

Coordination Mechanisms for Agent-Based Smart Grids

Tesis presentada por

Radu-Casian Mihailescu

Departamento de Arquitectura y Tecnología de Computadores,
Ciencias de la Computación e Inteligencia Artificial

para la obtención del título de
Doctor en Informática



2013

D. Sascha Ossowski, Profesor Catedrático de la Universidad Rey Juan Carlos,
con NIE X-1805919-M,

CERTIFICA: Que D. Radu-Casian Mihailescu, licenciado en Informática, ha realizado en el Dpto. de Arquitectura de Computadores y Ciencias de la Comunicación e Inteligencia Artificial, bajo su dirección, el trabajo correspondiente a la tesis doctoral titulada:

Coordination Mechanisms for Agent-Based Smart Grids

Revisado el presente trabajo, estima que puede ser presentado al tribunal que ha de juzgarlo. Y para que conste a efecto de lo establecido en la normativa reguladora del tercer ciclo de la Universidad Rey Juan Carlos, autoriza su presentación.

Móstoles, 2014.

Fdo.: Sascha Ossowski

Acknowledgments

Concluding the doctoral dissertation has been an exciting and challenging experience, which exceeded my expectations in many ways, including all the ups and downs along the way. While any list will inevitably be incomplete, I would like to mention here several principal actors.

First and foremost, I want to thank my thesis advisor, Sascha Ossowski, who guided and supported me throughout the PhD process, always providing me with insightful advice. Our discussions provided me with lots of valuable pieces of advice, while I have learned to pursue intuitive insights into mature results.

I would like to thank all the co-authors of the papers published, especially to Matteo Vasirani for benefiting from his comments and questions, to Matthias Klusch for his expert and valuable feedback and to René Schumann for taking such a fervent interest in my work.

I am also grateful that the course of my PhD research offered me the possibility of three research visits: I would like to thank Michael Schumacher who received me at HES-SO (Switzerland) during my research stay for his pertinent suggestions, Matthias Klusch who supervised my research visit at DFKI (Germany) and Jörg Müller for my time spent at TU-Clausthal (Germany). Many thanks also to Jürgen Dunkel and Moser Fagundes who accepted to provide the reviews of the thesis promptly and with due diligence. I would also like to thank my former advisers, respectively, Stefan Holban from UPT (Romania) and Weng Fai from NUS (Singapore) for supporting my early research endeavors and interest in AI.

I am especially thankful to all the past and present members of the Artificial Intelligence Group at URJC: to Alberto and Holger for many fruitful discussions of both scientific and non-scientific nature, to all the good friends and colleagues (Carlos, Cliff, Jose, Jota, Levan, Luis, Marin, Ramon, Ruben, Roberto, Dani, Victor) for their willingness to help on numerous occasions and for creating a pleasant working atmosphere. Special thanks to Carmen for her help with various administrative tasks.

Finally, special acknowledgement goes to all of my friends and family! To all of them, I say Thanks for being there and supporting me throughout this journey. I am particularly grateful to my parents for their unconditional love and for being a constant source of encouragement in all my endeavors.

Dedicated to my loving parents.

Thesis advisor
Prof. Sascha Ossowski

Author
Radu-Casian Mihailescu

Abstract

Governments around the globe are heavily investing in upgrading the ageing infrastructure of the electricity grid. The imperative for this is driven primarily by regulatory requirements and the high cost of inefficiently delivering energy. The infrastructure is continuously improving, as more and more smart meters are installed, coupled with the proliferation of controllable loads and distributed generation. However, network operators, utilities, as well as end-consumers and small-scale producers are struggling to extract value from such systems and are exploring new ways for optimizing the performance of their deployed assets.

This thesis introduces a multiagent approach for modelling the emerging complexity of the energy industry. The multiagent system paradigm is an ideal candidate for delivering a framework that captures the inherent distributed and dynamic nature of *smart grids*. While the traditionally centralized management of the system becomes less viable in the context of distributed generation and controllable loads, the underlying thread of this thesis advocates the design and implementation of *coordination mechanisms* capable to integrate and manage a large-scale integration of such devices via agent-based control.

We begin by proposing *dynamic micro-grids*, a new conceptual organization of the network, adequate to integrate today's traditional users into an interactive, internet-like system, in the sense that power flow will become bidirectional and energy management will become distributed in the grid due to the many actors involved in the operation of the system. The mechanisms proposed for micro-grid formation are oriented towards producing sub-systems of the grid that are exhibiting reduced transmission losses and an efficient utilization of renewables, as well as endowing the system with self-adaptation techniques for coping with dynamic environments.

We further aim to enhance the operation of the micro-grid formations by mainly focusing on two aspects. On one hand (*supply-side*) we are concerned with seamlessly integrating distributed generation to ensure a reliable service of energy supply comparable to what a large power plant delivers today. We first address the economic benefits of *virtual power plants* in a game-like setting and then go on to propose a DCOP-based formalism for solving the schedule generation problem, while accounting for the stochastic behavior of intermittent supply. On the other hand (*consumer-side*), we apply the use of game mechanics to drive the behaviour of *prosumers* towards efficient grid-wise use of energy. In order to cope with the challenges faced by current electricity networks, we propose a game layer on top of the electricity grid infrastructure and the use *coordination mechanisms* as a catalyst for change, encouraging participation of prosumers in the energy field towards reduced costs, lower carbon generation and increased grid resilience in the form of *demand response* and *demand-side management* solutions. Finally, we propose a *collusion detection* mechanism that complements the above-mentioned solutions in the sense of inspecting for patterns where agents tacitly cooperate through illicit monopoly tactics to manipulate energy markets.

Contents

1	Introduction	1
1.1	Objectives	4
1.2	Methodological remarks	7
1.3	Structure of the thesis	9
2	Background	11
2.1	The Vision behind Smart Electricity Grids	11
2.2	Open Multi-Agent Systems and Agentification of Smart Grids	19
2.3	Coordination methods in MAS	22
2.3.1	Emergent Coordination and Agents	23
2.3.2	Distributed Constraint Optimization Problem and Agents	25
2.3.3	Game theory and Agents	28
2.4	Coordination and control in the <i>smart grid</i>	33
2.4.1	Enhancing the control of the electricity grid	33
2.4.2	Coordination and market adaptation to support transitioning to a smarter grid	41
2.5	Simulator	52
2.6	Chapter Summary	56
3	Coordination Mechanisms for <i>Dynamic Micro-Grid Formation</i>	58
3.1	Microgrid Solution Approach	60
3.2	Problem representation	62
3.3	MAS-based algorithm for Microgrid formation	65
3.4	Experimental results	69
3.5	Coalitional mechanism for Microgrid self-adaptation	73
3.6	Experimental results	78
3.7	Discussion	82
4	Coordination Mechanisms for <i>Virtual Power Plants</i>	85
4.1	An Investigation of <i>Cooperation</i> in Smart Grids	87
4.1.1	Agent-based Model	88

4.1.2	Simulation Results	91
4.2	Stochastic Optimization for Virtual Power Plants	99
4.2.1	Problem definition & Model	101
4.2.2	Stochastic extension	104
4.2.3	Solving the StochDCOP formalization	107
4.2.4	Experimental setting	111
4.3	Discussion	115
5	Coordination Mechanisms for <i>Forward & Intraday Markets</i>	119
5.1	Adaptive Demand-Side Management	124
5.1.1	MAS Framework for Smart Distribution Grids	125
5.1.2	The Non-Cooperative Game Formulation	129
5.1.3	A MAS-based Control Mechanism	135
5.1.4	Numerical results	139
5.1.5	Discussion	142
5.2	Dynamic Coalition Formation for Power Regulation	146
5.2.1	A Coalitional Game Formulation for Intraday Power Regulation	148
5.2.2	BSV-Stable Payoff Distribution for Dynamic Environments . .	153
5.2.3	Privacy-Preserving Layer	157
5.2.4	<i>eCOOP</i> : Putting it all together	159
5.2.5	Empirical Evaluation	162
5.2.6	Discussion	169
5.3	Chapter Summary	173
6	Detection of <i>Collusive Behavior</i> in Energy Markets	174
6.1	Energy Markets	177
6.2	Collusion Detection Methodology	178
6.2.1	CPA based Behavioral Screening Phase	180
6.2.2	Verification Phase for Collusive Coalitions	183
6.2.3	Validation Phase via Principal Component Analysis	185
6.3	Experimental results	187
6.3.1	Data preparation	187
6.3.2	Case study	188
6.4	Chapter Summary	193
7	Conclusions	195
7.1	Contributions	195
7.2	Future Research Directions	198
A	Author Bibliography	200

B Resumen en castellano	203
B.1 Antecedentes	203
B.2 Objetivos	206
B.2.1 Estructura de la tesis	210
B.3 Metodología	211
B.4 Conclusiones	213
B.4.1 Contribuciones	214
B.4.2 Líneas de investigación futuras	217
References	219

List of Figures

1.1	Areas of contributions (with our contributions marked in light green)	4
1.2	Methodology	7
2.1	Challenges in smart grids	12
2.2	Electricity consumption growth 2007-2050 (Blue Map Scenario [2]) . .	14
2.3	Portion of variable generation of electricity by region 2007-2050 (Blue Map Scenario [2])	14
2.4	European Roadmap for Smart Grid deployment with associated costs [72]	17
2.5	Integrating devices into the Smart Home model	18
2.6	Approaches to smart grids	34
2.7	Factorial hidden Markov model representation for 2 devices; $x_t^{(i)}$ is the hidden state i of a device at time t ; \bar{y}_t is the aggregated power at time t ; $y_t^{(i)}$ the individual power of a device at time t ; shaded nodes denote observed variables and unshaded nodes represent hidden variables. . .	35
2.8	The IDAPS multi-agent architecture [138]	40
2.9	PowerMatcher architecture [87]	43
2.10	Schematic overview of MAS architecture for PEV scheduling [177] . .	45
2.11	The MAS levels of organization. [36]	47
2.12	Aggregated net-benefits for business stakeholders [176]	50
2.13	Context structure of the simulator.	53
3.1	Percentage gain of our approach against optimal	70
3.2	Evolution of performance over time of different communication network configurations	71
3.3	Social welfare of the system in number of steps. a) normal behavior ; b)induced variations	78
3.4	Histogram representation for the utilities of the coalitions	81
4.1	The evolution of the number of agents with respect to their strategy space for θ with $\mu=5$	91

4.2	Payoffs as a function of variable energy availability θ and number of agents in the system.	92
4.3	The evolution of the agents behavior with respect to the strategy selection rule for θ with $\mu=9$. a) <i>gradient ascent</i> ; b) <i>win-stay, lose-shift</i> ; c) <i>imitate best strategy</i> ; d) <i>regret minimization</i> ;	96
4.4	a) Population share in mixed-strategy repeated games b) Average pay-offs in repeated games	97
4.5	DCOP representation for a VPP scenario	102
4.6	Reformulating the DCOP representation of the VPP scenario	104
4.7	Revised DCOP representation of the VPP to include multiple time-slots	105
4.8	<i>StochDCOP</i> representation of the VPP scenario for a single time-slot t	106
4.9	Computational time for generating a schedule	113
5.1	Generation optimization to meet demand for scenario: <i>dispatch cost of generator 1 < cost generator 2 < cost generator 3</i>	120
5.2	Average daily consumption diagram. L represents the desired level of demand.	121
5.3	Market timeline	122
5.4	Agent-based Smart Distribution Grid	129
5.5	The sigmoid function f_i and the role of parameters: i) a determines the slope of the curve; ii) p determines the point around which the performance growth is the fastest.	131
5.6	Generic user utility as a function of incurred cost for different values of k	133
5.7	Best response curves and equilibrium of a two-player game	139
5.8	Realization of the scenario	140
5.9	Evolution of user (a) bids and (b) power allocation	141
5.10	Average utility vs. power supply Θ	141
5.11	Structure of the coalitional game	150
5.12	Example of generating payoff configuration through bilateral Shapley value	157
5.13	Example of homomorphic cryptosystem	158
5.14	Load curve for an individual household	163
5.15	Comparison over aggregated demand patterns	166
5.16	Evolution of (a) <i>load factor</i> , (b) <i>carbon emissions</i> and (c) <i>percentage savings</i> for different degrees of elasticity of demand	167
5.17	Simulation results for year-round load factor comparison	169
6.1	Mechanism flow-chart	178
6.2	a) Plot of the trading activity P_1^6 corresponding to market agent a_6 for January 2012. b) The associated CUSUM chart for the time-series of market agent a_6	188

6.3	Plot of the covariance matrix for the collusion candidate set of market agents using a grayscale range of colors. Correlations of 1.0 are plotted in white, while no correlation are plotted in black. The diagonal elements of the matrix represent self-correlations and are thus all 1.	190
6.4	a) Boxplot representation of the trading dataset; b) Scree Test.	191
6.5	Orthonormal principal component coefficients for each agent and principal component scores for each observation.	192
B.1	Areas of contributions (with our contributions marked in light green)	206
B.2	Methodology	211

List of Tables

0.1	Publications related to each chapter referenced in author's bibliography	xvii
5.1	Cost of stability as a function of network load	142
6.1	Results of the Change-Point Analysis with MSE Estimates for market agents with confidence level above or equal to 95%, representing the collusion candidate set.	189

List of Abbreviations

AEMO	Australian Energy Market Operator
AHAM	Association of Home Appliance Manufactures
CHP	Combined Heat and Power
CP	Change-Point
CPA	Change-Point Analysis
CUSUM	Cumulative Sum Charts
DCF	Dynamic Coalition Formation
DCOP	Distributed Constraint Optimization Problem
DCSP	Distributed Constraint Satisfaction Problem
DER	Distributed Energy Resource
DG	Distributed Generation
DLC	Direct Load Control
DNOA	Distribution Network Operator Agent
DSM	Demand Side Management
EV	Electric Vehicles
FHMM	Factorial Hidden Markov Model
GIV	Grid Integrated Vehicle
HCI	Human-Computer Interaction
HMM	Hidden Markov Model
HTTP	Hyper-Text Transfer Protocol
ICT	Information Communication Technology
IDAPS	Intelligent Distributed Autonomous Power Systems
ISO	International Organization for Standardization

ISP	Internet Service Provider
MADP	Mean Absolute Percentage Deviation
MAS	Multi-Agent System
MILP	Mixed-Integer Linear Programming
MIP	Mixed-Integer Programming
MSR	Mean Square Error
NAP	Network Access Point
NIALM	Non-Intrusive Appliance Load Monitoring
PCA	Principal Component Analysis
PEV	Plug-in Electric Vehicles
PHEV	Plug-in Hybrid Electric Vehicles
PJM	Pennsylvania - New Jersey - Maryland Interconnection
PLC	Power Line Communication
POP	Point of Presence
PSO	Particle Swarm Optimization
PSR	Power Supply Restoration
PV	Photovoltaic
REDD	Reference Energy Disaggregation Dataset
RES	Renewable Energy Resource
RTP	Real-Time Pricing
TSO	Transmission System Operator
TU	Transferable Utility
UI	User Interface
VPP	Virtual Power Plant
WESM	Wholesale Electricity Spot Market
XML	Extensible Markup Language

Thesis	Publications	Reference
Chapter 3	Dynamic coalition adaptation for efficient agent-based virtual power plants	(7)
	An Organizational approach to agent-based virtual power stations via coalitional games	(10)
	Dynamic coalition formation and adaptation for virtual power stations in smart grids	(12)
	Towards agent-based virtual power stations via multi-level coalition formation	(13)
Chapter 4	Stochastic multiobjective optimization for hybrid power generation systems	(6)
	An investigation of emergent collaboration under uncertainty and minimal information in energy domains	(9)
	An agent-based model of emergent collaboration under uncertainty in energy domains	(11)
Chapter 5	e COOP: privacy-preserving dynamic coalition formation for power regulation in the smart grid	(1)
	The gamification of agent-based smart grids	(2)
	Desiderata for agent-based power regulation protocols in smart grids	(3)
	A game-theoretic coordination framework for energy-efficient demand-side management in the smart grid	(8)
	A dynamic coalitional game approach to autonomous demand response in the smart grid	(15)
Chapter 6	A Framework for Fraud Discovery via Illicit Agreements in Energy Markets	(14)
	Detection of Collusive Behavior in Energy Markets	(4)
	Detection of Collusive Behavior in Energy Markets (Extended Abstract)	(5)

Table 0.1: Publications related to each chapter referenced in author's bibliography

Chapter 1

Introduction

Prediction is very difficult, especially about the future.

— *Niels Bohr* —

From its humble beginning in October 29, 1969 starting with four host computer systems in 3420 Boelter Hall at UCLA, when ARPANET its precursor gasped its first breaths, the Internet has grown nowadays to a scale of hundreds of millions. Among the best things about it is that no one can claim ownership of the Internet. In fact, the single entity we refer to as the Internet is a collection of networks of various types and sizes. Notably, there is no overall controlling network that regulates the functioning of this system. Its very name comes from the idea of networks being interconnected. The basic underlying structure of the Internet comprises several high-level networks connecting to each other through Network Access Points (NAPs). In order to reach this backbone, home users or companies organized into local area networks (LANs), firstly connect to an Internet Service Provider (ISP). The Point of Presence (POP) is where local users access the ISP's network and thus become part of this network. Of course, it may be the case that the ISP may need to connect to a larger network beforehand of reaching the NAP. In the end, Internet providers interconnect at NAPs agreeing to all intercommunications between them. The interconnection is achieved by high-speed backbones, which are typically fiber optic trunk lines, representing fiber optic cables combined together to increase the capacity. The actual machines on the

internet are distinguished into two categories: servers and clients. Servers are the machines that provide services to other machines. There are multiple kinds of servers such as database servers, file servers, mail servers, print servers, web servers, gaming servers, application servers, etc. Clients are the machines that request tasks to the servers. Accessing a particular service on a server by a client is done using specific protocols. For instance, a client machine running a web browser directs its requests to a specific software server running on the server machine. the server client interaction conforms to the hypertext transfer protocol (*HTTP*), which describes how the client and server will have their conversation.

In 1964, artificial intelligence pioneer Dr. Arthur L. Samuel wrote an article for *New Scientist* titled, "The Banishment of Paper-Work" where he imagined the future of the Internet for the year 1984 prior to its actual existence. Although his predictions proved to be significantly more optimistic in terms of their time to implementation, they were essentially correct:

"One will be able to browse through the fiction section of the central library, enjoy an evening's light entertainment viewing any movie that has ever been produced (for a suitable fee, of course, since Hollywood will still be commercial), or inquire as to the previous day's production figures for tin in Bolivia - all for the asking via one's remote terminal. Libraries for books will have ceased to exist in the more advanced countries except for a few which will be preserved at museums, and most of the world's knowledge will be in machine-readable form. Perhaps it would be more correct to say, all of the world's recorded knowledge will be in this form since the art of programming computers to read printed and handwritten material will have been fully developed. However, the storage problem will make it imperative that a more condensed form of recording be used, a form which will only be machine-readable, and which will be translated into human-readable form by one's computer on demand."

While looking back, we also look forward. By analogy, we envision the developments in the energy sector to transform this industry into an organization closely akin to what the Internet appears to be today. It is clear that the bandwidth (bits per second) did not diminish in time, quite the opposite, it reached unexpected unforeseeable heights. Similarly, we predict that power consumption (Joules per second)

we'll exhibit significant increases that are difficult to even conceive at this point, not least, largely due to undiscovered utilizations. With this assumption in mind the aspect of *efficiency* becomes a major focal point in the existence of future electricity grids. While nowadays a pseudo-efficient utilization of energy is still possible, this will no longer be the case in the future. Imagine if you will different types of generation resources based on wind, solar or tidal power playing the role of servers in the network and clients denoting consumer organized in LAN configurations connected through various ISP-like service providers to the grid. Protocols designed for an efficient resource allocation will regulate the consumer-producer interaction in a way much similar to what HTTP provides for the Internet today. Network actors could coordinate to ensure provision of certain services by entering organizations through the ISP's POP or participating independently by accessing the Smart Grid's NAPs. But most importantly, the experience (UI) of how we are now interacting with the grid is going to be transformed. I summarize below my subjective predictions as to what the so-called Smart Grid is going to look like by 2029:

Currently the grid handles well the problem of distributing energy. Moving energy about long distances is both inefficient and costly. The new grid will specialize in sharing and exchanging energy locally, from intermittent sources to intermittent sinks. With big central power plants largely removed from the system in favour of clean, distributed energy sources, such a configuration of the grid will benefit from openness, robustness and reliability, being less prone to major failures. Importantly, the local aspect is going to be prevalent, as the large scale optimization will result from the local level reconfigurations for increased energy efficiency. Moreover, in this highly dynamic and complex ecosystem of energy supply and consumption, utilities, as we know them today, will cease to exist. The grid will run on a big data platform that will enable to run applications by various stakeholders. Software agents will accurately predict consumption and generation patterns and be delegated much of the consumer planning duties. Agents will negotiate the usage of their devices, whose utilization will no longer be merely thought of as using energy but rather as consuming services through Google-like facilitators. Overall the electricity grid will be an infrastructure manageable by an overlaying network where all smart appliances, the car, the smart phone, the laptop, the washing machine etc. are collaborating transparently as part of a global intelligence assisting humans in optimizing their daily lives.

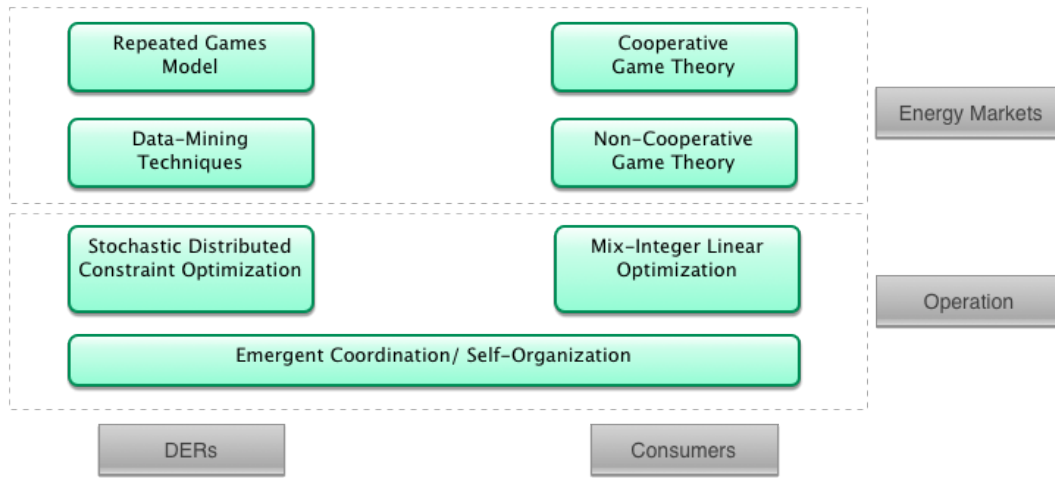


Figure 1.1: Areas of contributions (with our contributions marked in light green)

1.1 Objectives

Having these long-term predictions in mind, in this work we focus on studying distributed mechanisms for the control and management of future smart electricity grids. The concept of *smart grids* has been in use since at least 2003, when it appeared in the article "*Reliability demands will drive automation investments*", which described the necessity for a long term upgrade of the grid instead of short term fixes, regarding transmission capacity and network control systems [21]. Similarly, the US Department of Energy has been amongst the first advocates for such proposals [123].

From a core design principle point of view one can distinguish between two major possibilities in addressing the management problem of smart grids: *centralization* vs. *decentralization*. Centralized approaches are favoured primarily by *utilities*, which do not intend to relinquish their key participation in the energy value chain. Alternatively, a decentralized approach brings into question the role of the utility in future smart grids, bringing to the forefront the actual producers and consumers that make up the network, placing on them the burden of *coordination*. While it is true that democratizing the energy sector calls in for decentralized solutions, this is also what makes the problem hard. How can small-scale *producers* organize to match the reliability offered nowadays by large power plants? How can *consumers* become active

participants in the optimization of energy consumption? How can we minimize the loss of energy in the system and augment its reliability? What forms of organizing the network can we envision that avoid the added cost incurred by today's utilities?

In this thesis we take on the decentralized approach, addressing it from an agent-based perspective [191], where intelligent autonomous agents act on behalf of the actors in the grid, interacting between themselves and with the infrastructure, in order to optimize the state of the network. The objective of the thesis maps onto key milestones that we consider the development of the smart grid must undertake in bringing about the smart grid vision. Figure 1.1 provides a map of the research areas wherein our contributions lie. In light green we mark all those fields that are tackled in this thesis by providing models, solution concepts and algorithms, where coordination plays a key role in problem resolution, while the dark boxes represent the level at which the problem occurs. This work pursues the following set of objectives:

- (O1) **Develop a smart grid model for organizing the actors in the network based on *microgrid* principles of resilience and reduced loss of energy in transmission, that incorporates the dynamic nature of the environment.** The emergence of a complex, dynamic, heterogeneous and distributed system of energy production and consumption requires a radically different approach that can adapt adequately to these new conditions and ensure that energy can be utilized efficiently. The starting point of this work is to investigate a restructuring of the delivery infrastructure and to introduce an approach characterised by openness, robustness and reliability, which could exploit the arising efficiency potentials. Moreover, it is meant that the new setting will lay the ground forward towards providing a framework where negotiations and cooperation among all entities will lead to an increased energy efficiency.
- (O2) The second aspect addressed has to do with the integration of distributed energy resources, that are largely excluded from the wholesale market due to their perceived inefficiency and unreliability. **How can we better instrument the organization of these devices so that they can represent the equivalent distributed decomposition of a big centralized power plant ?** Recall the

Internet analogy in the beginning. Similarly to how the Internet evolved from the multi-user accessible mainframes into a distributed network of machines, the current monolithic, centralized power plants will be replaced by distributed energy generation. The vulnerability of grid infrastructure has called into question the viability of basing energy production on renewables.

(O2.1) **Analyse the evolution of individually rational agreements where agents can counter their individual inefficiencies by group formation techniques.**

(O2.2) **Develop models and algorithms capable to optimize the operation of such formations able to cope with the inherent stochastic environment.**

(O3) Thirdly, we believe firmly that consumers will have a much greater role to play in the future grid control and monitoring. We distinguish several aspects where consumer engagement can impact significantly the efficiency of the grid.

(O3.1) **Design market dynamics that provide the ability to manage demand in such a way that it does away with the need for expensive and inefficient stand-by generation.** Provide regular reductions in demand for certain periods of the day by reaching equilibriums that enable an efficient allocation of the available resources.

(O3.2) **Develop models, mechanisms and algorithms to dynamically control demand to be able to respond to sharp requirements in demand reduction, aiming to balance generation and consumption in near real-time conditions.**

(O3.3) **Develop general purpose and flexible enough model that an agent can deploy to automate the home energy usage by managing deferrable loads.**

(O4) Finally, as the underpinning of our approach is in deploying a decentralized, agent-based perspective to the smart grid challenges, where self-interested agents




Figure 1.2: Methodology

interact in various scenarios, it is important to **devise complementary mechanisms for collusion detection for the energy markets under consideration.**

1.2 Methodological remarks

Today's electric grids are huge structures that are managed in a largely centralised fashion, based on sensors and actuators that are placed on a few (but strategically chosen) parts of the grid. Control room management is usually supported and complemented by field engineers, who may report on the status of transmission lines and can manually perform repair actions when needed. Today's electricity markets are also highly regulated and at best loosely connected, and given the existing governance structures it is still unclear to what level and at what pace the power grid will actually become smarter.

In this thesis, we set out from an optimistic projection of the future to this regard. In particular, we assume an existing physical infrastructure apt to operate in the presence of a diffusion of electrical power generation devices based primarily on renewable resources. Also, the presence of smart nodes with communication and computation capabilities is assumed hereafter for all devices in the system. Moreover, we rely on the existence of a smart metering infrastructure. Finally, and most importantly, we

envision a grid where today's passive consumers of electricity are no longer present, but active players in energy consumption and micro-production populate the system. This has two major consequences for research. First, it becomes important to study and determine the inefficiencies of current electricity networks, along with drawing an understanding of how can we expect that technology deployment may improve on these findings. Second, we need to investigate mechanisms to tackle such problems, which will differ in important aspects since it is unlikely that some general solution can be used in a wide range of settings. As the actors in the grid are becoming active participants we are especially interested in the issues that require coordination in large, dynamic, complex systems. We take a multi-agent perspective for the creation of large-scale systems with predictable behaviour, capable to generate desirable global properties.

Starting off from this setting, our goal is to design, develop and validate *coordination mechanisms* in accordance to the objectives identified in the previous paragraph. The framework under which these objectives are addressed is synthesized in Figure 1.2. There are four main parts of the thesis. The starting point of this work involves model construction, in the sense of a novel reorganizations of the network based on dynamic microgrid formations. The microgrid structuring of the grid leads to an intractable problem formulation, which we address by taking an emergent coordination approach. As depicted in Figure 1.2, once microgrids are formed, we address the two aspects of producer and consumer coordination in a two-level market context by proposing in the following chapters *virtual power plant* and *demand-side management and response* techniques respectively, complemented by *collusion detection* procedures.

We evaluated the performance of our approach using our own custom simulator that seeks to bring together the areas of Figure 1.1 in the context of multiagent environments, using real-world data subject to its availability. Notably though, in reality, the possibilities of testing energy management systems are limited due to the lack of an existing advanced metering infrastructure. When real-world datasets are actually collected there are still many barriers in making them available for research purposes. They generally represent proprietary data own by utilities, which are hes-

itant to disclose this type of information either due to business policies or due to privacy regulation on consumer data. Obtaining reliable data is particularly difficult when the information is necessary for a large number of actors in the system, over long periods and with a high granularity (e.g. device usage at the household level). Hence, the data used for the experiments herein do reflect real-world scenarios to the extent they are available, while sometimes we resort to artificially generate individual profiles from aggregated data and occasionally generate synthetic data from scratch.

1.3 Structure of the thesis

This thesis is structured in the following way:

1. Chapter 2 describes the background that underlies this thesis with an emphasis on the cross-pollinating domains of future smart electricity grids and open multiagent systems. We then go on to revise the state of the art focusing on the approaches related to our abovementioned objectives.
2. Chapter 3 tackles the first objective of the thesis, analysing the state of the grid given a considerable penetration of distributed energy resources. In this new setting we argue that a different type of organization needs to be deployed for efficient grid management based on small-scale modular solutions that are at once both highly adaptable and flexible enough to cope with the stochastic and dynamic environment.
3. In Chapter 4 we introduce and evaluate our proposal for virtual power plant creation and operation for the integration of renewable and distributed energy sources. Given the uncertainty regarding renewable supply our approach devises a way where heterogeneous devices cooperatively mimic the reliability characteristics of a traditional power plant.
4. Next, in Chapter 5 we tackle the coordination gap between supply and demand and introduce our proposal for changing the demand patterns of end consumers, by shifting load to flatten the load profile and maximize the use of de-

ployed assets. We address this problem by differentiating between demand-side management techniques for peak-reduction in the following day and intra-day mechanisms designed to cope with supply-demand imbalances in near real-time scenarios.

5. Chapter 6 addresses the challenges of instituting a collusion free energy market environment. Integration of distributed generation all the way down to the household level is a major shift towards democratizing today's energy supply model. Nevertheless the liberalization of markets is known to be prone towards manipulations and illicit monopoly tactics. We introduce here a mechanism to facilitate the system in avoiding such undesirable states.
6. Finally, Chapter 7 draws the main conclusions of the thesis, discussing the main results achieved and pointing towards future research directions.

Chapter 2

Background

Every phase of evolution commences by being in a state of unstable force and proceeds through organization to equilibrium. Equilibrium having been achieved, no further development is possible without once more over-setting the stability and passing through a phase of contending forces.

— *Kabbalah* —

2.1 The Vision behind Smart Electricity Grids

It comes at no surprise that energy represents the fundamental building block of the modern consumer-based society we know today. At every point in history when the energy sector hit an inflection point, there has occurred an economic revolution. The last two major economic revolutions are no exception to this. In the 1800s, steam power was one of the most important technologies behind the industrial revolution. Likewise, in the 1900s, the mass usage of electricity along with the oil-powered combustion engine gave rise to a new economic paradigm.

Arguably, we are now preceding a third major paradigm shift, this time caused by the emergence of renewable, distributed energy resource. Albeit being a clean source of energy, renewables challenge the current organization of the electricity grid. They are intermittent in the sense that their output can vary significantly over short periods of time according to local environmental conditions. Also, due to their small-

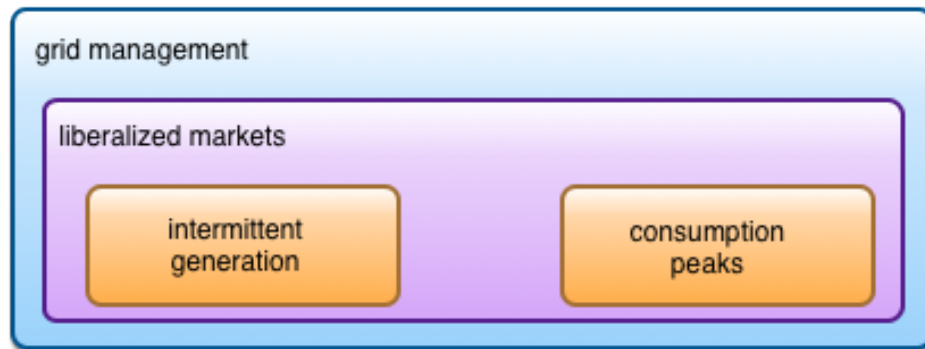


Figure 2.1: Challenges in smart grids

scale output, for a sustainable provision of energy, there is a need for hundreds of thousands or even millions of such devices embedded across both the transmission and distribution networks. Moreover, the presence of these devices implies a bidirectional flow of electricity in the network. Thus, the integration of sustainable forms of energy is transforming the electricity grid in a number of ways. Most importantly, the realization of the *smart grid* vision requires a great inter-disciplinary effort that spans many areas. In this context, Computer Science and predominantly Artificial Intelligence is expected to provide an algorithmic layer that will bring about a fundamental reorganization of the grid. We will return to the convergence of these two sectors after a brief review of what the smart grid vision entails.

According to ISO, the definition of *smart grid* is:

The application of technologies to all aspects of the energy transmission and delivery system that provide better monitoring, control and efficient use of the system. The ISOs goal is to enable and integrate all applicable smart technologies while operating the grid reliably, securely and efficiently, and facilitate effective, open markets that engage and empower consumers while meeting environmental and energy policies.

Although electrification was the most significant achievement of the 20th Century, the grid still operates the way it did almost 100 years ago. That is, having central power plants feed the energy flow over the grid to consumers. Given the expected growth in demand for electricity around the world (see Figure 2.2) averaged at a global 115%, the ageing infrastructure could affect system reliability, stability and

security. Fortunately, the same study, BLUE Map Scenario [2], estimates that during this time horizon (2007-2050) the penetration rates of renewable, variable generation will also increase between 15% and 20% (see Figure 2.3).

In Figure 2.1 we bring together the challenges, that we believe, the transition to a smarter grid must address and differentiate them into separate levels as follows.

1. The problem of accommodating all generation and storage capacities.

Aside from conventional generation methods, there is a growing array of micro-generation devices that are becoming affordable to residential, commercial and industrial customers. In order to boost efficiency, the widespread adoption of these technologies will require that owners will be able to sell their excess energy to utilities, thus moving from a passive to an active role in the system.

Enhancements:

- *Local energy management.*
- *Reduction in transmission losses.*
- *Reduced Greenhouse Gas Emissions.*
- *Increased interest in electricity market opportunities for customers with in-house generation capabilities.*
- *Capability to meet increasing consumer demand without adding infrastructure that impacts CO₂ emissions (new power plants).*

2. The problem with peaks. During periods of peak energy usage grid operators are required to bring peaker plants online so that they can ensure to meet peak demand. Moreover, they do so without a precise information about when demand will peak and how high it will go. These peaks, some of which are registered for only a limited number of hours per year are yet responsible for setting the capacity requirements of the whole distribution grid. Also the usage of peaker plants imply higher operational costs as well as higher greenhouse gas emissions.

Enhancements:




Figure 2.2: Electricity consumption growth 2007-2050 (Blue Map Scenario [2])



Figure 2.3: Portion of variable generation of electricity by region 2007-2050 (Blue Map Scenario [2])

- *The total grid capacity could be significantly lowered, thus great investments in extending the network infrastructure to accommodate the increasing demand could be avoided.*
- *Encouraging customers to smooth or reduce overall energy consumption leads to lower emissions and costs. For instance, in the US, a 5% increase in energy efficiency would equvalate eliminating the greenhouse gas emissions and need for fuel from 53 million cars. Again, in the US, McKinsey estimated that demand managements could provide \$59 billion in annual benefits by 2019 [73].*
- *Increasing customer awareness.*
- *Reduce costs due to the need for building new power plants.*

3. **The problem of introducing liberalised markets.** It has to do with the state withdrawing its involvement in the energy industry. Liberalization however requires putting into place a market structure within which effective competition could be achieved, where for instance, any restrictions on customers from changing their supplier are removed. This is set to affect the business models of companies in a major way.

Enhancements:

- *Free choice of power supplier for electricity customers. Customers will be able to switch between suppliers or to reenter contract negotiations.*
- *In a liberalised market economy the pursuit of profit by private owners will lead to efficiency improvement and cost saving. Benefits are to be passed to end customers in the form of lower prices.*
- *Enabling new products and new services (e.g. Platforms to solve the electricity tariff selection problem, DSM, VPP).*
- *Liberalization separates the responsibility of security and maintenance of transmission and distribution networks from the generation business.*

4. **The problem of grid surveillance and advanced metering infrastructure.** To this point, monitoring has been mainly addressed at the level of high-voltage transmission grids. Increasing the degree of automation for better quality of service is an important step, through installation of sensors and controls all the way down to distribution. Using near-real time solution could be crucial in isolating outages, while the rest of the system is restored to normal operation. More hardware solution require automated re-closers, switches and capacitors.

Enhancements:

- *Grid resiliency & interoperability of electricity networks. As monitoring ensures localization of unexpected contingencies, the following phase is in automating the ability to react and produce self-healing actions in order to enable prompt solutions to events that impact the security of supply and power quality. A classic example of a massive blackout is the one that occurred in 2003, the Northeast blackout, which affected more than 50 million people and caused a complete standstill from communications to traffic to banking, etc. In the US only, power outages and interruptions cost at least \$150 billions each year [73].*
- *Remote metering eliminates the costs associated with manual meter readings. More than 50 million smart meters are planned to be installed by 2015 in the US [73].*

The model envisioned for the deployment of the smart grid is in many ways analogous to the internet, in terms of distributed decision making throughout the system and bidirectional flows of electricity and information. This places important emphasis on the role of Computer Science to provide for internet-like protocols in operating a network of heterogeneous components without the traditional monopoly-based regulation of supply. Consumers would be no longer bind to a certain supplier of power, while the network use could also vary, as the grid could self-determine its configuration.

Smart Grids Functionalities	Functional Project	YEAR										Costs (M€)	
		2010	2011	2012	2013	2014	2015	2016	2017	2018	2019		2020
Active Demand Response and integration with Smart Homes	D1	ADDRESS			Active Demand Response								190
	D2	BEWARE Smart Homes/Smart Grids			Integration with Smart Homes								120
Smart Metering Infrastructure & Data Processing	D3	OPEN METER	Smart Metering Infrastructure										150
	D4	Smart Metering Data Processing											20
Integration of RES, storage and EV	D5	Active Distribution Network	Integration of small DER										90
	D6	Active Distribution Network	Integration of medium DER										150
	D7	STORAGE TECHNOLOGY		Integration of storage technologies									60
	D8	ELECTRIC VEHICLES		Integration of Electric Vehicles									100
Planning, monitoring and control	D9	Active Distribution Network	Monitoring and control of LV networks										100
	D10	Active Distribution Network	Automation and Control of MV networks										90
	D11		New methods and systems support										80
Integrated communication Infrastructure	D12	Active Distribution Network	Integrated Communications Solution										50
Total											1.200		

Figure 2.4: European Roadmap for Smart Grid deployment with associated costs [72]

After defining the main smart grid functionalities, we illustrate the European targets to meet these goals in Figure 2.4, which is the detailed roadmap with budget allocations. As the table shows, most of these technologies are supposed to become implemented ahead of 2018. The R&D activities have been organized into 12 functional projects, which leverage on relevant running projects and investments already done at national and European level. The vision includes a broad spectrum of stakeholders which range from consumers, utility companies and network operators to new businesses for advanced electricity services and various solution providers (e.g. microgrids, on-site generation, virtual utilities, etc).

Bringing about the smart grid vision is highly dependent on the mass adoption of new technological advances. We classify these driving factors into three categories detailed below. The integration of this type of technologies down to the individual household level has been referred to as the *smart home* - an interconnection an automated control over electronic devices, renewable generation and sensing equipment - depicted in Figure 2.5.




Figure 2.5: Integrating devices into the Smart Home model

- a) *Smart Appliances.* According to the Association of Home Appliance Manufacturers (AHAM), smart appliance refers to a modernization of the electricity usage system of a home appliance so that it monitors, protects and automatically adjusts its operation to the needs of its owner [124]. The product segmentation of these devices generally includes clothes dryers, clothes washers, refrigerators, freezers, dishwashers, ovens, microwaves, coffee makers as well as other devices. The goal is to have each home ultimately equipped with a home energy management system that connects all these appliances and automatically operates an energy optimization for consumers at the level of the total home energy usage. In this same category we also include electric/hybrid vehicles or battery storages capabilities.
- b) *Smart Metering.* Importantly, in order to quantify the home usage profile of users, digital electricity meters with bi-directional communication capabilities, that link end-users to utilities need to be installed. Additionally other devices that act as sensors, such as smart plugs, enables consumers to track electricity usage by the outlet.
- c) *Distributed Generation Technologies.* Here we include small power generators typically located at consumers' sites. The portfolio of such devices includes: micro-turbines, fuel cells, photovoltaic cells, solar thermal arrays, wind turbines, recip-

rocating engines, combined heat and power units (CHP).

2.2 Open Multi-Agent Systems and Agentification of Smart Grids

The field of multi-agent systems has emerged as an overarching software paradigm for developing architectures that contain a large number of dynamically interacting components, each with its separate control and which are engaged in complex interaction protocols. Applications and possible uses of multi-agent systems cover a wide spectrum. For instance, applications of multi-agent systems include e-commerce trading, network and resource management, production control, supply chain management, complex system modelling and simulations, telecommunications, decision support systems, manufacturing systems, traffic monitoring, aircraft maintenance, military logistics planning, simulation of real world, video games, power systems etc.

In very simple terms, an agent is a computer system that performs independent actions on behalf of its user. Then, a multi-agent environment is a system where a number of agents representing users, with different goals and motivations, perform actions to fulfil their tasks through various ways of interaction such as cooperation, coordination and negotiation. Hence, the multi-agent paradigm dwells upon two key questions. The first is referred to as the *micro* design and has to do with modelling and building autonomous agents that can be delegated user goals and that are capable to achieve them. The second aspect regards the *macro* design, that is engineering the interaction rules by which system-level goals can be achieved, while agents with different goals can strive to fulfil theirs through various forms of coordination, cooperation and negotiation. In other words, given a society of self-interested agents, can cooperation yield as an emergent phenomena? Or, how could agents resolve conflicts by reaching agreements and coordinate to resolve group goals? As multi-agent systems are getting more complex, with many dynamically interacting components, the best way to control agent activity becomes an important question. Moreover these systems are *large-scale* and *open* in the sense that the components themselves nor

their behaviour cannot be predetermined ahead of execution.

There is not a universally accepted consensus as to what an intelligent autonomous agent can be defined as, however there is a general framework that encapsulates the list of capabilities that such an agent is expected to have [191]. Agents are characterised by *proactiveness*. They are to be designed so that they reflect the user's goals and are able to exhibit a goal-oriented behavior, by producing a plan, whose execution ensure satisfying its objectives. Agents are *reactive*. They are able to perceive the environment and as there are changes occurring, agents are expected to respond in order to satisfy their user goals. Agents are endowed with a *social ability*. This capability allows agents to perform complex decision-making procedures where they reason about the goals of other agents and engage in interactions, much like in real-world settings, where in order to achieve their goals they need to cooperate, coordinate and negotiate.

The benefits of applying multi-agent system to the smart grid domain are essentially twofold. Firstly, MAS provides a *modeling* approach to represent real-world situations. Secondly, MAS provides a way for *building* such systems.

In terms of modelling, according to [191], MAS builds upon the object-oriented approach, where entities in a system are represented as objects. This means that it maintains the benefits of data encapsulation. The data structure allows for associating object attributes that are only accessible to the extent that the object methods allow it. The agents paradigm augments this approach by adding another layer of abstraction, where direct access to these methods is not allowed. Instead, the only available action is communicating between agents via message passing. The autonomy of an agent can thus be guaranteed. While in object-oriented modelling an external object can call methods used by other objects, which have no option but to execute them, an external agent can at most request another agent to execute a certain function. Agents only receive requests, but retain autonomy to decide and schedule their own actions. Hence deploying the multi-agent system as a modelling approach is beneficial for simulating complex systems such as the smart grid, where actors have certain attributes and a certain degree of autonomy given by a set of possible, executable actions. Especially they can be related to economic encounters.

Consider for instance modelling market-like mechanisms where different parties compete for the allocation of scarce resources or cooperate to achieve certain tasks based on an economic rationale. Agents cannot manipulate directly the actions of other agents, but have however an indirect access through negotiation.

In terms of building systems, MAS can be thought of as a platform for distributed environments. Based on the same notion of autonomy and strategic behavior, deploying MAS provides a distributed control in the sense of enabling prompt, adaptable local control. Moreover MAS enables building robust, flexible and extensible systems. Robustness is an important property of distributed systems, which need to ensure a certain level of fault-tolerance, so that the system as a whole can still perform its functionality even in the event of partial failure. Flexibility can be understood both at the agents level, as the choice of executing different actions, as well as at the system's level, in providing agents with flexible coordination and communication methods. Due to its modularity, in the sense of object-oriented programming MAS are easily extensible, allowing to add new functionality without the need to reimplement the existing functionality. Because the focus is placed on agent interaction via communication, MAS are open architectures. Thus, as long as agents adhere to messaging standards, there is a clear decoupling between the agents and the environment. This allows for a heterogeneous population of agents to interact regardless of the programming languages used in their implementation. Even more importantly, openness, in a multiagent context, can be understood as providing a framework which can accommodate agents designed with different interests, preferences and behaviours in mind.

Exploring the application of open MAS to modern power system is justifiable due to:

- the existence of a large number of interacting actors with private goal-driven behavior, whether we are talking about consumers of energy, producers of energy or network operators; explicitly modelling the overall system behavior is impractical;
- the heterogeneous nature of actors that are required to interact, coordinate,

cooperate and negotiate;

- the requirement of an open system that places no restriction on actors entering or exiting the system;
- the requirement of appending system functionality on an ongoing basis;
- the requirement of managing and controlling the system via local decision-making, according to the organization of the grid in different subsystem;

2.3 Coordination methods in MAS

Arguably, the defining problem in multiagent systems is that of *coordination*. Similar to rules and regulations in human societies, multiagent systems require coordination mechanisms in order to work efficiently.

In [191], Wooldridge defines MAS coordination as the process of managing interdependencies between agents' activities. In contrast to traditional computer systems, in MAS, coordination is achieved at run time rather than design time. The decentralized nature of MAS allows for adaptability, so that each agent can decide its appropriate course of actions at run time in order to produce the best possible result, given their available processing, communication and information resources. Coordination is required for guiding the agents' behavior towards an efficient usage of these resources. A comprehensive overview of coordination models is given in [125, 159], where the authors categorize the many diverse approaches to agent coordination in two main classes, the *subjective* and *objective* approaches, depending on whether they adopt the agent's or the engineer's viewpoint, respectively.

A taxonomy for coordination relationships is provided by von Martial in [103]. Essentially, it distinguishes between positive and negative relationships between agents' activities. Positive relationships denote situations where by combining the actions of two or more agents some benefit can be derived for at least one of the agents. Negative relationships appear when a certain resource is involved, be it either consumable or non-consumable. When such interdependencies occur, the choice of methods used

to resolve a coordination problem will significantly affect the overall system performance. An extensive outlook on coordination relationships in the context of artificial agent societies is introduced in [127]. The authors propose the notion of social structure and formalize it as an external factor in order to derive an operational model of social coordination. Its applicability is demonstrated in a road traffic management scenario [62].

Given the complex domain of smart grids, which we are interested in, in this thesis we consider that a combination of coordination techniques needs to be used, with different methods appropriate at different levels of abstraction. We go on to detail several approaches relevant for our work, that can be found in the literature and tackle the problem of coordination.

2.3.1 Emergent Coordination and Agents

On one hand MAS can be engineered as a *closed* computer system that solves predefined problems in a distributed way, with an explicit coordination procedure. On the other hand, at the opposite end of the spectrum are *open* MAS that assume a high level of interaction between a large number of highly stochastic, heterogeneous agents operating within a dynamic environment. In [128] the authors introduce the notion of *emergent coordination* and discuss its potential for efficiently handling coordination in open environments. An efficient way to coordinate the more complex case of open systems is to design a limited set of rules, a common communication protocol and a specific role distribution. A careful intentional design of the above should produce the desired emergent property, which is a higher level property, caused by the interaction of its lower level components. This type of coordination is also known as *self-organization*. A general characterization of self-organizing behaviour is given in [34] for the following properties:

- *absence of explicit external control* - This is a mandatory property that states that the system is autonomous; it imposes and changes its organization based solely on internal decisions and without following any explicit external reorganization command. This property refers to the self-part of the self-organization

definition.

