grid pays each KWh delivered using the predetermined price $p(t)$ but only for the number of KWh that was agreed to deliver. If the delivered flexibility is greater than the one estimated, then the rest flexibility is paid using the standard price $p_{standard}(t)$. The final payment $V_{Grid,i}$ is calculated using Equation 3.19, re-stated here as Equation 5.4 for convenience.

$$V_{Grid,i} = \begin{cases} flex_i(t) \cdot p(t) & flex_{diff}(t) \leq 0 \\ \widetilde{flex}_i(t) \cdot p(t) + flex_{diff}(t) \cdot p_{standard}(t) & flex_{diff}(t) > 0 \end{cases}$$

(5.4)

In the rest of this section, we will demonstrate the experimental results of our Aggregator framework when the Grid uses “Simple Payments” to pay the delivered flexibility upon the agreed KWh price.

5.2.1 Static LFE accuracy

The first use case we are going to examine is when LFEs have static prediction accuracy through the course of the simulations. Unlike the previous section, where CRPS was used to calculate the Grid-to-Aggregator/DER payments, here, LFEs don’t get penalized for being inaccurate; hence the optimal bidding strategy from the side of the LFEs would be to bid for much larger amounts of flexibility than the ones they can deliver, so their contribution, even if less than the one promised, is awarded through the pre-agreed price $p(t)$ which is usually higher than the current electricity price $p_{standard}(t)$. However, in our simulations, we have made the assumption that the Aggregator and the LFEs that directly trade with the Grid will never exploit this feature since we haven’t developed any penalty mechanism for this edge case.

Total Profits(€) of LFEs via the:	Static LFE accuracy	
	Aggregator	Grid
Experiment 1: Simple Selection	61 K	41K
Experiment 2: Using CRPS only	61 K	43 K
Experiment 3: DQN Reward 1	75 K	7 K
Experiment 3: DQN Reward 2	70 K	10 K
Experiment 4: Singleton LFEs	0	92 K
Experiment 5: All LFEs participate	81 K	0

TABLE 5.3: The total profits of LFEs via the Aggregator and the Grid at the end of the simulations for every experimental scenario

It should be noted that the total profits accumulated by LFEs in the previous section using the “CRPS based Grid-to-Aggregator” payment cannot be compared with the profits gained when using the “Simple” Grid-to-Aggregator

payment because of fundamental differences in the way they are stated. For example, the CRPS Grid-to-Aggregator payment has an additional logarithmic term multiplying the final payment intending to incentivize the formation of larger cooperatives. By contrast, the “Simple” does not have such a multiplier. Instead, we can compare the relative profits of each LFE when using different pricing mechanisms.

First, we can take a look at Table 5.3 showing the total profits of the LFEs for this static LFE accuracy use case and the way they earned their money in every experimental scenario. We can observe that the first two experimental scenarios reward the LFEs similarly, in detail, they have accumulated 61 thousand euros by tradings through the Aggregator and around 41 thousand euros by trading directly with the Grid. Similarly, the rewards are distributed for the two DQN scenarios, with around 75K euros gained from the Aggregator trades while only around 10K euros via the Grid. However, the most striking feature of this table is the difference in the profits when all LFEs trade alone (Experiment four) and when all LFEs participate in the Aggregator (Experiment 5). In particular, the total profits of the singleton LFEs stood at 92K euros, which is 11K euros higher than the total payments when every LFE trades through the Aggregator. The latter clearly shows that traditional Aggregator schemes, where every DER asset participates in the Aggregator, are not cost-efficient when using simple usage payments that do not individually penalize LFEs for their inaccurate estimations. On the contrary, our aggregation framework manages to accumulate higher amounts of profit for the simple selection and the CRPS-only methods, summing to a total of 102K€ (61K via the Aggregator and 41K via the Grid) and 104K (61K + 43K) euros, respectively.




FIGURE 5.13: “Static LFE accuracy”: Total profits (€) of the LFEs in every experimental scenario

Figure 5.13 illustrates the total profits of every LFE after the end of the simulations. In this use case, we can infer that the usage of the simple selection and CRPS methods guarantees the highest total payments when compared to every other selection method. In particular, the CRPS-only selection method manages to differentiate from the simple selection method by a few thousand euros in most cases, such as for LFE_1 ; while in the worse cases, it leads to equal total payments for LFEs. Furthermore, we can observe that when using simple usage payments, the singleton LFEs scenario achieves greater overall performance when compared to DQN methods and the second baseline method.

Figure 5.14 presents the average payment per KWh for every LFE in the six tested scenarios. This figure verifies the previous findings, clearly depicting the superiority of the CRPS-only method that is deployed by this framework. The simple selection method is the runner-up accumulating an average payment of a few more cents for LFE_3 and LFE_4 , while in every other case, it scores a few euros less than the CRPS. The “singleton LFEs” baseline scenario is the third most efficient option having a difference in its average payments fluctuating from 1 euro up to 3 euros less, which is significant since each LFE sells many MWh. In the final place come the remaining selection methods having a much worse performance varying up to 5 euros less than the CRPS-only method.