- *decentralized control* - A self-organizing system can work under decentralized control. In this case, there is no internal central authority or centralized information flow. As a result, access to global information is limited by the locality of interactions, which is governed by simple rules. This property is generally not mandatory, since we can observe internal central control in many natural self-organizing systems. However, in the context of MAS, the existence of decentralized control is also considered a mandatory property of self-organizing MAS.
- *dynamic operation* - This mandatory property is linked to the system evolution in time. Since the organization evolves independently of any external control, this property implies continuity in the self-organization process.

Predominantly, the approaches proposed in this space introduce nature-inspired techniques. To exemplify several representative cases, the notion of stigmergy describes the use of pheromones as environmental markers to drive individual and social behaviors in ants. The authors in [39, 38] adapted this concept to a coordination mechanism for agent population. Similarly, a field-based coordination solution is given in [101], inspired by how mass particles move and self-organize according to electromagnetic fields. The chemical metaphor is exploited by [182, 183] in so called chemical tuple spaces. A more general approach is proposed in [152], based on tuples which represent ordered collections of information items, accessible through a set of operation primitives such as put, browse and retrieve. Overall, this work considers an information-driven coordination medium that regulates the agent interaction. The application of such approaches range from the control of unmanned vehicles, to news management, to the coordination of e-health systems [39, 102, 116].

Self-organization implies that agents follow a predefined set of rules that govern the space of interaction forming particular organizational structures, which facilitate achieving individual and system goals. Organizations allow simple agents to exhibit complex group behavior, as well as for more sophisticated agents to reduce

the complexity of their reasoning [66]. For instance, hierarchies are simple forms of organization, which represent agents that are conceptually arranged in a tree-like structure, limiting agent interaction based on their connectivity. Here, control flows downward from the upper levels to the ones below. Holarchies are specific cases of hierarchies denoting nested, self-similar grouped hierarchies. The basic unit of organization is called a holon. Another case of agent structures are teams. Commonly the type and pattern of interactions can be quite arbitrary here, however, generally, each agent will take on one or more roles required by the team goal. Federations are characterised by having a group of agents delegate certain tasks to a group representative. Congregations represent a more flat, long-lived type of organization comprised by heterogeneous agents.

2.3.2 Distributed Constraint Optimization Problem and Agents

Consider for instance the problem of traffic light control. Given a road network, our goal is to coordinate a set of traffic lights so that the vehicles' average travel times is minimized, by ensuring that as many vehicles as possible can keep a specific speed, without stopping. If each agent actuates a traffic light, it means that the agent's decision about the state of its assigned traffic light depends on the state of nearby ones. In other words, the agent is faced with a decision problem, where its optimal choice depends simultaneously on the decisions made by others. This situation occurs in many real-world decision problems. Supposing that the common goal of the agents is to find a joint choice of decisions that maximizes a certain global function, in this case the average travel time, they need to exchange information about their respective constraints in order to coordinate their decisions.

The framework of *distributed constraint satisfaction* has been developed in order to mathematically formalize such distributed decision problems. *Decision variables* are used to model each decision that a given agent must take. Thus, the agent assigns a value to a decision variable from the set of possible values corresponding to the possible choices for that decision. Additionally, *constraints* on subsets of decision variables are used to express which combinations of value assignments to

these decision variables are allowed or disallowed. By generalizing constraints to represent costs or utilities for particular variable-value assignments for the agents, we fall under the category of *distributed constraint optimization*. Notably, the distributed nature of this formalization is evident firstly, *i*) in terms of the agents' decision power in selecting the values for the decision variables under their control and secondly, *ii*) in the sense that knowledge is distributed, having each agent knowing only the constraints that involve the decision variables he is controlling.

Formally, a *discrete distributed constraint satisfaction problem (DCSP)* [194] is a tuple $\langle A, X, D, M, C \rangle$ such that:

- $A = \{a_1, \dots, a_k\}$ is a set of agents;
- $X = \{x_1, \dots, x_n\}$ is a set of decision variables
- $D = \{D_1, \dots, D_n\}$ is a set of finite variable domains, such that each variable x_i takes values in D_i ;
- $M : X \rightarrow A$ an ownership mapping that assigns each variables to the agent that owns it;
- $C = \{c_1, \dots, c_m\}$ is a set of hard constraints, where each c_i is $s(c_i)$ -ary function of scope $(x_{i_1}, \dots, x_{i_{s(c_i)}})$ and $c_i : D_{i_1} \times \dots \times D_{i_{s(c_i)}} \rightarrow \{false, true\}$ gives the allowed value combinations for the corresponding constraint c_i , assigning false to infeasible tuples, and true to feasible ones.

A solution is a complete assignment consistent with all constraints: $\bigwedge_{c_i \in C} c_i = true$.

The *distributed constraint optimization problem (DCOP)* is a generalization of the above, defined as a tuple $\langle A, X, D, M, C \rangle$ such that A, X, D and M are the same and:

- $C = \{c_1, \dots, c_m\}$ is a set of soft constraints, where each constraint is a function $c_i : D_{i_1} \times \dots \times D_{i_{|c_i|}} \rightarrow \mathbb{R} \cup \{+\infty\}$ that assigns a cost $c_i(x_{i_1}, \dots, x_{i_{|c_i|}})$ to combinations of assignments to a subset of decision variables, infinite costs corresponding to infeasible assignments.

A solution is an assignment to all decision variables that minimizes the sum of all costs:

$$(x_1^*, x_2^*, \dots, x_n^*) = \arg \min_{x_1, \dots, x_n} \sum_i c_i$$

The above formulation denotes constraints through costs, thus it is a cost minimization problem. Similarly, in situation where constraints define utilities, DCOPs can be expressed as maximization problems.

The algorithms proposed for solving DCOPs can be classified into:

- *complete algorithms*, which guaranteed to find a feasible solution when there exists one
- *incomplete algorithms*, which sacrifice optimality to obtain fast (any-time) solutions.

Complete algorithms adopt two main techniques: search and dynamic programming. The optimality guarantees requires an exponentially increasing coordination overhead for both approaches. Although search-based algorithms only require linear-size messages, they need an exponential number of messages. Alternatively, dynamic programming algorithms only require a linear number of messages, but their complexity lies on the message size, which may be very large.

The simplest search-based complete DCOP algorithm introduced was *synchronous branch and bound* (SynchBB) [64]. It represents a straightforward distributed adaptation of the well-known centralized *branch and bound* mechanism. The first algorithm proposed which performed a decentralized search, allowing asynchronous operations was *asynchronous distributed optimization* (ADOPT) [112]. The novelty in ADOPT was in compiling a DCOP into a pseudo-tree structure, which was then used as a hierarchy to communicate among agents. A recent extension of ADOPT is BnB-ADOPT [193], the difference consisting in replacing the best-first search strategy for a depth-first branch-and-bound search. Importantly, BnB-ADOPT has been shown to provide optimal solutions up to one order of magnitude faster than ADOPT.

The class of dynamic programming approaches is marked by the *dynamic programming optimization protocol* (DPOP) algorithm [94], which also operates along a pseudo-tree variable ordering. DPOP is an instance of the general bucket elimination scheme [32], which has been adapted for the distributed case. Given a pseudo-tree arrangement of the variables, the algorithm requires two stages. First, each variable, starting from the leaves of the tree, *joins* the constraints involving itself as well as its descendants, parent and pseudo-parents and then projects itself out of the join, sending the result to its parent. During the second phase the root receives the aggregated result and can thus select the optimal value assignment for its variable. Next, optimal decisions are propagated down the pseudo-tree, until all variables have been assigned optimal values. Most of the algorithms subsequently proposed, based on dynamic programming, have been defined as extensions to DPOP with the aim of providing different trade-offs: M-DPOP, MB-DPOP, A-DPOP [136, 135, 134].

Incomplete DCOP algorithms are characterized by some form of local search, where agents begin from an initial (arbitrary) solution and attempt to improve it based on local information available from neighbouring agents. They are commonly differentiated based on the strategy they apply for escaping local minima. For example the *distributed stochastic algorithm* (DSA) [45] prescribes that agents will perform, with some probability, non-strictly improving changes in order to avoid such situations. Another example, *distributed asynchronous local optimization* (DALO) [81] is based on the idea of optimising for groups of k agents. The *max-sum* algorithm [43] introduces a modified representation of the DCOP by reasoning on a factor graph, which is a bipartite graph where function nodes stand for constraints and variable nodes stand for the decision variables in the DCOP.

2.3.3 Game theory and Agents

Throughout the multiagent literature, *game theory* proves to be predominantly the theoretical tool in use for the analysis of such systems, providing the mathematical foundation for studying interactions among self-interested agents [126, 114, 13]. As the name suggests, the agents' interaction is structured in the form of so-called *games*,

which generally speaking, represent protocols that regulate the space of interaction. A key distinction in game theory is made between games that are cooperative and games that are non-cooperative.

Cooperative Games

Of particular interest for the multiagent domain are *coalitional* or *cooperative games*, studying how groups of self interested players interact to accomplish more together than they could individually achieve. A coalitional game is modelled based on a population of agents (the set of players in the game), a set of actions available to each agent and a preference profile over the joint outcomes, particular again to every agent. The preference of an agent over the set of possible outcomes is formally captured by means of a *utility function*, which assigns to every outcome a real number. The game is abstract in that no indication of how the utility function should be derived is given, so that details can be omitted and the framework remains general. Instances when such a modelling approach is applicable are situations where utility can be gained or costs can be reduced by agents working together, forming a *coalition*. Here, each subset of agents represents a potential coalition. It is important to highlight that although cooperation is required in order to attain some desirable outcome, the agents always seek the actions most likely to bring the highest utility for themselves. Clear examples for such scenarios are task allocation problems, where groups of agents, potentially heterogeneous in their capabilities, are required for performing the tasks, given that the ability to solve a task is greater than any single agent can offer [169, 170, 162].

Under the assumption of a particular type of reasoning employed by the players (e.g. rationality), there is a need to predict the outcomes of these games. *Solution concepts* in coalitional games describe sets of outcomes that explicitly take cooperation into account and attempt to ensure some form of stability. It is often the case that desirable outcomes, from the point of view of the system's social welfare cannot be enforced in a game because several agents may deviate in order to maximize their utility. This tension between social welfare and stability is often evident. Thus, game

theory proposes solution concepts, defining rational outcomes. Also, in the game theoretic literature, solution concepts are usually defined under the assumption of superadditive games where the *grand coalition* (the set of all players) forms. The problem is then restricted to solving the division of payoffs among the players. Many real-world problems are however non-superadditive such that negative interactions prevent the players from forming the grand coalition. This leads to the more interesting setting of games with *coalition structure* (partitions of the set of players). The common problem with solution concepts is that they may not offer guarantees nor on the existence of a rational outcome, nor on the uniqueness of a rational outcome. Another implicit assumption often used is to consider that the coalitional utility can be divided amongst the members of the coalition in any way that they choose. Formally, games with this property are said to be *transferable utility games* (TU games).

The possible outcomes of a game can be evaluated broadly according to two criteria: *fairness* and *stability*, which lead to two classes of solution concepts.

i) Representatively, *the core* is one of the most well-known *stability* solution concepts in coalitional games and can be understood in the sense that an outcome is stable if no deviations are profitable [50]. Formally, for a game $G = (N, v)$ *the core* is the set of all outcomes (CS, \mathbf{x}) such that $x(C) \geq v(C)$ for every $C \subseteq N$, where

- $N = \{1, \dots, n\}$ is a finite, non-empty set of agents
- $v : 2^N \rightarrow \mathbb{R}$ is a characteristic function, which maps each coalition $C \subseteq N$ to a real number $v(C)$
- $CS = \{C^1, \dots, C^k\}$ is a coalition structure; $\bigcup_{j=1}^k C^j = N$ and $C^i \cap C^j = \emptyset$
- $\mathbf{x} = (x_1, \dots, x_n)$ is a payoff vector for a coalition structure $CS = \{C^1, \dots, C^k\}$ over N where $x_i \geq 0$ for all $i \in N$ and $\sum_{i \in C^j} x_i \leq v(C^j)$ for any $j \in \{1, \dots, k\}$.

ii) Another axiomatic solution concept, representative for solutions that aim to capture the notion of *fairness* is the *Shapley value* [161]. It represents the expected

marginal contribution that a player brings to the set of players preceding him in a coalition, while considering each coalition equally likely to form, as well as the size of the coalitions. The intuition behind the *Shapley value* is that the payment that each agent receives should be proportional to his contribution averaged over all possible orderings, or permutations, of the players. Formally, for a game $G = (N, v)$, with $|N| = n$, the *Shapley value* of a player $i \in N$ is denoted by:

$$\sigma_i(G) = \frac{1}{n!} \sum_{\pi \in \Pi_N} v(S_\pi(i) \cup \{i\}) - v(S_\pi(i))$$

where $S_\pi(i)$ is the set of all predecessors of i in a given a permutation π from the set of all possible permutations Π_N of N .

We have introduced two solution concepts that are perhaps the most convincing and representative equilibrium concepts in coalitional games for the two classes previously mentioned. In addition, numerous other proposals, that exhibit similar properties and drawbacks, have been made in the literature, such as *the bargaining set*, *stable set*, *nucleolus*, and *kernel* [126].

Non-Cooperative Games

Put simply, *non-cooperative games* are distinct in the sense that *binding agreements* are not possible here. As we have seen, cooperative games unfold under the assumption that players can agree about the distribution of payoffs, although agreements are not explicitly specified by the rules of the game [133]. Moreover, while in coalitional games, the basic modelling unit is the coalition, in non-cooperative games the basic modelling unit is the individual player. In the non-cooperative setting the background assumption is that binding agreements are not possible and players cannot trust one another, but all are trying to maximize their own utility taking decisions based solely on the information they have about the possible choices and corresponding utilities.

Modelling this type of games requires the following specifications: *i*) the number of players, *ii*) the set of possible actions available to each player, *iii*) the utility function of each player which he attempts to optimize. It is further important to mention that players will simultaneously choose an action to perform, resulting in a particular outcome. Thus, it is clear that the actual outcome that will result will depend on the particular combination of the actions performed, such that every player can influence the outcome. In line with the game-theory literature, actions are commonly referred to as *strategies*. The question that arises in a non-cooperative game, from the player's perspective is: what to do in any given scenario?

Now, suppose that for some strategy s_1 and s_2 of player i there will be a set of possible outcomes ω_1 and ω_2 respectively. It is said that strategy s_1 *dominates* s_2 if *every* outcome in ω_1 is preferred by i over *every* outcome in ω_2 . Therefore, in the presence of a dominant strategy among all possible strategies, the player's decision in a game is straightforward. The dominant strategy will guarantee the best outcome. Of course, computing the dominant strategy is a considerably complicated task, especially when there are many utility-maximizing players, whose actions can affect each other's utilities.

Arguably, the most influential solution concept in game theory is the *Nash equilibrium* [126], which determines the outcome of a game provided that every player is rational in adopting its best strategy. Intuitively, the *Nash equilibrium* is a stable strategy profile, such that no player would have any incentive to change his strategy. Nevertheless, there are several associated negative results regarding this solution concept, namely: *i*) not every game has a Nash equilibrium and *ii*) some games have more than one Nash equilibrium. For an interesting application of non-cooperative games, and particularly the Nash solution we refer to [53], where the authors propose a novel game approach to the problem of decision-making with respect to information security investments.

Overall, game theory brings about extremely important concepts to the multi-agent domain, however its applicability needs to be addressed in a careful manner and several aspects need to be acknowledged [56, 150]. Importantly, computational complexity is largely overlooked by game theory, limiting itself to offer descriptive

concepts, emphasizing properties of optimal solutions, though disregarding the fact that usually computing them turns out to be computationally hard and that we are normally dealing with resource-bounded players. Moreover, if in a game there are multiple equilibrium states, players have no way of knowing which one will be played. Furthermore, game theory puts forward an assumption that is unreasonable for most multi-agent systems. Namely, it models games with perfect information, where all aspects of the game are considered common knowledge for all players. Quite the opposite, most real-world scenarios represent games with minimal information, where not all actions and payoffs of other agents are directly observable. Sometimes, it may well be the case that not even all the possible moves are known for a player in every situation of the game. To conclude, there are multiple ways to go beyond standard game-theoretic concepts in order to ensure their applicability to software settings, which requires considerable attention.

2.4 Coordination and control in the *smart grid*

The literature on smart grid solution revolves around several key topics previously identified. This part of the chapter is reserved for an overview of existing work both from an academic and a commercial perspective. We provide hereafter a review on smart grid approaches structured according to Figure 2.6. Noticeably, the classification intertwines synergically the challenges and the driving factor of the smart grid, highlighted in section 2.1.

2.4.1 Enhancing the control of the electricity grid

Delivering the advanced metering infrastructure

One thread of work closely related to the roll-out of smart meters focuses on the usage of autonomous agents to automatically model and predict the use of energy at the household level. An accurate understanding of this information could help householders to better manage their energy consumption. Given that smart meters can only provide an indication of the total consumption of a household, extracting

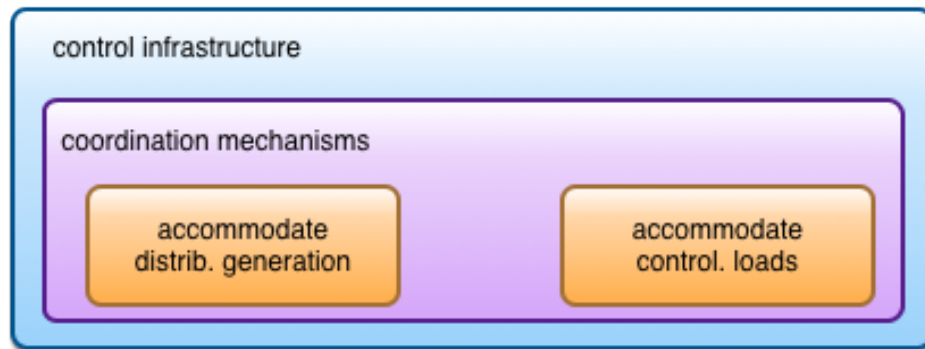


Figure 2.6: Approaches to smart grids

meaningful information implies undergoing an energy disaggregation procedure. Also known as non-intrusive appliance load monitoring (NIALM), energy disaggregation aims to break down a household's total energy usage into individual appliances. The initial solution, based on finite state machines, was proposed by Hart [60], who showed that different classes of appliances produce distinct power consumption signatures. More recent approaches apply various machine learning techniques to perform energy disaggregation. In general, these approaches difference themselves by making different assumptions about the prior information available for their considered setting.

The class of algorithms that base themselves on supervised learning methods make the assumption of having an available set of data for training. This is a limiting factor since it requires installing sub-metering devices for all appliances instead of the usual whole-home power monitor. Hence, it is often impractical to assume consumers to perform time-consuming, inconvenient and expensive procedures. Nevertheless several publicly available data for such domain start to emerge. Such is the case of the reference energy disaggregation data set (REDD) [89]. It comprises data sampled every minute for a period of several months for six households, with a total of 268 unique monitors that have recorded more than 1 terabyte of raw data. In [89] Kolter et al. also provide a result benchmarked on the REDD data, based on their approach using the factorial hidden Markov model (FHMM) technique. The idea behind this work is modelling each device as a hidden Markov model. The operation of each device is a sequence of one or more (possibly different) *on* states and *off* states. A




Figure 2.7: Factorial hidden Markov model representation for 2 devices; $x_t^{(i)}$ is the hidden state i of a device at time t ; \bar{y}_t is the aggregated power at time t ; $y_t^{(i)}$ the individual power of a device at time t ; shaded nodes denote observed variables and unshaded nodes represent hidden variables.

FHMM [49] can model multiple hidden state sequences. Factorial learning algorithms infer the posterior probability of individual device consumption, given an observed sequence of aggregated consumption. A general graphical representation of FHMM is showed in Figure 2.7.

At the other end of the spectrum of machine learning methods are the unsupervised techniques. No prior datasets of recorded appliance usage are required here, however there is often the assumption of knowing the number of appliances or that the appliances are labelled manually after the disaggregation procedure is performed. In [82] Kim et al. apply a variant of a conditional factorial hidden semi-Markov model. Their approach is intended for low sampling rates, which makes it more widely applicable. Under the same category of unsupervised learning and hidden Markov models variants, Parson et al. [131] introduces an iterative method for load disaggregation benchmarked against the abovementioned REDD dataset. They frame the problem under a different assumption, that of having prior knowledge of the generic appliance types under investigation. From here, without data on individual appliance consumption patterns, they tune to specific instances of appliances using aggregated data only. Basically, they reduce the necessary information used in the training process by tak-

ing generic models of appliance types and training them to specific instances. After having learned the parameters for each appliance, the disaggregation procedure uses an extension of the Viterbi algorithm. This further means that there is also prior knowledge about the existing type and number of appliances under investigation, which can be considered as a hard limitation. Another unsupervised approach benchmarked against REDD, which yields the highest performance, is proposed by Kolter and Jaakkola [88]. Again, they employ a variant of HMM, namely additive FHMMs and develop an approximate inference procedure, as the accurate inference is not computationally tractable. The drawback here is that the algorithm proposed requires sampling of data that is several orders of magnitude more frequent than the solution introduced in [131].

In [175], the authors propose a compromise between the two classes of algorithms presented thus far. They introduce a semi-supervised technique for identifying the minimal subset of appliances that need to be individually monitored to maximize accuracy. The idea behind their approach is to provide a trade-off between cost of monitoring and accuracy. The modelling technique used in this work is closely akin to the ones already presented, being based on FHMMs. Their methodology is structured in three stages. First, they train a FHMM model relying only on aggregated household-level information. The result is a possible number of Markov chains, each corresponding to a different device in the monitored space. Second, they select a subset of appliances to be individually monitored using HMMs. Third, using the HMM chains from step two, they retrain the FHMM model. Then, steps two and three are reiterated until an acceptable accuracy is attained. From here the authors define a heuristic approach based on which the subset of appliances is selected, based on practical considerations (e.g. the appliance with the highest consumption is expected to improve the disaggregation accuracy the most). Again they run the experiments on the REDD dataset and conclude that determining the proper subset of devices to be monitored can increase significantly the accuracy of their algorithm. Specifically, they show that the percentage of the total energy correctly classified does not have a high improvement after increasing the number of monitored appliances above a cer-

tain value¹. Of note is recognizing some of the harder assumptions present in their approach, which include knowing the number of appliances at each household, as well as their states and their average and peak power consumption.

Although employing HMMs to model the time-series of aggregated power readings has been the most prevalent approach used in NIALM and delivered the most interesting result, there are several different techniques that deserve to be mentioned here. While HMMs allows for a convenient abstraction from appliance-specific characteristics, in doing so, certain valuable information may be overlooked. Thus, if we can assume a higher frequency of sampling, new feature such as harmonics and signal waveforms could be also considered. Except for the high cost of deploying such sensors, there arise other issues such as the transmission and storage of data.

An extension of the initial NIALM algorithm based on edge detection to incorporate harmonics is presented in [93] and [99, 97, 190]. Performing continuous calculation of signal harmonics, the authors in [166] propose an appliance detection techniques that uses neural networks. Although they report a good detection accuracy, the training of the neural network requires datasets for all possible combinations of the appliances. Under the assumption that the average household includes up to 50 appliances, the combinatorial explosion makes this approach impractical.

Another category of algorithms introduce a noise-based technique for NIALM. In [132] the authors assume the electric noise monitoring of a single socket of the household. The appliance signatures though, have been shown to depend on the household wiring [54]. In this latter work, high frequency electromagnetic interference (EMI) is used to detect devices with a higher accuracy. Other features used to identify different appliances include, wavelet transform features [25] and I-V curve properties [98, 91]. Also, several techniques from the pattern recognition domain have been applied to this problem that consider standard classifiers such as the nearest-neighbour [54, 11, 12] or the Bayes classifier [37]. As can be observed, the majority of these works are concerned with residential settings. The investigation of the industrial

¹This value is shown to correspond on average to a number of three appliances in the context of a usual residential setting of seven appliances per household.

sector was considered in [119], where the focus is on large loads. It is usual in these type of settings to design appliance-specific rules for detecting devices [90]. Following this approach though is clearly impractical for consumer devices, which must include a vast number of models.

Conclusively, there is no one solution developed to this point that can perform NIALM for all type of devices and settings, but a set of different techniques that base their approach on different assumptions. Unfortunately, most approaches require expensive sensors and large datasets for apriori training of the models. Even so, there remains significant room for accuracy improvement. One obvious way to enhance performance is the concomitant usage of several features and algorithms. We speculate that a viable approach is the design of a meta-algorithm, where different algorithms are good at different parts of the input space. The interesting question then becomes combining these partial results into a single solution that outperforms any of the initial ones.

Delivering grid resiliency

Also known as self-healing capabilities, network resilience deals with the problems of detecting faulty states in the system and autonomously carrying out emergency controls and restoration procedures by managing the flow of power to resume normal operation. The current procedures used in handling such events are unfortunately limited to having human operators perform these tasks, due to a low level of automation especially in distribution systems. Power Supply Restoration (PSR) is the problem of reconfiguring a network in such a way as to isolate the faults and resupply as many customers as possible in the shortest amount of time. Hence, PSR is subject to computational runtime requirements that can generally be extended to minutes at most. Additionally, a significant body of work has been investigating the phenomenon of power outages and cascading blackouts [1, 23, 5]. A multiagent system solution for power system disturbance diagnosis is proposed in [67], where through a process of inter-agent communication, agents collaborate to provide a comprehensive disturbance diagnosis.

A straightforward multiagent approach for PSR is given in [121]. The system consists of two layers. The lower level includes agents for consumers, DERs and monitors that have the ability to sense, communicate and act. The upper layer is represented by a management agent that carries out an optimization procedure based on integer linear programming models. The scenario assumes the occurrence of disruptions in energy availability due to inoperative lines or generators. Thus, the optimization goal is to identify the optimal alternatives of supply sources capable to serve the unbalanced demand sites in the network.

The limitation of such an approach has to do with the complexity and scale of the network under surveillance. Searching for optimal solutions, moreover, in a centralized manner, within acceptable time constraints becomes infeasible as the size of the network considered grows. In [115], Nagata et al. propose a multiagent restoration scheme using a similar two layer approach based on bus agents and a single facilitator agent. The idea is to reach a efficient suboptimal configuration, by having bus agents use a set of simple rules for restoring the system, while the facilitator agent guides the process. The solution is unfortunately capable to handle only single faults, as opposed to the more realistic scenario of having multiple faults simultaneously in the system.

In [172] Thiebaut et al. take a step forward in their approach. On one hand their work deals with level- k plans for arbitrary values of k . Current approaches are often restricted to level-1 plans, which means that only switches located on faulty feeders are considered for remote operation. A level- k plan can operate switches with fault-distance² less than k . On the other hand, the paper distinguishes between two different aspects of the PSR problem. First, they address the typical problem of finding an optimal final network configuration. A flexible mixed-integer programming (MIP) framework is introduced for solving this problem, which benefits from allowing general network configurations to be modelled. Second, the authors address a sequencing problem to determine how best to transition the network into the op-

²Fault-distance is defined as the minimum number of open switches that must be traversed to reach a faulty element from the switch in the state immediately following the fault.




Figure 2.8: The IDAPS multi-agent architecture [138]

timal configuration obtained in the previous step. Similar solutions [26] based on MIP optimization fail to consider the critical issue of sequencing switching operations. In simulation the authors argue that solving the two problems jointly does not satisfy the runtime requirements of PSR, even for modest network sizes and number of faults. The two-step approach advocated does however produce suboptimal plans. Empirically though, the solution proves to produce good enough results.

In terms of existing infrastructure to support this type of automation, of note is the undergoing project proposed by a team at Advanced Research Institute of Virginia Tech [139, 144, 138]. They support an architecture called Intelligent Distributed Autonomous Power Systems (IDAPS), whose goal is to detect upstream outages and react accordingly to allow the micro-grid to operate autonomously in an islanded mode. In essence IDAPS offers a software alternative for islanding a microgrid instead of the traditional hardware-based zonal protection system. This means that the microgrid is able to autonomously disconnect itself from the local distribution utility and maintain the integrity of the system, while continuing operation. The authors advocate for a setting that requires a reduced number of exchanged messages along with an overall reduced complexity. Agents assume one of the four roles depicted in Fig. 2.8 and work in collaboration to secure critical loads within the microgrid during outages. Specifically, the control agent is responsible for detecting faults in the system and operating circuit breakers in order for transitioning into island mode, as well as

informing the user agent and DER agent. The scheme assumes a pre-defined load priority. The user agent and DER agent exchange messages to determine the amount of energy internally available, so that non-critical loads are disconnected if needed in order to stabilize the grid. IDAPS claims to facilitate seamless transition from grid connected to an island mode when upstream outages are detected.

Although, altogether there appears a vast body of work on PSR and grid resilience techniques at large, existing approaches as the abovementioned ones, rely on simplifications such as ignoring power flows and capacity constraints (e.g. [40]) or restrictions to particular network topologies (e.g. [15],[55]), which rule out a direct application of these solutions. Moreover, most papers consider that the faulty network elements are exactly located, in comparison to the real scenarios where there is an uncertainty in fault location, which considerably complicates the PSR problem and which makes it an important limitation. Overall the solutions proposed for this category of smart grid challenges appear still to be in an early phase and preparedness for emergency situations remains lacking.

2.4.2 Coordination and market adaptation to support transitioning to a smarter grid

Encouraging consumer engagement in grid management

The forthcoming of smart appliances is opening up a new research line, providing consumers with a more active role in power balancing in the grid. The benefits that result from actively controlling such appliances in order to avoid peaks in demand, to balance supply from intermittent resources or to avoid costly investments in expanding the current infrastructure, have been investigated in a number of studies. Synergies between controlling demand and solar photovoltaic generators have been addressed in [33]. Similarly, evaluating the benefits from balancing demand with intermittent wind power generation is presented in [163]. The authors also give an estimation of the benefits derived from applying demand control on the issue of network congestions. Some hints in the sense of consumer acceptance of such techniques is provided in [107].

In [168] the authors estimate a cost comparison between the investments required in transmission and distribution networks opposed to the investments required in order to increase the utilization of the grid. All these works suggest a consensus regarding the need and opportunity of encouraging consumer engagement in grid operation.

In the following, we touch on several notable proposals. Gottwalt et al. [52] presents a simulation-based approach assuming households equipped with smart metering and intelligent appliances. In their experiments they replace the widespread adoption of flat tariffs with time-based electricity prices. These are more often known as time-of-use tariffs. Essentially they represent a fix day-ahead pricing for each time interval of the following day. The authors assume a reduced penetration rate of shiftable appliances that resembles the current typical appliance set in Germany. For simplicity they also limit the usage of appliances to three operational modes only, which set the finishing time per device to 5h, 10h and until 6 a.m. the next day, respectively. In their simulations they report avalanche effects produced by shifting appliances to hours with cheaper price. That is, they manage to eliminate the initial peaks, only at the expense of producing new peak loads. Hence the savings incurred by consumers are low.

In [113], Molderink et al. propose a three-step control methodology focused on domestic energy streams within a micro-grid. The first step is a local prediction procedure achieved through an 'in-house' system, based on a neural network model. Although predicting the consumption profile is run locally the mechanism as a whole follows a centralized approach. This is shown in step two, which ensures the global planning. Thus, given an objective of the global demand pattern the algorithm decides at which time appliances are switched on/off. The planning horizon covers a one day interval and is shown to be NP-complete in [14]. Hence the authors propose a heuristic that aims to generate a "good enough" plan through an iterative distributed dynamic programming procedure. Finally, step three is a method for applying real-time adjustments to controlling the devices in the house.

Another approach based on optimization algorithms is given in [70]. Here, Hubert et al. focus on the problem of energy optimization at residential level. The model each house as endowed with both generating units as well as controlable loads.




Figure 2.9: PowerMatcher architecture [87]

They further assume the access to real-time pricing information, placing the question onto determining the power exchange with the grid, while satisfying certain user requirements.

Automation of load control for sustainable buildings is addressed by Wang et al. in [185]. Here, a hierarchical multi-agent based control is utilized for energy and comfort management. Agents are classified into four layers, that include *i*) agents monitoring flow between the building and the grid, *ii*) a central coordinator agent, *iii*) multiple local controller agents and *iiii*) load agents. The 'intelligence' of the system lies primarily at the level of the central agents that is responsible for determining the power dispatch to the agents situated at the lower levels, which are directly controlling devices within the building. Automatically controllable devices are differentiated into three categories. The task of the central agent is to optimize an objective function that represents user comfort and which represents a prioritization over three user desirable set points linked to temperature, illumination and air quality respectively. To solve this problem the authors apply a heuristic based on particle swarm optimization (PSO), which inherently offers no guarantees for reaching the global optimum solution, but that is empirically shown to return good enough results after 10 iterations of the PSO procedure.

A thread of work that is gaining more credibility in the smart grid domain is

focused on distributed decision making based on microeconomics. The key feature of such an approach is in modelling individual economic agents that interact strategically, in order for pursuing their private interests. The main premise of these approaches is that the large majority of consumers do not want central coordinators, such as utility companies, to control the system in their home, regardless of the potential for savings. Kok et al. [86, 87] implemented an agent architecture that runs an electricity market on a micro-grid setting. They coin the system the *PowerMatcher* and importantly, provide field experiments to validate their approach. The problem of striving for an equal supply and demand is directly transformed into a typical price-based market-based control problem, through a hierarchical agents structure. Namely, they consider a local device agent that controls the operation of every device and concentrator agents that represent a number of device agents in the communication with the auctioneer agent, which performs the price-forming process. An representation of the *PowerMatcher* agent architecture is given in Figure 2.10.

Ramchurn et al. [147] propose an approach inspired from the principle of homeostasis, which is based on three recurring actions: sensing, sending a control signal and feedback. Sensing has to do with exploiting external conditions, such as weather forecast to predict renewable generation and load profiles. For signalling, the authors introduce a carbon-based pricing scheme that matches the carbon intensity of the grid in real-time. Finally, the communication between consumer and supplier close the loop and aims to ensure that the aggregated demand from all the consumers is as close as possible to the real-time supply of producers. The signal provided by suppliers to consumers is a target increase or decrease in consumption, with regard to the previous day. In case consumers achieve this target consumption they are guaranteed a fixed lower price for electricity, while they are set to pay a more expensive one. Obviously, the underlying assumption is that consumers would actually contract such a plan where they need to continuously adapt demand to these signals in order to maintain a fixed minimal cost of electricity.

Vinyals et al. discusses in [181] the concept of a virtual energy consumer, that comprises a collection of consumers with complementary needs. The aim is to represent such groups as single entities that can act directly in the market. Distinguishing

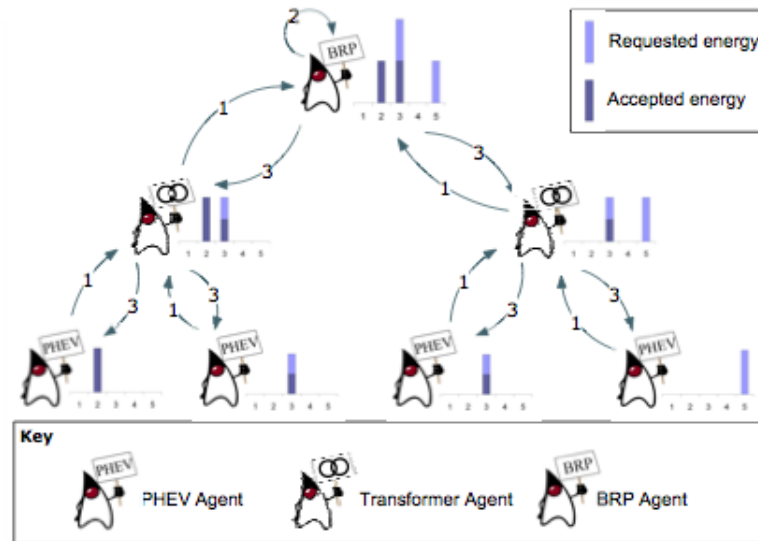


Figure 2.10: Schematic overview of MAS architecture for PEV scheduling [177]

between long-term and short-term energy markets, the goal for a group of consumers is to determine the quantities of energy to be bought in each of these markets so that the needs of the group are met, while minimising costs. The experiments prove that such an approach is viable only when there is a significant difference in the prices of the two markets (e.g. one is double than the other).

A category of deferrable loads with significant impact, assuming a major expansion in the near future, are plug-in (hybrid) electric vehicles (PHEV). In [27] the authors claim that a penetration of over 10% of PEVs will cause power losses and voltage excursions that would be unacceptable under current norms. In their approach the authors devise a classic centralised solver to determine the charging schedule.

Vandael et al. [177] propose a hierarchical structure for solving the problem of scheduling PEVs. The paper introduces a decentralized, multi-agent system solution for coordinated charging. Thus, the initial problem is decomposed into several sub-problems that are locally solved by intelligent agents residing at the charging substations. The owners of PEVs are represent their charging preferences as an 'intention graph'. Through simulation they manage to show a 14% improvement in cost.

Gerding et al. address in [48] the considerable strain on electricity distribution net-

works that is expected to be generated by PHEVs. They propose an online auction protocol, where agents representing owners of PHEVs place bids for power allocation. The authors model the setting as an online mechanism design problem, where agents are incentivized to report truthfully their demand. In terms of scalability, the algorithm is said to handle up to hundreds of agents. Experiments show that in comparison to well known scheduling heuristics, which assume non-strategic behavior, the algorithm performs well, achieving similar results. The downside of this approach is that it requires occasionally to discharge units of electricity that have been previously stored in the PHEV's battery so that the property of incentive compatibility is maintained. Moreover they make the assumption of an instantaneous 'burning' of the allocated power.

In [179], Vasirani et al. introduce a mechanism inspired from lottery scheduling, a randomised resource allocation mechanism that has been developed for operating systems. Their approach allows to purposefully determine the efficiency-fairness trade-off of the mechanism. In terms of efficiency, the goal is to maximise the allocation of power to the agent that values it the most, while fairness stands for an egalitarian allocation. The underlying assumption here is that electricity is priced according to a fixed, per-unit price plan, while the possibility of any other type of dynamic pricing scheme is not considered. Notably, they do not assume that the preference of PEVs (e.g. the preferred time to charge, the quantity of electricity that is needed) are necessarily truthfully revealed.

Integration of generation with profile uncertainty

The growing pressure of increasing DER adoption driven by environmental concerns is faced with transitioning from the passive operation of transmission and primarily, distribution networks, to an active one. This means a shift from traditional central control to a new distributed controllability. Failing to do so will lead to system instabilities and inefficiencies due to over-capacity issues and under-utilisation of the assets, respectively. To achieve this, a new technical architecture along with changes in the commercial and regulatory structure of liberalising electricity industries is re-

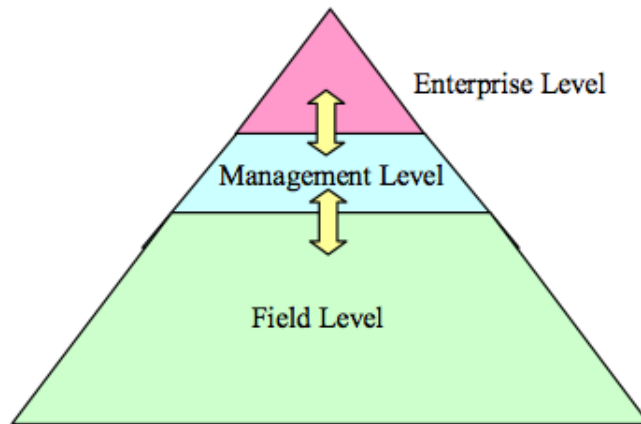


Figure 2.11: The MAS levels of organization. [36]

quired. In this subsection we review some of the prominent solutions investigated to be able to facilitate a smooth integration of DERs from an ICT perspective, capable to cope with the increasing complexity of interaction required to perform decentralised system management.

The concept of a *Virtual Power Plant* was developed to enhance the visibility and control of DERs in the system. As there is no strictly recognised understanding of what a VPP must ensue [35], in practice every VPP-like solution addresses different aspects and takes on a different view on its purpose and functionality. In principle though, the VPP is responsible for aggregating and managing the combined operation of various production units. Also, as the name suggests, a VPP is comparable to a transmission connected generating plant. Therefore, it must be characterized in a similar manner by operating costs, generation limits, generation schedule, etc. creating a single operating profile from a composite of the parameters characterising each DER.

Again, the multiagent paradigm appears to be a very useful tool for the operation and control of a power system. In [36] the authors present the advantages of using agents for VPP control by giving a high level description of the system (see Fig 2.11). On the field level are all agents associated to the production units. The management level has the responsibility to communicate and coordinate among the existing

agents. Lastly, commercial considerations for the VPP, such as market participation are deliberated at the enterprise level. Beyond this perspective though, the paper fails to give any concrete approach towards implementing such a system.

A significant number of approaches focus on a restrictive set of DER devices or on particular combinations of DER and controllable loads. For instance, in [158] the authors identify the combined heat and power technology as the one which has got the highest potential for the integration into a virtual power plant. In particular they look into the connection of combined heat and power micro-units. The study tries to estimate whether a CHP VPP is viable from an economical standpoint instead of the conventional power plant. The conclusion is that in terms of economical feasibility, the price for a produced electrical kilowatthour is higher than the price from of a conventional power plant.

Costa et al. [29] analyse the idea of coupling generation from wind power with storage facilities. For operating the VPP in an electricity market scenario, the authors propose a method based on dynamic programming techniques. Expectedly, the simulation shows that using energy storage improves the baseline scenario (no storage) by an approximate 5% under certain conditions. Although not proposing a solution to cope with the situation, the experiments also show that the uncertainty in wind power output is the cause for a 18% income loss for the VPP. A similar economical analysis is performed for the usage of PEV batteries by extending their life-cycle and installing them in buildings. In [10] the authors consider a microgrid scenario and investigate the optimal equipment combination for the task of ancillary services provision for the setting of California's electricity market.

In [24] the authors take a different perspective on VPPs, analysing the profitability of rational autonomous DER-agents, representing small-to-medium size renewable electricity producers, to group together into cooperatives and sell their energy to the electricity grid. The paper proposes a pricing mechanism with certain desirable properties. In particular, it introduces a scheme to allocate payments within the cooperative that is in the core and as such, no subset of members has a financial incentive to break away from the cooperative. The pricing mechanism is however dependent on the extent to which the grid operator is able to quantify the reliability

of DER estimation. A critical limitation of the approach is in using point estimations, which given the uncertainty around production of renewable energy has the risk that estimations can be widely off. Moreover, distributing rewards for cooperative members turns out to be an intractable problem for large cooperatives. The experiments performed restrict the number of participating agents to 24.

An interesting study addresses the energy inefficiencies that occur especially in off-grid, remote villages in the developing world [3]. The paper focuses on a particular scenario where households equipped with renewable units and electric batteries are operated in isolation. The authors propose an energy exchange between homes so that overall, the system benefits from a reduction in energy loss and prolongs the life of batteries. Under particular circumstances detailed in the paper, based on real world data, the need for energy charging is said to be reduced by up to 65%, while the use of energy can be improved by up to 80%. The solution is based on autonomous coordination of such resources. In particular the ingenuity consists of constructing a scheme that does not rely on monetary payments, which would not represent a feasible solution due to a serious hindrance in economical and social development of such communities. The agent utility is defined in terms of total battery charging and used as transferable utility among agents, given that the efficiency of electric batteries degrades with usage and time. Of note, is the fact that calculating an optimal allocation beyond 16 agents is not feasible in reasonable time, even when the exchange takes place over a single day. Thus the paper proposes a heuristic to approximate the optimal allocation, which is empirically shown to perform reasonably well. The approach is also suitable for the problem of unit sizing [78, 47], that is the procedure of determining the optimal size for renewables and storage units in the system.

In a more practical approach, the European project FENIX [17, 18] aimed to conceptualise, design and demonstrate a technical architecture and commercial framework that would enable distributed generation in a cost efficient, secure and sustainable fashion. The project consortium consists of 20 European partners and a large laboratory at ISET used for demonstrating the capabilities of the system. Their variant of Virtual Power Plants represents an aggregation of distributed generators

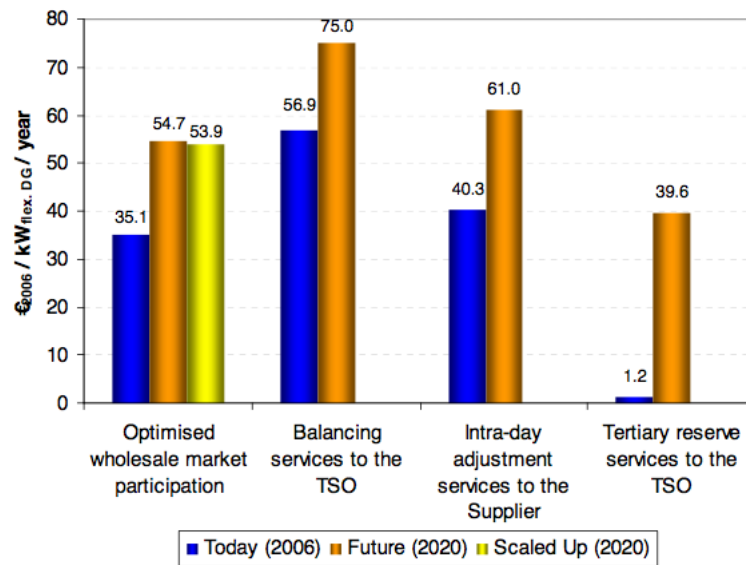


Figure 2.12: Aggregated net-benefits for business stakeholders [176]

through a centralised control architecture. The objective is to include thousands to millions of DER units in a single VPP. While the conceptual aspects of the architecture are kept simple and do not stand out in any particular way, the project attracts attention in terms of the demonstrations that were performed on real networks. The experiments were split into two scenarios: the Northern and the Southern scenario. The aim of the former scenario was to demonstrate the value of market participation in a VPP, using a cluster of small scale generators linked to a common low voltage network. The centralised control is responsible for how to best to dispatch the generation in its portfolio. This implies an optimisation procedure where the VPP offers the possible reactive power capacity and the corresponding costs and receive set points from the network operator. The latter scenario was focused on demonstrating the value of having distributed generation connected to medium voltage networks to deliver ancillary services. Overall, the experiments claim that the operational costs can be reduced between 10% and 48%. An indication of net-benefits for business stakeholders is given in Fig. 2.12, which shows significant returns for investors. In the future, returns are expected to increase mainly due to higher electricity prices.

Electric vehicles (EVs) essentially represent storage units, thus can also be thought

of as generation in a grid integrated vehicle (GIV) scenario. However, using EVs to provide power to the grid at various points in time differs from regular storage units in that it has to deal with the uncertainty in generation caused by the driving behaviors of the EV users. Vandael et al. report in a recent study [178] an evaluation of a real-world GIV scenario performed at the University of Delaware in collaboration with PJM (Pennsylvania - New Jersey - Maryland) Interconnection³. The test environment considers an infrastructure of 15 EVs owned by Delaware University interacting with PJM. The evaluation demonstrates that even for a small number of EVs the ancillary services can be provided in a reliable manner, while preserving the driving requirements of the user. Unfortunately, there is no insight into the financial benefits of the EVs owners, that would expect to be rewarded according to their contribution. Also the authors make key assumptions in terms of the predictability of driving behaviors. Instead of proposing a mechanism where an agent learns the driving patterns of its user in order to infer the charging flexibility, this is assumed to be given, which in fact avoids the problem of dealing with this uncertainty.

³PJM is the largest TSO in the world, servicing 13 states in the Northern and Midwestern US and operating several power markets.

2.5 Simulator

In order to empirically evaluate the different *coordination mechanisms* proposed in this thesis we developed a custom simulator implemented using the Java Platform Standard Edition (Java SE)⁴. Also, in the development process we have been using the Repast software package⁵.

Repast Symphony is pure Java extended portfolio for developing agent-based models. Essentially, it provides a number of modelling components for constructing complex software-intensive systems, as well as supporting the usage of external statistical programs, like Matlab. Noticeably, it has also been used as the groundwork for developing several reference simulators such as PowerTac, which is a simulation of future retail electric power markets, where brokers compete with each other in terms of profitability [80] and AMES, the wholesale power market test bed for the systematic experimental study of new market designs and policies [100]. Moreover, as we go on to describe in the following, the components introduced by Repast fit well to the nature of the simulations implemented in this thesis.

Model Building

Following the structure of the thesis, we can differentiate three main areas for evaluating the coordination mechanisms proposed herein. Firstly, in Chapter 3 we are concerned with the evaluation of mechanisms designed for dynamic microgrid formation through the interaction of consumer and producer agents in a distributed manner. Secondly, in Chapter 4, we address the problem of coordinating producer agents, starting with investigating a game setting w.r.t economical benefits of VPPs and then evaluating the efficiency of our proposed mechanism for VPP coordination. Also, we scrutinize in Chapter 6 the behavior of producer agents in market environments with respect to the problem of collusion. Thirdly, in Chapter 5 we simulate mechanisms for coordinating consumer agents for optimizing demand for

⁴<http://www.oracle.com/us/technologies/java/overview/index.html>

⁵<http://repast.sourceforge.net/download.php>

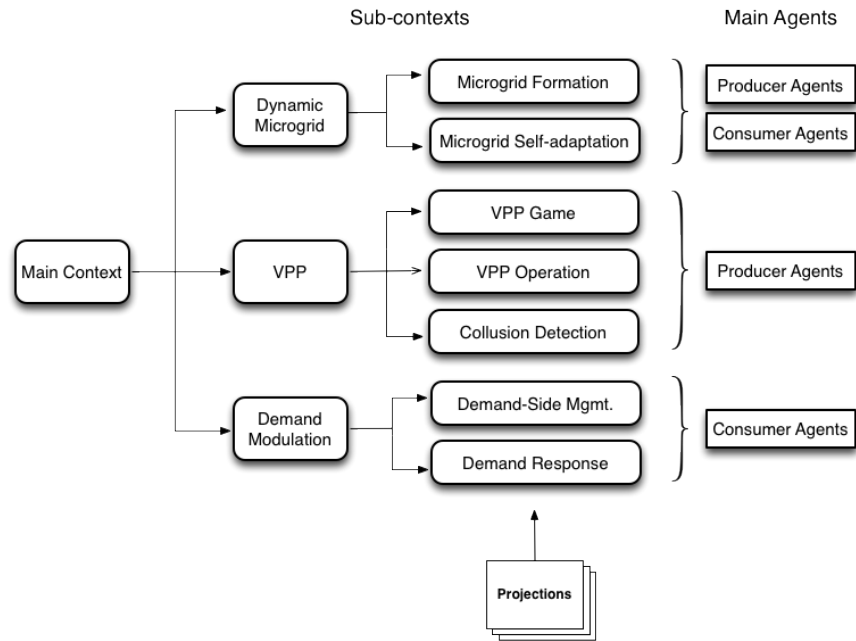


Figure 2.13: Context structure of the simulator.

the day-ahead and intraday operations of the grid. Although these challenges, from a simulation environment standpoint are diverse, we have engineered the simulator based on several *loosely coupled* components. In this regard, the Repast package has introduced several convenient modalities.

Namely, the main concepts introduced in Repast are *contexts* and *projections*. A context can be understood in the sense of a named container that holds model components (any type of Java object), in our case, agents. Additionally, contexts can be hierarchically nested and agents can be present in multiple contexts, as well as change them dynamically. In Figure 2.13 we represent the context structure of the simulator. The main context, also the core object in Repast, holds all sub-contexts and their corresponding agents. We include three main sub-contexts, which correspond respectively to the three main areas identified above, which are then further contextualized. Also, in Figure 2.13 the agent types for each sub-context are shown.

The control of relationships among the agents is supported by means of projections, which represent structures defined upon the agents in the context. Projections

are created for specific contexts and automatically contain every agent within that context. A context may and usually does include more than one projection, which allows for flexibility in testing protocols and agents in different environments by interchanging projections. For instance, in the *microgrid formation* context we define different communication topologies between agents, in order to evaluate the efficiency of our protocol given different instantiations of the network. Agents can then make use of these projections to access or provide information by querying and prompting respectively their neighbourhood according to the protocols prescribed by our mechanism.

Agent Behavior and Simulation Engine

At the simulator level, managing the execution of the agents' actions is done in a synchronous cyclic fashion, where at each time-step the schedule iterates through the set of agents, executing actions following the given agent behaviors. The agent behavior describes the actions performed by the agent during simulation either based on *time-driven* processes or *event-driven* methods. In principle, a mixture of both approaches are used throughout the simulations, still there are clear instances where one or the other are more suitable and sufficient. For example, in the *VPP Game* context we implement several agent behaviors that correspond to particular strategies. Then, we run a repeated game scenario, where at each round of the game, which comprises a number of predefined time-steps (following the game protocol), the participating agents are required to perform actions based on their strategies. What happens in fact is that agents, implemented as objects, cause actions to occur by registering them with a *Scheduler*, implemented as a discrete event clock manager, that runs the simulation.

```
//Define params for action start at time-step 1 and repeat every 10 steps
ScheduleParameters params = ScheduleParameters.createRepeating(1, 10);

//Schedule agent PlayerGradAsc to execute method choose given params
schedule.schedule(params, PlayerGradAsc, "choose");
```


Alternatively, in the *Microgrid Formation* context we simulate essentially peer-to-peer algorithms where the agents' behaviors are centred primarily on the interactions with the agents in their vicinity. More appropriately, we use here dynamic scheduling instead of using static times for the schedule parameters. The idea is that agents can schedule actions to the main *Scheduler* by specifying triggering conditions based on querying particular types of neighbours.

```
@Watch(watcheeClassName = "PlayerAgent", watcheeFieldName = "committed",
query = "linked_from", whenToTrigger = WatcherTriggerSchedule.LATER,
scheduleTriggerDelta = 1, scheduleTriggerPriority = 0)
public void takeAction(){
}
```

For the remaining contexts we use a combination of the abovementioned approaches.

Configuration

Provided that the different contexts correspond to the different sub-problems addressed in this thesis, we devise batch simulations for the separate scenarios considered. On the one hand there are the high level configuration files, such as *context.xml* which specifies the hierarchical organization of the sub-contexts, while on the other hand are the context specific configuration files which comprise the necessary parameter initializations (projections, no. of runs, no of agents, etc.). Additionally, for agent initializations, XML *parameter sweep files* are described. Defining a sweep can also be done by means of specifying nested parameter setters, so that before incrementing the values of the parent parameter, the space of the child parameter will be swept. *Object Loggers* are used for recording the state of objects (agents) or sets of objects, which are sometimes exported, subject to post-processing and charting using the Matlab plug-in.

2.6 Chapter Summary

This chapter was devoted for laying the foundation on *smart electricity grids* from different perspectives and introducing several domain-dependent concepts. We then reviewed and classified the main challenges of this domain sketching out general solutions to tackle these issues. In Section 2.2 we briefly discussed the open multiagent systems paradigm emphasizing its applicability to this domain, as well as visiting some of the relevant coordination mechanisms in Section 2.3.

The existing literature on *smart grids*, discussed in Section 2.4, distinguishes between two main lines of work.

On the one hand, there is a considerable body of work dedicated towards enhancing the energy infrastructure by providing software support that can either perform interpretations and processing of smart meter data, or utility software for managing grid resilience. In this thesis we are **not** looking into the various ways to infer end-user consumption profiling, nor other related forecasting approaches, however these aspects are addresses throughout this work to the extend that they represent prerequisites for the mechanisms proposed herein. With regard to the problem of grid resilience depicted in the second part of Section 2.4.1, where the focus is placed on fault detection and system restoration, we advocate for a novel preemptive approach based on a dynamic reorganization of the actors in the grid into resilient clusters in accordance to their profile variations. In doing so, we introduce a novel model for *smart grids* that exploits organizational and power loss inefficiencies of current systems.

On the other hand, the transformations occurring in the energy sector, that are increasing the autonomy and flexibility of the actors involved, require a shift from current centralized control over the system towards a coordination framework that can deal both with the unpredictability at the supply-side as well as at the demand-side. As seen in the first part of Section 2.4.2, there are many challenges in shaping the energy consumption, making consumer demand adaptable to energy supply. The solutions proposed so far fall either in the category of *i)* centralized solutions that exert direct load control, overwriting consumer autonomy by not allowing consumer

to fully retain control of their own appliances, or *ii*) enable consumer autonomy at the expense of undesirable effects like herding (creating new peaks), which oppose the initial goal in the first place. In this work we focus on the design of mechanisms where agents coordinate in such a way that we can mitigate the apparent contradiction between individual agents' and the system's goals. We take an integrative approach by formulating solutions in a game-like manner that allows us to structure the rules of the game in such a way that the desired global properties of the system emerge.

In the second part of Section 2.4.2 we have reviewed related work on the problem of integrating intermittent supply. However, none of the approaches are comprehensive and general enough to provide a framework to coordinate a heterogeneous group of distributed energy generators. Moreover, we augment our approach by tackling the important problem of coping with the inherent stochastic environment, while determining the optimal dispatch in a real-time manner.

In Section 2.5 we have presented a general description of the simulator developed to investigate the specific problems addressed in this thesis.

To sum up, the aim of this work is that of designing a full-fledged framework that extends the state of the art with an efficient scheme for the creation of a robust, intelligent electricity supply network and agent-based coordination algorithms for the problems of active consumer and producer integration.

Chapter 3

Coordination Mechanisms for *Dynamic Micro-Grid Formation*

The grid, which is made up of everything from power lines to generators to the meters in your home, still runs on century-old technology. It wastes too much energy, it costs us too much money, and it is too susceptible to outages and blackouts. To meet the energy challenge and create a 21st century energy economy, we need a 21st century grid.

— *President Obama* —

In today's energy value chain *retailers* play a central role. They bare the responsibility of energy supply and hence, deliver the billing system to businesses and household customers alike. Retailers are the last value adding party before energy is actually delivered. Electricity retailing holds the function of estimating the amount of energy used by (domestic and industrial) consumers and committing large power plants to this amount (often located at considerable distances), which is then carried over the transmission network to various pools of consumers. It is essentially an intermediary servicing the balance between electricity production and consumption for a certain profit margin. The coming of smart grids, bringing about the opportunities of small-scale, distributed generation, requires to drastically rethink the fundamentals of the energy ecosystem. Owing to the transition of traditional customers to their

new role of energy *prosumers*, new processes and interactions between these grid participants need to be designed. In this chapter we advocate an approach that places prosumers at the very heart of the smart grid.

We envision a setting where participants cooperate to ensure the grid functions efficiently, while the security of supply is ensured. Essentially by placing *intelligence* at the prosumer level, we *remove retailers from the loop of supply-demand-matching* and substitute their role through *decentralized peer-to-peer mechanism for dynamic microgrid-formation*. To achieve this, we deconstruct the retailer managed macrogrids through conception of regional configurations including a feature of locality into the market that supports a distributed, sustainable provision of supply.

- By *microgrid* we denote a cluster of prosumers that distribute power locally.
- They service small geographical regions so as to relieve the demand on the macrogrid.
- They are self-sufficient, as they do not rely on external power inputs, allowing them to operate even when other parts of the grid are dysfunctional (so-called *islanded mode*).
- They *minimize the power loss* occurring due to the transfer of power over the network lines.
- They allow the instantiation of distribution-level energy markets where participants can agilely react to ongoing, local levels of production and consumption.

Additionally, it is important to say that in contrast to the traditional microgrid, the prosumer configurations we introduce here are not static in nature, but are required to adapt in accordance with the forecasted profiles of energy consumption and production. Also, in today's market structures small-scale producers are prevented from direct participation because of capacity-related barriers to entry and unpredictability. By moving to a distribution-scale energy market we aim to provide a platform where a more efficient coordination of production for DERs is possible, as well as consumer-side mechanisms to target high peak consumption.

The remaining of this chapter is organized as follows. We begin in Section 3.1 by generally discussing the background for our approach. In Section 3.2 we introduce the formalization of the problem. Our agent-based organizational model is introduced in Section 3.3, with emphasis on the coalition self-adaptation scheme proposed in Section 3.5. In Sections 3.4 and 3.6 we present experimental results, while Section 3.7 concludes, relating our solution to existing techniques for team formation in MAS.

3.1 Microgrid Solution Approach

In this chapter, an agent-based organizational model for a smart energy system is introduced relying on a dynamic coalition formation mechanism for microgrid configuration. Central to this mechanism we propose a notion of stability that stems from the existent solution concepts in coalitional games. The process is intended as an open-ended organizational adaptation, concerned with achieving stable configurations that meet the desired functionalities within stochastic scenarios. We deploy the mechanism in distributed environments populated by negotiating agents and give empirical results that prove a significant improvement of organizational efficiency.

In recent years, there has been an increasing interest in the integration of distributed, small-scale, renewable generation into the power system. An efficient use of distributed energy resources may increase the flexibility and the resilience of the power system at the distribution level. Furthermore, it is possible to reduce the dependence from large-scale, non-renewable, power plants and therefore to contribute to a sensible reduction of CO₂ emissions. According to the US department of Energy, a 5% increase in grid efficiency is equivalent to the fuel and CO₂ emission of 53 million cars.

The potential allure of the multiagent system paradigm to the power industry has been extensively documented so far [104]. To this respect, several management systems have been proposed for the organization of the grid. On the one hand microgrids [61] have been advocated as a subsystem of the distribution network, formed by generation, storage and load devices, interconnected at the electrical and the informational level. Micro-grids can be intended as a systemic approach to realize

the emerging potential of distributed generation.

Setting aside from this approach that aims at imposing a static architectural control, whether centralized or not, on already predefined micro-grids, our vision is intended at proposing a method for congregating the smart-grid actors (DERs and consumers alike) to dynamically approximate optimal micro-grid configurations. To this end, the procedure is designed such that it develops a new notion of microgrids by integrating prosumers in the form of *dynamic configurations*. A dynamic microgrid is conceived as a bundle of prosumer that are connected through an informational infrastructure and act in a coordinated way as a single entity for a determined period of time. The challenging problem related to the implementation of this concept is the distributed control of the DERs, mainly due to the stochastic behaviour of the system and the heterogeneity of the devices involved.

The aim in this chapter is modelling the coordination of dynamic microgrids in the sense of dynamic coalition formation. Instead of considering centralized architectures [143], we claim that a dynamic, bottom-up, approximation of optimal configurations is more effective to ensure flexibility and robustness to the system.

Variable	Description	Variable	Description
\mathcal{A}	set of agents a_i	e	excess energy profile of a coalition
\mathcal{P}	set of provider agents	σ	power loss in a coalition
\mathcal{L}	set of consumer agents	c	cost of a coalition
\mathcal{G}	communication graph	ψ	target energy profile
\mathcal{N}_i	set of neighbours for agent a_i	s	state of an agent
E	set of edges	p	probability for coalition initiation
β_i	energy profile for agent a_i	φ	association coefficient
Φ	set of constraints	x, y	cost vectors
S, T	coalitions	μ, τ	thresholds
\mathcal{CS}	coalition structure	u	coalition utility
f	efficiency	k	dissipation rate
$P_{a_i \rightarrow a_j}$	power exchanged between a_i, a_j	l	line length

3.2 Problem representation

The problem we address in our approach is one where the coalition formation procedure is projected on an underlying network topology. Moreover, the domain is non-superadditive, in the sense that gains resulting from forming coalitions are limited by the actual cost of coalition formation¹ and coordination, thus the grand coalition is seldom the optimal structure. Additionally, the problem is subject to the dynamism of the environment. In contrast to static coalition formation, dynamic settings represent a more complex issue since the focus is not merely on analyzing the coalitional structure, but the main aspect under investigation is how the formation of the coalitional structure takes place through the players' interactions and its adaptability to environmental variations or externalities, which denote the evolution of the system. The challenge is to develop mechanisms that permit large numbers of autonomous agents to collectively achieve a desired functionality by permanent adaptive dynamics. Therefore, it is desirable that the coalition formation process takes place in a distributed manner, leveraging on the autonomy of the agents, which spontaneously organize into topologies and functionalities to meet the desired objectives. We believe that in order to solve these issues, the problem must be understood in the context of self-organization by providing a minimum set of interaction rules that would lead to an efficient achievement of the underlined desiderata.

Returning to our initial setting, the algorithm we propose is illustrated in the context of smart energy systems. We set to investigate the integration of renewable energy resources to the grid in the form of microgrids by means of aggregating the power generating potential of various devices in a novel way in the context of MAS. As system designers, we choose to enable the autonomous agents with the basic coordination primitives, and leave to the agents to self-organize and coordinate as the situation may demand, in a fully distributed manner.

We model the problem as a dynamic coalition formation procedure with the following formalization:

¹For instance, similar scenarios appear when the cost of forming a coalition can be perceived through the negotiation process and information exchange which incur costs.

Let $M = \langle \mathcal{A}, \mathcal{G}, \beta_i, \mathcal{CS}, \Phi, v \rangle$ be a multi-agent system where:

- $\mathcal{A} = \{a_1, a_2, \dots, a_n\}$ represents the set of agents of a given portion of the distribution grid. We assume that each stakeholder that is connected to the grid is represented by a software agent that manages its corresponding devices (e.g., generators, storage devices, intelligent loads).
- $\mathcal{G} = (\mathcal{A}, E)$ is the underlying communication network denoted as an undirected graph where the set of vertices is the set of agents and edges are communication links. \mathcal{N}_i represents a_i 's set of neighbours s.t. $\forall a_i, a_j \in \mathcal{A}$, if $(a_i, a_j) \in E$ then $a_i \in \mathcal{N}_j$ and $a_j \in \mathcal{N}_i$.
- β_i is the forecast amount of electricity for the following day associated with agent a_i . If $\beta_i > 0$, then agent a_i is a *provider*, whilst if $\beta_i < 0$ then agent a_i is a *consumer* (or load). Let $\mathcal{P} \subseteq \mathcal{A}$ denote the set of providers, and $\mathcal{L} \subseteq \mathcal{A}$ the set of consumers. In this chapter we assume that an agent is either a provider or a load, and therefore $\mathcal{L} \cup \mathcal{P} = \mathcal{A}$, $\mathcal{L} \cap \mathcal{P} = \emptyset$. We will refer onwards generically, to an agent belonging to set \mathcal{P} as PA , and to an agent belonging to set \mathcal{L} as LA .
- A *coalition* S is a subset $S \subseteq \mathcal{A}$ that satisfies the set of constraints $\Phi = \{\phi_1, \phi_2\}$:

$$\phi_1 : \sum_{a_i \in S_j} \beta_i > \sigma(S_j)$$

$$\phi_2 : \sum_{a_i \in \mathcal{P}_j} \beta_i \in [\psi_1, \psi_2]$$

where $\mathcal{P}_j = \mathcal{P} \cap S_j$; $\psi_1, \psi_2 \in \mathbb{R}_+$ and

- the *excess* $e(S_j) = \sum_{a_i \in S_j} \beta_i$ represents the energetic balance within a coalition.
- $\sigma : 2^{\mathcal{A}} \rightarrow \mathbb{R}_+$ specifies for every coalition $S \subseteq \mathcal{A}$ the *power loss* over the distribution lines inside the coalition.

- ϕ_1 enforces that each coalition is able to supply electricity to all of its loads, so that the energetic balance e between generation and consumption remains positive within the coalition, while also accounting for the power loss in the system given by σ ; ϕ_2 prescribes that each coalition must realise a specific generation profile of electricity that would qualify them as microgrid, lower and upper bounded by ψ_1 and ψ_2 respectively.
- $c : S \rightarrow \mathbb{R}$ represents the *cost of a coalition*, which quantifies the power loss within coalition S and its deviation from an equilibrated energetic balance where supply matches demand:

$$c(S) = \sigma(S) + |e(S)|$$

With the assumptions and definitions above, we formulate the coalition formation problem as follows:

Determine a coalition structure $\mathcal{CS}^* = \{S_1, S_2, \dots, S_m\}$ from the set of possible coalition structures \mathcal{CS} , representing an exhaustive and disjoint partition of the set of agents \mathcal{A} , such that the total loss of power over the distribution lines is minimized:

$$\mathcal{CS}^* = \arg \min_{\mathcal{CS}} \sum_{S_j \in \mathcal{CS}} c(S_j)$$

where

$$\bigcup_{j=1}^m S_j = \mathcal{A}, \quad S_j \cap S_k = \emptyset, \forall j \neq k$$

Again, the goal of the coordination problem is obtaining a partitioning of \mathcal{A} into a coalition structure \mathcal{CS}^* , where each coalition complies with the set of constraints Φ and at the same time maximises the *social welfare* of the system in the sense of minimizing loss of power over \mathcal{CS}^* , without jeopardizing the functionality of any of the coalitions. We leave aside for now what this trade-off implies and further develop on this issue in Sections 3.5 and 3.4.

3.3 MAS-based algorithm for Microgrid formation

The first stage of the process concerns essentially the coalition formation mechanism, which itself consists of three phases and which would yield subsystems of the grid that represent a self-sufficient electrical and informational interconnection of generation, storage and load devices.

Coalition initiation

Initially, during the *coalition formation phase*, neighbouring PAs need to coalesce in order to attain a minimum threshold power that would qualify them as microgrid (ϕ_2). The actors can assume one of the following states: *committed* or *uncommitted* to a coalition. Function $s : \mathcal{A} \rightarrow \{\textit{committed}, \textit{uncommitted}\}$ returns the status of one agent. In the beginning, all agents are assumed to be uncommitted. Additionally, we presume that the agents' ability to communicate is bounded by topological constraints: a_i and a_j can communicate within the predefined distance Δ ; N_i will denote the set of neighbouring agents for a_i . The agent's decision of collaborating towards establishing a microgrid is based on aggregating information about agents in its proximity, as well as other domain dependent values such as association coefficients, as will be further detailed in the following sections. There are numerous possible procedures for initiating a coalition (e.g. leader election [171]), though in order of placing into focus our coalition formation mechanism we have chosen a straightforward approach for accomplishing this task. Thus, those PAs whose energy availability exceeds a predefined value (ψ^*) are entitled to establish themselves as microgrid initiators and will do so with a probability p inversely proportional to the number of the agent's uncommitted PA neighbours that are also set to do so.

$$p(a_j) = 1 - \frac{|\{a_i | a_i \in N_j, \mathcal{P}, s(a_i) = \textit{uncommitted}, \beta(a_i) \geq \psi^*\}|}{|\{a_i | a_i \in N_j, \mathcal{P}, s(a_i) = \textit{uncommitted}\}|}$$

Algorithm 1 Provider Aggregation; Initiator behavior.

```

for each  $a_i \in \mathcal{P}$  do
  if  $\beta(a_i) \geq \psi^*$  &  $s(a_i) = \text{uncommitted}$  then
    do with probability  $p$ 
       $\text{role}(a_i) = CA$ 
      /* coalition expansion proposal */
      for each  $a_j \in PA$  &  $a_j \in N(a_i)$  &  $s(a_j) = \text{uncommitted}$  do
        |  $\text{sendProposal}(a_j)$ 
      end
    end
  /* resolve joining coalition request */
  if  $\text{incomingRequestList} \neq \text{NULL}$  then
    for each  $a_j$  in  $\text{incomingRequestList}$  do
      |  $List \leftarrow \text{constructOrderedCandidateSet}(a_j)$ 
    end
    while  $e(S_i) < \psi_2$  do
      |  $a_j = \text{popUp}(List)$ 
      |  $s(a_j) = \text{committed}$ 
      |  $S_i \leftarrow S_i \cup \{a_j\}$ 
      |  $e(S_i) += \beta(a_j)$ 
    end
  end
end

```

Provider Aggregation phase

The next stage to be undergone regards the *aggregation of providers*. The initiator PA assumes the role of CA (coordinator agent) for the emergent coalition. This procedure is realized in a self-organizing fashion that iteratively constructs the coalition structure through a distributed mechanism.

Algorithm 2 Provider Aggregation; PA behavior.

```

for each  $a_i \in \mathcal{P}$  do
  if  $\beta(a_i) \geq \psi^*$  &&  $s(a_i) = uncommitted$  then
    if  $incomingRequestList \neq NULL$  then
      for each  $a_j$  in  $incomingRequestList$  do
         $List \leftarrow constructOrderedCandidateSet(a_j)$ ;
      end
    end
    /*initiate contact*/
    for each  $a_j \in N(a_i) \&\& (s(a_j) = committed \mid\mid (role(a_j) = CA))$  do
      if ( $profile(CA(a_j))$ ) then
         $List \leftarrow constructOrderedCandidateSet(CA(a_j))$ ;
      end
    end
     $sendProposal(popUp(List))$ ;
  end
end

```

As an indicator of the degree to which an agents a_i impacts the formation of a coalition S , we introduce the term *association coefficient*, that quantifies the power loss incurred by S if a_i joins the coalition.

$$\varphi_S(a_i) = \sigma(S \cup \{a_i\}) - \sigma(S) \quad (3.1)$$

With these considerations in mind, the mechanism proceeds in a Contract Net-like manner [164] with the following steps, as depicted by the *algorithms* 1 and 2: *i*) CAs send requests to their neighbouring PAs indicating the microgrid profile they want to realize, in terms of the energetic potential ψ (the constraint ϕ_2 imposed on the coalition formation process may vary according to the desired feature of the emerging microgrid); *ii*) based upon this specifications PAs evaluate the association

coefficients for joining different CAs and select the one that minimizes φ from their respective candidate set; *iii*) finally, CAs receive these responses and take the decision of committing PAs subject to achieving the energetic potential of the coalition and minimizing the power loss.

In the pseudo-code of the algorithm in table 2 the *popUp* function extracts and eliminates from the PA's list the CA that meets the criteria of being the most adequate candidate. Thus, in case the PA would not have been selected by the CA, the latter would not remain in the candidate set of the respective PA. To be noted that the decision of selecting the best candidate is carried without a complete representation of the environment, but rather based on local information. This is the underlying reason for the evolutionary nature of the algorithm that iteratively approximates a solution through refinement steps. It's worth mentioning that, while CAs are looking for suitable suppliers of energy, the PAs themselves are pursuing an active role, ongoing an exploration phase of expanding their candidate set. For instance, information regarding a coalition obtained from another PA already committed to that particular coalition could cause the agent to submit a joining request without being explicitly addressed by the CA. This is meant at increasing the convergence rate towards emerging coalitions.

Consumer Aggregation phase

Once the microgrid's energetic potential has been ensured, the only remaining phase for the coalition formation process requires the *aggregation of consumers* (LAs), which operates in a similar manner to the one previously explained. *i*) LAs proceed by submitting their forecasted demand to the CAs in their proximity; *ii*) for each such request, the CAs calculate the *association coefficient* that reflects the power loss of the potential coalition *i* assuming the joining of the particular LA, while ensuring the energetic balance of the coalition remains positive *iii*) LAs will conclude the procedure by selecting the coalition whose power loss σ is least increased by their commitment. The decision is unequivocally accepted by the CA since it comes as a response to its precedent proposal and it is exclusively addressed to the CA. The LAs' preference for

acting in this sense is justified by the fact that the utility of the coalition would have a direct effect on the price of energy being traded within the coalition as well as its reliability.

Thus, the organization of agents is fundamentally correlated with the association network amongst them that emerges from their interactions in superposition to the topological communication network. The resulting interdependence between the agents within each coalition strongly impacts the characteristics and outcome of the procedure.

3.4 Experimental results

In this section we first evaluate our mechanism for a restricted scenario, considering a small agent population. The problem of generating a coalition structure optimally is a computationally intense one, as the sheer number of combinations makes exhaustive search impossible as the number of agents increases (e.g. computing the optimal \mathcal{CS} for 20 agents requires around 3.4×10^9 operations). This demonstrates the applicability and advantages of our domain-driven heuristic method. In the first scenario we consider a small-scale setting such that we can provide an empirical comparison to the optimal solution, which is however inappropriate in real-world cases. The experiments presented in the second part focus on a large-scale scenario and aim at understanding how the underlying interaction topologies are able to foster, or else hinder, the organizational efficiency of our proposed multiagent model. We are interested in determining the effects of the communication structure on the system's behavior at large, in order for exploiting the network configurations to achieve certain goals.

We initially run our experiments for a population of $|\mathcal{A}| = 15$ agents ($|\mathcal{P}| = 5$ and $|\mathcal{L}| = 10$). The energy profiles of consumers are uniformly distributed with $\beta \in [10, 40]kWh$, while for producers we assume DER generators with $\beta \in [30, 70]kWh$. We generate synthetic topologies to resemble statistical properties of real power grids. A topology represents in fact a graph where nodes denote agents and edges are the transmission lines. Particularly, we use the findings in [63] to generate topologies




Figure 3.1: Percentage gain of our approach against optimal

that preserve the probability for node degree as in the IEEE 300 bus network. The lines' length are uniformly distributed with $l \in [5, 25]km$. For the sake of simplicity we consider the following power loss approximation:

$$\sigma(\{a_i, a_j\}) = k \cdot l(a_i, a_j) \cdot P_{a_i \rightarrow a_j}, \text{ where } a_i \in \mathcal{P}, a_j \in \mathcal{L} \quad (3.2)$$

such that the power loss σ between agents a_i and a_j is proportional to the amount of power exchanged $P_{a_i \rightarrow a_j}$ and the line length l between them, according to the dissipation rate k . An interesting aspect is how the communication between agents occurs. For now we will assume PLC communication through the existing power line communication infrastructure. This means that the communication graph \mathcal{G} has essentially the same structure as the topology. We report results averaged over 100 topologies. Figure 3.1 gives the relative percent gain obtained comparing the performance of our approach to the optimal solution as a function of the dissipation rate k , which yields an average worst-case result of 78%, while attaining a maximum 83% gain.

Secondly, we are concerned with the sensitivity of the coalition formation process with regard to the communication network between agents and provide three different network models. For our experiments we have chosen to set aside from the less realistic

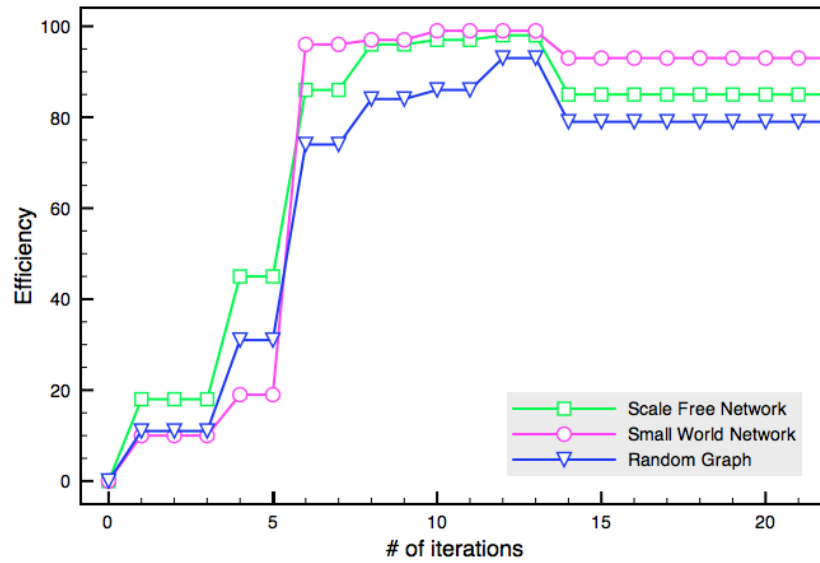


Figure 3.2: Evolution of performance over time of different communication network configurations

scenarios of regular networks and rather consider complex topologies of non-trivial connectivity patterns. As underlying topologies for our MAS we have focused on random graphs, small-world networks and scale-free networks. The choice for these particular configurations was justified by their proliferation in real-world, as well as in agent-based scenarios [19, 22]. In order to perform a relevant comparison between these classes of networks we have considered the same average connectivity per node equal to 4.

For generating the random graph structure we have used the model proposed in [41], where undirected edges are placed randomly between a fixed number of vertices, resulting in a network where each possible edge is present with equal probability. The second class under investigation addresses networks exhibiting the small-world effect. Basically, this means that the mean geodesic distance² scales logarithmically or slower with network size for fixed mean degree. Thus, such graphs are highly clustered and have small geodesic distance. We have implemented the model introduced by Watts

²Shortest path through the network from one vertex to another

and Strogatz [187], which is capable of obtaining small-world networks by applying the least amount of reconnecting transformations to regular graphs. The last category of networks considered are scale-free architectures, which are essentially characterized by the existence in the network of few nodes with many links. According to Barabasi, Albert and Jeong [4], constructing such topologies can be achieved by growing the network by means of preferential attachment to already existing nodes, proportional to their current degrees. It is important to say that this model generates a scale-free network that does not have the small-world property.

With the purpose of studying how the communication structure will affect the behavior of the system, we use as a measure of *efficiency* the percentage of agents that have committed to joining coalitions within a pre-set time frame:

$$f = \frac{|\{a_i | a_i \in \mathcal{A}, s(a_i) = \text{committed}\}|}{|\mathcal{A}|} \quad (3.3)$$

More precisely we simulate a system comprising 100 agents throughout a maximum of 30 rounds of interaction, according to the scheme proposed in Section 3.3. The energetic capacity for the microgrid is assumed at an average of 1MW (small-scale). For these experiments we have considered commercially available residential DERs of 5 capacity classes (in kW) varying uniformly over the set 50, 100, 150, 200 and 250. The results presented have been obtained by averaging over 20 realizations. The system proves to reach a stable organization in a short number of steps by applying the scheme proposed for all the topologies analyzed. However, in term of the efficiency metric abovementioned, there appears to be notable differences.

To begin with, one can observe from Figure 3.2 that the random graph topology is outperformed by both scale-free and small-world networks. In contrast, the small-world configuration is the one that dominates, exhibiting a significant improvement of approximately 15% over random networks and slightly under 10% over large-scale networks. Besides, it is the only complex network that is able to reach an efficiency of above 90%. This is intuitively reasonable according to the intrinsic properties of such networks, which enable rapid spread of information due to its short geodesic distance and high clusterization. Therefore, dynamical processes are facilitated, as

quick communication can be provided between distant parts of the system.

Scale-free networks appear to be an intermediate solution, as they outperform random graphs while being surpassed by small-world networks. In spite of this, it presents a functional advantage as opposed to the other ones if we take into account how the topology affects the robustness and stability of the system. Scale-free networks exhibit resilience to random failures because of the existence of a few number of hubs. On one hand this implies that in the event of a node failure, there is a high probability that the node has a small degree and such, its removal would have a little effect on the functioning of the system at large. Furthermore, while most nodes have a low degree, they happen to lie on few paths that connect other nodes and so, their removal does not affect drastically the overall communication. On the other hand though, deliberate attacks could cause significant vulnerability to the system. Moreover, the experiments indicate that scale-free networks are able of attaining a significantly higher efficiency during the first iteration steps in comparison to the other configurations considered. This is obviously an important aspect in case a good solution needs to be reached in a bounded amount of time.

Thus, results suggest that small-world configurations primarily and scale-free networks represent the best choices for an underlying communication topology for our MAS scenario. As the experiments have shown, the choice for a particular configuration needs though to take into account several trade-offs, according to the prioritization of the system's performance criteria.

Along these lines, we extend our model in the following section with an inter-coalition self-adaptation mechanism that enables for a reorganization of the coalitions in order to assure an enhanced coordination, required by the transient nature of the environment.

3.5 Coalitional mechanism for Microgrid self-adaptation

The mechanism presented hereafter proposes an organizational design for managing the smart grid actors, that operates at the level of the electricity distribution network. Namely, the foremost issue we address in this section regards the notion of

stability that a system of a given random coalition structure is capable to achieve, given the dynamism of the environment. This represents an adaptive process, providing much needed flexibility and functional scalability.

The mechanism proceeds as follows and is synthesized in Algorithm 3. Each coalition designates a coordinator agent (*CA*) which shall be performing the inter-coalition negotiations. There are numerous possible procedures for leader election [171], though in order of placing into focus our mechanism we have chosen a straightforward approach of assigning this role for each coalition to the *PA* with the highest value of β .

We note in advance that the process executes asynchronously and in parallel for all agents. Communication amongst agents assumes the use of time-outs by means of which agents place upper bounds, specifying the amount of time allocated for receiving a reply. In case no reply is received in due time, the particular agent is simply disregarded from being considered as a candidate for coalition reorganization. Here coalition adaptation is achieved in a self-organizing fashion by opportunistic aggregation of agents, while maximizing coalitional benefits by means of taking advantage of local network resources.

As it had been described in Section 3.2, all agents a_i submit on a daily basis their forecasted profile β_i , which typically does not differ exceedingly from their previous one. Nevertheless these cumulative variations might entail a reorganization of the coalition structure *CS* for the following forecasted period in order to assure enhanced coordination at the coalition level. Therefore, consequent to calculating the cost c of the coalition given the existing *LAs* and *PAs* comprising it, it is to be determined the *PA* actors that would qualify to be signed-off, or the profile of the actors that would be eligible to be signed-in to the coalition. Interconnected coalitions should incorporate a control mechanism for achieving a basic energy optimization for the entire system via a close coordination with neighbouring coalitions. Otherwise, the security and stability of the main grid could be threatened severely. The mechanism is in such a way designed that it proves to be consistent with our proposed solution concept introduced hereafter.

Thus, the problem we are facing in open organizations requires a modification of

the coalition structure due to the variations occurring within the system. With these considerations in mind we seek a notion of equilibrium and a corresponding negotiation scheme, which allows for a reorganization of the coalition structure. Furthermore, the solution concept should reflect the decentralization outlook of our scenario, while minimizing the structural adaptations by providing a minimum set of interaction rules in order of attaining the desired stability properties amongst negotiating agents.

Algorithm 3 Adaptation Algorithm. CA behavior.

t

```

if  $a_i \in \mathcal{P}$  &  $\beta(a_i) > \beta(a_j), \forall a_j \in S, i \neq j$  then
  Set  $\text{role}(a_i) = \text{CA}$ 
  \ \ construct objection
   $a_k = \arg \min_{a \in S} \varphi_S(a)$ 
  for each  $T \in \mathcal{N}_S$  do
    | Send ( $\text{Objection}(a_k), T$ )
  end
  \ \ resolve response
   $\text{Extract}(\text{acceptList}, \text{incomingRespList}); \text{Sort}(\text{acceptList})$ 
   $a_k = \text{popUp}(\text{acceptList})$ 
   $S \leftarrow S \setminus \{a_k\}; \text{Send}('ack', T)$ 
  \ \ resolve objection
  for each  $\text{Obj}(a_k) \in \text{Sort}(\text{incomingObjList})$  do
    | if  $\exists$  counterObj to Obj then
      | Send ( $\text{CounterObj}, T$ )
    | end
    | else
      | Send ( $'accept', T$ );  $S \leftarrow S \cup \{a_k\}$ 
    | end
  end
end

```

end

The solution we propose here is inspired by game-theoretic approaches on issues of stability and negotiation. For further considerations on notions of stability, their strength, limitations and interrelations we refer the reader to [126]. In our scenario, of utmost importance is the agents' capability to coordinate and reorganize into groups or coalitions within dynamic environments. Moreover, we advocate for reorienting game-theory to accommodate situations where coordination is a more likely descriptor of the game rather than simply self-interested settings. As it is emphasized in [126], an equivalent formulation for solution concepts can be interpreted in terms of objections and counter-objections.

More formally, given a coalition structure $\mathcal{CS} = \{S_1, S_2, \dots, S_m\}$ we denote by $x = \{c(S_1), c(S_2), \dots, c(S_m)\}$ the associated cost vector. We define our negotiation scheme as follows:

- (S, y) is an objection of coalition S to x against coalition T if S excludes i and $c(S \cup \{i\}) > c(S)$
- coalition T counteracts to the objection of coalition S against accepting player i if $c(T \cup \{i\})/c(T) > 1 + \mu$ or $c(T \cup \{i\}) + c(S) > c(T) + c(S \cup \{i\}) - \tau$.

The objection (S, y) may be interpreted as an argument of coalition S for excluding i resulting in cost vector y where its cost is being decreased. Our solution models situations where such objections cause unstable outcomes only if coalition T to which the objection has been addressed fails to counterobject by asserting that S 's demand is not justified since T 's cost under y by accepting i would be larger than it was under x . Such a response would have hold if we simply presumed players to be self-interested and not mind the social welfare of the system. If on the contrary, players are concerned with the overall efficiency of the system, they would consider accepting the greater sacrifice of y in comparison to x only if this would account for an improvement of S that exceeds the deterioration of T 's performance by at least the margin τ . Thus, τ is the threshold gain required in order for justifying the deviation, whereas μ represents S 's tolerance to suboptimal gains.

Recalling our microgrid scenario, it becomes imperative, as system designers, to

endow the system with the possibility for relaxing standards of their individual performance in the interest of the social welfare. Our proposed mechanisms, thus aims at assessing how and to what extent this may be achieved in order to satisfy the desired system functionalities. We further elaborate on this matter in Section 3.6 based on the experiments performed.

For applying this solution concept to our setting, we additionally need to take into account the underlying topology and thus restrain the inter-coalition negotiation to the given network structure (representing a particularization of the more generic outline presented so far).

Thus, each coalition perceives a local solution with respect to its neighbourhood. Accordingly, from coalition S_i 's local view point at iteration l the local solution is:

$$\mathcal{CS}_i(l) = \{S_{i1}, S_{i2}, \dots, S_{ik}\}, S_{ik} \in \mathcal{N}_{S_i}$$

A potential argument of one coalition would trigger reactions in its vicinity and so, coalitions need to make local adaptive decisions. Therefore, the system is able to determine the most appropriate organizational structure at run-time in the absence of a central controller and in a scalable manner.

When multiple objections are being addressed to one coalition, its decision of considering one would be based on the criteria of maximizing parameter τ , while minimizing parameter μ . The distribution of PA agents amongst microgrids is repeated until there are no neighbour microgrids that would gain a lower cost value in terms of the interaction scheme described. Also worth remembering is that the procedure is ought to occur with the domain dependent constraints, that impose maintaining the profile of the coalition within certain limits (see Section 3.2). Finally, we stress that the aim for our proposed scheme is intended towards an open-ended organizational adaptation concerned with achieving stable configurations in dynamic environments where one-shot optimization procedure are inapplicable.




Figure 3.3: Social welfare of the system in number of steps. a) normal behavior ; b) induced variations

3.6 Experimental results

In this section we focus on emphasizing the results attained during the inter-coalition interaction phase, based upon the solution concept introduced and projected on arbitrary grid configurations. We use hereafter the notion of a coalition interchangeably with the term microgrid. We simulate a system comprising of up to 1000 coalitions deployed in arbitrary topological configurations, which have been generated according to the same considerations discussed in Section 3.4 The energetic capacity for the microgrid is assumed at an average of 6MW (small-scale microgrid). The simulations assume daily variations for the generated energy of each coalition, bounded to an extent of at most 20%. For these experiments we have considered commercially available residential DERs of 3 capacity classes uniformly distributed over the values 15, 20 and 25 (kW). We presume the distribution of DERs in the grid capable of matching the overall consumers' demand. The results presented have been obtained by averaging over 20 realisations (statistically significant for reducing the results' variance below 1%).

To begin with, we first evaluate the performance of our algorithm attained through

the negotiation scheme introduced. Given the cooperative scenario reflected by our chosen solution concept we have set aside from the Pareto optimal instance³ where self-interested agents agree to participate in a trade if and only if the contract increases the agent's immediate utility. This basic type of negotiations alone, outside the smart grid domain, has been proved to reach a local optima, with higher social welfare than others [6].

Alternatively, our chosen scheme for negotiation is primarily aimed at increasing the social welfare of the system and thus, avoiding some of the imbalances that could occur otherwise at the coalition level and which would severely affect the microgrid in our scenario. Such a situation would have corresponded to the undesired case of a microgrid unable to assure an acceptable match of supply with demand. In doing so, the mechanism proposed herein is able to improve on the quality of the local optima reached, while still employing a straightforward self-organizing scheme that avoids an otherwise exponential lookahead. The negotiation is based on the actor's local perspective, not assuming the configuration of the other coalitions to be known. The experiments performed reveal that the procedure leads to a local optimum rapidly, to a higher average social welfare and even more importantly, decreases the occurrence of coalitions far from equilibrium.

Figure 3.3 points out the average percent increase in social welfare, that the system manages to attain from an initial state to a stable one, achieved during the course of the adaptation phase, against the optimal allocation of DERs in the system. The system proves to reach a stable organization in approximately 50 steps by agreeing on reassigning DERs between coalitions according to the scheme proposed. As the graph in Figure 3.3a illustrates, a stable configuration of the system is abruptly reached, meaning that the agreements realized earlier improve the social welfare more than the ones performed later. Furthermore, the solution applied is an anytime algorithm that achieves a monotonic improvement of the global (social) welfare of the system, which can thus only improve at each time step. This is obviously an important aspect when the best solution needs to be reached in a bounded amount of time. Hence we comply

³Represented in the graphs as the individualistic approach

with our objective of converging abruptly to an efficient and stable configuration of the system.

Transferring agents individually between coalitions as opposed to bundles of actors, although more time consuming, avoids a known outcome, that of concentrating the agents among only a few coalitions [6]. Single transfers have been shown to have the tendency of diffusing the spread of actors into more evenly configurations. This is clearly a desired state for our organization structure. Another important aspect achieved as a result of this is as well, performing a minimization of the structural adaptation required.

Moreover, the organizational model proposed proves to be able to operate in open environments and dynamically stabilize the system while actors are being added or removed to the system. In fact its adaptable features leverage on the inherent variations in the system permitting it to escape local optima. In Figure 3.3b we plot a less usual instance, where the system undergoes considerable variations during the adaptation phase, as some of the actors of the coalitions are removed (possibly due to failures) and some have joined the system. We consider a rather extreme situation where the proportion of coalitions that experience such modifications is considered at 30%, while within each coalition up to 10% of the actors have been disconnected or alternatively, have been appended to the system. The system demonstrates a capability to reorganize and reach a stationary configuration as the spikes injected into the system are flattened in a small number of steps.

We have evaluated the performance of a coalition in terms of the incurred cost. An analogous way to think about the performance of a coalition is in terms of its *relative utility*, which we define in Equation 3.4. In other words, the utility of coalition S represents the percentual performance of S with regard to cost, relative to the best and worst performing coalitions in \mathcal{CS} .

$$u(S) = \left(1 - \frac{c(S)}{\max_{S_j \in \mathcal{CS}} c(S_j) - \min_{S_i \in \mathcal{CS}} c(S_i)} \right) \cdot 100 \quad (3.4)$$

Subsequently we perform a series of experiments to draw more insight to the solution concept introduced. On one hand, our negotiation scheme implies that deviations

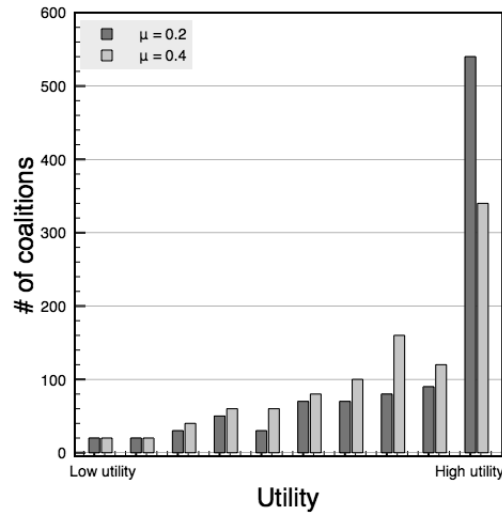


Figure 3.4: Histogram representation for the utilities of the coalitions

would only occur if a certain minimum gain τ has been achieved. On the other hand, the extent to which a coalition is willing to decrease its efficiency in detriment of the gain in social welfare is represented by a satisfactory parameter μ . This represents in effect a percentage, which defines what an acceptable performance would be and how tolerant is one coalition towards suboptimal performance. For our simulations we chose an initial value of 0.4 and considered a homogeneous population of actors in the system. Although this does not make the objective of our scenario, heterogeneity amongst the actors involved may as well be introduced.

Following, we analyzed the implications of the dependency on this parameter for a better understanding of its functionality. Thus, we have analyzed the stationary states the system falls into as a function of μ . For large values of μ , meaning that coalitions are willing to significantly decrease their utility with respect to the improvement of the global welfare of the system, we encountered an expected global increase in utility, but a considerable variation in the allocation of utilities in the system. Instead, when only lower decrease in performance is accepted by each coalition the results obtained are plotted in Figure 3.4. The diagrams of Figure 3.4 illustrate a histogram representation of the coalitions' utilities discretised in increasing order of their worth. It can be seen that a 20% increase of μ reduced significantly the number of coalitions

operating at high utility denoted by the first column of the histogram, while the number of coalitions operating at lower levels of utility has been increased. The results emphasized that the best performance⁴ of the system, in terms not only of social welfare but also relative to the distribution of utilities over \mathcal{CS} , was obtained for values of μ in the vicinity of 0.2 . Somewhat surprisingly, what the experiments show is that being willing to accept lower efficiencies in the benefit of the global performance is only advantageous to a certain extent. In actuality, there is a trade-off to be taken into account. Although the overall system utility increases, the ratio between the number of coalitions with low utility and those with high utility is increasing as well. So, for assessing the performance of the system not only should we be interested in the global utility, but also in having a uniform distribution of high utilities for the majority of the coalitions.

3.7 Discussion

In the following we give a brief review on related approaches, positioning our mechanism in the context of teamwork in agent organizations and more precisely, along the problem of structuring a set of individuals as a team of cooperative agents that pursue an institutional goal.

An important body of work has been devoted to the question of how to best partition a group of agents (in non-superadditive domains), which is essentially a combinatorial problem with an exponential search space [160]. The proposed solutions can be analysed according to different attributes of the solution method, such as optimality, centralisation and dynamism.

A first class of algorithms are those that are run centrally by an omniscient agent that tries to find an optimal solution, or at least a solution that is bounded from

⁴In our acceptance of best performance we restrict the results to a number of threshold values. In terms of utility distribution, namely we would like the number of coalitions pertained by the highest utility class (last column of Figure 3.4) to represent a minimum of 50% of all coalitions, while the remaining classes to be below the limit of 10%. Moreover what we have achieved is to maintain the inferior fraction of lower utility classes, each below 5% of the total number of coalitions. In terms of average percent increase in social welfare we impose an 80% improvement.

optimal. As a subclass, a common practice is that of employing dynamic programming solutions [153]. The complexity of such algorithms, although significantly better than an exhaustive enumeration of all coalition structures, is prohibitive in the number of agents, being usually suitable for situations of at most 20 agents. The second subclass of this category of algorithms is built upon interpreting the problem in the sense of a coalition structure graph, introduced by Sandholm in [155]. Extensions of this approach consider different pruning techniques in order to establish solutions bounded from optimal, such as the one proposed in [31]. For instance, using this algorithm one may compute a faster solution when smaller bounds are desirable. Still, this type of algorithms remain severely prohibitive in terms of scalability. Same is the case for the state-of-the-art algorithm [145], which divides the search space into partitions based on integer partition and performs branch and bound search.

The second class of algorithms are oriented towards providing solutions rapidly, however not guaranteeing optimality, nor worst-case bounds for the solution. This type of algorithms is known to be scalable and more applicable to real-world scenarios. Amongst them we mention several notable efforts employing, genetic algorithms [160], swarm optimization [195] or constraint satisfaction techniques [16]. However, the limitation of these algorithms lies in that they represent a centralized, one-shot optimization procedure.