FIGURE 5.14: “Static LFE accuracy”: Comparison of the average payment (€) per MWh sold of every LFE in all the experimental scenarios

We continue this extensive evaluation of the proposed DER aggregation framework by taking a deeper look at the way LFEs accumulated their profits (Figure 5.15). The bar charts illustrated here have certain similarities with the respective charts of the previous sections. We can observe that the first two

selection methods select only the best-predictor LFEs to participate in the Aggregator trades, while the DQN methods select most of the LFEs.




FIGURE 5.15: “Static LFE accuracy”: Comparison of the amount of money (€) LFEs got via the Aggregator and the Grid

Observing Figure 5.16, we can see that all LFEs traded approximately five GWh in total per simulation. The first two scenarios managed to split it almost equally, with around three GWh sold via the Aggregator and over two GWh directly to the Grid. Also, DQN methods delivered most of their energy via the Aggregator and traded only half a GWh directly to the Grid. However, it is interesting that in this “Simple Payment” results section, the

Singleton LFEs delivered less energy overall when compared to when every LFE participates in the Aggregator (Experiment 5).




FIGURE 5.16: “Static LFE accuracy”: Comparison of the Total LFE Energy sold(GWh) for every experimental scenario




FIGURE 5.17: Static LFE accuracy: Comparing the LFE participation percentages (in the Aggregator cooperative) for every experimental scenario

Furthermore, we can take a closer look at the participation percentages for every LFE during the simulations by observing Figure 5.17. Here, we can verify our initial expectations and previous observations that Simple Selection and CRPS-only methods select only the best LFE predictors; while DQN methods usually select most of the LFEs with other criteria, probably deriving from the way their RL models were trained.

Finally, Figure 5.18 compares the flexibility selling channel of the LFEs. Here, we can also see that every LFE sells a different amount of energy; more specifically, the first five LFEs, LFE_9 , and LFE_{12} sell over 0.5 GWh while the rest LFEs sell under 0.5 GWh. This heterogeneity allows us to have more realistic simulations trying to imitate the performance of real-life DERs in the Smart Grid.




FIGURE 5.18: “Static LFE accuracy”: Comparing the flexibility selling channel (directly to Grid or via the Aggregator) and the total flexibility (GWh) sold

5.2.2 Dynamic LFE accuracy

The last use case we study is the “Dynamic LFE-accuracy”. In this use case, the accuracy of the predictions of the LFEs fluctuates during the games, leading to unstable performances from the side of the predictors. We do not expect to see significant differences because, as mentioned earlier, the lack of a penalizing mechanism, such as CRPS, can result in unfair payments. This means that if an LFE_i trades a specific amount of flexibility $flex_{LFE_i}(t)$ that is *much less* (low-accuracy predictor) than the agreed upon the flexibility to deliver ($\widetilde{flex}_{LFE_i}(t)$) to the Grid, then an Aggregator that trades exactly (perfect predictor) the pre-agreed amount of flexibility $\widetilde{flex}_{Agg}(t)$ will be paid the exact same amount of money per KWh as the very inaccurate LFE. This underscores a significant problem of existing energy markets that greatly impacts the Grid’s stability. However, we can take a closer look at the experimental results of this use case when we use our novel DER aggregation framework and then compare it with other use cases, such as the CRPS payment used to calculate the reward for the Grid-to-Aggregator/LFEs transactions.

Total Profits(€) of LFEs via the:	Dynamic LFE accuracy	
	Aggregator	Grid
Experiment 1: Simple Selection	47 K	68 K
Experiment 2: Using CRPS only	46 K	78 K
Experiment 3: DQN Reward 1	72 K	24 K
Experiment 3: DQN Reward 2	72 K	18 K
Experiment 4: Singleton LFEs	0	122 K
Experiment 5: All LFEs participate	80 K	0

TABLE 5.4: The total profits of LFEs via the Aggregator and the Grid at the end of the simulations for every experimental scenario

Initially, we can observe Table 5.4, which compares the total amount of money LFEs have gathered by trading via the Aggregator and by trading directly with the Grid. First, Simple Selection and CRPS-only methods have accumulated similar amounts of money, the only difference being that the second had 10K euros more directly with the Grid than the first. Additionally, it seems like the first reward method performs better in this use case since it accumulated 6K euros more than the “DQN reward 2” method. Interestingly, we can

observe that the Singleton LFEs have gathered 122K euros which is equal to the total profits of the CRPS-only scenario, showing that Singleton LFEs are efficient when there are a lot of fluctuations in the prediction accuracy of the LFEs, and the Grid uses simple payments. However, we need to observe the other figures, too, to have a spherical view of the results so we can characterize how efficient our Aggregator framework is.