Thirdly, we identify the class of decentralized and dynamic methods. Here, quite the opposite, considering a multi-agent environment, it is desirable that computing the solution could be achieved in a decentralized manner, based on the agents' local utility computations, that seek to find feasible coalition structures through series of negotiations. Decentralized solutions are especially suitable for scenarios where coalitions have to adapt their structure due to the dynamic nature of the environment. Along these lines, Klusch et al. introduced in [83] a distributed and completely decentralized process for coalition formation, able of operating in open systems. Another relevant instance of such algorithms is the one proposed by Apt et al in [8], though limiting, in the sense of allowing transformations only to make use of simple split and merge rules. A satisfying coalition formation algorithms is proposed in [165]. Here, agents have an incomplete view of the world, time and computational constraints,

the coalition formation goal being one of meeting minimum requirements rather than achieving maximum performance. The algorithms aim to cope with the dynamism of the system by engaging in opportunistic negotiations and of high-risk, while the success of the formation of coalitions cannot be guaranteed.

In this chapter, we have modelled the domain-dependent constraints posed by the electricity domain and propose a coordination mechanism that falls in the third category of coalition formation solutions designed for dynamic environments.

As a proof of concept, our work has introduced a dynamic coalition-based model deployed in distributed environments of negotiating agents. The adaptation mechanism introduced performs an open-ended adaptation of groups of organizational agents, converging towards stable configurations. In particular, we have highlighted the applicability of this approach through the design of a distributed adaptive scheme for the problem of microgrid configuration. This process resulted in partitions of the grid that would be able to commit to a steady and robust generation profile requiring less or no energy from traditional power plants.

In terms of a full-fledged operational deployment of this solution in a real setting, complementary techniques that enable to incorporate the electrical features of the power system in a more factual form, such as load-flow computation analyses, that verifies for contingencies and maintain the system within its operational limits, need to be considered. Also, taking into account that at present the proliferation of DERs in the grid is still yet to achieve an adequate level, an interesting future line of work would be to address, at a more granular level, the most suitable techniques for efficiently deploying such devices throughout the topology of the grid.

Chapter 4

Coordination Mechanisms for *Virtual Power Plants*

It is the long history of humankind (and animal kind, too), those who learned to collaborate and improvise most effectively have prevailed.

— *Charles Darwin* —

In this Chapter we tackle the implications of an increasing embedded generation in the grid. We start off by addressing the organizational models that may emerge from an ecosystem comprised of a mixed bag of distributed energy generation facilities, in Section 4.1. We evaluate the economic viability of VPPs, which take the form of energy aggregators bringing together a portfolio of smaller generators and operating them as a unified and flexible resource on the day-ahead energy market or selling their power as system reserve. Secondly, in Section 4.2 we look closely into mechanisms capable to couple and co-optimize energy resources in real-time in order to accommodate a given load shape, facilitating the orderly integration of intermittent renewable resources into the grid's operations.

Variable Sect. 4.1	Description	Variable Sect. 4.2	Description
\mathcal{A}	action space of each agent	A	set of agents
s_1, s_2	agent strategies	V, R	set of decision and random variables
$\mathcal{P}_1, \mathcal{P}_1$	probabilities for strategy selection	D, Δ	set of domains decision, random var
N	set of agents	F	set of cost functions
n	number agents	α	objective function
\mathcal{N}_{a_i}	set of neighbours for agent a_i	T	set of time-slots
θ	stochastic variable of the system	P_{loss}	energy loss
\mathcal{R}	agent's payoff/reward	E_{CO_2}	carbon emissions
p	energy unit price	C_g	production cost
α	learning rate	β	penalty factor
$D(\mu, \sigma)$	normal distribution	L_j	demand a time-slot j
\mathcal{G}	communication graph	e	evaluation function
\mathcal{V}	set of vertices	C	set of soft constraints
\mathcal{E}	set of edges	n	number of decision variables
L_i	denotes the average regret of a_i	m	number of random variables

4.1 An Investigation of *Cooperation* in Smart Grids

In this Section we study the phenomenon of evolution of cooperation in the electricity domain, where self-interested agents representing distributed energy resources (DERs) strategize for maximizing payoff. From the system's viewpoint cooperation represents a solution capable to cope with the increasing complexity, generated by the introduction of DERs to the grid. The problem domain is modelled from a multi-agent system high-level perspective. We report on experiments with this model, giving the underlying understanding for the emergent behavior, in order to determine if and under what conditions such a collaborative behavior would hold. Finally we suggest how insights from this model can inspire mechanisms to instill cooperation as the dominant strategy.

Conventional methods for energy generation, transmission and distribution are about to experience a radical change. This is on one hand due to the ever increasing demand in energy consumption (e.g. electric vehicles) and secondly, due to the proliferation of distributed generators (e.g. renewable energy) to be connected to the grid. Current power networks will no longer be able to provide the required level of reliability and robustness and thus there is a need for a more flexible connection and management of the system.

Various approaches for the advent of agent technologies to this domain have been proposed thus far, that range from micro-grid architectures [61, 104], demand-side management [57, 174] and micro-storage solutions [184] to plug-in hybrid vehicles coordination [77]. The benefits of applying the multi-agent systems paradigm as an approach for distributed control of the Grid entails primarily: autonomy, scalability, flexibility, extensibility, fault tolerance and reduced maintenance [104]. The actors existing in the grid (i.e. consumer loads, distributed generators) represent different owners with particular user behaviors, hence deploying an agent-based distributed control over the system becomes highly suitable for this scenario. Moreover, decentralization increases the system's reliability in case of failures, enables local adaptability to dynamic situations at runtime and allows coordination as opposed to the more complex task of centralised management.

The outline for the rest of this Section is as follows. Section 4.1.1 describes the agent-based framework used for modelling the problem domain. Then, in Section 4.1.2 we discuss implementation details and outline a series of experiments that exhibit the phenomenon of emergent cooperation.

4.1.1 Agent-based Model

Given that the penetration of distributed energy resources is expected to increase significantly [122], integrating these devices to the grid poses difficult challenges. As a solution for reducing the complexity of managing the system at large, the aggregation of DERs as virtual power plants has been proposed. A VPP is conceived as a bundle of distributed energy generators that are connected through an informational infrastructure and act in a coordinated way as a single entity, being represented as one resource to the system operator. Generally, they represent renewable energy resources such as wind or solar power, which accounts for high variability depending on environmental conditions. We conceive an economic encounter where we synthesize the main decision factors that influence the way DER devices are operated in relation to today's market structures and regulations in place. Our goal is to draw an understanding about what an efficient administration of resources means under different regimes of the system which are detailed hereafter.

Now, considering the distributed nature of DERs and their selfishly driven behavior it comes natural representing them as autonomous agents interacting in an open and highly dynamic environment. Also, we consider agents to be controlling identical DERs in terms of their capacity profile. Agents are thus interested in maximizing the payoff obtained by selling their available energy. There are several day-ahead power markets where agents may choose to bid: (i) *Baseload Power Market*: typically this power is currently provided by large-scale power plants round-the-clock and at low costs per kWh; (ii) *Peak Power Market*: this represents the additional power necessary during high-demand intervals of time; (iii) *Spinning Reserves Market*: designed for ensuring the reliability of the grid in case of transmission line failures or similar contingencies, in practice they are rarely used but are being paid for the duration

they are available; (iv) *Regulation Markets*: required in order to regulate the frequency and voltage in the grid, must be capable to respond to frequent real-time signals from the grid operator.

Taking into account the profile of DERs, in our model we consider for the agents the options of participating in either the *Baseload* or *Peak Power Market*, reflecting two opposing strategies, which denote exposure to risk. There is clearly an uncertainty regarding energy availability for the following forecasted day. On one hand selecting to bid in the *Baseload Market* ensures a lower profit (lower kWh rates) regardless of the amount of energy provided, whilst the *Peak Power Market* would guarantee higher profits (during high-demand intervals) given sufficient energy availability. Thus, the former option represents risk-aversion, as the latter denotes a risk-seeking behavior.

Secondly, agents need to decide whether they prefer cooperation, which takes the form of a VPP, where agents reallocate energy in order to mitigate day-ahead predictions and fulfil bids such that the payoff of the VPP is maximized or rather, choose a non-cooperative behavior, where the agent's payoff depends solely on its own choice with regard to risk.

Therefore, essentially we modelled a game where the action space for each agent is a two dimensional binary state space $\mathcal{A} = s_1 \times s_2$, characterizing willingness to cooperate as the strategy set $s_1 = \{cooperate(1), defect(-1)\}$ and aversion to risk by $s_2 = \{risk-seeking(1), risk-averse(-1)\}$. Moreover, the stochastic and dynamic nature of the system is due to the uncertainty of available energy and the varying number of agents participating in the system at each iteration of the game, respectively. Specifically, the n -player game proceeds in rounds, each consisting of three phases:

- (i) the *game playing phase*, where agents commit to a particular choice of strategies
 - s_1 : bid for the *Baseload* or *Peak Power Market*
 - s_2 : participate in VPP or act alone in the market
- (ii) *output of the game*, where payoffs are computed for each agent based on their particular choice of strategies and the actual amount of energy available (θ)

- (iii) the *strategy update phase*, where based on the result agents may change their choices for the next round

Now, assuming our previous considerations about participation in different energy markets, we represent the agents' payoff as a function of the energy that the agents are able to provide, which is in turn the stochastic variable of the system θ . The availability of the energy provide through renewable resources fluctuates randomly in time and is denoted here by θ . Recall that for simplicity we assume in our model all agents to be controlling identical DER devices, meaning all agents will have the same amount of energy available. Specifically, by adopting a *risk-averse strategy* the energy unit price p (referred in terms of kWh rates) is moderate, however the payoff is received regardless of the amount of energy available: $\mathcal{R} = p \cdot \theta$. On the contrary, a risk-seeking strategy would return a considerably higher payoff in case of high energy availability but, a significantly lower one otherwise.

$$\mathcal{R} = \begin{cases} 2 \cdot p \cdot \theta & \text{if } \theta \geq \varkappa \\ 0 & \text{if } \theta < \varkappa \end{cases}$$

The payoff functions for the two strategies are defined to approximate current pricing conditions in the renewable energy sector, representing a close-to-flat rate for the former strategy, as for the latter generating twice the profit for large values of θ and close to zero otherwise.

Moreover the choice over strategy s_1 influences the agents' payoff in the following way. If they choose *not* to cooperate their payoff depends solely on their choice over strategy s_2 . If they select to cooperate as VPP, all agents adopting the cooperative strategy will pull together the total amount of energy available and redistribute it in a way to cover the most profitable bids that the agents in the VPP have committed to.

Choices are modelled probabilistically, each agent is characterised at each time-step by two probabilities. Let \mathcal{P}_1 represent for each agent the probability for choosing cooperation, whilst the complementary probability corresponds to defecting. Similarly, we associate \mathcal{P}_2 with the probability of risk-seeking behavior and $1 - \mathcal{P}_2$ with risk-aversion. The update rule computes new probabilities for the next iteration for

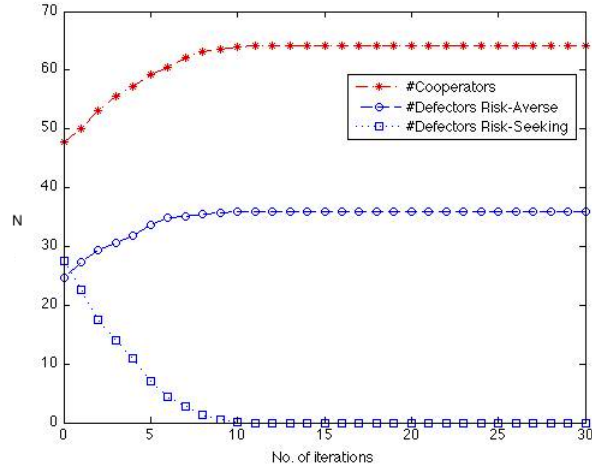


Figure 4.1: The evolution of the number of agents with respect to their strategy space for θ with $\mu=5$.

each agent a_i of the set of all agents N based upon previous round results, by applying a *satisfying gradient ascent rule*:

$$\mathcal{P}_i^{t+1} \leftarrow \mathcal{P}_i^t + \alpha \cdot s_i^t \cdot (\mathcal{R}_t(\bar{\mathcal{A}}_{a_i}) / \operatorname{argmax}_{a_j \in \mathcal{N}} \mathcal{R}_t(\bar{\mathcal{A}}_{a_j}) - \varepsilon)$$

where \mathcal{P}_i^t is the current probability associated with strategy s_i ; \mathcal{P}_i^{t+1} is the probability associated with strategy s_i for the following iteration; α represents the learning rate; $\mathcal{R}_t(\bar{\mathcal{A}}_{a_j})$ is the payoff of agent a_j at iteration t corresponding to its chosen strategy set $\bar{\mathcal{A}}_{a_j}$. Intuitively, the probability of a particular strategy is updated according to the percent payoff difference between the agent's current strategy and the best performing strategy in the system, minus the margin acceptability ε . Parameter ε represents the trade-off between the payoff difference and the probability to switch.

4.1.2 Simulation Results

Complete Information Games

The purpose of our experiments is analysing the conditions under which there is evolution of cooperation for our domain specific setting and understand whether the




Figure 4.2: Payoffs as a function of variable energy availability θ and number of agents in the system.

system-level goals and those of our selfishly driven agent population align or on the contrary, they are conflicting.

We start under the assumption of complete information, where all aspects of the game are considered common knowledge for all players and the payoffs of the other agents are directly observable. We simulated a system consisting of a population of at most 100 agents. The initial probabilities, based on which agents are determining their strategies are allocated randomly with uniform probability, while the learning rate is set to $\alpha = 0.01$ and $\varepsilon = 0.5$. The variability of the environment is represented as the stochastic value θ , which determines the amount of energy available for each agent. At each iteration, the value for $\theta(t)$ is generated randomly from a normal distribution $D(\mu, \sigma)$. The mean value μ is kept fixed for each realization of the game, while $\sigma = 1$. The simulation then proceeds for each round according with the three phases previously described in Section 4.1.1.

What the experiments show is that for each mean value of θ the system converges to stable configurations, in terms of the ratios with which strategies are selected. Convergence is reached in approximately 20 rounds. Figure 4.1 highlights for a population of 100 agents and a particular mean value of θ the fraction of agents that have selected to cooperate and for those that have defected, their strategy regarding risk. We take this last stationary state of the system as the final result. Following, we plot this data in Figure 4.2, representing agents' payoff as a function of θ , by averaging each data point over 1000 runs.

It appears that a cooperative strategy yields the highest payoff for the agents in a situation where neither a risk-seeking nor a risk-averse strategy is clearly dominant. Thus, in case of uncertainty cooperation proves to be the optimal choice, accounting for a minimization of risk. However, when there is a higher level of certainty with regard to strategy s_2 for selecting a suitable energy market, defecting outperforms cooperation. The underlying reason for this counter-intuitive result is the fact that the reallocation of energy between cooperative agents produces also better results for the agents with an incorrect choice for strategy s_2 , than would defecting. This misguided feedback is causing them to react suboptimal, and thus, in detriment of the cooperative VPP at large. In comparison defecting agents show better adaptability.

Moreover, we take into account a scenario where the number of agents is varying. In Figure 4.2 the y -axis shows how the number of agents is influencing the outcome of the game. Consistent with the abovementioned explanation, in conditions of higher certainty about risk, a lesser number of cooperating agents are capable to better adapt their strategies, rather than a larger coalition. Therefore we can conclude that the limitation of the number of cooperating agents proves here to represent a solution for promoting cooperation as a dominant strategy. This would allow for members of a coalition to better respond to variations and adapt their respective strategies, while maintaining a dynamic equilibrium with the environmental changes.

Limited Information Environments

In this section we set to investigate the emergence of cooperation for games played with limited information, wherein the global knowledge assumption of observable payoffs for all players (as opposed to complete information games, see Section 4.1.2) does no longer hold.

Thus, we consider the agents representing DERs deployed over a spatial distribution and restrict observability to their vicinity. Formally, agents are connected via an underlying network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, denoted as an undirected graph, where the set of vertices \mathcal{V} is the set of agents and edges \mathcal{E} are the set of links in \mathcal{G} , which connect the agents in \mathcal{V} . An agent a_i may only directly observe agent a_j 's payoff if they are neighbours. \mathcal{N}_{a_i} represents the set of a_i 's neighbours, i.e. $\mathcal{N}_{a_i} = \{a_j | (a_i, a_j) \in \mathcal{E}\} \subset \mathcal{V}$.

For generating the random graph structure we have used the model proposed in [41], where undirected edges are placed randomly between a fixed number of vertices, resulting in a network where each possible edge is present with equal probability p . Similarly to our previous experimental setting, we consider a population of $n = 100$ agents. We set $p = \log(n)/n$ to ensure the graph connectedness and an average connectivity per node equal to 4, corresponding to topological configurations for generic meshed suburban network models¹.

Strategy Selection and Convergence.

The strategy selection rule determines which strategy to play from the agent's strategy space. Considering the given limited information scenario, we use the following types of strategy selection rules²:

To maintain normalization, these two probabilities will be restrained to the interval $[0; 1]$, meaning that they will be set to 0 or 1 when they grow beyond these limits.

- *The gradient ascent strategy rule* performs an identical estimation to the one

¹Identified as the most suitable setting for the deployment of medium-scale VPPs, thus the particular spatial distribution considered.

²Note that probabilities are restrained to the interval $[0, 1]$ by setting their value to 0 or 1 accordingly in case they exceed this bound.

detailed in Section 4.1.1, revising its strategy selection probabilities \mathcal{P}_i according to the reward gradient, with one important difference. Namely, each agent can only perceive the local highest payoff in its vicinity, as opposed to the global highest payoff in the system:

$$\mathcal{P}_i^{t+1} \leftarrow \mathcal{P}_i^t + \alpha \cdot s_i^t \cdot (\mathcal{R}_t(\bar{\mathcal{A}}_{a_i}) / \operatorname{argmax}_{a_j \in \mathcal{N}_{a_i}} \mathcal{R}_t(\bar{\mathcal{A}}_{a_j}) - \varepsilon)$$

- *Win-stay, lose-shift rule* [120] maintains the current strategy selection probabilities only if the current payoff is at least as high as in the previous iteration round. Otherwise, revises \mathcal{P}_i proportional to the difference of its current and last payoff:

$$\Delta \mathcal{R}_t = \mathcal{R}_t - \mathcal{R}_{t-1}; \mathcal{P}_i^{t+1} \leftarrow \mathcal{P}_i^t + \alpha \cdot s_i^t \cdot \Delta \mathcal{R}_t$$

This approach is highly suitable for limited information environments as it only requires keeping track of short-term previous payoffs.

- *Imitate best strategy rule* [7] identifies the agent with the highest payoff in its neighbourhood and adopts the same probabilities for strategy selection.

$$\mathcal{P}_i^{t+1}(a_i) \leftarrow \mathcal{P}_i^t(a_j), \text{ where } \mathcal{R}_{t-1}(\bar{\mathcal{A}}_{a_j}) = \operatorname{argmax}_{a_k \in \mathcal{N}_{a_i}} \mathcal{R}_{t-1}(\bar{\mathcal{A}}_{a_k})$$

- *Regret minimization strategy rule* [157] consists of playing the strategy that minimizes the average overall regret.

$$L_i^T(k) = \frac{1}{T} \sum_{t=1}^T (\operatorname{argmax}_{a_j \in \mathcal{N}_{a_i}} \mathcal{R}_t(\bar{\mathcal{A}}_{a_j}) - \mathcal{R}_t(\bar{\mathcal{A}}_{a_i}))$$

where $L_i^T(k)$ denotes the average overall regret of agent a_i for strategy k over the previous T iterations.

Evaluation.

In the previous section we ran our simulation under the assumption of global system knowledge. However, in a practical scenario, it is more likely that agents are capable to acquire only partial knowledge of the system. With these considerations in

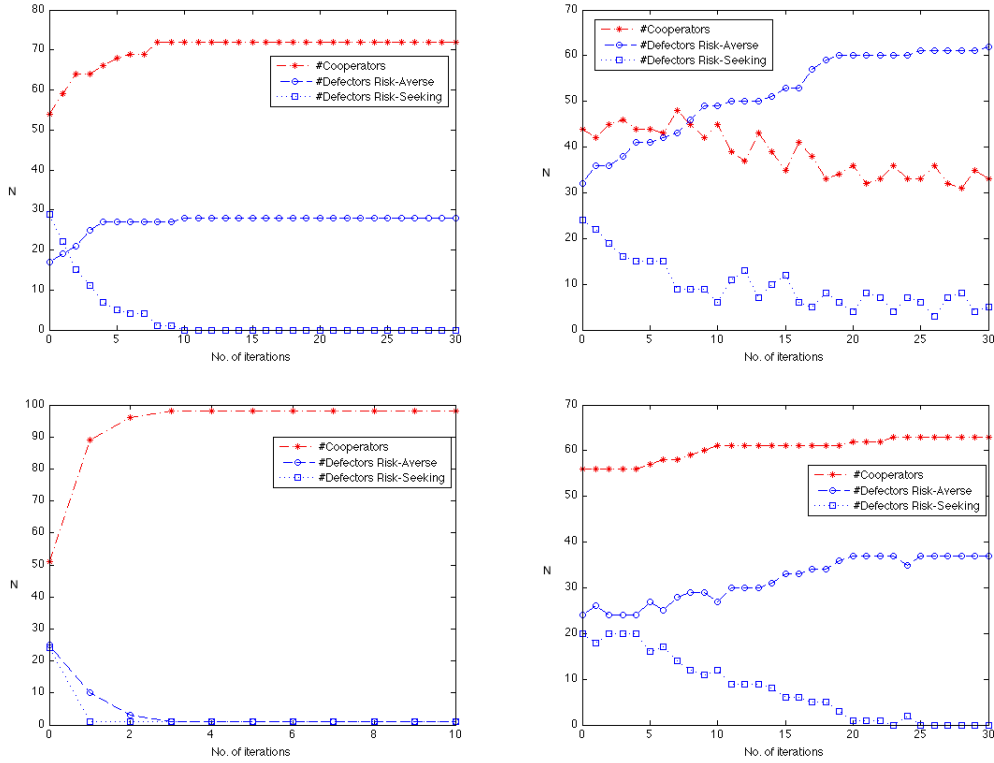


Figure 4.3: The evolution of the agents behavior with respect to the strategy selection rule for θ with $\mu=9$. a) *gradient ascent* ; b) *win-stay, lose-shift* ; c) *imitate best strategy* ; d) *regret minimization* ;

mind we deploy agents over a spatial distribution and investigate to what extent this may affect results. Moreover we derived several strategy selection rules for limited information environments and performed a comparative analysis. The value used for θ was chosen to reflect highly uncertain scenarios where cooperation is ought to emerge as the dominant strategy.

It is interesting to observe from the plots in Figure 4.3 that the system converges to similar results, obtained for games with minimal information. Particularly, the *gradient ascent* strategy attains just about the same stationary states in terms of the fraction of cooperating agents via local estimations only, though convergence occurs after more iterations of the game.

For the *win-stay, lose-shift* rule we find a lower fraction of cooperators, denoting a

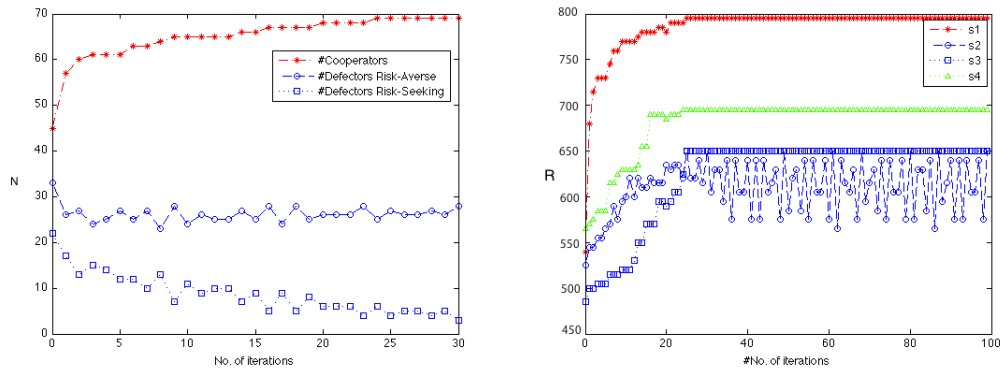


Figure 4.4: a) Population share in mixed-strategy repeated games b) Average payoffs in repeated games

lower efficiency for adopting the optimal strategy. Here strategy selection solely relies on the agent’s own past results. This enables applying it to settings where payoffs of other agents are unobservable, while still producing satisfactory outcomes. We found that the *win-stay, lose-shift rule* was outperformed by the *gradient ascent* strategy by a margin of approximately 20%.

In contrast, the *imitate best strategy rule* proves a rapid convergence to optimal for the majority of the agent population. However, note that the success of this rule depends on the agents’ capability of perceiving the internal states of neighbouring agents for copying their probabilities of strategy selection, which may not always hold as a reasonable assumption.

Finally, the *regret minimization strategy rule* shows a foreseeable result. Given the highly uncertain conditions chosen for this experiment, the number of agents that adopt a *risk-seeking* behavior is dramatically reduced, as *risk-averse defectors* and *cooperators* become majoritarian.

Also, it is interesting to analyse the impact on performance considering a heterogeneous population of agents. We conducted experiments for mixed-strategy repeated games, by having each of the previously detailed strategy selection rules equally represented in the population. Similarly, we evaluated the percentage of the population that converged to optimal behavior for high uncertainty conditions. Figure 4.4 a) shows that a heterogeneous population of agents achieves almost about the same

level of cooperation for limited information domains as opposed to complete information scenarios. Additionally, it is interesting to see which strategies have been the top performers of the game in terms of the payoff obtained against the different types of participating agents. As depicted in Figure 4.4 b), the highest average payoff after 50 rounds of the game was achieved by the *imitate best strategy rule*, where all agents employing it converged to the highest payoff possible (which was also the case for the self play game). The *win-stay, lose-shift rule* population share returned the second highest payoff, while the *gradient ascent* and *regret minimization strategy rules* were respectively significantly outperformed. These results suggest that even in limited information settings agents can learn optimal behavior under the more strong assumption of local complete information, while moderate performance can still be achieved relying solely on the agent's own past results.

In the following section we address closely the inner mechanisms of operating VPPs, in order to relax some of the underlying assumptions. In particular we intend to experiment with a heterogeneous population of agents, in the sense of the type of DERs and their output capacity in order to reflect more accurately realistic scenarios.

4.2 Stochastic Optimization for Virtual Power Plants

A wide variety of renewable distributed energy technologies are now commercially available. However, due to their small-scale capacity and intermittent nature, integrating such devices to the grid may degrade the security and reliability of the distribution systems. Virtual power plants representing clusters of distributed generators may alleviate these drawbacks by aggregating a reliable energy supply. In this section, we introduce a methodology for efficient demand-response in the context of hybrid power generation systems that combine different sources of energy. Specifically, the planning problem is formulated as a *stochastic* DCOP for determining the optimal dispatch. We then go on to propose a new distributed algorithm for solving the scheduling problem and provide an empirical analysis of our approach for a *smart grid* scenario.

The vision of the *smart grid* aims to ensure an increased amount of energy supply via clean energy generators [123]. However, this type of devices pose serious challenges to the efficiency and reliability of the grid due to their volatile nature and dependence on external factors (e.g. wind power, solar power, tidal power, plug-in hybrid vehicles, etc.). Moreover, renewable generators are typically small-scale, heterogeneous and distributed within the electricity grid, which makes the coordination problem significantly more difficult. Along this line, the concept of VPP has been introduced as a means to aggregate various distributed energy generators (DERs). A VPP is defined as a bundle of DERs that are connected through an informational infrastructure and act in a coordinated way as a single virtual entity [36, 143, 109]. Multi-agent systems have been regarded as natural approach for VPP coordination and various attempts have been made to provide an efficient coordination scheme. Centralized structures for managing VPPs are proposed in [137, 143]. The shortcomings of these approaches are essentially twofold. Firstly, the organizational structure proposed fails to address an open system setting, where DERs may be dynamically appended to the system, or (temporarily) removed. Secondly, the inherent stochastic nature of DERs is not explicitly factored into the model and thus, the extent to which the system can withstand disturbances in terms of power generation remains

unknown. In this work we are particularly concerned with the need for a more flexible coordination model, where the VPP can readily adapt to different operational modes while being capable to cope with and reason about the uncertainty in the system.

In [24] the authors propose a novel pricing mechanism, that addresses the question of allocating payoff amongst the VPP members, so that it can guarantee that no subset of agents has an incentive to break away. Similarly, [151] introduces a payment mechanism based on proper scoring rules that incentivise agents to report private probabilistic predictions truthfully. Although complementary to our work, pricing mechanisms are out of scope here. Here, our goal is restricted to providing a solution for generating a schedule for a given hybrid power generation system under uncertainty conditions.

The contribution of this Section is threefold: *i)* Firstly, we provide a new representation of the VPP scheduling problem by formalizing it in the context of stochastic DCOP; *ii)* Secondly, we propose a distributed algorithm for solving this optimization problem by means of exploiting domain-dependent characteristics of our model; *iii)* Thirdly, we present an empirical evaluation of our proposed approach and highlight its superior performance in terms of computational time. Although the deployment of DERs is expected to increase considerably in the following years, they are currently overlooked in production planning by the distribution network operator for being categorized as unreliable. Therefore, providing a mechanism that can ensure an efficient, large-scale coordination of such devices is of crucial importance for replacing conventional, carbon-intensive power stations.

The remainder of Chapter 4.2 is organized as follows. In Section 4.2.1 we present a DCOP formalization for a hybrid power generation system. Section 4.2.2 extends our formalism to incorporate the inherent stochasticity of the problem to our representation. Section 4.2.3 dwells on the usability of existing solutions for solving the *StochDCOP* and proposes a the new resolution, together with a theoretical analysis of our algorithm. An empirical evaluation of our model is given in Section 4.2.4.

4.2.1 Problem definition & Model

Distributed Constraint Optimization (DCOP) has been proposed in modeling a wide variety of multiagent coordination problems [106]. In this section we use this technique to investigate the problem of coordinating generation from intermittent resources, constituted in the form of a VPP. Thus, given a demand profile that needs to be satisfied, distributed energy resources (DERs) must coordinate their power output selection so that a global objective function is optimized. The global objective function arises as a distributed reasoning problem over a set of constraints, where each DER knows about the constraint in which its output is involved.

Formally, a DCOP is defined as $\langle A, V, D, F, \alpha, \sigma \rangle$ with

- $A = \{a_1, a_2, \dots, a_m\}$ a set of agents
- $V = \{x_1, x_2, \dots, x_n\}$ a set of variables
- $D = \{D_1, D_2, \dots, D_n\}$ a set of domains, variable x_i taking values in D_i
- F a set of cost functions $f_i : D_{i1} \times \dots \times D_{ij} \rightarrow \mathbb{N}_0 \cup \infty$
- α an objective function, defines the aggregation of F
- $\sigma : V \rightarrow A$ a distribution function of variables to agents

The goal is to find an assignment to all variables that minimizes the result of the α function. As we have introduced in Section 2.3.2, it is often the case that α is defined as the summation over the set of cost functions F , which we will continue to assume henceforth.

The DCOP representation enables us to capture the *plug-and-play* characteristics of DERs, which may recurrently switch between online and offline modes, according to exogenous operating conditions. Notably, our model ensures adaptiveness to this dynamic setting by allowing DER units to be placed (or appended) seamlessly at any point within the grid, without re-engineering the control logic of the system and thus preserving its *plug-and-play* functionality. Hence, by means of the decomposability property of the DCOP formalism we can guarantee that DERs could optimize their




Figure 4.5: DCOP representation for a VPP scenario

joint performance in a flexible manner. In our representation we assume each agent represents a DER. Consequently, each agent is assigned a single variable denoting its power output, which it must select from the discrete set of possible outputs, corresponding to the domain of the device. Because of this assumption, the terms *agent* and *variable* will be used interchangeably. Note, in later research we can extend our model towards complex decision problems within each DER, by formalizing this as a set of local constraints. This can be captured with the used model, and will result in a DCOP with complex local problems. Only the agent who is assigned a variable has control of its value and knowledge of its domain.

As depicted in Figure 4.5, the representation of our DCOP corresponds to a hypergraph. A hypergraph is a generalization of a graph, where an edge can connect any number of vertices. Formally, a hypergraph H is a pair $H = (X, E)$, where X is a set of elements, called nodes or vertices, and E is a set of non-empty subsets of X called hyperedges or links. Therefore, E is a subset of $\mathcal{P}(X) \setminus \{0\}$, where $\mathcal{P}(X)$ is the power set of X . Here, vertices correspond to agents and edges to cost functions. Then, $x_{i1}, x_{i2}, \dots, x_{ik}$ are neighbors if there is a cost function that binds them.

Figure 4.5 gives a simple example of a constraint hypergraph for our application domain. We consider a VPP consisting of n distributed generators $G = \{g_1, g_2, \dots, g_n\}$, where each generator g_i controls its power output variable v_i via an agent a_i . The discrete domain of power output values, where variable v_i can take values, is given by the set $D_i = \{d_{i1}, d_{i2}, \dots, d_{ik}\}$. The self-loop edges of the graph, denoted as unary constraints f_i defines the costs for generating the selected power output v_i by g_i .

Thereby these costs have to be defined as a linear combination of several factors, as shown in Equation 4.1. This approach allows the VPP planner to decide on the best compromise between energy loss P_{loss} , carbon emissions E_{CO_2} and production cost C_g .

$$f_i(x_i) = w_1 P_{loss}(x_i) + w_2 E_{CO_2}(x_i) + w_3 C_g(x_i) \quad (4.1)$$

Now, the coordination problem spans over a discretized nonempty and finite set of distinct and successive time periods $T = \{t_1, \dots, t_m\}$. Suppose that for each time slot t_j the VPP has committed to generate a specific amount of energy, denoted as L_j . To ensure that for a specific time slot the correct amount of energy is produced we have introduced a constraint in the DCOP by \tilde{f} , a hyperedge of n -cardinality that links all g_i vertices. Thus, for each time slot t_j we can define a constraint, like \tilde{f} , that guarantees to find assignments that satisfies the demand of L_j , if such an assignment exists. Otherwise, the optimization problem will aim to minimize the cost incurred for not complying with the designated demand in the market, denoted by the penalty factor β .

$$\tilde{f}(x_1, x_2, \dots, x_n) = \beta |L_j - \sum_{1 \leq i \leq n} x_i|, \forall t_j \in T \quad (4.2)$$

Furthermore, given the specifications of the generators, additional constraints need to be factored into the model. Concretely, the power output of a generator g_i can be revised between two consecutive time-slots within a predefined limit l_{max_i} . Thus we have to duplicate the variables and constraints, shown in Figure 4.5, for different time slots and link them together with constraints, that represents these generator-specific characteristics (Figure 4.7). Constraints that define the output behavior of a generator in time will be of the form as shown in Equation 4.3, restricting the values of x_i for the next time slot $t + 1$, based on the values at preceding time slots.

$$\hat{f}_i(x_{i,t-k}, \dots, x_{i,t}) \rightarrow x_{i,t+1} \quad (4.3)$$




Figure 4.6: Reformulating the DCOP representation of the VPP scenario

Basic DCOP preprocessing and optimization

Now, the first step towards solving the optimal dispatch planning problem requires a transformation of our constraint graph so that standard DCOP algorithms could be efficiently applicable. Thus, we substitute the n -ary constraint by introducing a *complex local problem* at the VPP agent level, as depicted in Figure 4.6. This requires introducing the decision variables g_i' , which mirror the value choices of each generator g_i respectively. In turn, \tilde{f} becomes a unary constraint to be addressed locally by the VPP agent, reducing the message passing overhead.

4.2.2 Stochastic extension

Renewable resources are intrinsically prone to inaccurate estimations of their generating output. The challenge is then to account for this stochastic information and optimize the variables with respect to a probabilistic distribution of their expected domain, reflecting decisions under uncertainty.

In [95] Leaute and Faltings introduce *StochDCOP*, an extension of the traditional




Figure 4.7: Revised DCOP representation of the VPP to include multiple time-slots

DCOP formalism, which includes sources of uncertainty in the form of random, uncontrollable variables with known probability distributions. We start from the definition in [95] and proceed to show how our VPP scenario maps on this stochastic adaptations of the DCOP formalism.

A *StochDCOP* is defined as a tuple $\langle A, V, D, \alpha, \sigma, R, \Delta, P, C, e \rangle$, where:

- A, X, D, α, σ are defined as in standard DCOP;
- $R = \{r_1, \dots, r_q\}$ is a set of random variables modeling future, uncontrollable events;
- $\Delta = \{\Delta_1, \dots, \Delta_q\}$ is a set of (not necessarily finite) domains for the random variables such that r_i takes values in Δ_i ;
- $P = \{\pi_1, \dots, \pi_q\}$ is a set of probability distributions for the random variables; with each distribution $\pi_i : \Delta_i \rightarrow [0, 1]$ defines the probability law for random variable r_i , where the value of π_i sum up to 1;




Figure 4.8: *StochDCOP* representation of the VPP scenario for a single time-slot t

- $C = \{c_1, \dots, c_p\}$ is a set of soft constraints over mixed subsets of decision and random variables;
- e is an evaluation function, which, given a constraint $c_i(x_{i1}, \dots, x_{il}, r'_{i1}, \dots, r'_{ik})$, associates, to each assignment of values to c_i 's decision variables, a single cost $e_{c_i}(x_{i1}, \dots, x_{il}) \in \mathbb{R}$ that is independent of the r_i 's.

A solution is an assignment of values to all decision variables that is independent of the random variables, such that:

$$(x_1^*, x_2^*, \dots, x_n^*) = \arg \min_{x_1, \dots, x_n} \{e_{\sum_i c_i}(x_1, x_2, \dots, x_n)\}$$

In our scenario weather conditions are modeled as random variables, that represent exogenous sources of uncertainty. This formulation enables to represent multiple agents' costs as depending on a common source of uncertainty. For instance consider the scenario in Figure 4.8. Supposing generators g_1 and g_2 are wind turbines, each with particular characteristics, controlled via agents a_1 and a_2 respectively. We then introduce the random variable r_1 , which denotes the forecasted wind speed through

an a priori known probability distribution π_i . The impact of random variable r_1 on a_1 is expressed by the constraints c_1 , having the constraint return a cost distribution instead of just single cost value. Similarly is the case for the constraint c_2 , that binds a_2 to r_1 . Hence, the respective cost distributions of a_1 and a_2 are correlated by the fact that they both depend on the uncontrollable variable r_1 . Furthermore, the probability distributions for the random variables are assumed independent of each other and of the decision variables. The knowledge of each random variable, its domain and its probability distribution is only relevant to the agents in control of neighboring decision variables and thus shared only among these agents. Consequently, solving a *StochDCOP* involves choosing values for the decision variables only, with respect to the given probability distributions of the random variables. Also note that by defining the evaluation function e of the *StochDCOP* we essentially decide how to evaluate the solution quality using a single criterion.

4.2.3 Solving the *StochDCOP* formalization

Given the *StochDCOP* formalism of our *hybrid power generation system*, we now come to the question of determining a solution. The general approach for solving a *StochDCOPs* was introduced by Leaute and Faltings in [95]. Essentially, it consists of three phases. Firstly, it abstracts away from the random variables and generates a pseudo-tree³ arrangement for the decision variables only, via a depth-first traversal of the constraint graph. Secondly, random variables are assigned to decision variables. Thirdly, once the constrained graph has been generated, it simply applies a standard DCOP algorithm, namely DPOP⁴. As we have detailed previously in Section 4.2.1 and Section 4.2.2, the problem representation for the *hybrid power generation system* defines a particular DCOP configuration, as depicted in Figure 4.8. The general representation constitutes a two-level tree, when considering only decision variables.

³By definition, a *pseudo-tree* is a generalization of a tree, in which a node is allowed to have additional links, called back-edges, with remote ancestors, called pseudo-parents and with remote descendants, called pseudo-children, but never with nodes in other branches of the tree.

⁴The approach could be applied to other DCOP algorithms as well.

Algorithm 4Distributed reconfiguration of constraint graph, for each decision variable x

```

1: procedure PROPAGATE CONSTRAINTS PHASE
2:   for all constraints  $c_i(x, x_i)$  do
3:     if  $x_i$  is a random variable then
4:       append( $DepList, c_i$ )
5:     end if
6:   end for
7:   for all incoming messages  $M_j = \{DepList_j\}$  from child  $x_j$  do
8:      $DepList \leftarrow DepList \cup M_j$ 
9:   end for
10:  send  $DepList$  to parent of  $x$  (root)
11: end procedure

12: procedure CONSTRUCT NEW CONSTRAINT GRAPH PHASE
13:  if  $x$  is root then sort  $DepList$  and go to line 18
14:    wait until receive messages  $M = \{DepList\}$  from new parent
15:    while  $\exists r \in DepList$  such that  $\forall x_i \in DepList \nexists c(x_i, r)$  do
16:      assign random variable  $r$  to decision variable  $x$ 
17:    end while
18:    while  $\exists$  subset  $S \in DepList$  so that:
19:     $\exists r_i$  and  $\forall x_i \in S \exists c(x_i, r_i)$  and  $\forall x_j \in DepList \setminus S \nexists c(x_j, r_i)$  do
20:      split  $DepList$  into  $S$  and  $DepList \setminus S$ 
21:       $y \leftarrow popUp(DepList)$ 
22:       $create(c_{new}(x, y))$ 
23:      send  $DepList$  to new child  $y$ 
24:    end while
25:  end if
26:  for all  $r \in DepList$  do
27:    if  $lca(r)$  is not assigned and  $x == lca(r)$  then
28:      mark decision variable  $x$  as  $lca(r)$ 
29:    end if
30:  end for
31: end procedure
=0

```

Additionally, the presence of random variables introduces constraints between themselves and any number of decision variables situated at the second level of the tree. Thus, according to the approach in [95], it appears that we only need to be concerned with the assignment of random variables, since the tree-structure of our DCOP is already in place.

However, simply using the existing DCOP is ought to result in poor performance. The inefficiency is because all of the dependency information is to be centralized at the root agent, which quickly becomes a bottleneck for our scenario. Therefore, we look to adopt a pseudo-tree ordering of the decision variables that can better exploit the topological structure of the problem. Consequently, we introduce a new distributed algorithm (Algorithm 1) for *pseudo-tree generation* that better suits our scenario. Moreover, we go on to show that the assignment of random variables can be performed during the pseudo-tree generation, rather than through an additional phase of tree traversals, which would be inefficient. The goal is then to retain as much parallelism as possible and hereby improve performance. To address this, we decouple the problem by generating a separate branch in the pseudo-tree for each random variable, that includes all decision variables adjacent to a respective random variable, on which computations will be performed in parallel.

Initially, the algorithm propagates up the constraint graph the dependencies on all random variables (lines 1 to 11). This means that each decision variable waits for the reception of a message from each of its children (if any), specifying dependency on random variables and then joins them all together with its own constraints. Next, the result is sent to its parent. As we have described before, largely, for DER devices, it is the case that their operation relies only on one prediction factor, represented via a random variable. Nevertheless, for the sake of generality we do include potential dependencies for a decision variable on any number of random variables. Hence, once the root variable has obtained messages from all its children it computes maximal, ordered lists of decision variables that depend on the same random variable(s) (line 13-24). Each such list will represent a separate sub-tree in our new configuration. The lists of decision variables are sorted based on the number of dependencies on random variables, in decreasing order.

Now, the top-down phase of the algorithm is actually responsible for constructing the new tree topology (lines 12-31). During this procedure, each list is essentially interpreted as a stack. The current node will extract the head of the list and insert a link⁵ to this node, then propagate the list to this new child. However, if the current node identifies that the list consists of subsets of decision variables that depend on distinct random variables, the list is split and the procedure continues recursively (lines 18-24). Note that using this lists' ordering we can further parallelize the computation by branching out the tree structure once a remaining subset of decision variables in the list no longer depend on one or more of the initial random variables from that list.

Completeness and Complexity Analysis

Another important consequence of using Algorithm 1 is that we can guarantee *consistency* for the resulting constraint graph. Consistency means that the pseudo-tree is chosen such that each random variable is positioned in the graph lower than any other decision variable, with whom it shares a constraint. Thus, the modification of the constraint graph to ensure consistency guarantees *completeness*, regardless of the properties of the *evaluation function* e (linear/non-linear) [95]. This is reflected in the algorithm (lines 15-17) by checking, at each iteration during the propagation of the dependency list *DepList*, whether there is a random variable, that none of the remaining decision variables shares a dependency with. If this is the case, the respective random variable is assigned to the current decision variable. It follows, that the information regarding dependencies on all random variables is propagated up the pseudo-tree. Again, we take advantage of the new graph configuration we constructed and stop the propagation of this type of information for each random variable r as soon as no additional decision variable depends on r . This corresponds in effect to the lowest common ancestor of r in the pseudo-tree. In our algorithm, given the list

⁵Similarly to Figure 4.6, the edges inserted between the root and the first level nodes represent equality constraints, while the remaining edges that are to be inserted represent additive constraints. Thus, the complex local problem at the root, given by Equation 4.2, essentially remains the same while the local variables now mirror the value choices for each sub-tree component.

of dependencies *DepList* that is propagated down the tree, determining the decision variables where we can stop the upward transmission of information concerning each random variable is in fact straightforward (lines 26-30).

Looking at the amount of information exchange, the algorithm requires a total of $2 \times (n - 1)$ exchanged messages, one per each edge of the graph for the upwards and downwards propagation phases respectively. Given that the messages contain information about constraints on random variables, this implies that the algorithm yields a complexity of $O(nm)$, where m is the number of random variables and n is the number of decision variables. Finally, the newly resulted problem generated via our proposed algorithm is solvable using any traditional DCOP algorithm. As shown in [94], for instance, DPOP introduces an additional worst-case complexity of $O(nD_{max}^n \Delta_{max}^m)$ in terms of information exchange (where the maximum domain size of the decision and random variables are represented by D_{max}^n and Δ_{max}^m respectively). We can thus conclude that against DPOP, the added complexity for our algorithm for distributively generating a consistent pseudo-tree can be considered largely negligible.

4.2.4 Experimental setting

Mathematical Model of DERs

This section describes the calculation of the output power for some of the most prevalent distributed energy generators commercially available.

Wind Turbine

The output power of a wind turbine is a function of the area swept out by its rotor (A), the air density (ρ) and most notably, the cubic power of the wind speed (v) [44]. It is calculated using the following equation:

$$P_w = 0.5\alpha\rho Av^3$$

PV modules

The output power of a PV system is dependent on the solar irradiance (I_T) as well as the characteristics of the module itself, which include the surface area of the cells (A) and the system's efficiency (η) [44].

$$P_{pv} = I_T A \eta$$

where the solar radiation can be estimated as $I_T = I_b R_b + I_d R_d + (I_b + I_d) R_r$

I_b and I_d denote the direct normal and diffuse solar radiations;

R_d and R_r are the tilt factors for the diffuse and reflected solar radiation.

CHP unit

In contrast to the devices abovementioned whose energy availability is dependent on meteorological conditions, a CHP system is driven by natural gas and is used to meet the electrical demand. Hence, CHP units can easily be adapted to various power output characteristics, while the running cost for CHP is given by the fuel cost.

$$C_{CHP} = E_{CHP} \frac{C_{gas}}{\eta + H_R}$$

where E_{CHP} is the electricity load (kW);

C_{gas} is the price of gas;

η represents the electricity efficiency (%);

H_R is the heat rate (kWh/m^3)

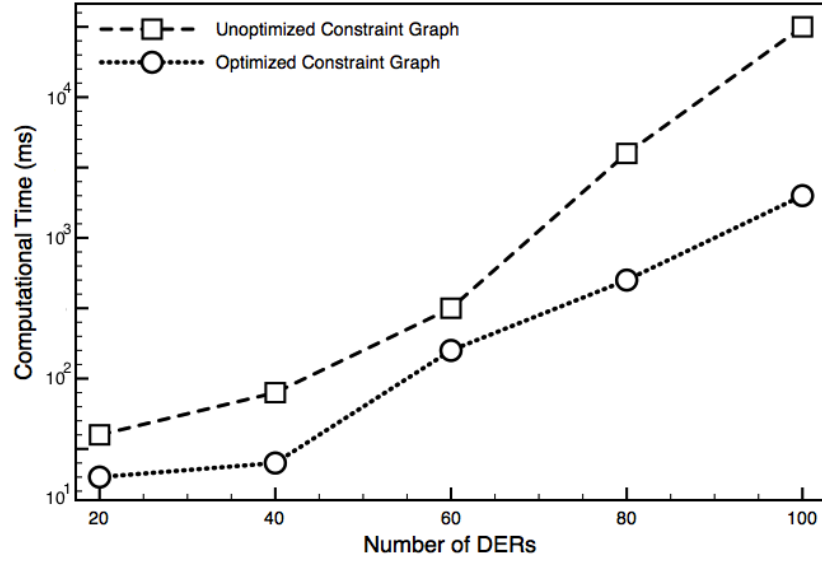


Figure 4.9: Computational time for generating a schedule

Micro-storage system

For our scenario we further assume that VPPs may also include storage devices enabling them to coordinate with renewable energy generation facilities to maximise the usage of the energy produced from such sources. Micro-storage devices are characterized by the available storage capacity c_i and their storage efficiency α_i . Consequently, if a certain q amount of energy is stored, then at most $\alpha_i q$ may be discharged, discretizing the output domain of device g_i within the interval $D_i = [0, \alpha_i q]$.