FIGURE 5.19: “Dynamic LFE accuracy”: Comparison of the amount of money (€) LFEs got via the Aggregator and the Grid

With the help of Figure 5.19, we can get more insights into the performance of each selection method and of our framework in general. Figure 5.19 depicts the amount of money the LFEs accumulated by trading via the Grid and directly with the Grid. As it is natural, the first two selection methods have similar outcomes, however, the CRPS-only method results in slightly higher payments overall, we will extract more detailed information about the exact

amount of money from Figure 5.20. The DQN selection methods resulted in less total profits for every LFE when compared to the LFEs of Simple Selection, CRPS-only, and Singleton LFEs method. Furthermore, another notable observation is that when every LFE participates in the Aggregator, the total accumulated money is much less.




FIGURE 5.20: “Dynamic LFE accuracy”: Total profits (€) of the LFEs in every experimental scenario

In detail, by observing the total profits of each LFE in every scenario (Figure 5.20), we can see that the CRPS-only selection method is the most profitable in total for every LFE without any exception, showing the efficiency of our novel aggregation framework when combined with strictly proper rules, such as CRPS. To our surprise, the Singleton LFEs method follows quite closely with only a few hundred euros less than the first method. The third place belongs to the Simple selection method, which achieved solid performance of only a few hundred euros less than the Singleton LFEs. The two DQN methods follow with a much larger difference of a few thousand euros in total, showing that the dynamic reinforcement learning models failed to learn the optimal policy that maximizes the total profits of each LFE. Finally, it seems that it is quite wasteful—because potential profit is lost—to participate in traditional Aggregators (that have no selection methods and every DER joins) when the Grid pays the Aggregator using simple payments.

Now that we have seen the total profits, we can also observe another important metric showing the average payment of euros per energy sold (MWh), as depicted in Figure 5.21. We can notice that the CRPS-only selection method resulted in the highest average payment per MWh for every LFE. In particular, it follows a similar trend to that of the total profits of each LFE (Figure 5.20), meaning that the Singleton LFEs method is the runner-up with a slight difference varying from 0 euro (for LFE_1) up to 2 euros for LFE_2 . The




FIGURE 5.21: “Dynamic LFE accuracy”: Comparison of the average payment(€) per MWh sold of every LFE

Simple Selection method is quite profitable as well, rewarding LFEs on average with 1 euro less per MWh than the CRPS-only method. On the other hand, the DQN selection methods and the traditional Aggregator scheme (“All LFEs”) resulted in an average of more than five euros lost per MWh, rendering these methods quite wasteful.




FIGURE 5.22: “Dynamic LFE accuracy”: Comparison of the Total LFE Flexibility sold(GWh) for every experimental scenario

The lack of a penalizing mechanism in the “Simple payments” when calculating the payment from the Grid to the Aggregator has created an exciting result. More specifically, Figure 5.22 demonstrates the total flexibility traded

by all LFEs in every experimental scenario. Combining this figure with Figure 5.22 about the average payments, we can notice that even when the Simple Selection method and the CRPS-only method delivered the same amount of flexibility, as shown in Figure 5.22, their average payments had a notable difference of 1 euro per MWh less for the Simple Selection. By contrast, the Singleton LFEs scenario, where LFEs traded flexibility directly with the Grid, resulted in similar average payments per MWh with the CRPS-only scenario, which was the most profitable, as verified in Figure 5.21. Based on the previous observation about the difference in average prices in different experimental scenarios, we could intuitively state that there are potentially other selection patterns based on our Aggregator framework that could possibly increase even more the average payments of every LFE. Indeed, there is still a lot of space for additional experimentation when using other LFE selection methods or payment mechanisms.

Figure 5.23 depicts the participation percentages of the LFEs in the Aggregator's flexibility trades. The selection patterns of the first two methods are quite familiar to us since they are similar to the corresponding ones from the CRPS Grid-to-Aggregator payments. We can see that the threshold τ we have set for these methods has managed to create dynamic selection methods that are also cost-efficient. On the other hand, the DQN methods selected at all times only the best LFEs according to their evaluations, 7 LFEs for DQN $reward_1$ and 8 LFEs for DQN $reward_2$.




FIGURE 5.23: “Dynamic LFE accuracy”: Comparing the LFE participation percentages (in the Aggregator cooperative) for every experimental scenario

Finally, Figure 5.24 depicts the total energy sold by each LFE and their selling channel at the end of the simulations. The results depicted here are similar to the previous CRPS-based Grid-to-Aggregator payments.



FIGURE 5.24: “Dynamic LFE accuracy”: Comparing the flexibility selling channel (directly to Grid or via the Aggregator) and the total flexibility (GWh) sold

5.3 Summarizing the Results

Concluding, we can notice that in every case, the LFEs sold the same amount of energy regardless of the experimental scenario, the static or dynamic LFE accuracy, and the payment method used by the Grid. Generating a similar amount of energy in every simulation was an important condition that we made sure to ensure. Therefore, we could get statistically significant results on how our aggregation framework impacts the profits the LFEs make, especially when compared to a traditional Aggregator without any selection methods(Experiment 5) and when no Aggregator is gathering the DERs (Experiment 4).