For simplicity, we assume no maintenance costs of DERs, thus for renewables as well as micro-storage units the generation cost incurred is considered negligible. The carbon emissions for each generator g_i is calculated as a function of its output x_i and its carbon intensity $I_i \in \mathbb{R}^+ kgCO_2/kWh$, as follows: $E_{CO_2} = I_i x_i$. Additionally, the output variable of a generator g_i can take values in the domain $D_i = \{d_1, \dots, d_r\}$, $0 \leq d_i \leq P_{max}$, where d_i denotes a possible operating mode of the device and P_{max} is its maximum generating power.

Performance evaluation

In the following, we give a performance evaluation of our proposed algorithm in the *Smart Grid* domain. We consider an instantiation of a VPP scenario consisting of a combination of at most 100 DERs, that operate according to the specifications aforementioned. Namely, we simulate a VPP including a PV installation⁶ of at most 10MW and a number of at most 24 wind turbines⁷, with a maximum capacity of 17.56 MW. Additionally, we assume CHP units with variable outputs ranging between 100-200kW and a total storage capacity of 300kWh, represented by batteries, with 25kW and 30kW as charging and discharging rates respectively. For the target consumption profile that the VPP is ought to secure we have used datasets provided by the Australian Energy Market Operator⁸. A piece of Java software was implemented to synthesise this scenario and further integrated with the FRODO 2 platform [96], which has been released to the public under the Affero GPL license and has been extensively developed by numerous contributors from the DCSP and DCOP community.

The empirical results are summarized in Figure 4.9 by computing a generation schedule based on the given scenario. Firstly, the experiments we conducted show that our model adheres to the real-time constraints of the VPP scheduling problem, where for grid regulation the response time is constrained within minutes, or to even a couple of seconds. Secondly, our proposed algorithm delivers a superior performance in terms of computational time. Results show that applying our algorithm for a reconfiguration of the constraint graph produces a moderate decrease of computational time in the number of DERs, as opposed to standard approach used for solving the *StochDCOP* formulation. Specifically, as we approach the upper bound for our considered scenario, in terms of number of coordinated DERs, we improve the scheduling time by a factor of 10 (notice the *log* scale in Fig. 4.9). This becomes particularly important as the scale of the VPP increases and a larger portion of the consumption profile is

⁶<https://openpv.nrel.gov/>

⁷<http://www.sotaventogalicia.com/>

⁸<http://www.aemo.com.au/Electricity/Data>

accommodated via renewable resources.

4.3 Discussion

Generically, these intelligent electricity network technologies support the vision of a *Smart Grid* that aims at reducing the carbon footprint, while increasing energy efficiency and clean energy usage. A key approach in doing so, focuses on decreasing the high energy costs and emissions during peak demand periods by means of intelligently coordinating the variable output of wind or solar energy generators. In particular the cooperative concept of virtual power plants has been advocated as a viable organizational model [24, 110], from the grid operator's viewpoint, allowing to cope efficiently with the integration of the many distributed energy resources.

In [24], the authors start from the assumption of an existing VPP and propose a novel pricing mechanism, that addresses the question of allocating payoff amongst the VPP members, so that it can guarantee that no subset of agents has an incentive to break away. Additionally, the payoff scheme elicits truthful behavior from the VPP members, while their production estimates are evaluated by means of statistical methods. Similarly, the PowerMatcher described in [84] is a market-based multi-agent tree-architecture for balancing supply and demand within clusters of DERs.

From an organizational perspective earlier works [137, 143] suggested centralized structures for managing the grid, that come short in addressing an open system setting and the inherent stochastic nature of DERs. These issues have been captured in Chapter 3, where a dynamic coalition formation mechanism is proposed for the creation of dynamic microgrids, considering a scenario that introduces a decentralized algorithm according to which, agents efficiently self-organizing into coalitions representing the actual microgrid configurations.

Alternatively, in Section 4.1, we have taken a different perspective and question the emergence of such a phenomenon of evolution of cooperation itself in a population of self-interested agents in the context of VPPs. We modelled the problem as a repeated game and conducted an analysis in order to determine the context under which collaborative behavior could result. Gaining an understanding of what drives

cooperation under the assumption of rational agents is particularly important in designing system-level mechanisms and policies that could incentivize efficient resource allocation.

The emergence of cooperation amongst self-interested agents has received a lot of attention from various research areas including social sciences, behavioral economics or evolutionary biology [9, 156]. Several notable efforts have looked at different variations of the Prisoner's Dilemma game⁹ and studied the influences of parameters such as underlying network topology, interaction rules or updated rules [196, 65]. In Section 4.1 we have taken a different outlook on this issue and addressed the problem of stochastic environments, specifically represented here by the electricity domain, showing how collaborative behavior can emerge as an adaptive strategy for handling uncertainty. We have described an agent-based model for a smart electricity grid that assumes the presence of large number of distributed energy resources in the system and studied the phenomenon from the perspective of self-interested agents, that look to maximize their payoff, in order to determine if and under what conditions such a collaborative behavior might emerge.

Experiments have shown that cooperation amongst rational agents is an emergent phenomenon in this setting for a large fraction of the agent population, including for games played with limited information. In fact cooperation is the optimal strategy in situations of high uncertainty, where agents adopt it as an adaptation mechanism to variable environmental conditions. However, when uncertainty is decreased defecting may produce better results for good predicting agents. Whereas collaboration becomes more susceptible to suboptimal gains as some cooperating agents may still return acceptable payoffs although selecting suboptimal strategies, due to the redistribution of available energy. Therefore, in such instances, in order to instil cooperation as the dominant strategy further mechanisms need to be implemented at the VPP level in order to ensure optimality.

The work presented in Section 4.2 is intended to provide a flexible approach for

⁹Well-known two-player game-theoretic framework where agents have to chose between two strategies: cooperating or defecting. While defecting is the dominant strategy and the only Nash Equilibrium, it is also Pareto-inefficient as cooperation would make both players better off.

coordinating a hybrid power generation system. Here, our goal is restricted to providing a solution for generating a schedule for a given hybrid power generation system under uncertainty conditions. Firstly we have provided a new formulation of the scheduling problem in terms of a distributed constraint optimization. We then go on to extend our formalism to capture the inherent stochasticity of the domain. A novel distributed algorithm is introduced for solving this optimization problem by means of exploiting domain-dependent characteristics of our model. Finally, we provide an empirical evaluation of our approach, which shows its practical applicability and the performance improvements in terms of computational time speed-ups.

In Section 4.1 we have addressed specific energy markets suitable for DER participation. However, assuming reliable VPPs, there are multiple markets available for trading energy (e.g. baseload power market, peak power market, spinning reserves market, regulation market). A further interesting perspective on this work would be to investigate how would the VPP be able to maximize profits gained from participating in these different markets concurrently and what would the impact of such behavior be on the grid at large.

To sum up, the role of traditional power plants is expected to diminish considerably in the near future. Virtual Power Plants represent a key concept for the future smart grids, leveraging distributed generation resources. Importantly, the majority of DERs provide clean energy, significantly reducing the total carbon emissions. However, this also makes their power output highly weather dependent. Laying the foundations for DER integration needs therefore to provide solutions for systems that remotely and automatically dispatch and optimize generation from intermittent resources. Hence, there are many interesting questions to be investigated.

In this chapter, in Section 4.1 we proved the financial benefits of DER owners that come through VPP cooperation, under varying profit maximizing strategies and production conditions. Also, it is imperative that VPPs can package a wide variety of generation devices and storage units to ensure a reliable service of energy supply. Section 4.2 addresses this very aspect of generating a schedule for a given hybrid power generation system under uncertainty conditions. In this way, we essentially reconstruct the power plant functionality by aggregating its distributed equivalent.

Conclusively, in this chapter we align our VPP solution along the same underlying principle of tackling the emerging complexity of the system by cooperation and modularity.

Chapter 5

Coordination Mechanisms for *Forward & Intraday Markets*

*The behaviour of individuals is the tool with which the organisation
achieves its targets.*

— *Herbert Simon* —

A central pillar behind the vision for a new electricity grid is the *smart-meter*. Essentially, the smart-meter opens up an entirely uncharted territory in the sector of energy-efficiency, by empowering consumers to become an active participant towards improving the performance of the grid. It allows them to gain a better understanding of their consumption and therefore take better informed decisions about the impact of their usage. Against this background, the European Directive EU 2006/32/EC legislated to provide a complete smart-meter coverage until 2018 across Europe. Similar initiatives are currently running around the world.

In this chapter the focal point of our research is to provide a *consumer-centric* approach for balancing supply and demand by designing mechanisms that tie consumers into the loop and enable automated decision making on their behalf. We begin by setting the context and answering the question of how supply and demand are maintained in balance today and highlight the inherent inefficiencies.

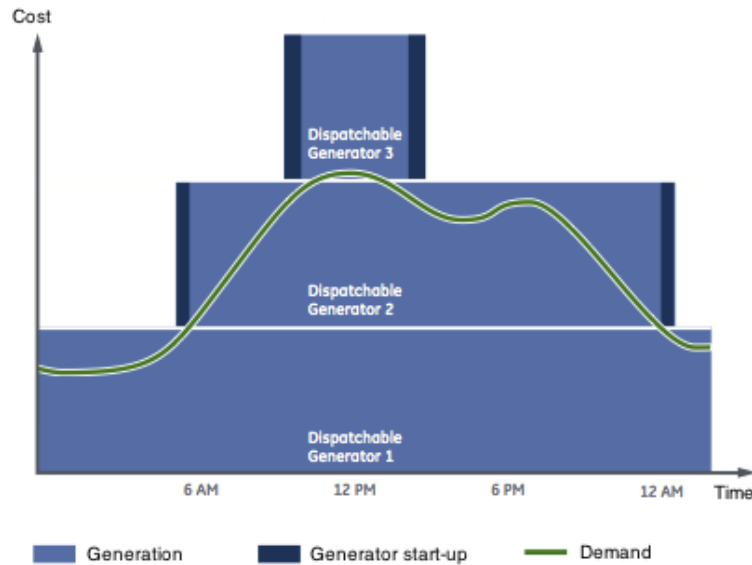


Figure 5.1: Generation optimization to meet demand for scenario:
 $\text{dispatch cost of generator 1} < \text{cost generator 2} < \text{cost generator 3}$

Conventionally, electricity is produced to meet the instantaneous demand, which is dependent on the current aggregated load of consumers. Electricity generation is characterized by a non-decreasing marginal cost of production. Why is this the case? As we have described in the previous chapter 4.2, a utility or virtual power plant operates its generators such that the load profile is met with minimum costs. This means that generators will be activated in the order of their economic efficiency (see Fig 5.1). Due to this allocation procedure, the marginal cost of electricity is non-decreasing. Thus, we can conclude that during intervals of higher demand the marginal cost of producing energy is higher, while during periods of lower demand the marginal cost is lower.

In Figure 5.2 we plot a generic, average representation of the load curve for a standard residential area. This pattern is the cause for a number of severe inefficiencies:

- Firstly, there is the hindrance of *cost-inefficiency*. Due to the abovementioned explanation, peaking power to supplement short-falls during peak demand can double or quadruple the marginal cost of electricity.




Figure 5.2: Average daily consumption diagram. L represents the desired level of demand.

- A second concern is the *eco-inefficiency*. Peakload and contingency periods are handled by firing carbon-intensive, peaking plant generators, which impact heavily the carbon footprint.
- Thirdly, as the demand of electricity is increasing steadily, during peakload intervals the power system will be operated near to its limits. This *network-inefficiency* is on one hand a threat for brownouts¹ or potentially more severe cascading blackouts. On the other hand, for the current situation, in order to cope with the increasing demand, high capital investments into the bulk power infrastructure are required.

It follows then that, if the load curve could be influenced so that peak-periods could be shifted to off-peak, the overall operating cost could be reduced. As represented in Figure 5.2, *electricity curtailment* means avoiding surplus periods denoted with 'C', as well as the times when energy is in deficiency, denoted with 'A'. The traditional operational mode of the grid where generation must follow demand needs

¹A *brownout* is a drop in voltage in an electrical power supply. Brownouts may also occur intentionally in order to reduce the load during emergency conditions, to prevent a total supply failure (blackout).

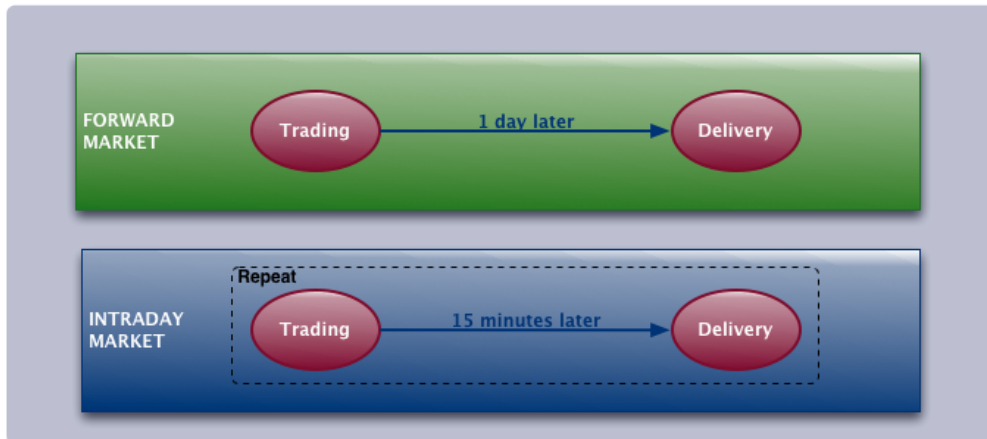


Figure 5.3: Market timeline

to be reversed to a business model where demand can follow generation. Aside for alleviating the enumerated inefficiencies, the flexibility of demand will foster the integration of renewable resources.

The general structure for trading energy distinguishes between two market timelines (see Figure 5.3). On one hand, the base-load energy requirements are traded in the day-ahead market, also referred as the *forward market*, which settles the price and volumes of electricity for each time-slot for the following day. On the other hand, the *intra-day market* is responsible for a real-time balancing of supply and demand, whenever there is an imbalance between the contracted energy supply and demand within a given time-slot. It is important to note that the difference in terms of time-horizon, means that different approaches are required to determine a viable provision of energy solution for these two markets. Against this background, consumers can be engaged to provide two types of services.

- Firstly, we shall address the forward market, which plans the day-ahead schedule. Here, based on the estimated demand a number of generators are committed for dispatch during the following day. This sets a threshold in terms of the available amount of energy to be distributed among consumers. The challenge here is to design a mechanism where consumer can mostly benefit from operating under the given constraints by adapting there consumption. This class of

services fall under the term of *demand-side management*, which we dwell upon in the first half of this chapter.

- Secondly, unlike the forward market, where the goal is flattening the energy consumption curve for the day ahead, the intraday market aims to balance generation and consumption in the event of a contingency. Here, consumers need to react quickly to provide *demand response* services, which are requested with only a few minutes notice. This topic is addressed in the second part. The key point here is that there is a real economic value proposition that concerns utilities, regulators and customers altogether, in contrast to other "green" initiatives that rely on financial subsidies or other intangible societal benefits. The application of autonomous consumer response for *active power regulation* could enable the system to respond dynamically to time-varying grid conditions through adaptive, localized, and continuous power regulation.

5.1 Adaptive Demand-Side Management

The concept of the Smart Grid envisions deploying intelligent agent technologies for managing the future electricity networks. One key challenge in this regard, generically known as Demand-Side Management (DSM), aims at eliciting a desired consumer behavior for increasing energy efficiency, via reducing peak energy demand and capitalizing on available, intermittent energy resources. To this end, game theory proves to be a particularly suitable tool for analysing such setting, given that consumers can be modelled as individual and independent decision makers whose behavior impacts all other consumers and the grid at large. Hence, in this section, we introduce a game-theoretic approach, which focuses on consumer coordination for energy efficiency. A non-cooperative game is proposed where agents representing consumers adapt their loads in order to maximize their user satisfaction, represented as a utility function. Next we address the existence and uniqueness of Nash Equilibrium for the proposed game. We then put forward under this framework a novel DSM algorithm for enabling decentralized control of the grid, which we analytically prove to converge to the Nash Equilibrium. Finally, the performance of the proposed mechanism is evaluated by simulation.

The outline for the rest of chapter 5.1 is as follows. Section 5.1.1 describes the agent-based framework used for modelling the problem domain. Then, in Section 5.1.2 we formulate the DSM problem by defining a game in the analytical setting of non-cooperative game-theory. The description of our proposed algorithm is given in Section 5.1.3. Numerical results are presented in Section 5.1.4. In Section 5.1.5 we provide a discussion of the solution proposed, relating it to existing work in this area.

Variable	Description	Variable	Description
\mathcal{N}	set of players	u	agent utility
\mathcal{L}	subset of consumer agents l_i	\mathcal{A}	n -tuple of action variables
\mathcal{T}	set of time-slots t_k	C_{max}	cost constrain per agent
β	agent profile function	Γ	non-cooperative game
\mathcal{D}	set of deferrable loads	Θ	desired consumption level
l	set of deferrable interruptibles	CoS	cost of stability
\bar{l}	set of deferrable non-interruptibles	k	iteration count
Φ	maximum power constrain	\mathbf{c}	cost vector
s	initial starting time slot per device	δ	duration of loads
ζ_1	earliest time slot for the load to start	r	preference rating per device
ζ_2	latest time slot for the load to start	d	device power rate (in kW)
φ	determines active periods of devices	Δ	deferment per device

5.1.1 MAS Framework for Smart Distribution Grids

We represent the set of consumers as the set of self-interested agents $\mathcal{L} = \{l_i \mid 0 < i \leq n\}$ that always aim at minimising their incurred costs. Each agent $l \in \mathcal{L}$ is characterized by its estimated load for the fixed set of time intervals $t_k \in \mathcal{T}$ via a profile function $\beta^{l_i}(t_k)$. The day-ahead aggregate of all profile functions for all agents represents the system's operational mode:

$$\beta = \begin{bmatrix} \beta^{l_1}(t_1) & \beta^{l_1}(t_2) & \dots & \beta^{l_1}(t_m) \\ \beta^{l_2}(t_1) & \beta^{l_2}(t_2) & \dots & \beta^{l_2}(t_m) \\ \vdots & \vdots & \vdots & \\ \beta^{l_n}(t_1) & \beta^{l_n}(t_2) & \dots & \beta^{l_n}(t_m) \end{bmatrix}$$

where each line of the matrix, $\bar{\beta}^{l_i}(t_k), \forall t_k \in \mathcal{T}, \forall l_i \in \mathcal{L}$, represents the operational mode vector for agent l_i .

Now, each agent l_i is responsible for computing $\bar{\beta}^{l_i}(t_k)$ according to its consumer's preferences. Generally, home appliances can be classified into deferrable and non-deferrable loads. While the latter are assumed to be fixed requirements of the user, the former category will be managed by the agent program. Each agent seeks to optimise energy usage in order to reduce costs as well as applying the least amount of rescheduling, required in order to meet budget constraints.

Within the domestic energy domain, for the category of non-deferrables, we exemplify: entertainment devices, lighting, computer usage, etc. Alternatively, the set of deferrables \mathcal{D} can further be differentiated into two specific categories: *i*) non-interruptible \bar{I} (e.g. washing machines, dishwashers, etc.) and *ii*) interruptible loads I (e.g. electric vehicles, heat pumps, radiators, etc.), where $\mathcal{D} = I \cup \bar{I}$. In the following we build specific models for the two classes of deferrable devices.

We associate with agent l_i the set of *deferrable non-interruptible* loads $d_j^{l_i} \in \bar{I}$ (in kW) and their initial starting time slots $s_j^{l_i}$ set by the user, as well as the duration $\delta_j^{l_i}$ and the descending preference rating $r_j^{l_i}$ for the respective device. In doing so, we can represent the objective of the agent in terms of a MILP (mix-integer linear program) optimization problem of determining the optimal deferments Δ_j for minimising loss of comfort and maintaining demand within certain specified constraints:

$$\operatorname{argmin}_{\Delta_j} \sum_{d_j^{l_i} \in \bar{I}} r_j^{l_i} |\Delta_j| \quad (5.1)$$

$$\text{subject to: } \sum_{d_j^{l_i} \in \bar{I}} d_j^{l_i} \varphi_j(t_k) \leq \Phi_i(t_k)$$

$$\zeta_1^j \leq s_j + \Delta_j + \delta_j \leq \zeta_2^j$$

$$\text{where: } \varphi_j(t_k) = \begin{cases} 1 & \text{if } t_k \in [s_j + \Delta_j, s_j + \Delta_j + \delta_j] \\ 0 & \text{otherwise} \end{cases}$$

such that $\mathcal{T} = \{1, \dots, 48\}$ discretizes the daily schedule over half-hourly time slots $t_k \in \mathcal{T}$; interval $[\zeta_1^j, \zeta_2^j]$ represents a hard constraint for scheduling load j ; $\Delta_j^{l_i}$ specifies the time deferment of load j for agent l_i ; φ_j determines the active periods for each load according to the revised schedule; $\Phi_i(t_k)$ denotes the maximum power constraints of the user for the respective time period t_k ; Deferments from the preset time for operating a particular device may be required in order to comply with the maximum power constrain Φ . The more the starting time for a device is deferred, the more discomfort it will cause to the user. The minimization in Equation 5.1.1 strives to reduce the user's discomfort by ensuring least deviations, following the preference rating ordering over the set of devices.

We now turn to modelling *deferrable interruptible* loads, by following a similar approach. For the previous category we minimized the deferment of each load with respect to the user's preferred starting time, in the order of their preference rating denoted by $r_j^{l_i}$. An interruptible load has an additional degree of freedom because it can further be interrupted and resumed as long as the scheduling happens within the hard constrained interval $[\zeta_1^j, \zeta_2^j]$. Hence, we are concerned here with minimising the deferment Δ_j from the earliest preferred finishing time $\zeta_1^j + \delta_j$. Let ζ_f^j denote the finishing time-slot for device $d_j^{l_i}$, while the remaining parameters preserve the same meanings. Then, the optimization problem for deferrable interruptible loads boils down to:

$$\operatorname{argmin}_{\Delta_j} \sum_{d_j^{l_i} \in I} r_j^{l_i} |\Delta_j| \quad (5.2)$$

$$\text{subject to: } \sum_{d_j^{l_i} \in I} d_j^{l_i} \varphi_j(t_k) \leq \Phi_i(t_k)$$

$$\sum_{t_k \in \mathcal{T}} \varphi_j(t_k) = \delta_j$$

$$\text{where: } \Delta_j = \zeta_f^j - (\zeta_1^j + \delta_j)$$

$$\text{and } \varphi_j(t_k) = \begin{cases} 0 & \text{if } t_k < \zeta_1^j \text{ or } t_k > \zeta_2^j \text{ or } t_k > \zeta_f^j \text{ or device is } \textit{inactive} \\ 1 & \text{device is } \textit{active} \end{cases}$$

We are now in the position to bring it all together and summarize the overall optimization procedure for both categories of deferrable loads, denoted respectively with subscripts I and \bar{I} :

$$\operatorname{argmin}_{\Delta_j^I, \Delta_j^{\bar{I}}} \left(\sum_{d_j^{l_i} \in I} r_j^{l_i} |\Delta_j^I| + \sum_{d_j^{l_i} \in \bar{I}} r_j^{l_i} |\Delta_j^{\bar{I}}| \right) \quad (5.3)$$

$$\text{subject to: } \sum_{d_j^{l_i} \in \bar{I}} d_j^{l_i} \varphi_j^{\bar{I}}(t_k) + \sum_{d_j^{l_i} \in I} d_j^{l_i} \varphi_j^I(t_k) \leq \Phi_i(t_k)$$

$$\zeta_1^j \leq s_j + \Delta_j^{\bar{I}} + \delta_j \leq \zeta_2^j$$

$$\sum_{t_k \in \mathcal{T}} \varphi_j^I(t_k) = \delta_j$$

$$\text{where: } \Delta_j^I = \zeta_f^j - (\zeta_1^j + \delta_j)$$

$$\text{and } \varphi_j^I(t_k) = \begin{cases} 0 & \text{if } t_k < \zeta_1^j \text{ or } t_k > \zeta_2^j \text{ or } t_k > \zeta_f^j \text{ or device is } \textit{inactive} \\ 1 & \text{device is } \textit{active} \end{cases}$$

$$\text{and } \varphi_j^{\bar{I}}(t_k) = \begin{cases} 1 & \text{if } t_k \in [s_j + \Delta_j^{\bar{I}}, s_j + \Delta_j^{\bar{I}} + \delta_j] \\ 0 & \text{otherwise} \end{cases}$$

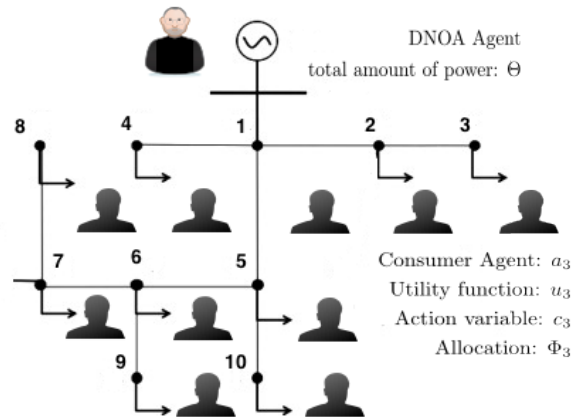


Figure 5.4: Agent-based Smart Distribution Grid

5.1.2 The Non-Cooperative Game Formulation

The scenario under consideration pertains to a typical setting for a residential area, where consumers optimise their energy usage based on the procedure detailed in the previous section. According to the discussion in the introduction of this chapter, there are several aspects against which the efficiency of the grid could be evaluated. We termed these: cost-efficiency, eco-efficiency and network-efficiency. For the remainder of Chapter 5.1 it is important to acknowledge that they come as a direct effect from setting the constrain on the load curve. Aside from this, there is of course, a limited amount of power that a distribution substation in a residential area can safely provide. The problem is thus closely akin to a resource allocation problem where consumers represented by software agents compete for a scarce resource. The setting assumes the presence of a distribution network operator agent (DNOA) that has the task of allocating the available amount of energy. Additionally, each consumer provides an initial set of information regarding preferred consumption patterns and delegates control to a software agents that interacts with the DNOA. We use a game-theoretic approach in order to better understand the expected behavior of the agents for the outlined scenario and consequently, to design a distributed mechanism capable to induce a socially desirable equilibrium. Moreover, we intend to do so without revealing private information about the agents' preferences.

We begin by giving the basic intuition behind the coordination mechanism proposed. The problem of efficient power allocation to consumers is modelled in the form of a *game*. Consumers are modelled as utility maximizing agents (or *players* in the game) that choose the cost that they are willing to pay such that they can maximize their utilities. The choice of an agents will influence the performance of other agents in the system. The general picture is given in Figure 5.4 representing consumer agents that have the capability to behave "selfishly" and strive to optimize their own utility unilaterally. The utility of an agents denotes a trade-off between the satisfaction of the allocated energy amount and the cost of this allocation. We dwell upon this in greater detail in the next paragraphs. The outcome of the game provides an allocation of power to agents from which, according to Section 5.1.1 the optimal schedule of devices is determined. Based on this premise we investigate: What strategy should an agent chose to maximize its utility? If agents selfishly select their utility-maximizing strategy, will there be a stable state where no agent can unilaterally improve its utility and where each agent attains its own optimum coincidentally?

Formally, let $\Gamma = \langle \mathcal{N}, \mathcal{A}, \{u_i\} \rangle$ be a n -person non-cooperative game that models the interaction between individual rational decision-makers with potentially conflicting objectives. $\mathcal{N} = \{1, 2, \dots, n\}$ represents the set of players (decision-makers); \mathcal{A} denotes the n -tuple of decision or action variables of all players, $\mathcal{A} = \times_{i=1}^n c_i$. The action variable of player j is denoted by the finite strategy set $c_j = [1, \dots, C_{(j)max}]$, where each player j is characterized by its upper cost bound $C_{(j)max}$. $u_i : \mathcal{A} \rightarrow \mathbb{R}$ is the set of utility functions that each player i wants to maximize. The utility of user j , defined by $u_j(c_j, \mathbf{c}_{-j})$, is a function of the action chosen by the j th player, c_j , and the actions chosen by all the players in the game other than player j , denoted as \mathbf{c}_{-j} . Thus, together c_j and \mathbf{c}_{-j} make up an action tuple \mathbf{c} , which is a unique choice of actions by each player. We will refer onwards to $\mathbf{c} = (c_1, c_2, \dots, c_n) \in \mathcal{A}$ as the cost vector, where \mathcal{A} is now the set of all cost vectors.

The outcome for the non-cooperative strategic game consists of the selection of cost vector \mathbf{c} , in which the players select their strategies without knowing the other players choices (moving simultaneously). Given this scenario, there may occur a set of Nash equilibrium points exhibiting robustness or local optimality in the sense




Figure 5.5: The sigmoid function f_i and the role of parameters: i) a determines the slope of the curve; ii) p determines the point around which the performance growth is the fastest.

that no player has any incentive to choose a different strategy. Furthermore, the Nash equilibrium is fair among all players in the sense that it is achieved by the fair competition among agents.

The Nash equilibrium (or non-cooperative equilibrium) is an action profile $\mathcal{A} = (c_1, \dots, c_n)$ at which no user may gain by unilaterally deviating. Formally, the action tuple $\mathbf{c}^* = (c_1^*, c_2^*, \dots, c_n^*)$ is a Nash equilibrium of the game iff $u_i(c_i^*, \mathbf{c}_{-j}^*) \geq u_i(c_i, \mathbf{c}_{-j}^*), \forall c_i \in \mathcal{A}_i$ and $\forall i \in \mathcal{N}$.

Utility Function

For determining the payoff of the game for each player we further need to specify the system-dependent utility functions u_i , based upon the agent model introduced in the previous Section 5.1.1. The utility function essentially measures the satisfaction experienced by each player, by capturing the trade-off between the cost of energy and the level of user requirements (corresponding to the initial profile function β^{l_i}), that are being actually met:

$$u_i = k \frac{f_i}{c_i} \quad (5.4)$$

We represent the first component of u_i as a function of the allocated energy x (which can be also be translated as the number of terminated scheduled loads) using a normalized sigmoid function:

$$f_i(x) = \frac{1}{1 + e^{-a(x-p)}} \quad (5.5)$$

In doing so, parametrizing the function allows the specification of different degrees of preference for following the initial user requirements as depicted in Figure 5.5, where p represents the sigmoid center, while a is the sigmoid width. The user can control, on one hand the preferable amount of terminated loads² via the parameter p and on the other hand, the steepness of the curve via a which corresponds to the user's sensitivity within its load range. Thus, for values below p we can observe a very low utility, while as we approach p , it increases gradually according to parameter a . This can be understood as low level utility and speed of growth when only a small number of load devices from the agent l_i 's set $d_j^{l_i}$ can be terminated, since a low number of activated devices would not yield much improvement in utility. However, as the number of activated devices increases, they result in a greater increase in user utility. We consider a utility of 1 if all user requirements have been achieved, meaning all $d_j^{l_i}$ have been scheduled. Additionally, as shown in Figure 5.5, the increase in utility decreases near the upper limit of the sigmoid function f_i .

Now, the overall performances perceived by the user (the a_i utility, u_i) is given by equation 5.4 and is represented in Figure 5.6 as the trade-off between user costs c_i and the level of requirements being met, denoted by f_i . It can be observed that as costs exceed the threshold budget constraints, utility is decreasing proportionally inverse to costs. Here the parameter k is additionally used to specify the maximum user utility as a function of incurred costs, which occurs when the partial derivative of u_i with respect to the cost c_i is zero:

²We remind that the optimization algorithm running at the agent level ensures that the deferment of loads is considered in an ascending order according to the preference rating $r_j^{l_i}$ specified by each user l_i for his device set $d_j^{l_i}$




Figure 5.6: Generic user utility as a function of incurred cost for different values of k .

$$\frac{\partial u_i}{\partial c_i} = 0 \quad (5.6)$$

Control Strategies

Before addressing the problem of analysing the Nash equilibria of the game, we define a player's *best response* as the strategy that maximizes her utility function for a given action tuple of the other players. Formally, the best response of player j to \mathbf{c}_{-j} is represented by:

$$\bar{c} \in \left\{ \underset{c_j \in \mathcal{A}_j}{\operatorname{argmax}} u_j(c_j, \mathbf{c}_{-j}) \right\} \quad (5.7)$$

Specifically, for our scenario, determining the best response strategy for player j corresponds to the selection of the optimal cost of energy for its associated day-ahead schedule. Essentially, the cost vector \mathbf{c} comprised of the choices of each agent, determines the actual amount of energy Φ awarded to each player given below, with respect to the grid operator's desired consumption level Θ :

$$\Phi_j = \frac{c_j}{\sum_{i=1, i \neq j}^n c_i + c} \Theta \quad (5.8)$$

As system designers, the particular choice for the energy allocation function Φ_j enables to validate certain properties of the game. Constant c is used here for normalization, with $c \leq C_{(j)max}, \forall j \in \mathcal{N}$. We further define the cost of stability (CoS) for the proposed game as the difference between the desired consumption bound Θ and the actual allocated energy:

$$CoS = \Theta - \sum_{i=1}^n \Phi_i \quad (5.9)$$

Equation 5.8 indicates that for a fixed bid \mathbf{c}_{-j} of all other players in the game except j , Φ_j has a linear growth as a function of c_j . Based on the designated amount of energy Φ_j , each agent adapts its energy usage by applying a deferment set $\Delta_i^{l_j}$ to meet its energy constraints. This in turn sets the number of scheduled loads $d_i^{l_j}$, resulted from computing a new profile function $\beta^{l_j}(t)$ obtained by solving the optimization problem in Section 5.1.1.

Existence of Equilibrium

Proposition 1. The game $\Gamma = \langle \mathcal{N}, \mathcal{A}, \{u_i\} \rangle$ has at least one Nash equilibrium point.

Theorem 1 (*Glicksberg-Fan [51, 42]*). A Nash equilibrium exists in a game $\Gamma = \langle \mathcal{I}, (\mathcal{S}_i)_{i \in \mathcal{I}}, \{u_i\}_{i \in \mathcal{I}} \rangle$ if for each $i \in \mathcal{I}$:

- (a) \mathcal{S}_i is a nonempty, convex and compact subset of a Euclidean space \mathcal{R}^n
- (b) $u_i(s_i, s_{-i})$ is continuous in s_{-i}
- (c) $u_i(s_i, s_{-i})$ is quasi-concave in s_i

Proof. In order to analyse the existence of equilibrium in the game Γ we can firstly directly infer condition (a), given that the action set \mathcal{A}_i is by definition non-empty and convex. Similarly, it is easy to show that \mathcal{A}_i is compact, given that it is on one hand bounded by $C_{(i)max}$ and secondly, it is closed as it includes the boundary points 1 and $C_{(i)max}$. Again, by definition u_i is continuous (b), remaining to demonstrate only that u_i is quasi-concave on \mathcal{A}_i (c).

A function $g : \mathcal{X} \rightarrow \mathbb{R}$ is quasi-concave iff either g is monotonic or it is single-peaked, meaning that $\exists x_0 \in \mathcal{X}$ so that $g(x)$ is non-decreasing on $\mathcal{X} \cap [0, x_0)$ and non-increasing on $\mathcal{X} \cap (x_0, \infty)$. Then, clearly this implies that for those values for which the first derivative g' is zero, the second derivative g'' must be negative. For notational convenience and without loss of generality, let $u_i = \frac{h(x)}{x}$, where $x \in [1, \infty)^3$ and $h(x)$ is obtained by substituting x in equation 5.5 with Φ_j from equation 5.8, representing utility as a function of cost c_j . Then the second derivative of the utility function is:

$$u_i''(x) = \frac{\partial^2}{\partial x^2} \left(\frac{h(x)}{x} \right) = \frac{h''(x)}{x} - \frac{2(h'(x)x - h(x))}{x^3}$$

where substituting $u_i'(x) = 0$ simplifies to:

$$\frac{\partial^2}{\partial x^2} \left(\frac{h(x)}{x} \right) = \frac{h''(x)}{x}$$

which takes only negative values, $\forall x \in [1, \infty)$ which includes any finite strategy set $[1, C_{(j)max}]$. Hence, the conditions for quasi-concavity are satisfied. For a more detailed discussion on quasi-concavity we refer the reader to [141]. \square

In the following we derive an instance of a standard fixed-point algorithm, by means of which such an equilibrium can be reached.

5.1.3 A MAS-based Control Mechanism

In this section we introduce a distributed algorithm designed for achieving decentralized control over the grid. The challenge of DSM mechanisms is to elicit a desired

³Hence we are assuming that the players' budget constraints $C_{(j)max}$ are sufficiently large s.t. each player can achieve c^* .

consumer behavior for increasing energy efficiency, via reducing peak energy demand and capitalizing on available, intermittent energy resources. Such mechanisms can essentially be classified into two categories: *i) cooperative* and *ii) non-cooperative* scenarios. The former assumes a collaborative settings where user preferences are subordinated to the grid operator's objective and can thus be overwritten. Contrary to this, our algorithm considers a non-cooperative environment where agents selfishly aim to maximise their utility. Several attempts have looked into applying price signals, as opposed to fixed energy prices, in order to regulate consumer behavior. This implies either using time-of-use pricing (different costs for peak and off-peak intervals) or real-time pricing schemes (energy price is specified for the following half-hour period). Unfortunately, although providing a better performance in terms of lowering peaks, both time-of-use and real-time pricing mainly display a reactive behavior shifting peaks and so, simply inducing other peaks during different time periods.

To cope with these drawbacks, we propose a decentralised multi-agent system algorithm for coordinating an adaptive demand side-management (ADDSM), resulting in an allocation where agents attain their own optimum coincidentally, while concealing specifications regarding their user preferences and valuations (utility functions). We build upon the non-cooperative game framework previously introduced, where players follow the rules detailed in Section 5.1.2. The outcome of the iterative game Γ considered, represents the system's day-ahead operational mode. The description of the ADDSM algorithm is as follows.

1. Set iteration count to $k = 0$. Players initialize their action variables c_i generating $\mathbf{c}(0)$, over each time slot $\mathcal{T} = \{1, \dots, 48\}$.
2. Increment iteration $k \leftarrow k + 1$.
3. Given the action variables of the other users $\mathbf{c}_{-j}(k - 1)$, each agent derives its utility based on the allocation provided by the DNOA according to Equation 5.8 .
4. Each player updates strategy by computing the best response w.r.t the previous iteration:

$$\bar{c}_j(k) = \operatorname{argmax}_{c_j \in \mathcal{A}_j} u_j(c_j, \mathbf{c}_{-j}(k-1)), \forall j \in \mathcal{N} \quad (5.10)$$

5. Stop if algorithm has reached convergence or if $k > \text{MaxIterations}$; otherwise go to Step 2.

The algorithm defines an iterative procedure where players adapt the cost of energy bid by applying a best strategy response with regard to all the other users' actions during the previous iteration, presuming players update their action variables simultaneously for every time instants $k = \{0, 1, 2, \dots\}$. The ADDSM algorithm falls in the class of *minimal information environments* as each player needs not be aware of the exact strategy selected by all other players, since their aggregated bids will actually determine the designated amount of available energy corresponding to each player. Based upon this information each player revises its action variable c_j . The procedure is terminated once all players have converged to a state where their changes in the profile function become less than a predefined bound, or an upper limit in the number of iterations has been reached, thus the convergence criterion. It is evident that in case of convergence, the algorithm will reach the Nash equilibrium, however it is still to be determined whether the algorithm always converges given that the Nash equilibrium exists, as has been demonstrated in Section 5.1.2.

Proposition 3. The ADDSM algorithm converges to the Nash equilibrium for the proposed non-cooperative game Γ .

Proof. We begin by defining the set-valued function $B : \mathcal{A} \rightarrow \mathcal{A}$ as $B(c) = \times_{i=1}^n \bar{c}_i$, where \bar{c}_i represents the best-response function of player i as defined in eq. 5.10, which results from solving eq. 5.6:

$$\bar{c}_i = \bar{\Phi}_i\left(\sum_{j=1, j \neq i}^n c_j + c\right)$$

where $\bar{\Phi}_i$ denotes the unique solution of the same eq.:

$$\frac{\partial}{\partial c_i} \left(\frac{f(\Phi_i)}{c_i} \right) = 0 \Leftrightarrow \Phi f'_i(\Phi) - f_i(\Phi) = 0.$$

Theorem 2 (*Kakutani fixed-point theorem [76]*). Let X be a compact convex subset of \mathcal{R}^n and let $f : X \rightarrow X$ be a set-valued function for which:

- (a) $\forall x \in X$ the set $f(x)$ is non-empty and convex
- (b) the graph of f is closed

Then there exists $x^* \in X$ such that $x^* \in f(x^*)$.

Now, the conditions of (a) have been proved in the demonstration of Proposition 1. Secondly, our set-valued function B has a closed graph since each best response function $\bar{c}_i(k)$ is continuous. Thus, according to Theorem 2, B has a fixed point, which is the Nash Equilibrium of the game.

Theorem 3. $M(\cdot) : X \rightarrow X$ is a standard function if it satisfies

- (a) Positivity: $X \in \mathcal{R}_+^n$
- (b) Monotonicity: $\forall x, y \in X$, if $x \leq y$ then $M(x) \leq M(y)$
- (c) Scalability: $\forall x \in X, \forall \alpha > 1, M(\alpha x) < \alpha M(x)$

Notice that verifying conditions (a)-(c) of Theorem 3 for function B is trivial given the formula for \bar{c}_i , however we will need this result to demonstrate convergence.

Theorem 4. For a standard function $B(\cdot)$, if the iteration $x(t+1) = B(x(t))$ has a fixed point, then:

- (a) the fixed point is unique
- (b) the sequence $\{x(t)\}$ converges to the fixed point

Hence, given that B is a standard function according to Theorem 3 and that it has a fixed point (Theorem 2), we can conclude that the ADDSM algorithm converges to the Nash Equilibrium for game Γ . \square

Moreover, property (a) of Theorem 4 also implies:

Proposition 2. The non-cooperative game Γ has a unique Nash equilibrium.




Figure 5.7: Best response curves and equilibrium of a two-player game

5.1.4 Numerical results

An Example

In Figure 5.7 we illustrate graphically a simplified instance of the problem for a 2-player fictitious game $\Gamma = \langle \{P_1, P_2\}, \mathcal{A}, \{u_1, u_2\} \rangle$. In order to find the Nash equilibria of the game we need to analyze first the best response curves for the two players. That is, to associate for every action of player P_2 , player P_1 's best response. On the horizontal axis of Fig. 5.7 are plotted the actions of player P_1 , while on the vertical axis we plot the actions of player P_2 . For constructing the best response curve $P_1(P_2)$ of player P_1 for a given bid of player P_2 , c_2 , we determine \bar{c}_1 by solving equation 5.6, which translates to finding the bid c_1 that yields maximum utility for player P_1 . The point of intersection of the two best response curves represents the Nash equilibrium of the game, which boils down to solving the two best response equations jointly. The intersection is a single point, thus there is a unique Nash Equilibrium that specifies for both P_1 and P_2 their optimal cost (w.r.t. u_i), to which the players converge via iterative strategy updates following the ADDSM algorithm. This implies that by selfishly maximizing her utility function and without need for revealing it, in a distributed manner each player settles to an equilibrium point where it can no further gain an increase in utility through individual effort.




Figure 5.8: Realization of the scenario

Performance Evaluation

In this section we report on numerical results for the analysis presented in the previous sections. To begin with, for the simulation we have considered an illustrative scenario for accommodating 10 consumer agents, where Figure 5.8 shows one particular realization of the system. We assume consumers to represent a load in the system, each varying over the interval $[0,10]$ kW for every time slot in \mathcal{T} . As described in section 5.1.2, the sigmoid curves, respectively corresponding to each user, allow us to represent explicitly the user preference as the trade-off between the cost of energy and the satisfaction level experienced from the allocated amount of energy. For this particular realization of the scenario we run the ADDSM algorithm proposed and we looked firstly at the stationary states that the system reaches.

As expected, the numerical experiments demonstrate that within the context of a non-cooperative game, when each agent operates independently to maximize its utility, the solution, representing the system's day-ahead operational mode (the set of energy allocation to users), converges to a unique stationary state - the Nash Equilibrium. In Figures 5.9a and 5.9b we represent the evolution of user bids and the evolution of allocated energy for user 1 to 6 respectively. The algorithm proves to converge to an equilibrium point in a short number of iterations. Going back to

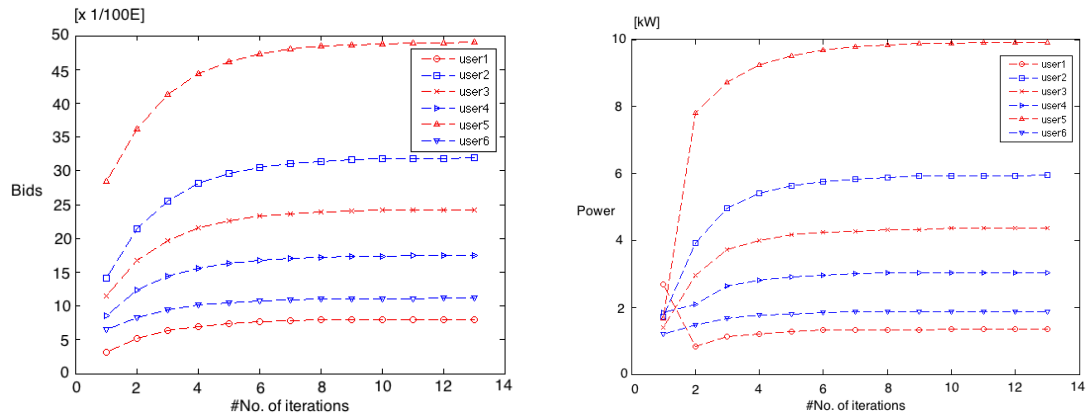


Figure 5.9: Evolution of user (a) bids and (b) power allocation


Figure 5.10: Average utility vs. power supply Θ

Figure 5.8 we can observe the relation between these results. It can be seen that the user with the highest demand is *user 5*. Then, it makes sense that at the equilibrium state, which the system reaches after running the algorithm, a significantly higher demand corresponds to a higher energy allocation as well as higher bids.

We now look at correlating the impact of the existing power supply versus the aggregated user utility. Gaining such an understanding is critical for the system operator, which can thus choose the power supply such that it may balance the user preferences with network load and cost of energy. Figure 5.10 shows the average utility of the system as a function of the energy availability in the system. The averaging is done over 1000 realizations for different initializations of the iterative algorithm.

Table 5.1: Cost of stability as a function of network load

$\Theta(kW)$	10	20	30	40	50	60	70	80	90
CoS (%)	1.7502	1.7505	1.7505	1.7507	1.7537	1.7543	1.7537	1.7528	1.7532

Additionally, for each run we determine the Nash equilibria of the game for values of the total amount of available energy ranging in the interval $[10, 100]$ kW. Finally, we plot the data of these last stationary states of the system as the final result. As the results show, the average utility of the system improves proportionally with the allocated energy. Hence the grid operator may directly derive an estimation of the energy supply's effect on the users' utility.

Finally, we address the notion of cost of stability (CoS) introduced in section 5.1.2. As previously described, the robustness of the game, in terms of a fair energy allocation amongst users, is a desirable property that comes with a price. Hence, we examine the worst-case loss of performance for the proposed ADDSM algorithm, by conducting a similar analysis repeated over 1000 realizations for each network load Θ considered. The results in Table 5.1 show that this allocation scheme returns worst-case solution values, which are very close to a complete energy distribution (within a 2% bound). Moreover, in practice, this can be translated in fact as a favourable property for our mechanism, as it provides a tolerance margin, considering the highly dynamic nature of the various intermittent and distributed energy resource to be integrated to the grid [140].

5.1.5 Discussion

A key challenge in the context of the *Smart Grid*, generically known as Demand-Side Management (DSM), aims at inducing a specific consumer behavior depending on available intermittent energy resources, in order to reduce peak energy demand and to efficiently integrate consumption from renewable sources. The pitfalls of having intervals of high peak demand include higher energy costs and increased carbon

emissions, apart from potentially destabilizing the grid and exposing it to the risk of cascading outages.

We classify existing DSM techniques according to two criteria of the solutions proposed: *centralization* and *cooperativeness*. Most of the work reported on DSM has adopted a centralized approach coupled with considering agents as participating in a cooperative environment [174, 57]. We regard such an approach as a simplifying assumption that overlooks consumer preferences and fails to represent the self-interested nature of consumers that strategize to minimize costs while maximizing their user satisfaction level. Moreover we advocate for fully distributed mechanisms that go beyond existing central control systems, which represent an impractical approach for coordinating complex large-scale settings.

In more recent work [184, 148] the authors are addressing these important issues. However, even though they avoid the need for a central controller that imposes a schedule on the system, they hardcode the decision mechanism at the agents' level, limiting severely the agents autonomy with regard to user preferences. Other related efforts that are confined to either centralized or cooperative environments have been proposed for managing demand. Amongst, we note [85], where the authors propose a statical hierarchical design for matching supply, which is centrally organized. Similarly, in [142] the authors introduce a dynamic tree structure where agents collaboratively perform local adaptations that propagate through each level of the tree, resulting in locally optimal solutions.

In Chapter 5.1, we have positioned our work at the opposite pole and only consider the case of non-cooperative users operating in a decentralized fashion. In order to do so, we design a non-cooperative day-ahead game where agents compete with each other to maximize their utility, which represents a specific balance of comfort levels over price. The outcome of the game results in an equilibrium state of the system, that is to be reached without implying any control over other agent's actions, nor requiring information regarding their preferences, that might expose them to strategic privacy invasions.

Noticeably, the underlying assumption of modelling the problem from a game-theoretic standpoint is that of considering that the utility of a consumer can be

revealed accurately in a function form that pertains certain properties. In Section 5.1.2 we have argued the versatility of our approach in characterizing the utility of the consumer through parameters that capture the user preference. However one may conceive it to be difficult for the average consumer to describe in a mathematical form the utility derived from using certain devices at various times. In this sense we point at possible extensions of this work, particularly in the machine learning domain. It would be especially useful to have a system capable to observe the consumption behavior of users and actually learn a utility functions without explicit human guidance.

To sum up chapter 5.1, we have applied game-theoretic concepts for one of the key arising challenges of the Smart Grid, Demand-Side Management. As the electricity demand is rapidly increasing, this puts increasing pressure on the power network requiring new mechanisms for accommodating peak usage. The key issue for enabling the network to work more efficiently is developing solutions, where consumers can mostly benefit from operating under a set of given constraints by adapting their load profile. We have represented the problem in the context of a multi-agent systems and formulated it in terms of a non-cooperative game, where agents optimize their individual objectives. Following, we derive the Nash equilibrium of the game proving the existence and uniqueness of such an operating point for our given setting. Finally, under this framework we introduce a decentralized algorithm for the proposed game that we analytically prove to converge to the Nash solution and give insights into a deployment of our mechanism through simulation.

NOTATIONS

Variable	Description	Variable	Description
\mathcal{A}	set of prosumer agents a_i	χ	set of agent actions
\mathcal{L}	set of deferrable loads l_i	w_{a_i}	discomfort cost
\mathcal{T}	set of time-slots t_k	val, h	thresholds
β	agent profile function	$\delta(t_k)$	changes in consumption
l_i	device power rate (in kW)	ν	coalition value
s	initial starting time slot per device	ν	coalition value
d	duration of load $lw_a(\mathcal{S})$	u	payoff distribution
α_c	corrective action	Δ	deferment
α	agent action	R	set of random variables r_i
α_S	joint action of coalition \mathcal{S}	Υ	elasticity of demand
\mathcal{S}	coalition	σ	bilateral Shapley value
p_c	probability for corrective action	\mathcal{G}	cooperative game
\mathcal{R}	reward function	$\mathcal{C}(\alpha_S)$	cost of coalition \mathcal{S}
\mathcal{P}	penalty function	$\mu(\mathcal{S})$	expected utility of \mathcal{S}
$\mu_a(\mathcal{S})$	expected utility of agent a in coalition \mathcal{S}	P	set of probabilities π_i

5.2 Dynamic Coalition Formation for Power Regulation

In this part of chapter 5 we focus on one particular area of the smart grid, namely, the challenges faced by distribution network operators in securing the balance between supply and demand in the *intraday market*, as a growing number of load controllable devices and small-scale, intermittent generators coming from renewables are expected to pervade the system. We introduce a multi-agent design to facilitate coordinating the various actors in the grid. The underpinning of our approach consists of an online cooperation scheme, *eCOOP*, where agents learn a prediction model regarding potential coalition partners and thus, can respond in an agile manner to situations that are occurring in the grid, by means of negotiating and formulating speculative solutions, with respect to the estimated behavior of the system. We provide a computational characterisation for our solution in terms of complexity, as well as an empirical analysis against the state-of-the-art mechanism, showing a performance improvement of about 17%.

Recent years have seen the advent of distributed energy resources (DERs) with particular emphasis for a cleaner generation of electricity, predominantly based on wind and solar power [140]. Albeit representing a sustainable form of energy, renewables pose a major challenge to current electricity networks due to their stochastic behavior. DERs are essentially characterised by small-scale, intermittent and highly unpredictable output. In this context, embedding such devices to the ageing infrastructure of distribution networks requires novel approaches for managing the grid efficiently [71, 57, 58]. Given this setting, the organization of the exchange electricity markets is also expected to change.

Currently, the majority of all power is being traded in what is known as the day-ahead spot market. Here, the following day is discretized over hourly time intervals and the market is cleared the day before, fixing the prices and volumes for the contracted amount of energy. In addition, shortages or excesses of energy are mitigated over the *intraday market*, which is cleared just before the actual power is delivered by

producers. Such circumstances may include (but are not limited to) compensating for errors in renewable energy forecasts, smoothing start-up ramps of conventional power plants, correct instantaneous mismatches between supply and demand and providing short-term contingency power in case of generator or transmission line failures.

Thus, as the network is becoming more reliant on the power generated by DERs, the role of the intraday market is expected to gain significant importance [59]. The goal is then to maximise the usage of clean energy upon its availability and maintain the delicate balance between supply and demand in real-time. In order to do so, demand should be able to adapt to the volatility in supply. This can be made possible assuming that consumers too can engage in an online, self-interested negotiation for shifting loads and thus adapting their demand. Moreover, the system ought to react in real-time to sudden changes of the aggregated generation profile in order to balance supply from intermittent renewable resources, while complying with consumer requirements. Here, we apply the multi-agent paradigm to devise a mechanism that enables local adaptability to dynamic situations at runtime and allows coordination, as opposed to the more complex task of centralised management [83]. We address the above-identified requirements by proposing a dynamic coalition formation (DCF) algorithm, where agents provide a bottom-up resolution for contingencies via a coordinated look-ahead response.

In more detail, the contribution of Chapter 5.2 is threefold:

1. Firstly, we provide a new representation of the power regulation problem by formalizing it in the context of dynamic coalitional games;
2. Secondly, we propose a distributed online protocol for solving this problem given its real-time constraints, where we integrate:
 - (a) a cooperation scheme that on one hand benefits from attractive economic properties and on the other hand is scalable and computationally tractable;
 - (b) prediction-based learning for reasoning about future interactions and states of the grid;
 - (c) privacy-preservation guarantees for non-intrusive negotiations;

3. Thirdly, we present an empirical evaluation of the approach against the state-of-the-art real-time pricing (RTP) mechanism.

The organization of the rest of chapter 5.2 is as follows. Section 5.2.1 introduces a new formalism for the intraday power regulation problem in terms of a dynamic coalition formation analysis. The challenge of an efficient payoff allocation procedure is addressed in Section 5.2.2, while Section 5.2.3 augments our approach in the context of privacy preservation. In Section 5.2.4 we put it all together and synthesize our coalition formation mechanism. Finally, Section 5.2.5 provides an empirical evaluation of our *eCOOP* scheme. Section 5.2.6 discusses our approach in the context of related work.

5.2.1 A Coalitional Game Formulation for Intraday Power Regulation

Currently, the grid operator is responsible for compiling the day-ahead schedule for power generation, that is explicitly passed to the actors in the grid. However, with the advent of renewable generation, these schedules are becoming volatile in nature, as they can be influenced by a wide variety of factors (e.g. wind speed, solar irradiance, consumer patterns, etc.), though their accuracy improves as the time-to-prediction elapses.

Henceforth, we take a standpoint where the grid operator, confronted with the uncertainty regarding both generation and consumption capacities, is running a continuous prediction of both supply and demand in the near future, in order to prepare for reductions in available supply or high-peak demand. We propose a mechanism owing to which, the grid operator can attempt to manipulate the behavior of these actors. Namely, once it determines that a *control action* needs to be executed, this information is published and becomes available to the actors in the respective region of the grid. Normally, due to the small capacity of individual actors, for obtaining a meaningful impact cooperation and coordination is required.

More formally, we represent *prosumers* as the set of self-interested agents $\mathcal{A} = \{a_i \mid 0 < i \leq n\}$ that always aim at maximizing their incurred gains. In doing so,

we associate with each consumer agent a_i the set of deferrable loads $l_j \in \mathcal{L}^i$ (in kW), operated over a discretized time schedule $\mathcal{T} = \{t_1, \dots, t_m\}$, by specifying their initial starting time slots s_j set by the user, their duration d_j , as well as the active periods for each load φ_j . For simplicity, we overload notation by denoting, for producer agents, with l_j the amount of energy to be generated during the interval d_j , starting at s_j . Against this background, we introduce the following definitions.

Definition 1. A *corrective action* is a tuple $\alpha_c = \langle t_i, t_j \rangle$, expressing the need to shift capacity from time slot t_i to t_j , without affecting the remaining time slots.

Definition 2. A *corrective action request* is a tuple $\langle \alpha_c, p_c, \mathcal{R}, \mathcal{P} \rangle$. The grid operator proposes *corrective action requests* by providing estimations that take the form of a probability distribution $\mathbb{P} : \mathcal{D} \rightarrow [0, 1]$, specifying the likelihood p_c of *corrective actions* $\alpha_c \in \mathcal{D}$ to be necessary. Additionally, functions $\mathcal{R}, \mathcal{P} : \mathbb{R} \rightarrow \mathbb{R}$ associate respectively, *monetary incentives* to be distributed amongst the members of the coalition that undertakes each task and *penalties* to be imposed for unfulfilled commitments, based on the capacity to be shifted.

Definition 3. Each agent $a \in \mathcal{A}$ is characterized by its baseline preferred consumption or generation, discretized over time slots $\mathcal{T} = \{t_1, \dots, t_m\}$ via the *profile function* β that aggregates its schedule:

$$\beta^{a_i}(t_k) = \sum_{l_j \in \mathcal{L}^i} l_j \varphi_j(t_k), \forall t_k \in \mathcal{T}$$

$$\varphi_j(t_k) = \begin{cases} 1 & \text{if } t_k \in [s_j, s_j + d_j] \\ 0 & \text{otherwise} \end{cases} \quad (5.11)$$

Now, we consider that each consumer agent a_i is characterized by a set of actions, which represent the shifting actions a_i is willing to take.

Definition 4. An *action* is a tuple $\alpha = \langle l, \Delta \rangle$ that specifies the potential deferment Δ of load l . For each agent a_i we denote its *flexibility domain* as the set of possible actions $\chi_{a_i} = \cup \alpha_j$.




Figure 5.11: Structure of the coalitional game

Essentially, an action produces an alteration to the initial profile β of the agent.

Definition 5. Function $\delta : \chi \rightarrow \mathbb{R}^T$ captures the changes in consumption for each time slot per action.

Definition 6. Let $\alpha_c = \langle t_i, t_j \rangle$ be a corrective action. Then an action α is *relevant* for α_c if the following holds:

$$\delta(\alpha, t) = \begin{cases} -q & \text{if } t = t_i \\ q & \text{if } t = t_j \\ 0 & \text{otherwise} \end{cases}$$

(5.12)

Definition 7. The discomfort cost $w_{a_i} : \chi_{a_i} \rightarrow \mathbb{R}$ quantifies the marginal loss of agent a_i in performing a particular action.

The business model (see Fig. 5.11) behind this approach implements a case-by-case monetary reward for each specific corrective action requested by the grid operator to consumers willing to participate. In principle, the reliability of the agents to carry out corrective actions should be the basis in committing the agents for such tasks. Importantly, in our approach the goal is in having the grid operator exempt from micromanaging the interactions with every agent individually. We address this issue by providing *reward* and *penalty* functions fit for this purpose. Explicitly, the reward function consists of two components (Equation 5.13), a *superadditive*⁴ function f and a *subadditive*⁵ function g . The threshold val specifies the point where increasing the capacity to be reduced by the agents is no longer desired by the grid operator.

$$\mathcal{R}(q) = \begin{cases} f(x) & \text{if } q < val \\ g(x) & \text{if } q \geq val \end{cases} \quad (5.13)$$

The penalty \mathcal{P} represents a superadditive function. Noticeably, while the reward incentivizes agents to perform joint actions in return for a higher pay, the penalty denotes a higher pay for failing to deliver this joint action. Thus, the problem of the grid operator in assessing the agents' reliability of actually delivering their actions is now being transferred to the agents that are incentivized to police themselves, with the scope of avoiding high penalties.

This models in effect a coalition game, where upon a corrective action request of a given probability (inline with Definition 2), agents can reallocate load usage over time schedule \mathcal{T} , in order to fulfil the corrective action and be eligible to collect the associated reward. Coalitions are formed based on the expected reward of the coalition and the individual costs that the agents incur in performing the actions. If the corrective action takes place and a coalition delivers the action α as promised,

⁴Suppose a_1 and a_2 can reduce demand with capacity q_1 and q_2 respectively. Superadditivity implies $\mathcal{R}(q_1 + q_2) > \mathcal{R}(q_1) + \mathcal{R}(q_2)$.

⁵Similarly, subadditivity implies $\mathcal{R}(q_1 + q_2) < \mathcal{R}(q_1) + \mathcal{R}(q_2)$

then the reward $\mathcal{R}(\alpha)$ is awarded to coalition. Contrary, if the coalition fails to deliver action α as promised the penalty $\mathcal{P}(\alpha)$ is to be imposed on the coalition.

Definition 8. A *coalition* is a subset of agents $\mathcal{S} \subseteq \mathcal{A}$ that agree on pursuing a set of actions $\alpha_{\mathcal{S}}$ called the *joint action* of coalition \mathcal{S} :

$$\alpha_{\mathcal{S}} \subseteq \bigcup_{a_i \in \mathcal{S}} \chi_{a_i}$$

Definition 9. Let $\alpha_c = \langle t_i, t_j \rangle$ be a corrective action. Let coalition \mathcal{S} *commit* joint action $\alpha_{\mathcal{S}}$ to a corrective action request α_C , producing a reduction of capacity q . Then, coalition \mathcal{S} is *compliant* if the following holds:

$$\sum_{\alpha_j \in \alpha_{\mathcal{S}}} \delta(\alpha_j, t) = \begin{cases} -q & \text{if } t = t_i \\ q & \text{if } t = t_j \\ 0 & \text{otherwise} \end{cases} \quad (5.14)$$

Definition 10. The *cost of coalition* \mathcal{S} sums up the discomfort costs for all actions performed by members of \mathcal{S} :

$$\mathcal{C}(\alpha_{\mathcal{S}}) = \sum_{\substack{\alpha_j \in \alpha_{\mathcal{S}} \\ a_i \in \mathcal{S}}} w_{a_i}(\alpha_j) \quad (5.15)$$

Definition 11. Let \mathcal{S} be a coalition with joint action $\alpha_{\mathcal{S}}$ that is compliant with $\alpha_c = \langle t_i, t_j \rangle$. The overall *coalition value* is computed based on whether the action $\alpha_{\mathcal{S}}$ has actually been delivered or not, by subtracting the discomfort cost of all coalition members from the given reward $\mathcal{R}(\alpha_c)$ or penalty $\mathcal{P}(\alpha_c)$, respectively:

$$\nu(\mathcal{S}) = \begin{cases} \mathcal{R}(\alpha_{\mathcal{S}}) - \mathcal{C}(\alpha_{\mathcal{S}}) & \text{if } \alpha_{\mathcal{S}} \text{ is delivered} \\ -\mathcal{P}(\alpha_{\mathcal{S}}) - \mathcal{C}(\alpha_{\mathcal{S}}) & \text{if } \alpha_{\mathcal{S}} \text{ is not delivered} \end{cases} \quad (5.16)$$

We are now in the position to define a number of key requirements for our power regulation protocol.

Requirements. Let $\langle \alpha_c, p_c, \mathcal{R}, \mathcal{P} \rangle$ be a corrective action requests dynamically initiated by the grid operator. Design a protocol where:

1. Agents self-organize to form a coalition structure \mathcal{CS} such that each coalition $\mathcal{S} \in \mathcal{CS}$ is *compliant* with the corrective action α_c
2. Determine a *payoff distribution* $u : \mathcal{A} \rightarrow \mathbb{R}$ that is:
 - (a) Individually rational *iff* $\forall a \in \mathcal{S} : u(a) \geq \nu(a)$
 - (b) Efficient *iff* $\sum_{a \in \mathcal{S}} u(a) = \nu(\mathcal{S})$
 - (c) Offers coalitional stability guaranties
3. Preserve data privacy w.r.t self valuation w_{a_i} of possible shifting action in χ_{a_i} during coalition negotiation

Example. Consider a 2-agent scenario, where a_1 's flexibility domain is represented by the action $\chi_{a_1} = \{\{\alpha_{a_1}^1 = \langle l_1, \Delta_1 \rangle, \alpha_{a_1}^2 = \langle l_2, \Delta_2 \rangle\}\}$, while for a_2 we denote χ_{a_2} as the actions $\chi_{a_2} = \{\alpha_{a_2}^1 = \langle l_3, \Delta_3 \rangle\}$. Function δ determines the modifications in consumption induced by these actions: $\delta(\alpha_{a_1}^1) = \{q_1, t_1 \rightarrow t_2\}$; $\delta(\alpha_{a_1}^2) = \{q_2, t_3 \rightarrow t_4\}$; $\delta(\alpha_{a_2}^1) = \{q_3, t_1 \rightarrow t_2\}$. For instances $\alpha_{a_1}^1$ means shifting the capacity q_1 from t_1 to t_2 . Suppose now the grid operator requires the corrective action $\alpha_c = \langle t_1, t_2 \rangle$. Consequently, the coalition of agents a_1 and a_2 could reduce consumption in t_1 with $q = q_1 + q_3$ and shift it to t_2 in compliance with the Grid's request.

5.2.2 BSV-Stable Payoff Distribution for Dynamic Environments

The starting point for this section gives further insight into modelling the rationale under which agents consider joining potential coalitions. It is important to realize that agents, representing both consumers and producers of energy in the grid, operate

within significant levels of uncertainty. We model a setting in which we consider the sources of uncertainty to be twofold. From the agent's perspective, on one hand the challenge is in accurately predicting its user's energy profile and preferences. On the other hand, in order to increase their coordination efficiency, agents need to build a prediction with regard to the expected behavior of potential coalition partners.

Definition 12. Given agent a 's estimation π of a joint action α_S actually occurring, the *expected utility of agent a in coalition S* is given by factoring in this probability:

$$\mu_a(\mathcal{S}) = \pi\mathcal{R}(\alpha_S) - (1 - \pi)\mathcal{P}(\alpha_S) - \mathcal{C}(\alpha_S) \quad (5.17)$$

Intuitively, the utility computation considers the expected coalitional reward, the expected penalty and the cost of performing the joint action.

Definition 13. Let \mathcal{S} be a coalition with joint action α_S . The *expected utility of \mathcal{S}* is the average over the individual utilities of the members $a \in \mathcal{S}$:

$$\mu(\mathcal{S}) = \frac{\sum_{a \in \mathcal{S}} \mu_a(\mathcal{S})}{|\mathcal{S}|} \quad (5.18)$$

Recall that we assume the grid operator to be providing estimations that take the form of a probability distribution $\mathbb{P} : \mathcal{D} \rightarrow [0, 1]$, that specifies the likelihood p_c of a corrective action $\alpha_c \in \mathcal{D}$ to be necessary. It is important to note that a corrective action will have different valuations for each agent. Agents will engage in a coalition formation procedure by playing the best response depending on their preferred strategy. Selecting a strategy, essentially boils down to a particular interpretation of the expected reward of pursuing a certain corrective action:

$$\alpha = \operatorname{argmax}_{\alpha_c \in \mathcal{D}} \mathbb{E}[\mathcal{R}(\alpha_c)] \quad (5.19)$$

Notice now that given the fact that corrective actions can only be estimated to occur, we have used for the strategy formulation the expected reward term, $\mathbb{E}[\mathcal{R}(\alpha_c)]$. Subsequently, each agent may adopt a different strategy according to its user's exposure to risk:

i) risk-neutral strategy: select the solution that maximizes the expected coalition reward: $\alpha = \operatorname{argmax}_{\alpha_c \in \mathcal{D}} p_c \mathcal{R}(\alpha_c)$

ii) risk-averse strategy: selects the solution over a restricted set of corrective actions with high probability for a given threshold h : $\alpha = \operatorname{argmax}_{\alpha_c \in \mathcal{D}} p_c \mathcal{R}(\alpha_c)$ if $p_c > h$

iii) risk-seeking strategy: selects the solution by favouring corrective actions with high monetary incentive, regardless of low probability of occurrence: $\alpha = \operatorname{argmax}_{\alpha_c \in \mathcal{D}} \mathcal{R}(\alpha_c)$

As previously detailed in Section 5.2.1, based on their expected utilities, agents engage in a coalitional game $\mathcal{G} = (\mathcal{A}, \mu)$. The solution of the game is a configuration $\langle \mathcal{CS}, u \rangle$ that specifies a *payoff distribution* $u : \mathcal{A} \rightarrow \mathbb{R}$ and a *coalition structure* \mathcal{CS} , which partitions the set of agents \mathcal{A} . According to the requirements of Section 5.2.1, the payoff distribution $u(a)$ is supposed to be individually rational, efficient and stable. This means that another aspect that needs to be addressed, concerns coming up with a payoff configuration that satisfies a notion of stability, implying that agents have an incentive for behaving in a certain way.

The payoff allocation scheme is resulting from running a negotiation procedure, where agents reschedule loads in order to meet the required constraints. Thus, considering the real-time constraints, for the payoff distribution, the protocol should minimize computational and communication demands. It is however well known that the classical stability concepts in coalitional game theory are of high computational complexity [126]. Consequently, for the payoff distribution we adopt an efficient version of Shapley value [161] introduced by Ketchpel in [79] and further developed in [28]:

Definition 14. The *bilateral Shapley value* $\sigma(\mathcal{S}_i, \mathcal{S}, \nu)$, $i \in \{1, 2\}$ in the bilateral coalition \mathcal{S} is equivalent to determining the Shapley value of \mathcal{S}_i in the game $(\{\mathcal{S}_1, \mathcal{S}_2\}, \nu)$:

$$\sigma(\mathcal{S}_i, \mathcal{S}, \nu) = \frac{1}{2}\nu(\mathcal{S}_i) + \frac{1}{2}(\nu(\mathcal{S}) - \nu(\mathcal{S}_k)) \quad (5.20)$$

with $k \in 1, 2$, $k \neq i$.

Notice that this can be rewritten such that the surplus of joining \mathcal{S}_1 and \mathcal{S}_2 into \mathcal{S} is distributed equally among \mathcal{S}_1 and \mathcal{S}_2 :

$$\sigma(\mathcal{S}_i, \mathcal{S}, \nu) = \nu(\mathcal{S}_i) - \frac{1}{2}(\nu(\mathcal{S}) - \nu(\mathcal{S}_1) - \nu(\mathcal{S}_2)) \quad (5.21)$$

Now, given two disjunct coalitions \mathcal{S}_1 and \mathcal{S}_2 , their union \mathcal{S} is called a *bilateral coalition*, while $\mathcal{S}_1, \mathcal{S}_2$ are subcoalitions of \mathcal{S} . In order for a bilateral coalition \mathcal{S} to be *recursively bilateral* it needs to represent the root node of a binary tree $T_{\mathcal{S}}$ for which *i)* every non-leaf node is a bilateral coalition and its subcoalitions are its children and *ii)* every leaf-node is a single-agent coalition. It follows then that a coalition structure \mathcal{CS} is *recursively bilateral* iff $\forall \mathcal{S} \in \mathcal{CS}$: \mathcal{S} is recursively bilateral or $\mathcal{S} = a, a \in \mathcal{A}$.

Definition 15. Given a game $\mathcal{G} = (\mathcal{A}, \nu)$ and a recursively bilateral coalition structure \mathcal{CS} , a payoff distribution u is called *recursively bilateral Shapley value stable* iff for each $\mathcal{S} \in \mathcal{CS}$, every non-leaf node \mathcal{S}^* in $T_{\mathcal{S}}$: $u(\mathcal{S}_i^*) = \sigma(\mathcal{S}_i^*, \mathcal{S}^*, \nu_{\mathcal{S}^*})$, $i \in \{1, 2\}$ with $\forall \mathcal{S}^{**} \subseteq \mathcal{A}$:

$$\nu_{\mathcal{S}^*}(\mathcal{S}^{**}) = \begin{cases} \sigma(\mathcal{S}_k^p, \mathcal{S}^p, \nu_{\mathcal{S}^p}) & \text{if } \mathcal{S}^p \in T_{\mathcal{S}}, \mathcal{S}^* = \mathcal{S}^{**} = \mathcal{S}_k^p, k \in \{1, 2\} \\ \nu(\mathcal{S}^{**}) & \text{otherwise} \end{cases} \quad (5.22)$$

Our aim is to find a recursively bilateral coalition structure \mathcal{CS} for game $\mathcal{G} = (\mathcal{A}, \mu)$, as well as a payoff distribution u that is recursively bilateral Shapley value stable. Notice that such a solution can be constructed incrementally through a bilateral merging process, where the intermediary coalition value is computed according to Equation 5.22.

Example. (see Fig. 5.12) Consider the following 3-agent scenario (\mathcal{A}, ν) with $\mathcal{A} = \{a_1, a_2, a_3\}$, where we demonstrate the calculations for the payoff distribution using the bilateral Shapley value:




Figure 5.12: Example of generating payoff configuration through bilateral Shapley value

- $\nu(a_1) = 1; \nu(a_2) = 0.5; \nu(a_3) = 0.5;$
- $\nu(\{a_1, a_2\}) = 4; \nu(\{a_2, a_3\}) = 2; \nu(\{a_1, a_3\}) = 2$
- $\nu(\{a_1, a_2, a_3\}) = 5.5$

It follows that merging into coalition $A = \{a_1, a_2\}$ and then into coalition $B = \{a_1, a_2, a_3\}$ yields the following payoff distributions: $\sigma(A, \{B\}, \nu) = 4 + 1/2(5.5 - 0.5 - 4) = 4.5$; $\sigma(\{a_3\}, \{B\}, \nu) = 0.5 + 1/2(5.5 - 0.5 - 4) = 1$. Similarly, the payoff of A is distributed recursively into $\sigma(\{a_1\}, A, \nu) = 2.25$ and $\sigma(\{a_2\}, A, \nu) = 1.75$.

5.2.3 Privacy-Preserving Layer

The cooperation scheme that we propose is run distributively among agents representing various actors in the grid, requiring that individual valuations, such as discomfort costs and the agents' utilities to form coalitions, need to be communicated between them. This implies that sensitive information will become distributed among numerous agents, without transmitting the data to a central (trusted) site. Thus, in order to avoid the possibility of malicious agents attempting to learn information about other agents, in particular, the discomfort costs of their actions, our scheme is to incorporate cryptographic primitives in order to perform *secure multi-party computations*. Evidently, there are numerous additional ways in which disclosure of energy consumption data may negatively impact consumers. As the granularity of the data collected and transmitted over the smart grid increases, privacy preservation is becoming an imperative concern [105, 149]. Primarily, protocols should prevent that

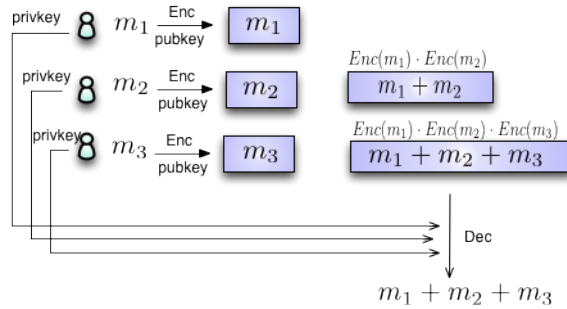


Figure 5.13: Example of homomorphic cryptosystem

private behavior could be derived, which may reflect personal routines, when a location is occupied, work schedules, or other information regarding occupant activities and lifestyle.

To address this, we look at *homomorphic* encryption schemes, which make it possible to perform operations directly on cyphertexts that correspond to operations on the initial cleartext messages. This means that despite that an agent cannot decrypt any of the individual messages received, he can however aggregate the messages using the homomorphic property and ask a subset of the sending agents to help it decrypt the result. Specifically, we are interested in applying an efficient additive homomorphic encryption scheme (see in Fig. 5.13).

Let $(pubkey, privkey)$ be a pair of public and matching private keys, $Enc(pubkey, m)$ a function that encrypts message m using the public key $pubkey$ and $m = Dec(privkey, Enc(pubkey, m))$ denotes the corresponding decryption function using the private key $privkey$. Then a public key cryptosystem with homomorphic property satisfies:

$$Enc(pubkey, m_1) \cdot Enc(pubkey, m_2) = Enc(pubkey, m_1 + m_2) \quad (5.23)$$

An efficient instantiation of such a scheme is the Paillier cryptosystem [129], which provides a fast encryption and decryption protocol, the details of which we leave aside due to space limitations.

Moreover, the public key cryptosystem is semantically secure, meaning that it is infeasible for a computationally bounded adversary (probabilistic polynomial-time adversary) to derive significant information about a message, when given only its

ciphertext and the corresponding public encryption key. More formally, this means that for any function f and plaintext m , the probability of guessing $f(m)$ does not increase, for a hypothetical adversary with polynomial resources, if he knows a ciphertext corresponding to m . Nevertheless, a coalition can compute the aggregate over these functions, which can then be revealed using the private shares of each coalition member.

5.2.4 *eCOOP*: Putting it all together

In the following we summarize the main steps outlined in the previous sections by giving the pseudocode representation of our online cooperation scheme. As detailed in Section 5.2.1, the grid operator holds the responsibility of monitoring the grid at large, in preparation for various instances of fluctuations, high-peaks, line overloads, reduced DER generation, etc. As a precautionary measure, the grid operator dynamically updates and publishes a list of corrective actions. The agents representing actors on the level of the low voltage grid need to coalesce in order to perform the actions indicated by the grid operator. According to the diagram in Fig. 5.11, that depicts the overall structure of the game in Section 5.2.1, the coalition formation procedure introduced hereafter corresponds to the second stage.

The *eCOOP* algorithm is run by every agent in the system. The starting point for each agent is in inspecting the global list of corrective actions provided by the grid operator, along with the probability of their occurrence. According to the user prescribed strategy of the agent, a set of target events in *EventQueue* is selected from *CorrectiveActionsList* (line 3), which induce a set of goal-oriented cooperative games that are solved concurrently. Next, for each target event the algorithm iteratively attempts to construct feasible coalitions starting from the initial set of singleton coalitions. A coalition represents an agreement between a group of agents for a successful resolution of a corrective action solicited by the grid operator. Based on the information exchange (line 30-34), each coalition computes internally the expected utility of a bilateral merger with a potential coalition partner, by including individual evaluations of past collaborations, which are reflected in the computation of the utility

μ of a coalition in a given coalition merger. Then, potential coalition formations are simulated via mergers of subcoalitions by computing the coalition value as the mean of the expected utilities of the merging coalitions (line 37). Following the assessment of potential coalition partners, for a particular candidate set, proposals are opportunistically advanced (line 41-52). The procedure terminates once the algorithm converges on a particular coalition structure. Finally, once the corrective action has been performed by the coalition the reward is distributed according to the BSV computation for that particular configuration (line 22), resulting in coalitions with stable payoff distributions. Specifically, once the event has elapsed, according to Equation 5.16, depending on the compliance or non-compliance with the corrective action, a reward or a penalty is determined respectively. The amount is then distributed down the coalitional tree based on the expected coalitional utilities μ (Definition 13) that were used in generating the tree structure. Additionally, agents update their probabilistic model (values of π) with the information inferred from the result of the coalition formation.

As we have established, we assume that the agents representing *prosumers* in the grid act selfishly, therefore, during the negotiation procedure for coalition formation, information about agents' profile must remain confidential. Firstly, this is achieved by communicating to potential coalition partners only a restricted set of actions that an agent is willing to take, instead of its complete profile. However, as this information represents the objective of negotiation, revealing it may expose agents to strategic behavior, in addition to the obvious risks of sharing detailed energy profiles (see Section 5.2.3). Our algorithm employs a homomorphic cryptosystem that allows agents to perform data aggregation without requiring that the data is decrypted beforehand. That is, agents can only determine the coalition value, instead of the individual preferences.

In the following we give the agent program of a leader agent a_i in a coalition, where a leader is determined by lexicographic order.

The complexity of the proposed DCF algorithm is given in the following propositions.

Algorithm 5**Data:** $\chi_{a_i}, \mu_{a_i}(S), \forall S \subset \mathcal{A}$

```

1: procedure ECOOP
2:   Update(CorrectiveActionsList)
3:   Select target event set EventQueue from CorrectiveActionsList according to agent strategy in Eq. 9
4:   for all target  $T_i \in \text{EventQueue}$  do
5:     INITIALIZE
6:     repeat
7:       if  $\exists S \in CS_{iter}$  so that  $a_i = \text{Lead}(S)$  then
8:         Det. relevant action set  $\alpha_S^{T_i} \subseteq \chi_S$  such that Eq. 2 holds  $\forall \alpha \in \alpha_S^{T_i}$ 
9:          $\alpha_S = \text{MergeActions}(\alpha_S^{T_i})$ ;  $\text{Candidate} = \emptyset$ 
10:        if  $\alpha_S$  not null then
11:          for all  $S' \in CS_i \setminus \{S\}$  do
12:             $\tilde{S} = S \cup S'$ 
13:            COMMUNICATE( $S', \tilde{S}$ )
14:            SIMULATE( $S', \tilde{S}$ )
15:          end for
16:          BILATERAL NEGOTIATION(Candidate)
17:        end if
18:         $iter := iter + 1$ 
19:      else break
20:    end if
21:  until Convergence( $CS$ ) or  $iter = \text{card}(\mathcal{A})$ 
22:  Compute recursively payoff vector  $u(C)$  for all  $C \in T_C$  as in Eq. 12 given  $\mu(C)$ 
23:  end for
24: end procedure
25:
26: function INITIALIZE
27:    $iter = 0$ ;  $CS_{iter} = \{\{a\} | a \in \mathcal{A}\}$ ;
28: end function
29:
30: function COMMUNICATE( $S', \tilde{S}$ )
31:   Send( $S, S', [\alpha_S; \text{Enc}(\mathcal{C}(\alpha_S))]$ ) Receive( $S, S', [\alpha_{S'}; \text{Enc}(\mathcal{C}(\alpha_{S'}))]$ )
32:   Aggregate  $\mu_S(\tilde{S})$  as in Eq. 8 based on  $\mu_a(\tilde{S})$ , for all  $a \in S$  using Homomorphic Scheme
33:   Send( $S, S', \text{Enc}(\mu_S(\tilde{S}))$ ) Receive( $S, S', \text{Enc}(\mu_{S'}(\tilde{S}))$ )
34: end function
35:
36: function SIMULATE( $S', \tilde{S}$ )
37:   Compute  $\mu(\tilde{S}) = \frac{\mu_S(\tilde{S}) + \mu_{S'}(\tilde{S})}{2}$ 
38:    $\text{Candidate} := \text{Append}(\text{Candidate}, S')$ 
39: end function
40:
41: function BILATERAL NEGOTIATION(Candidate)
42:    $S^* := \text{MaxValue}(\text{Candidate})$ 
43:   Send( $S, \text{Lead}(S^*), \text{MergeProposal}$ )
44:   if Receive( $S, \text{Lead}(S^*), \text{Agree}$ ) then
45:     Inform coalition members of forming  $S^* = S \cup S'$ 
46:   else  $\text{Candidate} := \text{Candidate} \setminus \{S^*\}$ ;
47:     if Candidate not null then
48:       BILATERAL NEGOTIATION(Candidate)
49:     else break
50:   end if
51: end if
52: end function

```

Proposition 1. *The computation complexity of the algorithm is $\mathcal{O}(pn^2m)$, where we denote with $n = |\mathcal{A}|$, $m = \max_{S \in CS} \{|\alpha_S|\}$, $p = \max\{|EventQueue|\}$.*

Proof. The number of iterations that the algorithm needs to cycle through is bounded by a) the maximum number of events in the global queue $\mathcal{O}(p)$ (line 6); b) the maximum number of coalition mergers that may occur $\mathcal{O}(n)$, which corresponds to the formation of the grand coalition (line 9); c) $\mathcal{O}(nm)$ the maximum number of operations required in order to construct the list *CandidateL*. Besides, the secure multi-party computation requires performing an encryption for every sent message, while the destination agent is needed to add the corresponding decryption. Hence, the overall complexity of the algorithm is $\mathcal{O}(p)\mathcal{O}(n)\mathcal{O}(nm) = \mathcal{O}(pn^2m)$. \square

Proposition 2. *The communication complexity of the algorithm in the number of messages per agent is $\mathcal{O}(mnp)$.*

Proof. During each run of the algorithm the number of messages sent by an agent is bounded by $\mathcal{O}(n) + \mathcal{O}(m)$ for the case of coalition representative agents, corresponding to inter-coalition negotiations and intra-coalition message passing respectively. Otherwise, a single message specifying μ_a is required to be sent to the coalition leader. In addition to this message, due to the usage of the cryptographic layer, an extra message per agent for every iteration is necessary for computing and sending the private shares of the coalition members to its leader. Thus, given at most pn rounds of the algorithm, the overall number of messages sent by an agent is $\mathcal{O}(mnp)$. \square

5.2.5 Empirical Evaluation

Experimental Set-up

In order to evaluate the performance of our proposed algorithm, experiments were conducted on real datasets obtained from the Australian Energy Market Operator (AEMO)⁶. It is important to note that AEMO centrally coordinates the dispatch

⁶<http://www.aemo.com.au/Electricity/Data/Price-and-Demand/Aggregated-Price-and-Demand-Data-Files>




Figure 5.14: Load curve for an individual household

procedure via a real-time pricing (RTP) scheme, by pooling the quantities of electricity required by consumers from available generators. Specifically, the dataset used for our first set of experiments archives price and aggregated demand data covering the month of September 2012 for each hourly slot, for the NSW region. While no detailed data was available on individual consumers, we infer this information and construct the agents' profile β by disaggregating the total demand. In doing so, we fix the number of agents to $N = 2252K$, derived from the number of households⁷ in the NSW region. In Figure 5.14 we plot the real consumption profile⁸ for a typical residential area. Next, based on this profile we generate stochastically, using a uniform distribution, new individual consumers that jointly match the initial aggregated demand of the AEMO. For our scenario, we used simulated consumption patterns for the number of N agents, where the consumption per agent per time slot is drawn from a uniform distribution $U(p_{min}, p_{max})$. We set the following parameters $p_{min}(t) = -0.15p(t)$ and $p_{max} = 0.15p(t)$, where $p(t)$ denotes the typical consumption at time slot t . We further assume that the strategies representing exposure to risk (Section 5.2.2) are equally represented in the consumer population and that the extent to which consumers are willing to reschedule demand by shifting loads is constrained to $\Upsilon = 25\%$, denoting the *elasticity of demand* as recent reports suggest [140]. For

⁷<http://www.abs.gov.au/AUSSTATS/abs@.nsf/DetailsPage/1338.1Dec%202010?OpenDocument>

⁸Available at UC Irvine Machine Learning Repository <http://archive.ics.uci.edu/ml/datasets.html>

associating shiftable loads to consumer profiles we generate loads with one time slot duration distributed uniformly over the set of time slots $\mathcal{T} = \{1, \dots, 24\}$, to match the given elasticity of demand. The number of time slots to which each load can be deferred is bounded to $\Delta \in [-5, 5]$.

Corrective actions are triggered by the grid operator for each timeslot where the ratio between the average consumption and the estimated consumption is outside the interval $[0.5, 1.5]$. Additionally, in order to give a measure of *robustness* for our approach, we factored into our simulation random variations in the power supply, accounting for fluctuations from renewable resources, which are estimated to cover about 13% of the total generation [140]. The mean absolute percentage deviation (MAPD) is bounded to an extent of at most 20%. These may also represent the cause for requesting corrective actions in case the abovementioned triggering condition is met. Also we consider that coalitions perform joint actions successfully with a 90% probability.

Predictive Model

The aspects of building an estimation model regarding potential coalition partners, based on previous encounters, as well as the agent's own estimated user behavior has been addressed in Section 5.2.1. In our experiments, we address both aspects in a unified approach by including sources of uncertainty in the form of random, uncontrollable variables with probability distributions, that each agent attempts to learn in an online fashion. Recall that for each agent $a \in \mathcal{A}$ there corresponds a set of (deferrable) loads \mathcal{L}^a . Essentially, the goal is in learning for a given action α^{l_j} , that shifts a load l_j , the likelihood that the shift occurs to a particular timeslot k . Suppose now agent a wants to determine the likelihood for the actions that constitute its flexibility set χ_a . Let $R = \{r_1, \dots, r_{|\chi_a|}\}$ denote the set of random variables modelling future, uncontrollable events and $\mathcal{D} = \{D_1, \dots, D_q\}$, a set of domains for the random variables such that r_i takes values in $D_i = \mathcal{T}$. Let $\sigma : R \rightarrow \chi_a$ be a distribution function of random variables to the agent's actions. Agent a learns $P = \{\pi_1, \dots, \pi_{|\chi_a|}\}$, which is a set of probability distributions for the random variables,

where each distribution $\pi_i : D_i \rightarrow [0, 1]$ defines the probability law for random variable r_i , so that the values of π_i sum up to 1.

Also, there is uncertainty regarding the expected behavior of potential coalition partners, which in turn needs to conform to their respective user demands in a timely fashion. Similarly, agent a tracks past encounters with other agents and builds a probability set P_i for each agent a_i . Consequently, we exploit the repeated game structure of the problem to learn a prediction model regarding future interactions and thus infer potential synergies between agents.

In order to compute the set of probabilities P , for the sake of clarity we adopt the *fictitious play* learning model⁹, where agents observe other agents', as well as their own user behavior. Concretely, for the latter case, the fictitious play requires that agent a models the set of random variables r_i by keeping, for each action of its user $\alpha^{lj} \in \chi_a$, a count $c_{\alpha_k}^j$ for each timeslot k :

$$\pi_{\alpha^{lj}}^k = \frac{c_{\alpha_k}^j}{\sum_i c_{\alpha_i}^j} \quad (5.24)$$

Of note is the fact that particular actions may be enforced by the user by setting the prior counts of the distribution. The same procedure holds for tracking agents that a has been previously exposed to during preceding runs of the algorithm. Moreover, for computing the probability of a joint action α_S , we average over the individual probabilities of each action $\alpha \in \alpha_S$:

$$\pi_{\alpha_S} = \frac{\sum_{\alpha \in \alpha_S} \pi_{\alpha}}{|\alpha_S|} \quad (5.25)$$

Results

Based on this data, we ran a comparison of the *eCOOP* algorithm against the existing RTP mechanism implemented in the Australian market. Results from these

⁹Of course, more complex functions could be considered, but this is beyond the scope of this work.




Figure 5.15: Comparison over aggregated demand patterns

experiments are shown in Figure 5.15, where we plot the average daily consumption patterns (in MW) for the given period. Based on our numerical experiments we can conclude that our coalition-based approach leads to a significant flattening of the energy consumption curve, as opposed to the RTP solution, although the overall consumption is maintained the same. Intuitively, Figure 5.15 clearly shows that by applying our proposed algorithm, ahead of critical peak periods, demand can efficiently adapt so that such instances are being prevented from occurring. In order to give a more quantitative measure for our results we consider the *load factor* metric [188], which represents the ratio of average power demand to the maximum (peak) demand. One of the key challenges behind bringing about the *smart grid* vision is particularly related to the improvements of load factors. Using this metric as an indicator of operational efficiency we can measure the disparity of the peak from average usage. Thus, the flattening of the demand curve corresponds to an increase of the load factor toward unity. For the one-month interval we have considered in our experiments, our approach produces a 14% increase of the load factor from 0.77 for RTP scheme to 0.91 when applying the *eCOOP* algorithm.

The second set of experiments are designed to speculate about future scenarios, when due to a wide adoption of plug-in electric vehicles (PEV) as well as electrification of heating, the proportion of shiftable consumption may increase significantly.

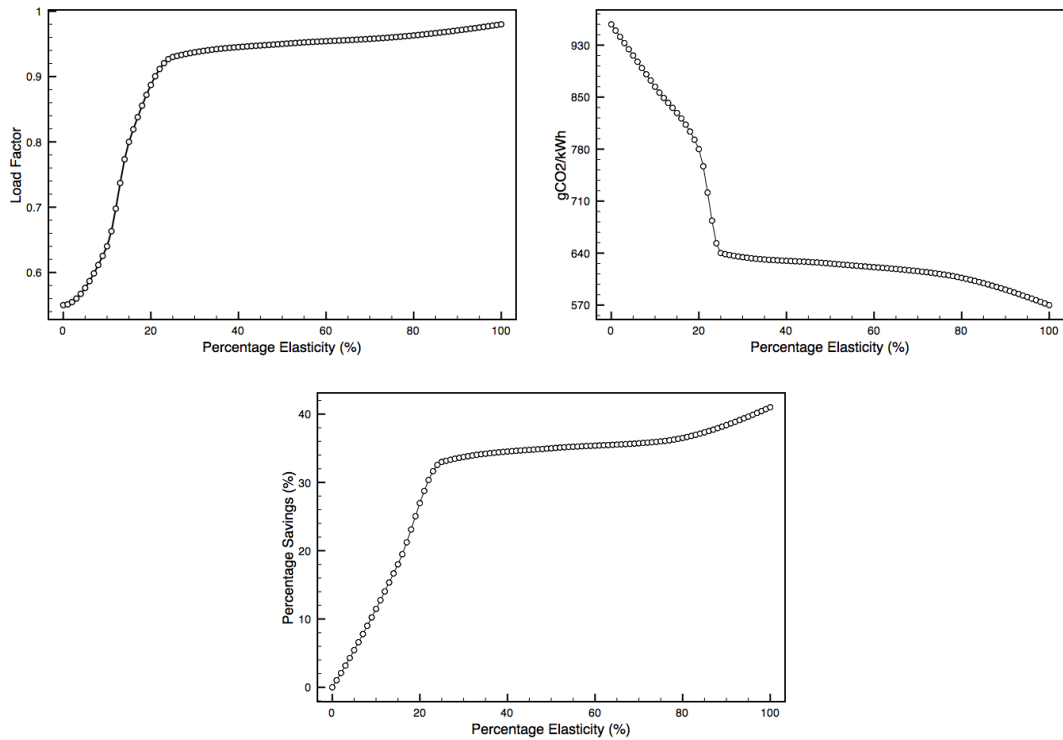


Figure 5.16: Evolution of (a) load factor, (b) carbon emissions and (c) percentage savings for different degrees of elasticity of demand

Specifically, we are interested to evaluate the impact of our mechanism for variations of the *elasticity of demand* Υ , by which we have denoted the percentage of energy a consumer is willing to defer upon an incoming request from the grid operator. In Figure 5.16 (a) we plot the effect of different values of Υ upon the load factor of the system. We have already seen from our initial setting very promising result regarding this parameter even for a moderate elasticity of 25%. A considerable increase in elasticity above this value is thus expected to provide negligible increases of the load factor, below 10%. Importantly though, the experiments show that against a consumer with zero elasticity, even a small increase can produce a large impact on the load factor, which is a very encouraging result.

Next, we investigate another criteria for evaluating the energetic efficiency of the system, in close relation to the previous. Hence, a high value of the load factor means a decrease usage of peaking plants and as a result, a lower carbon footprint. In other

words, a reduced consumption during peak intervals and an overall flatter demand means that energy can be generated from less polluting sources. The amount of emissions is in direct correspondence to the energy mix required in order to satisfy the aggregated demand. We use the AEMO dataset¹⁰ to determine the correspondence between a load factor value and the induced CO_2 emissions. Based on these findings, Figure 5.16 (b) shows the decrease in the amount of carbon emission per kWh by applying the *eCOOP* algorithm for various degrees of elasticity in demand. Not surprisingly, a similar pattern can be observed representing a steep reduction of emissions for small increases in elasticity for up to an approximate 20%.

Lastly, avoiding the need to deploy peaking plants can be directly translated into consumer savings. In Figure 5.15 we have represented the evolution of real-time pricing according to the given aggregated demand, provided by the AEMO dataset. Observably, off-peak intervals are correlated with lower prices, while higher prices correspond to peak periods. We now look into the peak timeslots, when corrective actions are requested by the grid operator and compute the aggregated costs of consumers for the RTP and *eCOOP* scheme respectively, based on the correlation between price and aggregated demand in our dataset. We plot the results in Figure 5.16 (c), where we give an estimation of percent savings incurred by consumers during peaking periods for varying degrees of elasticity. This is again an encouraging result, showing an approximate 30% reduction of kWh cost in return for an elasticity of 20%. Further flexibility in consumption can lead to a reduction of up to 40%.

Finally, we provide conclusive results for the performance of our algorithm, demonstrating how *eCOOP* outperforms the existing RTP scheme, evaluated over an extended year-round scenario. It is important to note that consumption patterns vary throughout the year. Specifically, winter and summer months are known to exhibit increased high-peak intervals due to an intense usage of electricity. We started our experiments investigating an average consumption month. For generality, we now give in Figure 5.17 the year-round results based on the load factor computation for the two approaches under consideration. It is interesting to observe that RTP produces differ-

¹⁰<http://www.aemo.com.au/Electricity/Settlements/Carbon-Dioxide-Equivalent-Intensity-Index>




Figure 5.17: Simulation results for year-round load factor comparison

ent outcomes depending on the particular period of the year, highly correlated to the expected consumption usage. For instance, on the one hand, the lowest efficiency is observed during January with a load factor value of 0.6 and on the other hand, April and November represent the highest efficiency months. In contrast, *eCOOP* consistently manages to attain a higher efficiency of an approximately 0.9 load factor value, invariantly of the period considered, while producing an average 17% improvement.

5.2.6 Discussion

Similar to our work, multi-agent systems have been proposed in the smart grid domain for the task of demand-side management in a number of studies [173, 184, 148]. Critical peak pricing or spot pricing mechanisms attempt to incentivize agents to adapt their demand, by reducing consumption during peak times [117]. Of course, this may end up in situations where peaks are only temporarily flattened and then shifted to different time intervals, as some of the research has shown, [167, 148]. More sophisticated solutions have proposed a game-theoretic framework, [111, 24], for a coordinated adaptation of the agents' behavior.