The aforementioned results clearly showcase the rise in the monetary gains of all stakeholders attributed to the use of our proposed DER aggregation framework. Interestingly, not only a single method but at least two, e.g., the Simple Selection method, CRPS-only selection method, and DQN selection in specific use cases, are performing better than the traditional Aggregator scheme where every DER is selected to participate in the collective Aggregator flexibility trades. Also, our DER aggregation framework outperforms, in terms of final profits, the other baseline scenario, that of “Singleton-LFEs”, when using both CRPS payments from Grid-to-Aggregator and simple payments from Grid-to-Aggregator, even though in the second environment, the difference is not as significant as in the first.

Chapter 6

Conclusions and Future Work

In this work, we put forward a novel flexibility aggregation framework for effectively integrating DERs in the Smart Grid. Our framework consists of a novel multiagent architecture along with various selection and pricing mechanisms found in the literature or introduced for the first time in this paper. Additionally, we presented a systematic experimental evaluation of our framework using data from the highly realistic PowerTAC simulator, which we extended to allow for the incorporation of flexibility aggregators and related entities and mechanisms.

We formulated various experimental scenarios to test our flexibility aggregation framework's performance in realistic situations. In particular, we have designed and tested two separate Grid-to-Aggregator payment scenarios. Therefore, we evaluated the performance of our proposed DER aggregation framework both in Smart Grid settings, using CRPS-based payments, and in the current Grid's settings, using "Simple" Payment.

Our results show that our framework can successfully contribute to the effective integration of DERs in the Grid, increasing their profits while also supporting the Grid's stability. The results show that LFEs that are accurately stating their flexibility predictions are always rewarded better when participating in our proposed DER aggregation framework than in baseline methods. Additionally, less-accurate LFEs managed by our framework are awarded better in most cases, even though their stated flexibility estimations can potentially have more significant errors.

The CRPS Selection mechanism delivered the highest average payments in most evaluation settings and for most LFEs, with only a few exceptions for less accurate LFE predictors. The second best method was the Simple Selection mechanism, being a close second to CRPS Selection when using CRPS Grid-to-Aggregator payment. By contrast, the singleton LFEs scenario was the second best method when using the "Simple" Grid-to-Aggregator payment. Finally, the DQN selection mechanisms were an interesting novel addition to our framework; however, their performance was not as consistently excellent as when the Aggregator utilized the CRPS Selection mechanism,

thus creating a need for further experimentation with this novel Reinforcement Learning approach for LFE cooperative formation.

6.1 Future Work

In terms of future work, there is still much space for experimentation with respect to cost-efficient DER Aggregation methods. In particular, we intend to experiment with alternative selection and pricing mechanisms by using other scoring rules that might apply better in some scenarios. For instance, we could experiment with another scenario that will use the “Simple Selection mechanism” (introduced in this thesis) to select which LFEs to participate in the aggregator flexibility trading when combined with the CRPS pricing mechanism to split the profits to the participating LFEs. Furthermore, there have been various scoring rules used in the literature for forecast evaluation in the past, such as Binary Scoring Rules [89], Surrogate Scoring Rules [90], Energy Score [20], etc., that could possibly have a remarkable impact on the way the new Selection mechanisms would perform. Additionally, many pricing mechanisms are used, either in Smart Grid settings [91], [92] or not [93], [94], that could be applied and tested on our framework to incentivize different DER behaviors.

Furthermore, we also plan to study methods (readily supported by our framework) that allow LFEs to replace inefficient DER assets after evaluating them. Specifically, we have not examined in this study how the LFE formation process—the process of selecting which DER to participate in the LFE—impacts the final prediction accuracy and the final profits of the other DERs in the LFE. Therefore, it could be reasonable to consider that some specific DERs could get better payments if they join another LFE with different characteristics. The results of this kind of experimentation would be most likely noteworthy.

Another research direction deriving from this work would be the study of the dynamic formation of multiple cooperatives with different properties in the same DER aggregation framework. This intuitive idea was generated naturally by observing the behavior of the LFEs in the experimental results of Chapter 5. In detail, it would be interesting to incorporate Aggregators that create more than one LFE cooperative, potentially using different mechanisms for each depending on its attributes, resulting in even higher profits for the LFEs.

Finally, it would be quite insightful to study the performance of our novel flexibility aggregation framework when placed in a larger Smart Grid environment with multiple competing aggregators. Therefore, enhancing our framework with the ability to include multiple aggregators competing for the representation of efficient LFEs, is also interesting future work that resembles with greater accuracy how the real DER markets will evolve.