Power regulation is however distinct in that the objective of a corrective action is well defined and localized to a particular region of the grid. Peakload and contingency periods are typically handled by means of adapting the power supply, by firing ex-

pensive, carbon-intensive, peaking plant generators. Here, the grid operator provides a request for a specific power regulation action that needs to be addressed in a timely fashion. While demand-side management may be regarded as a day-ahead scheduling problem, for grid regulation, the response time is constrained within minutes, or to even a couple of seconds.

Due to these challenges, the majority of current methods have been limited to propose solutions that can only be applied in the day-ahead market. For instance, confining the space possible actors, in [180], the authors explore the idea of coupling generation from wind farms with storage facilities, particularly batteries from electric vehicles. Furthermore, the approach follows the assumption of a hierarchical organization, where a group leader computes an optimised schedule to maximising profit, for a fixed number of given participants. In a similar approach, in [30], the authors consider a single owner for the entire system that allows the use of centralised control, based on dynamic programming scheduling.

More relevant to our context, in [77], the authors report some preliminary work on deploying electric vehicles (EVs) for power management in the grid. However, they restrict their study to a small-scale scenario, moreover, assuming centralized control over the set of EVs. This eludes some of the harder problems of operating within minimal information environments, where the assumptions of global knowledge and top-down control of centralization no longer hold.

More grounded approaches for the task of load management fall under the category of direct load control (DLC) as one of the best practices deployed so far to reduce system peak load by imposing a brute-force on/off strategy to control loads. In [189], the authors report on a study in cooperation with the Taiwan Power Company, where to achieve DLC, they use a multi-pass dynamic programming method to schedule the operation of air conditioner load in order to reach peak reduction and maximum cost savings. Another DLC scheduling solution is given in [118], where the goal is that of increasing the profit of the utility using a linear programming algorithm. Recognizing the importance of taking user preference and comfort into account, improved DLC solutions deploy a logic-based system to model by fuzzy variables the flexibility of interrupting the air-conditioner and electric water heater loads, in an attempt to

factor in user satisfaction [154, 192].

Clearly, the above-mentioned approaches pose a series of limitations. DLC methods, although practice assume full control over the consumer loads, which can be exercised at will. In a realistic scenario it is hard to imagine having consumers comply with such energy usage violations. Centralised solutions assuming a single owner of the system that has full control over the operation of all loads is again nonapplicable to instances where different stakeholders need to reach an agreement as to fulfil a given goal. Finally, applying various pricing schemes has also been shown to deliver poor results. For example, individual consumers may unilaterally decide to shift consumption from expensive time slots to cheap time slots, thus replacing peaks from one period to another. The problem here is due consumers *i)* not having a clear perception of the amount of energy that needs to be shifted, *ii)* having an interaction only with grid operator, while not being aware of the constraints and consumption preferences of other consumers and *iii)* not being able to opt in/out at will, dynamically, in the participation to energy management schemes. Moving towards a decentralized, agent-based setting of the electricity grid, we identified and resolved here a set of desiderata, that to best of our knowledge all current approaches fail to address.

To sum up, in Chapter 5.2 we were interested in mechanisms that can cope with an increasing amount of intermittent energy generated via renewable resources. We introduced the *eCOOP* agent-based algorithm, where look-ahead coalitional negotiations are run within minimal information environments in order to address the dynamism and uncertainty of the system. Furthermore, our protocol provides for computing an efficient payoff allocation scheme that guarantees stable coalitions, while satisfying privacy-preservation of sensitive data. Notably, our approach is guaranteed to achieve a security level of IND-CPA¹¹, which is the highest security level for homomorphic schemes. We have also provided an empirical evaluation of our approach based on real datasets and shown the advantages of using it in terms of increased grid efficiency.

¹¹IND stands for indistinguishability and CPA for chosen plaintext attacks [186]

It is important to point out that by design, the intervention of the grid operator addresses explicitly the shifting actions that consumer need to perform in order to collect the reward. In contrast to RTP schemes, this allows us to impose the necessary constraints so that by removing peaks we are not replacing them by new arising ones, which is also known as the *herding* effect. Moreover, the design of the reward function allows us to transfer the responsibility of determining the reliability of the agents to carry out corrective actions from the grid operator to the agents themselves.

In this work we have used a standard approach for computing the prediction model, namely fictitious play. In future work we plan to look into more complex models and asses their performance. Also, we are interested to evaluate our model in scenarios where consumers are not only willing to shift loads to different time intervals given monetary incentives, but may additionally be considering to reduce their total consumption given that a certain revenue could be attained. Expectedly, this is ought to further flatten demand and thus, increase the overall efficiency of the grid especially during periods when generation from renewables is highly fluctuating. Unfortunately, specifying this sort of parameters, such as the threshold in revenue to which consumers may react and the extent to which their consumption behavior may be altered remains an open question.

5.3 Chapter Summary

This Chapter introduces a *consumer-centric* framework for operating the smart grid. The problem is twofold and needs to be addressed in correspondence to the organization of the electricity market. Thus, we provide an inclusive approach by dwelling in Chapter 5.1 on the problem of *Demand-side management*, while Chapter 5.2 is focused on *Demand Response*. The former is formulated in the context of the *Forward market*, where consumers adapt their consumption pattern to adhere to the requirements of the DNO by providing regular reductions in demand for certain periods. The latter is concerned with mechanisms applicable in the *Intraday market*, where consumers have to wait for critical periods and cooperate in real-time to provide demand reduction services, while DSM aims to prevent these periods from occurring in the first place. Overall, we are providing a novel framework for managing the grid efficiently by envisioning a game layer on top of the electricity grid infrastructure that allows us to implement a consumer-centric approach as one key driving factor for a new vision of the grid via increasing participation of customers in the energy field.

Chapter 6

Detection of *Collusive Behavior* in Energy Markets

Singularity is almost invariably a clue. The more featureless and commonplace a crime is, the more difficult it is to bring it home.

— *Sherlock Holmes* —

Fundamental changes in the electrical energy sector are drawing on serious implications. One key arising challenge regards current energy markets, which are undergoing a transformation towards accommodating a more decentralized and sustainable provision of energy. As the number of traders in the market is increasing steadily and the trading activities are becoming more complex, the energy markets are becoming more exposed to potential fraud. In this chapter we address the problem of detecting collusive behavior, where a group of individual traders act together, inconsistently with the competitive model, to artificially manipulate the market and elicit illegal profits. We investigate collusion attacks in the energy market and propose a novel mechanism, showing the effectiveness and practical applicability of our method to real scenarios.

Market surveillance represents a serious challenge and it refers broadly to the detection of abnormal market behaviors, which are known to be predominant especially in the emerging markets. In practice, an efficient way for influencing the market and

gaining illegal profits is represented by the class of collusion-based malpractices. Conceptually, collusive behavior represents an attempt of a group of individual traders, that act together, to artificially manipulate the market (e.g. through price or market share) for maximizing their gains, in a manner inconsistent with a competitive model and in detriment to the other participants in the market.

While collusion has been reported in various market domains, the damages discovered were averaged at about a 25% increase in costs incurred by costumers [92]. Unfortunately, most of the collusion cases that have been discovered thus far came as a result of investigations triggered not through economic analysis, but rather due to customers' complaints or suspicious competition complaints (e.g. stainless steel industry, graphite electrodes, facsimile paper).

Though there are many ways in which collusion could be discovered, we address in this work collusion detection via the analysis of economic data, freely available in the market. The set of laws that regulate the market require continuous surveillance of trading activities. This can essentially be achieved either through online or offline surveillance. The former is constrained to analysing short-term data as well as being restricted to a limited time-window, thus being prone to overlooking occurrences of more complex types of fraud. Alternatively, offline techniques can encompass a wider spectrum of illegal trading strategies, while having an anticipatory or retroactive character. Obviously though, both approaches remain dependent on the amount of input data available.

In this chapter, we propose an effective offline method to identify collusive groups with respect to the domain of energy markets, which are being reoriented towards replacing the traditional top-down energy supply with a decentralized, market-oriented provision [146, 108, 109].

Chapter 6 is organized as follows: in Section 6.1 we provide background details to this work. Section 6.2 introduces our novel mechanism and discusses how to apply it to detect collusive behavior in the energy market. In Section 6.3 we show the experimental results. Finally, Section 6.4 concludes and points directions for future work.

NOTATIONS

\mathcal{A}	the set of market agents a_i	\mathcal{P}_r	probability of CP
\mathcal{D}	the set of devices d_i	MSE	mean square error estimator
\mathcal{T}	the set of time-slots t_k	$C(X, Y)$	covariance
CS	coalition structure	$\rho(X, Y)$	correlation coefficient
b_i^j	bid of agent a_i timestamped at tmp_i^j	λ_i	eigenvalues
X_i^j	time-series, trading activity of agent a_i	v_i	eigenvectors
F_A, F_B	distributions	Λ, V	diagonal and orthonormal matrices
α	mean value of X	P	principal components
S_i	cumulative sum	p	price
r	change-point	v	volume

6.1 Energy Markets

Emerging financial markets are inherently exposed to malpractices, whereas a subset of traders collaborate tacitly to manipulate the market for maximizing their individual gains. In the course of the last years, new challenges in the *Electricity Market* have come to the forefront, due to the transformations occurring in the power supply infrastructure. As new distributed energy resources (DERs: e.g. wind plants, photovoltaics, combined heat and power units) are pervading the electricity grid, they provide for an increasing number of participants in the market. This trend has driven the liberalization of electricity markets and the creation of power exchanges with the emphasis on decentralized power provision.

More formally, the participants in the market, consumers and suppliers of power, can be denoted by the *market agent* set $\mathcal{A} = \{a_1, \dots, a_n\}$, which exists in a bijective relationship with the set of devices $\mathcal{D} = \{d_1, \dots, d_n\}$ connected to the grid. The day-ahead market is discretized over a nonempty and finite set of distinct and successive time periods $\mathcal{T} = \{t_1, \dots, t_m\}$. Under this setting, a bid of agent a_i timestamped at tmp_i^j is represented by the function $b_i^j : \mathcal{T} \rightarrow \mathbb{R} \times \mathbb{R}$, specifying respectively, for time slot t_i , the amount of electricity requested or offered and the intended price per unit, as $b_i^j(t_i) = (v_j, p_j)$. To summarize the trading activity for agent a_i at trading day tmp_i^j we associate the time-series $X_i^j = \{b_i^j(t_1), \dots, b_i^j(t_m)\}$, which captures the existing bids for each time-slot in \mathcal{T} , else being considered a null bid.

Several approaches have attempted to automate the process of collusion detection based on economic data. A common method would be to apply supervised learning techniques, given that proven fraudulent activities, previously detected, could be obtained. Of course such datasets for training are not usually available in significant amount. Thus, some work has looked into unsupervised learning, namely using graph clustering algorithms. In [74], the collusive marker that the authors base their detection mechanism on, addresses circular trading, where a group of market agents are trading heavily among themselves aiming to raise the price of their shares. A Markov clustering algorithm is introduced and applied to the stock flow graph, which summarizes a trading database. In [130], the authors have adopted cross trading as a

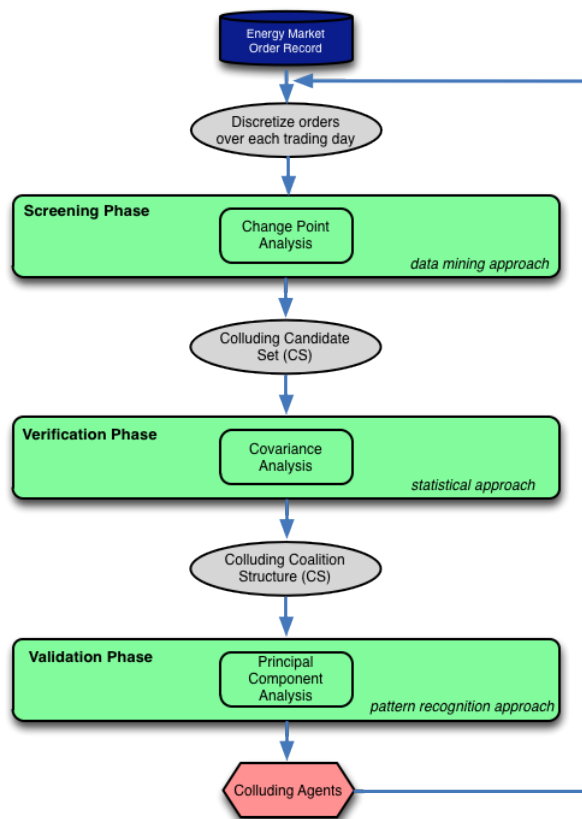


Figure 6.1: Mechanism flow-chart

collusive marker and compared the performance of different off-the-shelf clustering algorithms. A similar approach is introduced in [46], by means of employing a spectral clustering method.

In this chapter we address the phenomenon of collusion in the emerging energy markets where the collusive markers identified above do not hold, as novel domain driven markers need to be derived.

6.2 Collusion Detection Methodology

To start with, for discovering collusion in a market we need to be specific as to what behavioral patterns, that might be indicative of collusion, we will be looking for.

Our goal is to detect the presence of colluders that are coordinating their behavior at the expense of the rest of the market participants, by adopting a behavior inconsistent with what a competitive environment might entail. In this work, the *collusive marker* that we investigate in order to provide evidence of economic collusion consists of detecting colluders in the energy market based on the similar trading behaviors of agents. Colluders can generally be differentiated by similar trading patterns (which also depart from competition), as opposed to those outside their coalition. Thus, within a coalition of colluders their trading behavior should exhibit correlations when they should normally be independent.

Consequently, we design our collusion detection method as a three-stage process: a *screening phase*, a *verification phase* and a *validation phase*, as depicted in the diagram of Fig 6.1. The *screening phase* has the role of performing a triage through the *order record* of the market and identifying a set of market agents that are worthy of closer scrutiny. Due to the combinatorial explosion of possible colluder subsets and the data-intensive analysis, this phase is ought to output a set of candidates that will be further addressed during the *verification phase*.

Specifically, the *screening phase* is testing each market agent, looking for structural breaks in its behavior. Such behavioral breakpoint could be associated with the formation of a coalition of colluders (or with its dissolution). We approach this by running a *change-point analysis (CPA)* over the discretized order record of each market agent. If any behavioral breakpoints, which may be conducive to colluding coalitions, have been identified, they will represent the input for the second phase. Obviously such behavioral changes could be expected even when no collusion takes place. During *verification* we are essentially looking at two aspects: *i*) whether there is a price or volume correlation between the candidates' breakpoint and *ii*) if there appears inconsistencies with the competitive model. Finally, for the constructed solution, the *validation phase* seeks to reveal an underlying latent structure that can explain for the variability among observed, correlated behaviors.

While this methodology preserves a broad outlook into the realm of collusion detection in financial markets in general, further approaches can be directly derived for different contexts based upon the available data of the specific markets and agents.

Additional domain-specific insights (see case study in Section 6.3.2), such as inferring estimates of costs, may ease the distinction between collusion and competition.

6.2.1 CPA based Behavioral Screening Phase

This section demonstrates the potential usefulness of change-point analysis techniques for detection of colluding behavior in Energy markets. The approach undertaken in this work falls under the class of nonparametric change point detection methods [20], that do not rely on pre-specified parametric models and thus avoid strong model assumptions.

Change-point detection is the problem of discovering time points at which properties of time-series data change. Suppose that $X = \{x_1, \dots, x_n\}$ is a sequence of independent random variables, such that the first r observations $X^A = \{x_1^A, \dots, x_r^A\}$ are distributed as F_A and the remaining observations $X^B = \{x_1^B, \dots, x_{n-r}^B\}$ come from another unknown distribution F_B , where $F_A \neq F_B$. Hence, integer r is called change point. Representing the trading activity of each market agent $a_i \in \mathcal{A}$ in terms of time-series enables us to perform a change-point analysis (Algorithm 1).

So, when a behavioral breakpoint occurs this corresponds to a step change of the mean value of X at r from α to $\alpha + \delta$, where δ represents the minimum increase of the mean value of X . Then, the deviation from the average may indicate that a collusion attack is being launched, if the cumulative deviation is noticeably higher than the random fluctuations, lower-bounded by δ .

A popular method based on a recursive nonparametric change point detection scheme uses a combination of cumulative sum charts (CUSUM) and bootstrapping to detect the changes [75]. The analysis begins with the construction of the CUSUM chart by calculating and plotting a cumulative sum, based on the timeseries data X .

The cumulative sums can be recursively defined using a new sequence $\{S_n\}$:

$$\begin{cases} S_n = S_{n-1} + (x_j - \bar{X}) \\ S_0 = 0 \end{cases} \quad (6.1)$$

where by \bar{X} we denote the mean of the sample:

$$\bar{X} = \frac{\sum_{i=1}^n x_i}{n} \quad (6.2)$$

Thus, the cumulative sum series can be obtained iteratively by adding to the previous sum the difference between a current value and the sample mean. This means that the sequence $\{S_n\}$ always ends at zero ($S_n = 0$), as the differences computed at each iteration sum to zero. Based on these remarks a CUSUM chart can be interpreted as follows. An upward slope of the chart indicates that the corresponding values tend to be above the overall mean of the sample, while a downward slope indicates a period of time where the values tend to be below. When a sudden turn occurs this indicates that around this time, the mean has shifted, which represents a potential changepoint. Likewise, a relatively straight CUSUM represents a period where the average did not change.

In order to associate a confidence level with a changepoint occurrence, a bootstrap analysis can be performed. To start with, a number of bootstrap samples are generated by sampling without replacement, which essentially is a random reordering of the original sample values. Thus, considering the initial sample X of size n , a bootstrapping sample at iteration i will be obtained by permuting without replacing k elements, generating $X_k^i = \{x_1^i, x_2^i, \dots, x_n^i\}$. For each of these samples, the bootstrap CUSUM is computed similarly. Moreover, for each sample we need to determine the maximum, minimum and difference of the bootstrap CUSUM denoted respectively as $S_{max}^i = \max_{j=0,1,\dots,n} S_j$, $S_{min}^i = \min_{j=0,1,\dots,n} S_j$ and $S_{diff}^i = S_{max}^i - S_{min}^i$. A bootstrap analysis consists of performing a large number of bootstraps and counting the number of bootstraps for which S_{diff}^i is less than S_{diff} of the initial sample. Let N be the number of bootstrap samples performed. Then, the probability of a changepoint occurrence for a fixed point r is given by:

$$\mathcal{P}_r = \frac{\#\{j : S_{diff}^i \leq S_{diff}\}}{N} 100[\%] \quad (6.3)$$

The bootstrapping technique basically compares the S_{diff} value of the original data with the S_{diff}^i values from a number of bootstrap samples, which estimate how

much S_{diff} would vary if no change took place and checks whether these results are consistent. It is clear that a better estimate can be obtained by increasing the number of bootstrap samples, however statistically significant result can typically be obtained for a reasonable number of generations.

Now, if a changepoint has been detected, in order to determine when the change has occurred, different estimators can be employed. A straightforward approach would be to determine the changepoint r as the point furthest from zero in the CUSUM chart:

$$|S_r| = \operatorname{argmax}_{i=0,1,\dots,n} |S_i| \quad (6.4)$$

The point r estimates last point before the change occurred, while the point $r + 1$ estimates the first point after the change.

Alternatively, the changepoint occurrence can be estimated by applying the mean square error estimator (MSE). The idea behind this estimator is that of partitioning the data in two sequences $X_1 = \{x_1, \dots, x_r\}$ and $X_2 = \{x_{r+1}, \dots, x_n\}$ and estimating the mean of each sequence and compare this against the initial data:

$$MSE(r) = \sum_{i=1}^r (x_i - \bar{X}_1)^2 + \sum_{i=r+1}^n (x_i - \bar{X}_2)^2 \quad (6.5)$$

where $\bar{X}_1 = \frac{\sum_{i=1}^r x_i}{r}$ and $\bar{X}_2 = \frac{\sum_{i=r+1}^n x_i}{n-r}$.

Algorithm 4 Change point analysis

Data: X_n, N, k, δ **Result:** changepoint r (if any) $S_0 \leftarrow 0$ **for** $i \leftarrow 1$ to n **do**| $S_i \leftarrow S_{i-1} + (x_i - \bar{X})$ **end****for** $j \leftarrow 1$ to N **do**| generate bootstrap sample $X_k^j, S_0^j \leftarrow 0$ | **for** $i \leftarrow 1$ to n **do**| | $S_i^j \leftarrow S_{i-1}^j + (x_i^j - \bar{X}^j)$ | **end****end****for** $j \leftarrow 1$ to N **do**| $S_{diff}^j = S_{max}^j - S_{min}^j, cnt \leftarrow 0$ **if** $S_{diff}^j \leq S_{diff}$ **then**| | $cnt \leftarrow cnt + 1$ | **else**| **end****end** $\mathcal{P} = \frac{cnt}{N}$ **if** $\mathcal{P} \geq 1 - \delta$ **then**| apply estimator to determine changepoint r **else****end**

6.2.2 Verification Phase for Collusive Coalitions

Now, having identified a candidate set of potential colluders, which allows to narrow our search space, we proceed to the second phase. Here, the focus is on detecting correlations between any members of the candidate set, resulting in a coalition

structure¹. Considering two timeseries $X = \{x_1, x_2, \dots, x_n\}$ and $Y = \{y_1, y_2, \dots, y_n\}$ representing the trading activities of agents a_x and a_y respectively, we are interested to capture linear dependencies between the two variables: X and Y . As a measure of similarity, we evaluate the covariance, which determines how X and Y vary together:

$$C(X, Y) = \frac{1}{n} \sum_{i=1}^n (x(i) - \bar{x})(y(i) - \bar{y}) \quad (6.6)$$

where \bar{x} is the sample mean of the X values and \bar{y} is the sample mean of the Y values. The covariance measure ranges from 1 for perfectly correlated results, through 0 when there is no relation between X and Y , to -1 when the results are perfectly correlated negatively.

More generally, if we consider k variables we can construct the covariance matrix ($k \times k$), where an element (i, j) represents the covariance between the i th and j th variables. Removing the dependence of the covariance on the ranges of the variables can be done by standardization, dividing the result by the standard deviations of X and Y . The result is the correlation coefficient between X and Y :

$$\rho(X, Y) = \frac{\sum_{i=1}^n (x(i) - \bar{x})(y(i) - \bar{y})}{\left(\sum_{i=1}^n (x(i) - \bar{x})^2 \sum_{i=1}^n (y(i) - \bar{y})^2 \right)^{1/2}} \quad (6.7)$$

Algorithm 2 summarizes the verification phase according to the computations previously detailed.

¹If all members of the candidate set show correlations we refer to this as the grand colluding coalition, while no correlations corresponds to the empty set.

Algorithm 5 Collusion detection**Data:** order record set O from one trading day discretize over a fixed time slot se-quence $\tau = \{t_1, \dots, t_k\}$; set of market agents $A = \{a_1, \dots, a_n\}$ **Result:** collusion candidate set $\mathcal{C} = \{S_1, \dots, S_l\}$ $\mathcal{C} \leftarrow \emptyset$ **for** each market agent a_i **do**| extract order record X_i associated to agents a_i from O | run *change point analysis* for given X_i | **if** probability of change point occurrence $\geq 95\%$ **then**| | $\mathcal{C} \leftarrow \mathcal{C} \cup a_i$ determine change point r with estimator| **else****end**generating covariance matrix for elements $a_i \in \mathcal{C}$, $l = 1$ **for** $i \leftarrow 1$ to $\text{sizeof}(\mathcal{C})$ **do**| **for** $j \leftarrow i$ to $\text{sizeof}(\mathcal{C})$ **do**| | $M(i, j) \leftarrow \text{Cov}(X_i, Y_i)$ | | **if** $M(i, j) \geq 1 - \delta$ **then**| | | $S_l \leftarrow \{a_i, a_j\}$ **if** $a_i \in S_k$ or $a_j \in S_k$ **then**| | | | append S_l to S_k | | | **else**| | | | $l \leftarrow l + 1$ append S_l to \mathcal{C} | | | **end**| | **else**| | **end**| **end****end**

6.2.3 Validation Phase via Principal Component Analysis

At this point we have reached the *validation phase*, where we attempt to explain complex relations between the resulting set of colluders by inspecting for an underlying, unobservable, latent structure that governs their trading activities. The

interesting question that we investigate here is whether we can reduce the variance in the observed variables to an unspecified number of unknown, unrecorded variables that account for most of this variance. Our aim is thus to reveal the causal relation that influences the evolution of the multiple recorded variables. To this end, we base our method on a statistical approach known as principal component analysis (PCA) [68], which benefits from the property of not making any assumption about an underlying causal model.

Recall that for each collusive agent a_x , her trading activity is captured as a time-series $X = \{x_1, x_2, \dots, x_n\}$. Clearly, these variables can be directly observed and will be referred onwards as *manifest variables*. The *validation phase* aims to express this information in terms of a restricted set of new orthogonal variables that constitute a *latent variable model*. The previous *verification phase* provided the covariance matrix C based on which we assessed correlations between colluder candidates. According to the PCA approach the standard eigenvalue-decomposition of the covariance matrix yields:

$$Cv_i = \lambda v_i \quad (6.8)$$

where λ_i 's represent eigenvalues and v_i 's represent eigenvectors of C .

In the matrix form, we have:

$$C = V\Lambda V^T \quad (6.9)$$

where Λ is the diagonal matrix with the eigenvalues λ_i 's in the diagonal, arranged in descending order $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ and V is an orthonormal matrix. The principal components are then obtained by projecting the initial manifest variables X onto the reduced dimension space:

$$P = XV \quad (6.10)$$

Columns in P correspond to principal components, that is, standard normal variables describing patterns that are associated with causal relations in X . Note that only the first few components account for meaningful amounts of variance and therefore,

only these first few components are retained, interpreted, and used subsequently. The uncertainty regarding the number of components extracted is typically resolved by selecting the first few that explain a fraction greater than 90% of the total variance in the data. Hence, if this feature can be identified as a characteristic of the data, this in turn accounts for a validation of our initial findings in terms of collusion candidate sets. Essentially, it means that we can simplify the problem by replacing a group of variables with a few principal components, that are orthogonal to each other, so there is no redundant information, but which can give an explanation of the driving principal governing the behavior of the system.

6.3 Experimental results

6.3.1 Data preparation

Prior to running our collusion detection mechanism, the dataset needs to undergo a pre-processing phase. As discussed in Section 6.1, we retain from the order record of the Energy Market, for every agent, a time-series for each day consisting of their bids, with respect to the predefined time-slots \mathcal{T} of the day-ahead market. Therefore, agent a_i is characterized at day j by $X_i^j = \{b_i^j(t_1), \dots, b_i^j(t_m)\}$.

Now, in order to run a meaningful change-point analysis over this data, detecting relevant behavioral breakpoints, we need to span the investigation over a time-window of several days. Moreover, we need to relate the agents' trading patterns to the temporal organization of the day-ahead market, \mathcal{T} . Specifically, let's assume a time-window of length l days and a fixed discretization of the day-ahead market $\mathcal{T} = \{t_1, \dots, t_m\}$. This requires constructing for each agent a_i the set of time-series $X_k^i = \{b_1(t_k), \dots, b_l(t_k)\}$, $k = \overline{1, m}$. Next, recall that a bid, $b_i^j(t_i) = (v_j, p_j)$, from an energy supplier consists of the amount of electricity offered and the intended price per unit. This further implies that for each X_k^i there corresponds a time-series denoting price $P_k^i = \{p_1(t_k), \dots, p_l(t_k)\}$ and another for the intended trading volume $V_k^i = \{v_1(t_k), \dots, v_l(t_k)\}$. This representation of the data is used during the *screening phase* for generating the collusion candidate set, as well as during the *verification phase* for




Figure 6.2: a) Plot of the trading activity P_1^6 corresponding to market agent a_6 for January 2012. b) The associated CUSUM chart for the time-series of market agent a_6 .

detecting price or volume correlations.

6.3.2 Case study

In this section we report on results² obtained from applying our model to real datasets, collected from the Philippine Wholesale Electricity Spot Market (WESM)

²We remind the reader that the analysis of economic data only, has the role of discovering suspicious behavior and is not meant to provide conclusive evidence of collusion, nor substitute antitrust authorities, but rather to provide supporting evidence and triggers for deciding whether antitrust authorities should actively engage and further pursue such an investigation.

Table 6.1: Results of the Change-Point Analysis with MSE Estimates for market agents with confidence level above or equal to 95%, representing the collusion candidate set.

Market Agent	Change-Point	Confidence Level
a_1	3	95%
a_2	7	97%
a_3	6	98%
a_3	20	100%
a_4	9	96%
a_4	12	99%
a_5	2	99%
a_5	22	100%
a_6	24	100%
a_7	17	100%
a_7	26	95%
a_8	12	100%
a_9	10	100%
a_9	15	100%
a_9	28	98%
a_{10}	24	100%




Figure 6.3: Plot of the covariance matrix for the collusion candidate set of market agents using a grayscale range of colors. Correlations of 1.0 are plotted in white, while no correlation are plotted in black. The diagonal elements of the matrix represent self-correlations and are thus all 1.

[69]. We ran the analysis on the Market Bids submitted by the trading participants over a one month time-window (January 2012), for the Luzon region. The list of registered WESM market participants consists of 60 members; prices are listed in Pesos per MWh; the nominated energy quantity is given in MW.

We proceed with the *screening phase* by conducting an exhaustive change-point analysis over the bid records of each market participant. Table 6.1 summarizes the results obtained at this stage, outlining the *colluding candidate set* as input for the following phase. For the given scenario we have generated 1000 bootstrap samples for each run of the algorithm. The results indicate that out of the total number of market agents, behavioral breakpoints have been detected for 10 agents, some of which exhibiting multiple ones. An illustration of this process is given in Figure 6.2 for market agent a_6 . Figure 6.2a is a representation of a_6 's trading activity during the specified time-window. The result of performing the screening analysis over this data, using the MSE estimator, detects the occurrence of one change-point timestamped at

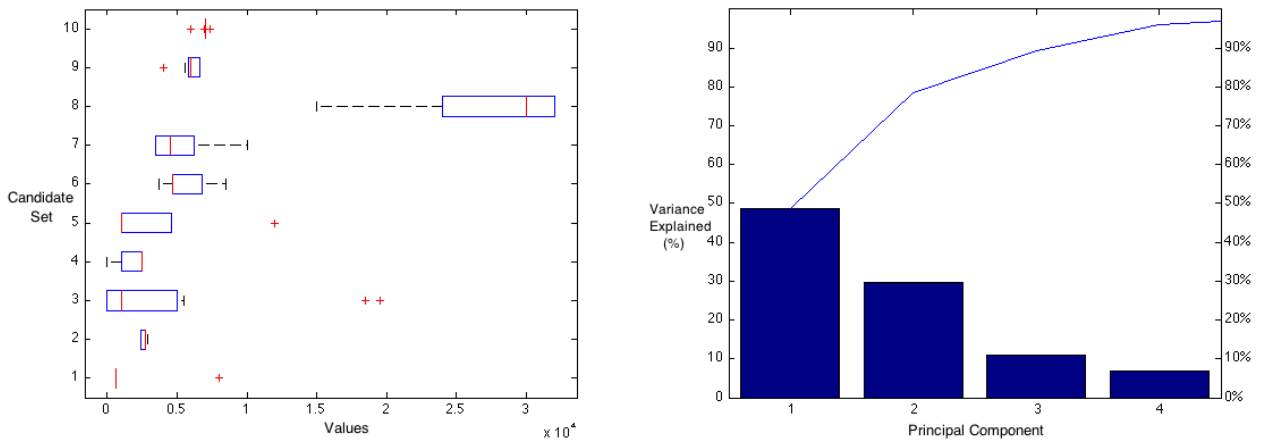


Figure 6.4: a) Boxplot representation of the trading dataset; b) Scree Test.

day 24. The confidence level associated, indicates a 100% accuracy. In Figure 6.2b we plot the corresponding CUSUM chart, highlighting a sudden shift in the average, associated with the change-point detection. Generally, it is not the case that change-points can be readily detected visually from the time series plot. A CUSUM chart however, can facilitate pinpointing shifts in the mean of the data, by identifying slope changes at the points where a change has occurred.

Next, the mechanism proceeds with the *verification phase*. As previously detailed, for the designated candidate set, generated during screening, we investigate further correlations between the market agents' trading patterns. This phase yields the covariance matrix, a representation of which is given in Figure 6.3. Here, we perform an exhaustive pairwise comparison of the candidate set. White squares denote a perfect correlation between the respective market agents, while black stands for no similarities. The color shading inbetween is indicative of the correlation strength. Thus, according to Algorithm 2 the coalition structure of colluders is determined. For the considered scenario, the coalition structure of potential colluders consists of solely one group: $CS = \{\{a_1, a_5, a_6, a_7, a_{10}\}\}$. For privacy concerns, we omit other direct reference to the suspected colluders. Note that the particular choice of the value of the correlation coefficient threshold δ is ought to impact the number of resulting colluders. On one hand, higher thresholds imply a higher level of confidence for the




Figure 6.5: Orthonormal principal component coefficients for each agent and principal component scores for each observation.

collusion detection process, but on the other hand it may as well reduce the number of possible suspects, disregarding certain abnormal trading behaviors. Alternatively, lower thresholds may result in including false colluders to the coalition structure and therefore reducing the accuracy of the mechanism. In this context, selecting a reliable correlation coefficient threshold is an important issue for the overall performance of the mechanism, which we plan to address in future work. Specifically, we intend to calibrate the system based on already proven cases of collusion and use such scenarios as training data. In addition to this, we plan to extend the model to integrate details regarding the devices \mathcal{D} , which are controlled by the market agents \mathcal{A} , such as DER type and geographical location. Adding this domain-dependent dimension is ought to provide further insight into differentiating between correlations that may come as a result of external conditions (e.g. weather conditions) and those that are irrespective to this regard.

Finally, we weight on our findings throughout the *validation phase*. We begin by plotting in Figure 6.4-a, comparatively, the distribution of the trading data for each member of the colluding candidate set in terms of the shape of the distribution, its

central value, and variability. Observably, there is more variability in the distributions of agents a_8 and a_3 than it appears for agents a_1 , a_2 and a_{10} . With respect to skewness, which measures the asymmetry of probability distributions, the data exhibits some common patterns, in the sense that the distribution of agents a_5 , a_6 and a_9 is skewed right, while for agents a_2 and a_4 is predominantly left skewed.

In Figure 6.5 the axes represent principal components, while the observed variables are represented as vectors. Figure 6.5 shows how the principal components form an orthogonal basis for the space of the data and how each observation is represented in terms of these components. Moreover, it captures the magnitude and sign of each variable's contribution to the first three principal components. Again, the data displays certain common patterns. For instance, the principal and third components have negative coefficients for all agents, while the second principal component has only positive coefficients.

Conclusively, Figure 6.4-b captures the *scree plot* representation of the first four components that explain 95% of the total variance. However, by itself, the first component appears to accounts for about 50% of the variance. By examining the plot we can conclude that altogether, the first three components explain to a large degree the total variability. Hence, this validates and supports the initial assumption of the existence of a reduced, unobservable number of factors that represent the driving forces that generated the original data.

6.4 Chapter Summary

In this chapter we have addressed the challenge of detecting collusive traders that collaborate illegally to increase their benefits at the expense of the other market participants. We have pose this question in the domain of the emerging energy markets, that are adapting to the integration of a diversity of distributed energy generators. Such contexts are especially susceptible to various trading malpractices.

The proposed method for discovering colluders consists of three phases. Firstly we apply a *screening phase* that performs a change-point analysis in order to detect behavioral breakpoint in the traders' activities, proposing a reduced candidate set of

possible colluders. Secondly, for the denominated group we run a *verification phase* aimed at revealing behavioral correlations. Thirdly, we seek to *validate* our results by gaining a deeper understanding of the driving factors that govern the behavior of the system. The procedure determines a potential coalition structure of colluders. We evaluate our mechanism on real datasets and show the effectiveness and practical applicability of our method, even for scenarios that are exploiting a minimal amount of data, that is freely available on the market.

As the space of *prosumer* engagement and participation in the operation of the *smart grid* is strongly augmented by this new paradigm, it is vital that mechanism for keeping the markets in check are developed in parallel. In this sense, an important part is played in the design of *collusion markers* which denote suspicious market behavior. This brings us to the limitation of applying automated collusion detection techniques. As the markets and the new procedures for efficient prosumer coordination mature, it is important that considerable attention is allocated for discovering new ways and patterns that threaten its operation in the form of non-benevolent participants that aim to 'game' the system. Investigating these issues and integrating new collusion markers, that may expose potential vulnerabilities in the energy market, represents new challenges for future work.

Chapter 7

Conclusions

The press, the machine, the railway, the telegraph are premises whose thousand-year conclusion no one has yet dared to draw.

— *Friedrich Nietzsche* —

In this chapter we summarize the contributions made in this thesis. Furthermore, we present some of the future lines that deserve to be addressed in future work.

7.1 Contributions

In the previous chapters we have taken a systemic approach to the problem of modernization of the electricity grid. In line with the mainstream view in the field of *smart grids*, we have identified two large groups of drivers that can catalyse the transformation of the grid which are occurring on the one hand at *(i)* the supply side through distributed energy generation at high penetration levels, and on the other hand, at *(ii)* the demand side through smart metering capabilities. Throughout the thesis we have applied the open multiagent paradigm to model a flexible representation of grid. Moreover, the central theme of this work is in framing the problems addressed as a coordinating effort where the self-interested nature of the agents is aligned with the system level requirements.

In particular, as the generation infrastructure is changing the operation and topology of the grid, our first objective in this work was to coordinate local optimization tasks that are shifted from central control systems towards subsystems of the grid driving a bottom-up efficiency in operation. In our approach, the dynamic distributed nature of the grid, as well as its large scale optimizations, will result at local level negotiations and cooperation among all the actors involved leading to energy efficiency. In this case, our objective as system designers was to tackle the emerging complexity by cooperation and modularity. By augmenting the prosumer model with computational and communicational capabilities we set the context for introducing a decentralized peer-to-peer mechanism for supply-demand-matching, proposing a new type of organization, *dynamic* microgrid, where the traditional role of the energy retailer becomes obsolete. The new setting proposed induces a macro-grid comprised of many grid-connected microgrids, that are exhibiting reduced transmission losses and an efficient utilization of renewables. Whenever the system deviates from its operating point, its components automatically reconfigure themselves adaptively to correct the problem.

In the same line of work, the second objective of this thesis was to address the problem of integration and controllability of DG systems, also referred to as VPPs. Here, coordination plays a key role for reconstructing the functionality of large power plants through large-scale deployment of DG. For this reason, we started off by modelling a game dynamic for underlying the economic effects of VPPs, by focusing on the formation of agent coalitions and analysing the conditions under which cooperation is beneficial, thus structuring the overall grid organization accordingly. We further went on to provided a new formulation of the scheduling problem in a VPP taking the form of distributed constraint optimization, extended to capture the inherent stochasticity of the domain. In doing so, we have designed algorithms capable to increase the potential for renewable power to contribute to the power supply in a reliable manner locally, by assessing the most viable combination of distribution automation and storage. The solution generates a plug-and-play environment where DERs can dynamically become VPP constituents. Importantly, uncertainty is addressed in a timely fashion allowing the system to reason reliably on its generation predictions.

The third major objective of this thesis lies in the area of consumer engagement as a future market segment that has not to this point been efficiently designed to bring this niche to a mass-market application. Increased market demand as well as climate change concerns are bringing a lot of pressure to the current ways in which consumption is managed. This new setting opened up interesting questions corresponding to two predominant energy markets. On the one hand, what is an efficient and fair way to determine consumers to not power-up their electricity-hungry devices at will, but make smart decisions in balancing costs and benefits in order to reschedule non-urgent activities to intervals when energy is abundant and cheaper. On the other hand, given this highly dynamic and complex ecosystem of production and consumption, can prosumers collaboratively service in an on-line manner stringent requirements coming from a grid operator that is responsible in monitoring the good operation of the network. We introduced a model where according to certain user specifications, an agent overlooks energy management at the household levels. We further take a game-theoretic standpoint to describe and analyse the agent interactions in the system, drawing attention on issues such as fairness and privacy-preservation. In relation to that, for the former case, we introduced the *ADDSM* protocol designed to capture the system goals, while enabling agents to adapt their demand profile based on economically rational decisions, with equilibrium guarantees. For the latter, we further propose the *eCOOP* protocol where agents can benefit from opportunities on timely reaction to market events. Again, we proposed a solution designed to cover a set of properties identified as key prerequisites for future smart grids, such as distributed agent coordination, stability, fairness, computational and communication simplicity, privacy-preservation, as well as the capability to operate in dynamic and stochastic environments.

Finally, it is clear that bringing about the decentralized and distributed reorganization of the electricity grid is unachievable without an intelligently engineered cooperation between the actors in the network. This in turn leaves room for ill-favored side-effects of cooperation and brings us to our fourth important objective, that of detecting for potential collusive behavior in the markets. We have introduced a mechanism for inspecting collusion patterns to tackle such practices based

on a combination of data mining techniques for change-point detection and principal component analysis. We expect such approaches to become a critical aspect in the monitorization of future energy markets that will need to cope with a significantly increased number of participation and the emergence of new mechanism, such as the ones proposed here for an efficient operation of the grid.

7.2 Future Research Directions

As we have been pointing out in the discussion of Chapter 2.6 on literature review there are several interesting topics that did not make the scope of this thesis, but on whose improvements our contribution could significantly benefit from. This also brings us to some of the limitations of this work.

The technology gap. The first obvious point is the limiting infrastructure that is currently in place, which forbids from a mass adoption of the solutions proposed in this work. Nevertheless, considering the spotlight that smart grid technologies has been under in the recent years as well as the impending climate concerns that are coming to the frontline, it is reasonable to believe that the required investments in deploying from smart meters to automated switching is going to also be supported in the future. Secondly we remark that the multiagent system domain is getting more traction in the smart grid environment. It is of course questionable the extent to which humans will choose to delegate their consumption preferences in favour of monetary incentives, however with the expected increase in the price of energy and carbon footprint, non-critical activities could easily be optimized to match costs benefits with lifestyle preferences of consumers.

An important gap that remains to be addressed is in perfecting techniques for residential and commercial energy profiling. In this work we make the assumption of having an accurate prediction of users' energy profiles on top of which our algorithms are run. It is difficult to imagine that without sound techniques that can perform accurate disaggregation of loads in the household or in the industrial sector, as well as predictions of demand profiles, that we would be able to successfully apply more sophisticated mechanisms as the ones described herein. In the beginning we have

pointed out briefly at some of the relevant approaches in this direction, emphasizing their current limitations. The same problem holds for deploying sensor technologies that can monitor the grid in a real-time fashion, estimating the future states of the grid and thus being able to perform adequate corrective actions. As we have been advocating thus far, policy engineering is a key aspect in engaging prosumers' participation in efficiency mechanisms. In order to be more realistic it is important that simulations rely, if not possible on actual real datasets, on artificially generated data that can approximate better the profiles of supply and demand.

The regulatory gap. Having to do with the assumptions considered herein, the smart grid faces the problem of lack of standardized modularity. Interoperability is key in enabling products or systems to interact with other products or systems. This means that bringing devices online, whether they are oriented to generation or consumption, should be done using service oriented architectures. This aspects leave a lot of room in applying existing, or developing new web service standards running on embedded devices. Thus, the heterogeneity can be abstracted away, allowing a straightforward way of accessing the functionality without focusing on the specific implementations. In this context, the vision for the smart home of the future prescribes a plug-and-play functionality of very device, so that the agent responsible for the home management system can easily incorporate these devices into its system and proceed with their utilization.

Appendix A

Author Bibliography

- (1) Mihailescu, R. C., M. Klusch, and S. Ossowski, "The COOP: Privacy-Preserving Dynamic Coalition Formation for Power Regulation in the Smart Grid.", *Proceedings of the 2nd International Conference on Agreement Technologies (AT)*, Beijing, China, August 1-2, 2013, 19-31, 2013.
- (2) Mihailescu, R. C., M. Vasirani, and S. Ossowski: The Gamification of Agent-Based Smart Grids. *ERCIM News* 2013(92), 31-32 (2013)
- (3) Mihailescu, R. C., M. Vasirani, and S. Ossowski, "Desiderata for agent-based power regulation protocols in smart grids", *Proceedings of the 1st International Conference on Agreement Technologies (AT)*, Dubrovnik, Croatia, 438-439, 2012.
- (4) Mihailescu, R. C., M. Vasirani, and S. Ossowski, "Detection of Collusive Behavior in Energy Markets", *Proceedings of the 1st International Conference on Agreement Technologies (AT)*, Dubrovnik, Croatia, 353-365, 2012.
- (5) Mihailescu, R. C., M. Vasirani, and S. Ossowski, "Detection of Collusive Behavior in Energy Markets (Extended Abstract)", *AAMAS Workshop on Agentes for Energy Systems (ATES)*, Valencia, Spain, 101-103, 2012.
- (6) Mihailescu, R. C., R. Schumann, and S. Ossowski, "Stochastic multiobjective optimization for hybrid power generation systems", *Proceedings of the 11th Eu-*

- ropean Workshop on Multiagent Systems (EUMAS), Ireland, Dublin, 224-238, 2012.
- (7) Mihailescu, R. C., M. Vasirani, and S. Ossowski, "Dynamic coalition adaptation for efficient agent-based virtual power plants", *Multi-Agent System Technologies - 9th German Conference (MATES)*, Berlin, Germany, 101-112, 2011.
 - (8) Mihailescu, R. C., M. Vasirani, and S. Ossowski, "A game-theoretic coordination framework for energy-efficient demand-side management in the smart grid", *Proceedings of the 10th European Workshop on Multi-agent Systems (EUMAS)*, 2011.
 - (9) Mihailescu, R. C., M. Vasirani, and S. Ossowski, "An investigation of emergent collaboration under uncertainty and minimal information in energy domains", *Agents in Principle, Agents in practice (PRIMA)*, Wollongong, Australia, 472-483, 2011.
 - (10) Mihailescu, R. C., M. Vasirani, and S. Ossowski, "An Organizational approach to agent-based virtual power stations via coalitional games", *Highlights in Practical Applications of Agents and Multiagent Systems-Advances in Intelligent and Soft Computing*, 125-134, 2011.
 - (11) Mihailescu, R. C., M. Vasirani, and S. Ossowski, "An agent-based model of emergent collaboration under uncertainty in energy domains", *Avances en Inteligencia Artificial (CAEPIA)*, San Cristóbal de la Laguna, Tenerife, Spain, Vol. 1, 2011.
 - (12) Mihailescu, R. C., M. Vasirani, and S. Ossowski, "Dynamic coalition formation and adaptation for virtual power stations in smart grids", *AAMAS Workshop on Agentes for Energy Systems (ATES)*, 85-88, 2011.
 - (13) Mihailescu, R. C., M. Vasirani, and S. Ossowski, "Towards agent-based virtual power stations via multi-level coalition formation", *The 1st Workshop on Agent Technologies for Energy System (ATES)*, 107-108, 2010.

- (14) Mihailescu, R. C. and S. Ossowski, "A Framework for Fraud Discovery via Illicit Agreements in Energy Markets", *Accepted for AI Communications Journal (To appear)*.
- (15) Mihailescu, R. C., M. Klusch, and S. Ossowski, "A dynamic coalitional game approach to autonomous demand response in the smart grid ", *Submitted to Journal of Information Systems Frontiers*.

Appendix B

Resumen en castellano

B.1 Antecedentes

Desde sus comienzos modestos del 29 de octubre de 1969, cuando, con cuatro sistemas host informáticos a 3420 Boelter Hall de UCLA, ARPANET, su precursor hizo los primeros pasos, el Internet ha crecido en presente, a un nivel de cientos de millones. Una de las mejores cosas sobre el Internet es que nadie puede pretender la propiedad sobre el mismo. De hecho, la entidad única llamada el Internet es una colección de redes de varios tipos y dimensiones. Sobre todo no existe una red global de control que reglamente la forma de funcionamiento de este sistema. La misma denominación proviene de la idea de redes interconectadas. La estructura fundamental de base del Internet incluye varias redes de alto nivel que se conectan unas a las otras a través de unos Network Access Points (Puntos de acceso a la red) (NAP). Para llegar a este backbone, los utilizadores caseros o las sociedades organizadas en Local Area Networks (redes locales) (LAN) se conectan al principio a un Internet Service Provider (Proveedor de servicios de Internet) (ISP). Point of Presence (punto de presencia) (POP) representa el lugar donde los utilizadores locales acceden la red del ISP, convirtiéndose en parte de esta red. Claro, es posible que el ISP deba conectarse a una red mayor antes de llegar a NAP. Por último, los

proveedores de Internet se interconectan al nivel de los NAP, conviniéndose sobre el cambio de comunicaciones entre ellos. La interconexión se realiza por los backbones de alta velocidad, normalmente backbones de fibra óptica, representando cables de fibra óptica combinados juntos para aumentar la capacidad. Las máquinas efectivas del Internet son divididas en dos categorías: servidores y clientes. Los servidores son esas máquinas que suministran servicios a otras máquinas. Hay varios tipos de servidores como por ejemplo: servidores de base de datos, servidores de archivos, servidores de correo electrónico, servidores de impresión, servidores web, servidores de juegos, servidores de aplicaciones etc. Los clientes son esas máquinas que solicitan procesos desde los servidores. Los clientes son esas máquinas que solicitan procesos de servidores. Acceder un cierto servicio en un servidor por un cliente se hace utilizando protocolos específicos. . Por ejemplo, una máquina cliente que utiliza un navegador web direcciona sus solicitudes hacia un servidor software específico que opera en la máquina servidor. La interacción cliente-servidor cumple con el Hypertext Transfer protocol (protocolo de transferencia de hipertexto) (HTTP), que describe la forma en la que el cliente y el servidor comunicarán.