Bibliography

- [1] P. Berrill, A. Arvesen, Y. Scholz, H. C. Gils, and E. G. Hertwich, "Environmental impacts of high penetration renewable energy scenarios for europe," *Environmental Research Letters*, vol. 11, no. 1, p. 014 012, Jan. 2016.
- [2] T. Adefarati and R. C. Bansal, "Integration of renewable distributed generators into the distribution system: A review," *Iet Renewable Power Generation*, vol. 10, pp. 873–884, 2016.
- [3] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings, "Putting the 'smarts' into the smart grid," *Communications of the ACM*, vol. 55, pp. 86–97, 2012.
- [4] M. A. Judge, A. Khan, A. Manzoor, and H. A. Khattak, "Overview of smart grid implementation: Frameworks, impact, performance and challenges," *Journal of Energy Storage*, vol. 49, p. 104 056, 2022.
- [5] M. L. Tuballa and M. L. S. Abundo, "A review of the development of smart grid technologies," *Renewable & Sustainable Energy Reviews*, vol. 59, pp. 710–725, 2016.
- [6] V. C. Gungor, D. Sahin, T. Koçak, S. Ergut, C. Buccella, C. Cecati, and G. P. Hancke, "A survey on smart grid potential applications and communication requirements," *IEEE Transactions on Industrial Informatics*, vol. 9, pp. 28–42, 2013.
- [7] I. C. Lampropoulos, T. Alskaif, J. Blom, and W. G. van Sark, "A framework for the provision of flexibility services at the transmission and distribution levels through aggregator companies," *Sustainable Energy, Grids and Networks*, vol. , pp. 175–191, 2019.
- [8] H. Sawall, A. Scheuriker, and D. Stetter, "Flexibility definition for smart grid cells in a decentralized energy system," in *Proceedings of the 7th International Conference on Smart Cities and Green ICT Systems*, Funchal, Madeira, Portugal: SCITEPRESS, 2018, pp. 130–139.
- [9] C. Iraklis, J. Smend, A. Almarzooqi, and A. Mnatsakanyan, "Flexibility forecast and resource composition methodology for virtual power plants," *ICECET*, vol. , pp. 1–7, 2021.
- [10] G. D. Zotti, S. A. Pourmousavi, J. M. Morales, H. Madsen, and N. K. Poulsen, "Consumers' flexibility estimation at the tso level for balancing services," *IEEE Transactions on Power Systems*, vol. 34, pp. 1918–1930, 2019.
- [11] L. Gkatzikis, I. Koutsopoulos, and T. Salonidis, "The role of aggregators in smart grid demand response markets," *IEEE Journal on Selected Areas in Communications*, vol. 31, pp. 1247–1257, 2013.

- [12] O. Okur, P. W. Heijnen, and Z. Lukszo, "Aggregator's business models in residential and service sectors: A review of operational and financial aspects," *Renewable and Sustainable Energy Reviews*, vol. 139, p. 110702, 2021.
- [13] O. Okur, "Aggregators' business models for flexibility from electricity consumers," Ph.D. dissertation, Delft University of Technology, 2021.
- [14] O. Valarezo, T. Gómez, J. P. Chaves-Ávila, L. Lind, M. Correa, D. U. Ziegler, and R. Escobar, "Analysis of new flexibility market models in europe," *Energies*, 2021.
- [15] Z. Broka and K. Baltputnis, "Handling of the rebound effect in independent aggregator framework," *2020 17th International Conference on the European Energy Market (EEM)*, vol. , pp. 1–5, 2020.
- [16] G. S. Wagh, S. Gupta, and S. Mishra, "A distributed privacy preserving framework for the smart grid," *2020 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)*, vol. , pp. 1–5, 2020.
- [17] N. Gai, K. Xue, P. He, B. Zhu, J. Liu, and D. He, "An efficient data aggregation scheme with local differential privacy in smart grid," *2020 16th International Conference on Mobility, Sensing and Networking (MSN)*, vol. , pp. 73–80, 2020.
- [18] C. Ibrahim, I. Mougharbel, H. Y. Kanaan, N. A. Daher, S. W. Georges, and M. Saad, "A review on the deployment of demand response programs with multiple aspects coexistence over smart grid platform," *Renewable and Sustainable Energy Reviews*, vol. 162, p. 112446, 2022.
- [19] T. Rawat, K. R. Niazi, N. Gupta, and S. Sharma, "A two stage interactive framework for demand side management in smart grid," *2019 IEEE 16th India Council International Conference (INDICON)*, vol. , pp. 1–4, 2019.
- [20] T. Gneiting and A. E. Raftery, "Strictly proper scoring rules, prediction, and estimation," *Journal of the American Statistical Association*, vol. 102, pp. 359–378, 2007.
- [21] R. S. Sutton and A. G. Barto, "Reinforcement learning: An introduction," *IEEE Transactions on Neural Networks*, vol. 16, pp. 285–286, 2005.
- [22] W. Ketter, J. Collins, and P. P. Reddy, "Power tac: A competitive economic simulation of the smart grid," *Energy Economics*, vol. 39, 2013.
- [23] V. Robu, R. Kota, G. Chalkiadakis, A. Rogers, and N. Jennings, "Cooperative virtual power plant formation using scoring rules," 1, vol. 26, Sep. 2021, pp. 370–376. DOI: [10.1609/aaai.v26i1.8172](https://doi.org/10.1609/aaai.v26i1.8172).
- [24] X. Fang, S. Misra, G. Xue, and D. Yang, "Smart grid — the new and improved power grid: A survey," *IEEE Communications Surveys & Tutorials*, vol. 14, pp. 944–980, 2012.
- [25] H. Farhangi, "The path of the smart grid," *IEEE Power and Energy Magazine*, vol. 8, 2010.
- [26] V. C. Gungor, D. Sahin, T. Koçak, S. Ergüt, C. Buccella, C. Cecati, and G. P. Hancke, "Smart grid technologies: Communication technologies and standards," *IEEE Transactions on Industrial Informatics*, vol. 7, pp. 529–539, 2011.
- [27] I. W. https://en.wikipedia.org/wiki/Smart_grid, *Smart grid*. 2022.

- [28] T. Hubert and S. Grijalva, "Realizing smart grid benefits requires energy optimization algorithms at residential level," *ISGT 2011*, pp. 1–8, 2011.
- [29] M. H. Rehmani, M. Reisslein, A. Rachedi, M. Erol-Kantarci, and M. Radenkovic, "Integrating renewable energy resources into the smart grid: Recent developments in information and communication technologies," *IEEE Transactions on Industrial Informatics*, vol. 14, pp. 2814–2825, 2018.
- [30] A. Nordling and S. Pädam, "Social costs and benefits of smart grid technologies 1," 2018.
- [31] 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*, vol. 16, pp. 1619–1635, 2015.
- [32] A.-M. Koufakis, E. S. Rigas, N. Bassiliades, and S. D. Ramchurn, "Towards an optimal ev charging scheduling scheme with v2g and v2v energy transfer," *SmartGridComm*, vol. , pp. 302–307, 2016.
- [33] M. M. R. Sahebi, E. A. Duki, M. Kia, A. Soroudi, and M. Ehsan, "Simultaneous emergency demand response programming and unit commitment programming in comparison with interruptible load contracts," *Iet Generation Transmission & Distribution*, vol. 6, pp. 605–611, 2012.
- [34] P. Vithayasrichareon, G. Mills, and I. Macgill, "Impact of electric vehicles and solar pv on future generation portfolio investment," *IEEE Transactions on Sustainable Energy*, vol. 6, pp. 899–908, 2015.
- [35] A. Ulbig and G. Andersson, "Analyzing operational flexibility of electric power systems," *2014 Power Systems Computation Conference*, pp. 1–8, 2014.
- [36] J. M. M. González, A. J. Conejo, H. Madsen, P. Pinson, and M. Zugno, "Integrating renewables in electricity markets: Operational problems," 2013.
- [37] J. Aghaei and M. I. Alizadeh, "Demand response in smart electricity grids equipped with renewable energy sources: A review," *Renewable & Sustainable Energy Reviews*, vol. 18, pp. 64–72, 2013.
- [38] H. Chen, T. N. Cong, W. Yang, C. Tan, Y. Li, and Y. Ding, "Progress in electrical energy storage system: A critical review," *Progress in Natural Science*, vol. 19, pp. 291–312, 2009.
- [39] C. Dang, X. Wang, C. Shao, and X. Wang, "Distributed generation planning for diversified participants in demand response to promote renewable energy integration," *Journal of Modern Power Systems and Clean Energy*, vol. , pp. 1–14, 2019.
- [40] M. Barbero, C. Corchero, L. C. Casals, L. Igualada, and F.-J. Heredia, "Critical evaluation of european balancing markets to enable the participation of demand aggregators," *Applied Energy*, vol. 264, 2020.
- [41] R. van der Veen and R. A. Hakvoort, "The electricity balancing market: Exploring the design challenge," *Utilities Policy*, vol. 43, pp. 186–194, 2016.

- [42] T. A. Muhaimin, "Electricity market of the future: Assessing economic feasibility and regulatory constraints for demand response aggregators in europe," 2015.
- [43] T. Schittekatte, V. Deschamps, and L. Meeus, "The regulatory framework for independent aggregators," *SSRN Electronic Journal*, vol. 34, 2021.
- [44] N. G. ESO, *Monthly balancing services summary 2018/19*, 2019.
- [45] A. P. Dawid, "The basic theory of the prequential approach to data analysis is described, and illustrated by means of both simulation experiments and applications to real data-sets.," 2016.
- [46] P. H. Garthwaite, J. B. Kadane, and A. O'Hagan, "Statistical methods for eliciting probability distributions," *Journal of the American Statistical Association*, vol. 100, pp. 680–701, 2005.
- [47] J. Vermorel, *Continuous ranked probability score (crps)*, 2016. [Online]. Available: <https://www.lokad.com/continuous-ranked-probability-score>.
- [48] *Basic diagram of reinforcement learning*. [Online]. Available: <https://www.kdnuggets.com/>.
- [49] A. Nair, B. McGrew, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Overcoming exploration in reinforcement learning with demonstrations," *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6292–6299, 2018.
- [50] C. J. C. H. Watkins and P. Dayan, "Q-learning," *Machine Learning*, vol. 8, pp. 279–292, 2004.
- [51] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. A. Riedmiller, A. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, pp. 529–533, 2015.
- [52] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. A. Riedmiller, "Playing atari with deep reinforcement learning," *ArXiv*, vol. abs/1312.5602, 2013.
- [53] H. Mao, M. Alizadeh, I. Menache, and S. Kandula, "Resource management with deep reinforcement learning," *Proceedings of the 15th ACM Workshop on Hot Topics in Networks*, 2016.
- [54] V. François-Lavet, P. Henderson, R. Islam, M. G. Bellemare, and J. Pineau, "An introduction to deep reinforcement learning," *Found. Trends Mach. Learn.*, vol. 11, pp. 219–354, 2018.
- [55] F. Zhang and Q. Yang, "Energy trading in smart grid: A deep reinforcement learning-based approach," *2020 Chinese Control And Decision Conference (CCDC)*, pp. 3677–3682, 2020.
- [56] P. Zhang, C. Wang, G. S. Aujla, and R. S. Batth, "Reledp: Reinforcement-learning-assisted dynamic pricing for wireless smart grid," *IEEE Wireless Communications*, vol. 28, pp. 62–69, 2021.
- [57] T. Lu, X. Chen, M. B. McElroy, C. P. Nielsen, Q. Wu, and Q. Ai, "A reinforcement learning-based decision system for electricity pricing plan

- selection by smart grid end users," *IEEE Transactions on Smart Grid*, vol. 12, pp. 2176–2187, 2021.
- [58] T. Qian, C. Shao, X. Wang, and M. Shahidehpour, "Deep reinforcement learning for ev charging navigation by coordinating smart grid and intelligent transportation system," *IEEE Transactions on Smart Grid*, vol. 11, pp. 1714–1723, 2020.
- [59] N. Naseri, S. Talari, W. Ketter, and J. Collins, "Dynamic retail market tariff design for an electricity aggregator using reinforcement learning," *Electric Power Systems Research*, 2022.
- [60] S. Orfanoudakis, "Developing an autonomous agent for automated electricity trading," Diploma thesis, School of Electrical and Computer Engineering, Technical University of Crete, Chania, Greece, 2021. DOI: <https://doi.org/10.26233/heallink.tuc.88411>.
- [61] J. Casazza and F. Delea, *Understanding Electric Power Systems: An Overview of Technology, the Marketplace, and Government Regulation*. Jun. 2010, pp. 1–320, ISBN: 9780470484180. DOI: [10.1002/9780470588475](https://doi.org/10.1002/9780470588475).
- [62] W. Ketter, J. Collins, and M. d. Weerd, "The 2020 power trading agent competition (march 30, 2020).," *ERIM Report Series Reference No. 2020-002*, 2020. DOI: <http://dx.doi.org/10.2139/ssrn.3564107>. [Online]. Available: <https://ssrn.com/abstract=3564107>.
- [63] S. Orfanoudakis, S. Kontos, C. Akasiadis, and G. Chalkiadakis, "Aiming for half gets you to the top: Winning powertac 2020," in *Multi-Agent Systems - 18th European Conference, EUMAS 2021, Virtual Event, June 28-29, 2021, Revised Selected Papers*, A. Rosenfeld and N. Talmon, Eds., ser. Lecture Notes in Computer Science, vol. 12802, Springer, 2021, pp. 144–159.
- [64] Y. Zheng, B. Celik, S. Suryanarayanan, A. A. Maciejewski, H. J. Siegel, and T. M. Hansen, "An aggregator-based resource allocation in the smart grid using an artificial neural network and sliding time window optimization," *IET Smart Grid*, vol. 4, no. 6, pp. 612–622, 2021.
- [65] F. Carducci, "Tools and methods to study the integration of flexibility assets in the future smart grid," Ph.D. dissertation, Università Politecnica delle Marche, 2017.
- [66] R. H. Jacobsen, D. Gabioud, G. Basso, P.-J. Alet, A. G. Azar, and E. S. M. Ebeid, "Semiah: An aggregator framework for european demand response programs," *2015 Euromicro Conference on Digital System Design*, vol. , pp. 470–477, 2015.
- [67] A. Castelo-Becerra, W. Zeng, and M.-Y. Chow, "Cooperative distributed aggregation algorithm for demand response using distributed energy storage devices," *2017 North American Power Symposium (NAPS)*, vol. , pp. 1–6, 2017.
- [68] Z. Elmrabet, N. Kaabouch, H. E. Ghazi, and H. E. Ghazi, "Cyber-security in smart grid: Survey and challenges," *ArXiv*, vol. abs/1809.02609, 2018.
- [69] C. Efthymiou and G. Kalogridis, "Smart grid privacy via anonymization of smart metering data," *2010 First IEEE International Conference on Smart Grid Communications*, pp. 238–243, 2010.

- [70] M. Shafie-khah, E. Heydarian-Forushani, G. J. Osório, F. A. S. Gil, J. Aghaei, M. Barani, and J. P. S. Catalão, "Optimal behavior of electric vehicle parking lots as demand response aggregation agents," *IEEE Transactions on Smart Grid*, vol. 7, pp. 2654–2665, 2016.
- [71] P. Goyal, A. Sharma, S. Vyas, and R. Kumar, "Customer and aggregator balanced dynamic electric vehicle charge scheduling in a smart grid framework," *ICEPES*, vol. , pp. 276–283, 2016.
- [72] S. I. Vagropoulos, D. K. Kyriazidis, and A. G. Bakirtzis, "Real-time charging management framework for electric vehicle aggregators in a market environment," *IEEE Transactions on Smart Grid*, vol. 7, 2016.
- [73] A. Perez-Diaz, E. H. Gerding, and F. McGroarty, "Coordination and payment mechanisms for electric vehicle aggregators," *Applied Energy*, vol. 212, pp. 185–195, 2018.
- [74] A. Perez-Diaz, E. Gerding, and F. McGroarty, "Coordination of electric vehicle aggregators: A coalitional approach," in *17th International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2018*, vol. 1, Southampton: International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS), 2018, pp. 676–684.
- [75] G. Chalkiadakis, E. Elkind, and M. Wooldridge, "Computational aspects of cooperative game theory," 6, vol. 5, 2011, pp. 1–168. DOI: [10.2200/S00355ED1V01Y201107AIM016](https://doi.org/10.2200/S00355ED1V01Y201107AIM016).
- [76] M. Wooldridge, "An introduction to multiagent systems, second edition," 2009.
- [77] G. Chalkiadakis, V. Robu, R. Kota, A. Rogers, and N. R. Jennings, "Co-operatives of distributed energy resources for efficient virtual power plants," in *The 10th International Conference on AAMAS - Volume 2*, ser. AAMAS '11, Taipei, Taiwan: IFAAMAS, 2011, pp. 787–794, ISBN: 0982657161.
- [78] C. Akasiadis and G. Chalkiadakis, "Decentralized large-scale electricity consumption shifting by prosumer cooperatives," in *Volume 285: ECAI 2016*, The Netherlands: IOS Press Ebooks, 2016, pp. 175–183.
- [79] M. Kubli and P. Canzi, "Business strategies for flexibility aggregators to steer clear of being "too small to bid"," *Renewable & Sustainable Energy Reviews*, vol. 143, 2021.
- [80] Y. Kabalci, "A survey on smart metering and smart grid communication," *Renewable & Sustainable Energy Reviews*, vol. 57, pp. 302–318, 2016.
- [81] F. Lezama, J. P. Soares, P. Hernandez-Leal, M. Kaisers, T. Pinto, and Z. A. Vale, "Local energy markets: Paving the path toward fully trans-active energy systems," *IEEE Transactions on Power Systems*, vol. 34, 2019.
- [82] K. Wang, Y. Wang, X. Hu, Y. Sun, D.-J. Deng, A. V. Vinel, and Y. Zhang, "Wireless big data computing in smart grid," *IEEE Wireless Communications*, vol. 24, pp. 58–64, 2017.
- [83] A. Mnatsakanyan, H. Albeshr, A. M. A. Marzooqi, and E. Bilbao, "Blockchain-integrated virtual power plant demonstration," *2nd International Conference on Smart Power & Internet Energy Systems*, vol. , pp. 172–175, 2020.

- [84] K. Kursawe, G. Danezis, and M. Kohlweiss, "Privacy-friendly aggregation for the smart-grid," in *PETS*, Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 175–191.
- [85] L. Bosman and S. Darling, "Difficulties and recommendations for more accurately predicting the performance of solar energy systems during the snow season," in *2016 IEEE International Conference on Renewable Energy Research and Applications (ICRERA)*, , IEEE, 2016, pp. 567–571.
- [86] V. Robu, G. Chalkiadakis, R. Kota, A. Rogers, and N. R. Jennings, "Rewarding cooperative virtual power plant formation using scoring rules," *Energy*, vol. 117, pp. 19–28, 2016.
- [87] *Flexible power*. [Online]. Available: <https://www.flexiblepower.co.uk/>.
- [88] *Flexible power payment mechanism*. [Online]. Available: <https://www.flexiblepower.co.uk/downloads/52>.
- [89] E. Neyman, G. Noarov, and S. M. Weinberg, "Binary scoring rules that incentivize precision," *Proceedings of the 22nd ACM Conference on Economics and Computation*, 2021.
- [90] Y. Liu, J. Wang, and Y. Chen, "Surrogate scoring rules," *Proceedings of the 21st ACM Conference on Economics and Computation*, 2020.
- [91] J. Ma, J. Deng, L. Song, and Z. Han, "Incentive mechanism for demand side management in smart grid using auction," *IEEE Transactions on Smart Grid*, vol. 5, pp. 1379–1388, 2014.
- [92] H. T. Rose, A. Rogers, and E. H. Gerding, "A scoring rule-based mechanism for aggregate demand prediction in the smart grid," in *AAMAS*, 2012.
- [93] C. Decker and R. Wattenhofer, "A fast and scalable payment network with bitcoin duplex micropayment channels," in *Safety-critical Systems Symposium*, 2015.
- [94] M. Assili and R. Ghazi, "An improved mechanism for capacity payment based on system dynamics modeling for investment planning in competitive electricity environment," *Energy Policy*, vol. 36, pp. 3703–3713, 2008.