En 1964, el pionero en inteligencia artificial, Dr. Arthur L. Samuel, escribió un artículo para New Scientist titulado The Banishment of Paper-Work (Eliminación del trabajo de oficina), donde imaginó el futuro del Internet para el año 1984, antes de su existencia efectiva. Aunque sus previsiones fueron aún más optimistas en cuanto al tiempo hasta la implementación, en esencia, fueron correctas:

"Una persona podrá navegar en la sección de literatura de la biblioteca central, podrá divertirse por la noche visionando cualquier de las películas producidas alguna vez (claro, pagando un coste conveniente, ya que en Hollywood todo es comercio) o podrá interesarse en las cifras de producción de hojalata del día anterior de Bolivia todas por una solicitud dirigida al terminal de la distancia propio. Las bibliotecas de libros dejan de existir en los países más desarrollados, excepto unas conservadas como en museos y la mayor parte de los conocimientos del mundo serán redactados de una forma legible en la máquina. Tal vez sería más correcto decir así: todos los conocimientos registrados del mundo serán redactados en esta forma, ya que el arte de la programación de los ordenadores que leerán el material editados y escrito a mano, hubiera sido totalmente desarrollado. A pesar de esto, el problema del almacenamiento hará imperativa la utilización de

una forma más condensada de registro, forma legible sólo a través de una máquina, que será traducida en una forma legible por el hombre a través del ordenador personal, a petición.”

Al mirar atrás también miramos hacia delante. Por analogía, preconizamos que los desarrollos del sector energético transformarán esta industria en una organización muy semejante a la que hoy parece el Internet. Está claro que la anchura de la banda (bits por segundo) no se ha reducido con el tiempo, sino al contrario, alcanzó alturas no preconizadas e inesperadas. Del mismo modo, tenemos preconizado que el consume de energía (julios por segundo) presentará aumentos significativos, difícil de imaginar en este momento, en gran parte gracias a unos usos no descubiertos. Con esta hipótesis en la mente, el aspecto de la eficiencia se convierte en un punto focal importante en la existencia de las futuras redes de electricidad. Mientras que hoy en día, la utilización pseudo-eficaz de la energía aún es posible, en el futuro ya no viene el caso. Imagínese varios tipos de recursos de generación basados en energía eólica, solar o de las mareas jugando el papel de unos servidores en la red y clientes representando a los consumidores organizados en configuraciones LAN conectados a la red por varios proveedores de servicios similares al ISP. Los protocolos concebidos para atribuirse eficazmente los recursos reglamentará la interacción consumidor-productor de una forma similar al modo de funcionamiento corriente del HTTP en el caso del Internet. Los participantes en la red (network actor) podría coordinar la red para asegurar el suministro de ciertos servicios por la entrada de las organizaciones a través del POP del ISP o a través de la participación de forma independiente al acceder los NAP Smart Grid. Pero, lo más importante, la experiencia (UI) de la forma en la que interaccionamos en el presente con la red cambiará. A continuación resumo mis previsiones subjetivas relacionada con la Red Inteligente (Smart Grid) hasta el año 2029:

Actualmente, la red trata bien el problema de la distribución de energía. El transporte de energía a largas distancias es ineficaz y caro. La nueva red se especializará en el cambio de energía a nivel local, desde las fuentes intermitentes a destinos intermitentes. Dado que las grandes centrales eléctricas centralizadas serán en gran parte eliminadas del sistema a favor de las fuentes de energía limpias, distribuidas, una configuración como

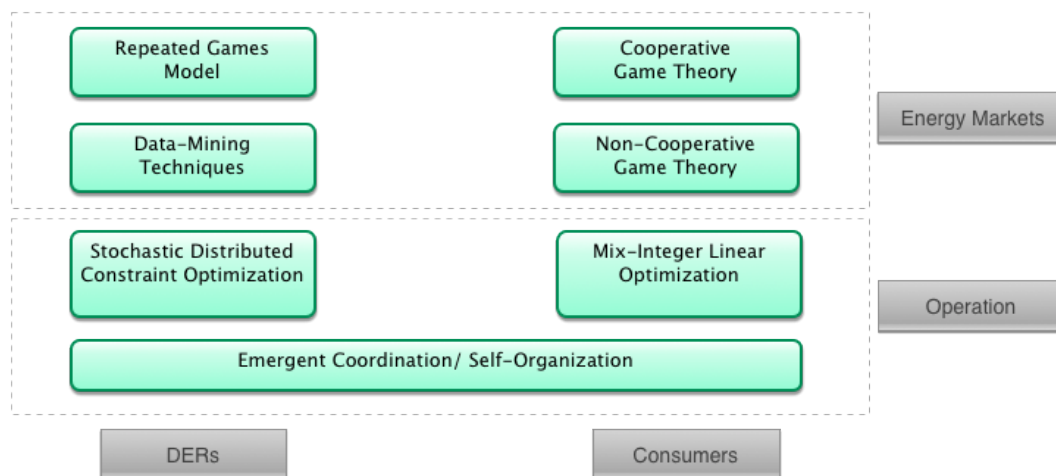


Figure B.1: Areas of contributions (with our contributions marked in light green)

esta de la red gozará de la apertura, robustez y fiabilidad, será menos predispuesta a molestias mayores. Lo importante es que el aspecto local predominará, ya que la optimización en gran escala resultará de las configuraciones al nivel local efectuadas para el aumento de la eficiencia energética. Es más, dentro de este ecosistema extremadamente dinámico y complejo de suministro y consumo de energía, las utilidades en la forma en las que las conocemos hoy, dejarán de existir. La red funcionará en base a una gran plataforma de datos que permitirá a las partes interesadas en utilizar las aplicaciones. Los agentes de software predecirán con exactitud los modelos de consumo y de generación y una gran parte de los deberes de planificación de los consumidores serán delegados a los mismos. Los agentes negociarán la utilización de sus dispositivos, cuya utilización ya no será considerada sólo una consumidora de energía, sino más bien consumidora de servicios a través de los facilitadores similares al servicio Google. Generalmente, la red de energía eléctrica será una infraestructura gestionable por una red de superposición en la que todos los aparatos inteligentes, el coche, el teléfono inteligente, el portátil, la lavadora etc colaboran de forma transparente como parte de una inteligencia global que asiste a la gente en la optimización de sus vidas diarias. ”

B.2 Objetivos

Teniendo en cuenta estas previsiones a largo plazo, en este trabajos nos concentramos en el estudio de los mecanismos distribuidos para el control y la gestión de las

futuras redes eléctricas inteligentes. El concepto de *redes inteligentes* (smart grids) se utiliza aún desde el año 2003, cuando apareció en el artículo "*Reliability demands will drive automation investments* (*"Demandas de fiabilidad llevarán a inversiones en automatización*), que describió la necesidad de una actualización a largo plazo de la red en lugar de efectuar las subsanaciones a corto plazo en cuanto a la capacidad de transmisión y los sistemas de control de la red [21]. Asimismo, el Departamento de Energía de los Estados Unidos fue uno de los primeros abogados de estas propuestas [123].

Desde el punto de vista del principio de proyección fundamental, hay que distinguir dos posibilidades mayores de enfoque del problema de gestión de las redes inteligentes: *centralización* versus *descentralización*. Los enfoques centralizados quedan favorecidos en primer lugar por las *utilidades*, que no tienen la intención de renunciar a su participación clave en la cadena energética de valores. Alternativamente, por un enfoque descentralizado se destaca la importancia de la utilidad en las futuras redes inteligentes, haciendo hincapié en primer lugar en los productores y consumidores reales que componen la red, delegándoles la carga de la *coordinación*.

Aunque es verdad que la democratización del sector energético solicita soluciones descentralizadas, este mismo es el que hace más difícil el problema. Cómo pueden organizarse los pequeños *productores* para igualar la fiabilidad ofrecida en presente por las grandes centrales eléctricas? Cómo pueden convertirse los *consumidores* en participantes activos en la optimización del consumo de energía? Cómo podemos reducir a mínimo las pérdidas de energía del sistema y aumentar la fiabilidad? Qué tipo de formas de organización de la red nos podemos imaginar que puedan evitar los costes adicionales a cargo de las utilidades corrientes?

En esta tesis, nos concentramos en el enfoque descentralizado al que trataremos desde una perspectiva basada en agentes [191], en la que los agentes autónomos inteligentes accionan a nombre de los participantes de la red, que interaccionan entre ellos y también con la infraestructura para optimizar el estado de la red. El objetivo de la tesis destaca las pautas a las que consideramos que el desarrollo de la red inteligente tiene que efectuar para realizar una visión de la red inteligente. Figura B.1 presenta un mapa de las zonas de investigación en las que se incorporan nuestras con-

tribuciones. Marcamos en verde claro todos esos campos que se enfocan en esta tesis suministrando modelos, conceptos de soluciones y algoritmos, donde la coordinación tiene una importancia clave en la resolución de problemas, mientras que las casillas negras representan el nivel donde aparece el problema. Este trabajo sigue esta serie de objetivos:

- (O1) **La elaboración de un modelo de red inteligente para la organización de los participantes en la red fundamental de principios *microgrid* de resistencia y de pérdida reducida de energía en transmisión, que incorpora la naturaleza dinámica del medio.** La aparición de un sistema complejo, dinámico, heterogéneo y distribuido de producción y consumo de energía necesita un enfoque radicalmente distinto, que se puede adaptar adecuadamente a estas nuevas condiciones y que puede asegurar que la energía se puede utilizar de forma eficaz. El punto de partida de este trabajo es la investigación de una reestructuración de la infraestructura de entrega y la introducción de un enfoque caracterizado por la apertura, robustez y fiabilidad, que podría explotar los rendimientos de eficiencia resultados. Además, se quiere que el nuevo contexto que facilite un marco en el que las negociaciones y la cooperación entre todas las entidades llevarán al aumento de la eficiencia energética.
- (O2) El segundo aspecto enfocado está relacionado con la integración de los recursos energéticos distribuidos, que son en gran parte excluidos del mercado al por mayor por su ineficiencia y falta de fiabilidad. **Cómo podemos realizar mejor la organización de estos dispositivos, de tal forma que pueda representar la descomposición distribuida equivalente de una gran central eléctrica centralizada?** No se olviden de la analogía con el Internet del principio. De forma similar al modo en el que el Internet evolucionó desde los ordenadores centrales accesibles por los usuarios múltiples a una red distribuida de máquinas, centrales eléctricas centralizadas monolíticas corrientes serán sustituidas por la generación energética distribuida. La vulnerabilidad de la infraestructura de la red pone en cuestión la viabilidad de la dependencia de la producción de energía en fuentes renovables.

- (O2.1) **El análisis de los acuerdos individuales racionales donde los agentes pueden equivaler a ineficiencias individuales a través de técnicas de formación de grupos.**
- (O2.2) **El desarrollo de los modelos y algoritmos capaces de optimizar la explotación de semejantes formaciones capaces de hacer frente al medio estocástico inherente.**
- (O3) En tercer lugar, confiamos que la importancia de los consumidores en el control y monitorización en el futuro de la red será mucho mayor. Distinguimos varios aspectos en los que la implicación del consumidor puede tener un impacto significativo sobre la eficiencia de la red.
- (O3.1) **Proyección de una dinámica de mercado que asegure la capacidad de gestionar la demanda de tal forma que elimine la necesidad de la generación de reserva cara e ineficiente.** Asegurar rebajas periódicas de demanda para ciertos períodos del día para alcanzar unos equilibrios que permiten una asignación eficaz de los recursos disponibles.
- (O3.2) **El desarrollo por modelos, mecanismos y algoritmos de control dinámico de la demanda para responder a los requisitos bruscos de reducción de la demanda, para equilibrar la producción y el consumo en condiciones casi en tiempo real.**
- (O3.3) **El desarrollo de un modelo de uso general y suficientemente flexible al que un agente pueda implementar para automatizar el consumo energético casero para la gestión de los deberes que se pueden posponer.**
- (O4) Por último, ya que la base de nuestro enfoque es la implementación de una perspectiva descentralizada, basada en agentes en retos de la red inteligente, en los que los agentes auto-interesados. **Interaccionan en distintos guión, es importante concebirse mecanismos adicionales para la detección de la colusión en el caso de los mercados energéticos bajo consideración.**

B.2.1 Estructura de la tesis

Esta tesis está estructurada de este modo:

1. El capítulo 2 describe las condiciones que están a base de esta tesis de doctorado, destacando los sectores transpolinizadores de las futuras redes de electricidad inteligentes y sistemas multiagente abiertos. Luego, pasamos a la revisión del estado de la técnica, concentrándonos en enfoques relacionados con los objetivos mencionados anteriormente.
2. El capítulo 3 enfoca el primer objetivo de la tesis, analizando el estado de la red en términos de una penetración considerable de los recursos energéticos distribuidos. En este nuevo contexto, consideramos que debe implementarse otro tipo de organización para la gestión eficiente de la red a base de unas soluciones modulares de pequeñas dimensiones que son, al mismo tiempo, extremadamente adaptables y suficientemente flexibles para poder hacer frente al entorno estocástico y dinámico.
3. En el capítulo 4 vamos a introducir y evaluar nuestra propuesta para la creación y explotación de una central eléctrica virtual para la integración de las fuentes de energía renovable distribuidas. Teniendo en cuenta la incertidumbre sobre las fuentes renovables, nuestro enfoque concibe una modalidad por la que los dispositivos heterogéneos imitan cooperativamente las características de confiabilidad de una central eléctrica tradicional.
4. Luego, en el capítulo 5, enfocamos la diferencia de coordinación entre la demanda y la oferta e introducimos nuestra propuesta para la modificación de los modelos de demanda de los consumidores finales, a través del traslado de la carga para aplanar la curva de carga y para maximizar la utilización de los activos implementados. Tratamos este problema a través de la diferenciación entre las técnicas de gestión en la parte de la demanda para rebaja de tope del día siguiente y los mecanismos durante el día, concebidos para hacer frente a los desequilibrios demanda-oferta en escenarios casi en tiempo real.

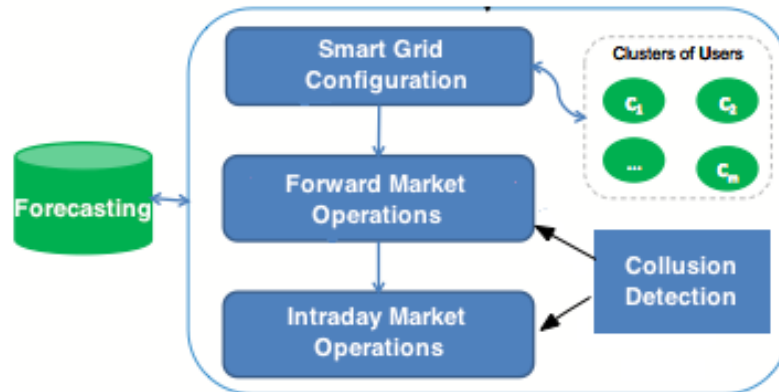


Figure B.2: Methodology

5. El Capítulo 6 enfoca los retos relacionadas con la institución de un entorno de mercado energético libre de colusiones. La integración de la generación distribuida hasta el nivel de hogar representa un cambio mayor hacia la democratización del modelo corriente de suministro de energía. A pesar de esto, se conoce el hecho de que la liberalización de los mercados está predispuesta a manipulaciones y tácticas ilegales de monopolio. Presentaremos aquí un mecanismo de facilitación por el que el sistema puede evitar este tipo de situaciones indeseadas.
6. Por lo ltimo, el capítulo 7 saca las principales conclusiones de la tesis, hablando sobre los principales resultados obtenidos e indicando futuras dirección de investigación.

B.3 Metodología

Las redes eléctricas actuales son estructuras gigantes gestionadas principalmente centralizado, basado en sensores y dispositivos de actuación ubicadas en algunas (pero estratégicamente elegidas) partes de reja. La gestión de la cámara de control es, de costumbre, sostenida y complementada por los ingenieros de terreno, que pueden informar sobre el estado de las líneas de transporte y pueden realizar manualmente

acciones de reparaciones, según el caso. Los mercados de energía eléctrica actuales están también reglamentadas y, en el mejor caso, conectados libremente y, teniendo en cuenta las estructuras de administración ya existentes, es todavía confuso a qué nivel y en qué ritmo la red eléctrica se convertirá en inteligente.

En esta tesis, partimos de una proyección optimista del futuro para este caso. Especialmente, suponemos una infraestructura física existente capaz de funcionar en presencia de una difusión de dispositivos de generación de energía eléctrica, basados, principalmente, en recursos renovables. También, se supone a continuación la presencia de algunos nudos inteligentes con capacidades de comunicación y de cálculo para todos los dispositivos del sistema. Más que esto, nos basamos en la existencia de una infraestructura de medida inteligente. Por último, lo más importante, preconizamos una cuadrícula en la que los consumidores pasivos corrientes de energía eléctrica no existen más, si son participantes activos en el consumo de energía y micro-producción que pueblan el sistema. Esto tiene dos consecuencias mayores para la investigación. Primero, es importante el estudio y la determinación de las ineficiencias de las redes actuales de energía eléctrica, junto al trazado del entendimiento del modo en que nos podemos esperar que la implementación de la tecnología se mejore a base de estas comprobaciones. Segundo, debemos investigar los mecanismos de lucha contra tales problemas, hecho que discordará en cuanto a los aspectos importantes, porque es poco probable que alguna solución general pueda ser utilizada en una gama larga de contextos. Porque los participantes de la red se convierten en participantes activos, estamos interesados especialmente en los problemas que necesitan coordinación en sistemas dinámicos y complejos. Enfocamos la perspectiva multiagente para la creación de algunos sistemas de grandes dimensiones, con un comportamiento previsible, capaces de generar propiedades deseables al nivel mundial.

Partiendo de este contexto, nuestro propósito es de diseñar, desarrollar y validar *mecanismos de coordinación* de acuerdo a los objetivos identificados en el apartado anterior. El marco en el que se enfocan estos objetivos se sintetiza en la Figura B.2. La tesis tiene cuatro partes principales. El punto de partida de esta obra implica la construcción del modelo, es decir una nueva reestructuración de la red a base de algunas formaciones microrred dinámica. La estructuración microrred de la red lleva

al formular un problema difícil que solucionar, que tratamos a través de un enfoque de coordinación emergente. Tal como se presenta en la Figura B.2, una vez que las microredes están formadas, tratamos los dos aspectos de la coordinación productor y consumidor en un contexto de mercado en dos niveles, proponiendo en los siguientes capítulos técnica *de gestión y de respuesta central eléctrica virtual* y, respectivamente, *en la parte de la demanda*, complementadas con procedimientos de *detección de la colusión*.

Hemos valorado el rendimiento de nuestro enfoque utilizando nuestro propio simulador personalizado, que desea reunir los campos de la Figura B.1 en el contexto de los entornos multiagente, utilizando datos del mundo real según la disponibilidad. Sin embargo, debe destacarse que, en realidad, las posibilidades de ensayo de los sistemas de gestión de la energía están limitadas por causa de la falta de una infraestructura de medida avanzada existente. Cuando los conjuntos de datos del mundo real están efectivamente recogidas, aún quedan muchos obstáculos para poderlos poner a disposición con fines de investigación. Estos representan, en general, datos de propiedad poseídas por los servicios, que hesitan revelar este tipo de informaciones por culpa de las políticas comerciales o como consecuencia de las reglamentaciones de confidencialidad de los datos de los consumidores. La obtención de datos viables es extremadamente difícil cuando las informaciones son necesarias para un gran número de participantes del sistema, por largos periodos de tiempo y con una alta granularidad (por ejemplo la utilización del aparato al nivel del hogar). Al efecto, los datos utilizados para los presentes experimentos reflectan escenarios del mundo real, si son disponibles, mientras que a veces recurrimos a la generación artificial de perfiles individuales de datos integrados y, ocasionalmente, generamos datos sintéticos desde cero.

B.4 Conclusiones

En este capítulo resumimos las aportaciones llevadas por esta tesis. Más que esto, presentaremos algunas de las futuras líneas que merecen ser enfocadas en una futura actividad.

B.4.1 Contribuciones

En los capítulos anteriores hemos enfocado sistémico el problema de la modernización de la red de energía eléctrica. De acuerdo al punto de vista principal en el sector de las *redes inteligentes*, hemos identificado dos grupos grandes de motores que pueden catalizar la transformación de la red, por una parte, en *(i)* la parte de la oferta por la generación de energía distribuida en altos niveles de penetración y, por otra parte, en *(ii)* la parte de demanda por las capacidades de medición inteligentes. A lo largo de la tesis de doctorado he aplicado el paradigma multiagente abierto para modelar una representación flexible de la red. Más que esto, el tema central de esta obra es la definición de los problemas enfocados como un esfuerzo de coordinación donde la naturaleza auto-interesada de los agentes está alineada a los requisitos al nivel de sistema.

Especialmente, a medida que la infraestructura de generación cambia la explotación y la topología de la red, nuestro primer objetivo en esta obra ha sido la coordinación de cargas de optimización local, pasadas de los sistemas de control centrales a los subsistemas de la red, llevando a un rendimiento de abajo hacia arriba de la explotación. En nuestro trámite, la naturaleza dinámica distribuida de la red y también sus optimizaciones de grandes dimensiones, llevará a negociación al nivel local y cooperación entre todos los participantes implicados, llevando a una eficiencia energética. En este caso, nuestro objetivo como diseñadores de sistema ha sido el enfoque de la complejidad emergente por cooperación y modularidad. A través de la aumentación del modelo de prosumidor con capacidades de cálculo y de comunicación, hemos fijado el marco para la introducción de un mecanismo descentralizado peer-to-peer para la concordancia demanda-oferta, proponiendo un nuevo tipo de organización, microred *dinámica*, donde el papel tradicional del proveedor de energía se convierte en caduco. El nuevo contexto propuesto induce una macrored formada de varios microredes conectadas a la red, que presentan pérdidas reducidas de transmisión y una utilización eficiente de las fuentes renovables de energía. Cada vez que el sistema se desvía de su punto de explotación, sus componentes se reconfiguran automáticamente y adaptivamente para corregir el problema.

En la misma actividad, el segundo objetivo de esta tesis ha sido el enfoque del problema de integración y controlabilidad de los sistemas DG, denominadas VPP. En este caso, la coordinación juega un papel clave para la reconstrucción de la funcionalidad de las grandes centrales eléctricas a través de la implementación de DG a gran escala. Por este motivo, hemos partido de la modelación de una dinámica de juego para la fundamentación de los efectos económicos de los VPP, enfocándose en la formación de una coalición de agentes y el análisis de las condiciones en las que la cooperación es benéfica, estructurando de este modo la organización general de dicha red. A continuación, hemos recurrido al suministro de una nueva fórmula del problema de programación en un VPP bajo la forma de una optimización de apremio distribuido, extendidas para captar la estocasticidad inherente del sector. Al efecto, hemos concebido algoritmos capaces de aumentar el potencial de la energía renovable de contribuir al suministro de energía eléctrica de modo viable al nivel local a través de la evaluación de la más fiable combinación entre la automatización de distribución y almacenamiento. La solución genera un entorno plug-and-play, donde los DER pueden convertirse dinámicamente en constituyentes VPP. Es importante retener que la incertidumbre está enfocada en tiempo útil, permitiéndole al sistema que piense viablemente sus previsiones de generación.

El tercer objetivo mayor de esta tesis reside en la zona de implicación del consumidor como futuro segmento de mercado que, hasta el día de hoy, no ha sido concebido de modo eficiente para llevar este nicho en una extensa aplicación de mercado. La demanda aumentada del mercado y también las preocupaciones relativas a los cambios climáticos ponen mucha presión en las modalidades corrientes en las que se gestiona el consumo. Este nuevo contexto crea preguntas interesantes para los dos mercados energéticos predominantes. Por una parte, ¿pe de o parte, cual es el modo eficiente y correcto por el que se pueden determinar a los consumidores a no alimentar sus dispositivos hambrientos para la energía eléctrica a sus anchas, sino tomar decisiones inteligentes en cuanto al equilibrar los costes y beneficios, para reprogramar actividades no urgentes en periodos de tiempo cuando hay energía mucha y barata? Por otra parte, teniendo en cuenta este ecosistema extremadamente dinámico y complejo de producción y consumo, pueden alimentar a los prosumidores en colaboración, en

línea, requisitos estrictos que vienen de un operador de red responsable para la monitorización de un buen funcionamiento de la red? Hemos introducido un modelo que, de acuerdo a algunas especificaciones del usuario, un agente vigila la gestión energética al nivel de los hogares. A continuación, enfocamos un punto de vista teórico de juego para describir y analizar las interacciones del agente del sistema, llamando la atención sobre los problemas y también sobre la corrección y conservación de la vida privada. En cuanto a este hecho, en el primer caso, hemos introducido el protocolo *ADDSM*, concebido para captar los objetivos del sistema, permitiéndoles a los agentes a adaptar su perfil de demanda a base de decisiones racionales desde punto de vista económico, con garantías de equilibrio. Para estas últimas, proponemos a continuación el protocolo *eCOOP* en el que los agentes pueden beneficiar de oportunidades de reacción en tiempo útil a los eventos del mercado. De nuevo, hemos propuesto una solución concebida para cubrir un conjunto de propiedades como premisas clave para las futuras redes inteligentes, como por ejemplo la coordinación distribuida de los agentes, estabilidad, corrección, sencillez de cálculo y de comunicación, conservación de la vida privada y también capacidad de operar en entornos dinámicos y estocásticos.

Por último, está claro que la realización de la reorganización descentralizada y distribuida de la red de energía eléctrica es imposible de realizar sin una cooperación inteligentemente arreglada entre los participantes de la red. Esto deja, a su turno, los efectos secundarios inoportunos de cooperación y nos lleva a nuestro cuarto objetivo, él de detectar el comportamiento posiblemente colusorio en el mercado. Hemos introducido un mecanismo para la inspección de los modelos de colusión para enfocar tales prácticas a base de una combinación de técnicas de minería de datos para la detección del punto de cambio y análisis de los componentes principales. Esperamos que tales enfoques se conviertan en un aspecto crítico en la monitorización de los futuros mercados energéticos que deberán hacer frente a una cantidad significativamente aumentada de participación y aparición de los nuevos mecanismos y también las propuestas aquí para un funcionamiento eficiente de la red.

B.4.2 Líneas de investigación futuras

Tal como hemos subrayado en la discusión del Capítulo 2.6 sobre la literatura de especialidad, existen varios temas interesantes que no son objeto de esta tesis, pero cuyas mejoras podrían llevar beneficios significativos a nuestra contribución. Esto nos lleva también a algunas de las limitaciones de esta obra.

Desfase tecnológico. El primer aspecto evidente es la infraestructura limitadora utilizada actualmente, que impide la adopción extendida de las soluciones propuestas en esta obra. Con todo esto, teniendo en cuenta que las tecnologías de redes inteligentes se encuentran desde varios años bajo la luz de los reflectores y que las preocupaciones climáticas inminentes pasa en primer plano, es razonable creer que las inversiones necesarias para la implementación de contadores inteligentes a la conmutación automática serán también apoyadas en el futuro. Segundo, destacamos que el sector del sistema multiagente recibe cada vez más tracción en el entorno de la red inteligente. Sin embargo, es cuestionable en qué medida la gente elija delegar sus preferencias de consumo a favor de los incentivos monetarios, pero, con el aumento preconizado del precio de la energía y de la huella de carbón, las actividades no críticas podrían ser fácilmente optimizadas para igualar los beneficios de costes con las preferencias de estilo de vida de los consumidores.

Un desfase importante que queda por ser enfocado es la mejora de las técnicas para la realización del perfil energético residencial y comercial. En esta obra, suponemos que tenemos una valoración exacta de los perfiles energéticos de los usuarios en virtud a los que funcionan nuestros algoritmos. Nos es difícil imaginarnos que, sin técnicas sólidas que puedan realizar la desagregación correcta de las cargas en el hogar o en el sector industrial y también las previsiones de perfiles de demanda, podrían aplicar exitosamente mecanismos más sofisticados, como el descrito aquí. Al principio, hemos indicado brevemente algunos de los enfoques relevantes al efecto, subrayando sus límites actuales. El mismo problema existe también para la implementación de tecnologías de sensores que pueden monitorizar la red en tiempo real, estimando los futuros estados de la red y, al efecto, teniendo la posibilidad de realizar acciones adecuadas de corrección. Tal como hemos sostenido hasta ahora, la creación de políticas

representa un aspecto clave en la implicación de la participación de los prosumidores en los mecanismos de eficiencia. Para ser más realista, las simulaciones deben basarse, si no en conjuntos de datos reales efectivos, en datos generados artificialmente, que pueden aproximar mejor los perfiles de demanda y oferta.

Faltas de reglamentación. En relación con las hipótesis consideradas aquí, la red inteligente se confronta con el problema de la falta de modularidad estandarizada. La interoperabilidad es un elemento clave en cuanto a la permisión de la interacción de los productos o sistemas con otros productos o sistemas. Esto significa que llevar los dispositivos en línea, indistintamente si esta están orientadas hacia la producción o consumo, debería hacerse utilizando arquitecturas orientadas en los servicios. Estos aspectos dejan mucho espacio en cuanto a la aplicación de los estándares de servicios web existentes o la elaboración de nuevos estándares de servicios web que funcionan en dispositivos integrados. De este modo, la heterogeneidad puede ser eliminada, permitiendo un modo simple para acceder la funcionalidad sin la concentración en implicaciones específicas. En este contexto, la visión para la casa inteligente del futuro prevé la funcionalidad del plug-and-play de cada dispositivo, de modo que el agente responsable para el sistema de gestión de la casa pueda incorporar fácilmente estos dispositivos en su sistema y continuar utilizarlas.

References

- [1] R. B. Adler, Jr. Daniel, S.L., C.R. Heising, M.G. Lauby, R.P. Ludorf, and T. S. White. An IEEE survey of us and canadian overhead transmission outages at 230 kv and above. *Power Delivery, IEEE Transactions on*, 9(1):21–39, 1994.
- [2] International Energy Agency. Energy technology perspectives. Technical report, IEA, 2010. <http://www.iea.org/publications/freepublications/publication/etp2010.pdf>.
- [3] Muddasser Alam, Sarvapali Ramchurn, and Alex Rogers. Cooperative energy exchange for the efficient use of energy and resources in remote communities. In *Twelfth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, pages 731–738, May 2013.
- [4] R. Albert and A.-L. Barabási. Topology of evolving networks: local events and universality. In *Phys. Rev. Lett.* 85, pages 5234–5237, 2000.
- [5] M. Amin and P. F. Schewe. Preventing blackouts. *Scientific American*, 296(5):60–67, 2007.
- [6] M. Andersson and T. Sandholm. Contract types for satisficing task allocation: experimental results. In *AAAI Spring Symposium Series: Satisficing Models, Stanford University, CA*, pages 1–7, March 1998.
- [7] Jose Apesteguia, Steffen Huck, and Jörg Oechssler. Imitation - theory and experimental evidence. *Journal of Economic Theory*, 136(1):217 – 235, 2007.
- [8] K. Apt and A. Witzel. A generic approach to coalition formation. *International Game Theory Review*, pages 347–367, 2009.
- [9] R. Axelrod. *The Evolution of Cooperation*. Basic Books, New York, 1984.
- [10] S. Beer, T. Gomez, D. Dallinger, I. Momber, C. Marnay, M. Stadler, and J. Lai. An economic analysis of used electric vehicle batteries integrated into commercial building microgrids. *Smart Grid, IEEE Transactions on*, 3(1):517–525, 2012.

- [11] Mario Berges, Ethan Goldman, H. Scott Matthews, and Lucio Soibelman. Learning systems for electric consumption of buildings. In *ASCE International Workshop on Computing in Civil Engineering*, pages 1–10. ASCE, 2009.
- [12] Mario Berges, Ethan Goldman, H. Scott Matthews, and Lucio Soibelman. Enhancing electricity audits in residential buildings with nonintrusive load monitoring. *Journal of Industrial Ecology*, 14(5):844–858, 2010.
- [13] K. Binmore. *Fun and Games: A Text on Game Theory*. Heath, Lexington, MA, 1992.
- [14] M.G.C. Bosman, V. Bakker, A. Molderink, J.L. Hurink, and G. J M Smit. On the microchp scheduling problem. In *3rd Global Conf. Power Control Optimization (PCO 2010)*, pages 367–374, 2010.
- [15] A. Botea, J. Rintanen, and D. Banerjee. Optimal reconfiguration for supply restoration with informed a search. *Smart Grid, IEEE Transactions on*, 3(2):583–593, 2012.
- [16] S. J. Brams, M. A. Jones, and D. Marc. Dynamic models of coalition formation: fallback vs. build-up. In *Theoretical Aspects of Rationality and Knowledge archive*, pages 187–200. Proceedings of the 9th conference of Theoretical aspects of rationality and knowledge, 2003.
- [17] M. Braun. Reactive power supply by distributed generators. In *Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century, 2008 IEEE*, pages 1–8, 2008.
- [18] M. Braun. Virtual power plant functionalities demonstrations in a large laboratory for distributed energy resources. In *Electricity Distribution - Part 1, 2009. CIRED 2009. 20th International Conference and Exhibition on*, pages 1–4, 2009.
- [19] Andrei Broder, Ravi Kumar, Farzin Maghoul, Prabhakar Raghavan, Sridhar Rajagopalan, Raymie Stata, Andrew Tomkins, and Janet Wiener. Graph structure in the web. *Comput. Netw.*, 33(1-6):309–320, 2000.
- [20] Boris E. Brodsky and Boris S. Darkhovsky. *Nonparametric Methods in Change-Point Problems*. Kluwer Academic Publishers, 1993.
- [21] Michael T. Burr. Reliability demands will drive automation investments. *Public Utilities Fortnightly*, 2003. <http://www.fortnightly.com/fortnightly/2003/11/technology-corridor>.
- [22] K. M. Carley. *Smart agents and the organizations of the future*. The Handbook of New Media, Thousand Oaks, CA, 2002.

- [23] B. A. Carreras, V. E. Lynch, I. Dobson, and D. E. Newman. Critical points and transitions in an electric power transmission model for cascading failure black-outs. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 12(4):985–994, 2002.
- [24] G. Chalkiadakis, V. Robu, R. Kota, A. Rogers, and N. R. Jennings. Cooperatives of distributed energy resources for efficient virtual power plants. In *Proc. of The 10th International Conference on Autonomous Agents and Multiagent Systems AAMAS*, pages 787–794, 2011.
- [25] W.L. Chan, A. T P So, and L.L. Lai. Harmonics load signature recognition by wavelets transforms. In *Electric Utility Deregulation and Restructuring and Power Technologies, 2000. Proceedings. DRPT 2000. International Conference on*, pages 666–671, 2000.
- [26] R.M. Ciric and D.S. Popovic. Multi-objective distribution network restoration using heuristic approach and mix integer programming method. *International Journal of Electrical Power Energy Systems*, 22(7):497 – 505, 2000.
- [27] K. Clement-Nyns, E. Haesen, and J. Driesen. The impact of charging plug-in hybrid electric vehicles on a residential distribution grid. *Power Systems, IEEE Transactions on*, 25(1):371–380, 2010.
- [28] Javier Contreras, Felix Wu, Matthias Klusch, and Onn Shehory. Coalition formation in a power transmission planning environment. In *Proc. of 2nd Intl. Conference on Practical Applications of Multi-Agent Systems, PAAM*, pages 21–23, 1997.
- [29] L.M. Costa, F. Bourry, J. Juban, and G. Kariniotakis. Management of energy storage coordinated with wind power under electricity market conditions. In *Proceedings of the 10th International Conference on Probabilistic Methods Applied to Power Systems, 2008. PMAAPS '08*, pages 1–8, 2008.
- [30] L.M. Costa, F. Bourry, J. Juban, and G. Kariniotakis. Management of energy storage coordinated with wind power under electricity market conditions. In *Probabilistic Methods Applied to Power Systems, 2008. PMAAPS '08. Proceedings of the 10th International Conference on*, pages 1–8, 2008.
- [31] V. D. Dang and N. R. Jennings. Generating coalition structures with finite bound from the optimal guarantees. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 564–571, 2004.
- [32] Rina Dechter. *Constraint Processing*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2003.

- [33] Paul Denholm and Robert M. Margolis. Evaluating the limits of solar photovoltaics (PV) in electric power systems utilizing energy storage and other enabling technologies. *Energy Policy*, 35(9):4424 – 4433, 2007.
- [34] Giovanna Di Marzo Serugendo, Marie-Pierre Gleizes, and Anthony Karageorgos. Self-organization in multi-agent systems. *Knowl. Eng. Rev.*, 20(2):165–189, 2005.
- [35] K. Dielmann and A. van der Velden. Virtual power plants (VPP) - a new perspective for energy generation? In *Modern Techniques and Technologies, 2003. MTT 2003. Proceedings of the 9th International Scientific and Practical Conference of Students, Post-graduates and Young Scientists*, pages 18–20, 2003.
- [36] A. L. Dimeas and N. D. Hatziargyriou. Agent based control of virtual power plants. In *International Conference on Intelligent Systems Applications to Power Systems (ISAP)*, pages 1–6, 2007.
- [37] M. R. Durling, Z. Ren, N. Visnevski, and L. E. Ray. Cognitive electric power meter. US Patent 7,693,670, 2010.
- [38] H. Dyke Parunak. A survey of environments and mechanisms for human-human stigmergy. In Danny Weyns, H. Dyke Parunak, and Fabien Michel, editors, *Environments for Multi-Agent Systems II*, volume 3830 of *Lecture Notes in Computer Science*, pages 163–186. Springer Berlin Heidelberg, 2006.
- [39] H. Dyke Parunak, SvenA. Brueckner, and John Sauter. Digital pheromones for coordination of unmanned vehicles. In Danny Weyns, H. Dyke Parunak, and Fabien Michel, editors, *Environments for Multi-Agent Systems*, volume 3374 of *Lecture Notes in Computer Science*, pages 246–263. Springer Berlin Heidelberg, 2005.
- [40] Stefan Edelkamp, Roman Englert, Jörg Hoffmann, Frederico dos S. Liporace, Sylvie Thiébaux, and Sebastian Trüg. Engineering benchmarks for planning: the domains used in the deterministic Part of IPC-4. *The Computing Research Repository (CoRR)*, 2011.
- [41] P. Erdos and A. Renyi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci.*, 5:17–61, 1960.
- [42] Ky Fan. Fixed-point and minimax theorems in locally convex topological linear spaces. *Proceedings of the National Academy of Sciences*, 38(2):121–126, 1952.
- [43] Alessandro Farinelli, Alex Rogers, Adrian Petcu, and N. R. Jennings. Decentralised coordination of low-power embedded devices using the max-sum

- algorithm. In *Seventh International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-08)*, pages 639–646, May 2008.
- [44] F. A. Farret and M. G. Simoes. *Integration of Alternative Sources of Energy*. NJ: Wiley-IEEE Press, Jan. 2006.
- [45] Stephen Fitzpatrick and Lambert Meertens. An experimental assessment of a stochastic, anytime, decentralized, soft colourer for sparse graphs. In Kathleen Steinhöfel, editor, *Stochastic Algorithms: Foundations and Applications*, volume 2264 of *Lecture Notes in Computer Science*, pages 49–64. Springer Berlin Heidelberg, 2001.
- [46] Markus Franke, Bettina Hoser, and Jan Schrder. On the analysis of irregular stock market trading behavior. In Christine Preisach, Hans Burkhardt, Lars Schmidt-Thieme, and Reinhold Decker, editors, *Data Analysis, Machine Learning and Applications*, Studies in Classification, Data Analysis, and Knowledge Organization, pages 355–362. Springer Berlin Heidelberg, 2008.
- [47] Zong Woo Geem. Size optimization for a hybrid photovoltaic-wind energy system. *International Journal of Electrical Power Energy Systems*, 42(1):448 – 451, 2012.
- [48] Enrico Gerding, Valentin Robu, Sebastian Stein, David Parkes, Alex Rogers, and Nick Jennings. Online mechanism design for electric vehicle charging. In *The Tenth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2011)*, pages 811–818, May 2011.
- [49] Zoubin Ghahramani and Michael I Jordan. Factorial hidden markov models. *Machine Learning*, 29(2-3):245–273, 1997.
- [50] Donald B. Gillies. Solutions to general non-zero-sum games. In Albert W. Tucker and R. Duncan Luce, editors, *Contributions to the Theory of Games IV*, Annals of Mathematics Studies, pages 47–85. Princeton University Press, 1959.
- [51] I. L. Glicksberg. A Further Generalization of the Kakutani Fixed Point Theorem, with Application to Nash Equilibrium Points. *Proceedings of the American Mathematical Society*, 3(1):170–174, 1952.
- [52] Sebastian Gottwalt, Wolfgang Ketter, Carsten Block, John Collins, and Christof Weinhardt. Demand side management - A simulation of household behavior under variable prices. *Energy Policy*, 39(12):8163 – 8174, 2011.
- [53] Jens Grossklags, Nicolas Christin, and John Chuang. Secure or insure?: A game-theoretic analysis of information security games. In *Proceedings of the*

- 17th International Conference on World Wide Web, WWW '08*, pages 209–218, New York, NY, USA, 2008. ACM.
- [54] Sidhant Gupta, Matthew S. Reynolds, and Shwetak N. Patel. Electrisesense: single-point sensing using emi for electrical event detection and classification in the home. In *Proceedings of the 12th ACM international conference on Ubiquitous computing, Ubicomp '10*, pages 139–148, New York, NY, USA, 2010. ACM.
- [55] Tarik Hadžic, Andrzej Wasowski, and Henrik R. Andersen. Techniques for efficient interactive configuration of distribution networks. In *Proceedings of the 20th international joint conference on Artificial intelligence, IJCAI'07*, pages 100–105, 2007.
- [56] Joseph Y. Halpern. A computer scientist looks at game theory. *Games and Economic Behavior*, 45(1):114–131, 2003.
- [57] J. Hammerstrom, J. Brous, D. P. Chassin, G. R. Horst, R. Kajfasz, P. Michie, T. V. Oliver, T. A. Carlon, C. Eustis, O. M. Jarvegren, W. Marek, R. L. Munson, and R. G. Pratt. Pacific northwest gridwise testbed demonstration projects Part II grid friendly appliance project. Technical report, Pacific Northwest National Laboratory, Richland, WA, October 2007. Tech. Rep. PNNL-17079, http://gridwise.pnnl.gov/docs/op_project_final_report_pnnl17167.pdf.
- [58] D. Hammerstrom et al. Pacific northwest gridwise testbed demonstration projects; Part I. olympic peninsula project. Technical report, Pacific Northwest National Laboratory, 2008. Tech. Rep. PNNL-17167 http://www.pnl.gov/main/publications/external/technical_reports/PNNL-17167.pdf.
- [59] Chris Harris. *Electricity markets: Pricing, structures & economics*. Publisher: Wiley, John Sons, 2011.
- [60] G. W. Hart. Nonintrusive appliance load monitoring. *Proceedings of the IEEE*, 80(12):1870–1891, 2002.
- [61] Nikos D. Hatziargyriou and Aris L. Dimeas. Operation of a multiagent system for microgrid control. *IEEE Transactions on Power Systems*, 20(3):1447–1455, 2005.
- [62] J.Z. Hernandez, S. Ossowski, and A. Garcia-Serrano. On multiagent coordination architectures: a traffic management case study. In *Proceedings of the 34th Annual Hawaii International Conference on System Sciences*, pages 1–9, 2001.

- [63] P. Hines and S. Blumsack. A centrality measure for electrical networks. In *Hawaii International Conference on System Sciences, Proceedings of the 41st Annual*, pages 185–185, 2008.
- [64] Katsutoshi Hirayama and Makoto Yokoo. Distributed partial constraint satisfaction problem. In Gert Smolka, editor, *Principles and Practice of Constraint Programming-CP97*, volume 1330 of *Lecture Notes in Computer Science*, pages 222–236. Springer Berlin Heidelberg, 1997.
- [65] L. M. Hofmann, N. Chakraborty, and K. Sycara. The evolution of cooperation in self-interested agent societies: A critical study. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '11*, pages 685–692. International Foundation for Autonomous Agents and Multiagent Systems, 2011.
- [66] Bryan Horling and Victor Lesser. A Survey of Multi-Agent Organizational Paradigms. *The Knowledge Engineering Review*, 19(4):281–316, 2005.
- [67] J. Hossack, S. D J McArthur, J.R. McDonald, J. Stokoe, and T. Cumming. A multi-agent approach to power system disturbance diagnosis. In *Power System Management and Control, 2002. Fifth International Conference on (Conf. Publ. No. 488)*, pages 317–322, 2002.
- [68] H. Hotelling. Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology*, 24(6):417–441, 1933.
- [69] The Philippine Wholesale Electricity Spot Market (WESM): <http://www.wesm.ph/>. <http://www.wesm.ph/>.
- [70] T. Hubert and S. Grijalva. Realizing smart grid benefits requires energy optimization algorithms at residential level. In *Innovative Smart Grid Technologies (ISGT), 2011 IEEE PES*, pages 1–8, 2011.
- [71] D.G. Infield, J. Short, C. Home, and L.L. Freris. Potential for domestic dynamic demand-side management in the uk. In *Power Engineering Society General Meeting, 2007. IEEE*, pages 1 –6, june 2007.
- [72] The European Electricity Grid Initiative. Roadmap 2010-18 and detailed implementation plan 2010-12. Technical report, EEGI, 2010. http://ec.europa.eu/energy/technology/initiatives/doc/grid_implementation_plan_final.pdf.
- [73] McKinsey Global Institute. Disruptive technologies: Advances that will transform life, business, and the global economy. Technical report, MGI, 2013. http://www.mckinsey.com/insights/business_technology/disruptive_technologies.

- [74] M.N. Islam, S.M.R. Haque, K.M. Alam, and M. Tarikuzzaman. An approach to improve collusion set detection using MCL algorithm. In *Computers and Information Technology, 2009. ICCIT '09. 12th International Conference on*, pages 237–242, dec. 2009.
- [75] Antoch Jaromir, Hukov Marie, and Veraverbeke Noel. Change-point problem and bootstrap. *Journal of Nonparametric Statistics*, 5(2):123–144, 1995.
- [76] Shizuo Kakutani. A generalization of Brouwer’s fixed point theorem. *Duke Mathematical Journal*, 8:416–427, 1941.
- [77] S. Kamboj, W. Kempton, and K. S. Decker. Deploying power grid-integrated electric vehicles as a multi-agent system. In *Proceedings of the Tenth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, pages 13–20, May 2011.
- [78] W. D. Kellogg, M.H. Nehrir, G. Venkataramanan, and V. Gerez. Generation unit sizing and cost analysis for stand-alone wind, photovoltaic, and hybrid wind/PV systems. *Energy Conversion, IEEE Transactions on*, 13(1):70–75, 1998.
- [79] Steven P. Ketchpel. Coalition formation among autonomous agents. In *Proc. European Workshop Modeling Autonomous Agents in a Multi-agent World (MAAMAW)*, pages 73–88, 1993.
- [80] Wolfgang Ketter, John Collins, and Prashant Reddy. Power TAC: A competitive economic simulation of the smart grid. *Energy Economics*, 39(0):262 – 270, 2013.
- [81] Christopher Kiekintveld, Zhengyu Yin, Atul Kumar, and Milind Tambe. Asynchronous algorithms for approximate distributed constraint optimization with quality bounds. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1, AAMAS '10*, pages 133–140, Richland, SC, 2010. International Foundation for Autonomous Agents and Multiagent Systems.
- [82] Hyungsul Kim, Manish Marwah, Martin F. Arlitt, Geoff Lyon, and Jiawei Han. Unsupervised disaggregation of low frequency power measurements. In *Proceedings of the Eleventh SIAM International Conference on Data Mining, SDM 2011, April 28-30, 2011, Mesa, Arizona, USA*, pages 747–758, 2011.
- [83] Matthias Klusch and Andreas Gerber. Dynamic coalition formation among rational agents. *IEEE Intelligent Systems*, 17(3):42–47, 2002.

- [84] J. K. Kok, C. J. Warmer, and I. G. Kamphuis. Powermatcher: multiagent control in the electricity infrastructure. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, AAMAS '05, pages 75–82, New York, NY, USA, 2005. ACM.
- [85] K. Kok, M. Scheepers, and R. Kamphuis. Intelligence in electricity networks for embedding renewables and distributed generation. In *Intelligent Infrastructures*. Springer, 42:179–209, 2010.
- [86] Koen Kok, S. Karnouskos, David Nestle, Aris Dimeas, Anke Weidlich, C. Warmer, Philipp Strauss, B. Buchholz, Stefan Drenkard, N. Hatziaargyriou, and Vali Lioliou. Smart houses for a smart grid. In *Electricity Distribution - Part 1, 2009. CIRED 2009. 20th International Conference and Exhibition on*, pages 1–4, 2009.
- [87] Koen Kok, Gerben Venekamp, and Pamala Macdougall. Market-based control in decentralized electrical power systems. In *Proceedings of the International Workshop on Agent Technologies for Energy Systems (ATES@AAMAS)*, pages 61–66, 2010.
- [88] J. Zico Kolter and Tommi Jaakkola. Approximate inference in additive factorial HMMs with application to energy disaggregation. *Journal of Machine Learning Research - Proceedings Track*, 22:1472–1482, 2012.
- [89] J. Zico Kolter and Matthew J. Johnson. REDD: A public data set for energy disaggregation research. In *Workshop on Data Mining Applications in Sustainability (SustKDD)*, 2011.
- [90] Farinaccio L. and Zmeureanu R. Using a pattern recognition approach to disaggregate the total electricity consumption in a house into the major end-uses. *Energy and Buildings*, 30(3):245 – 259, 1999.
- [91] H. Y. Lam, G. S K Fung, and W. K. Lee. A novel method to construct taxonomy electrical appliances based on load signaturesof. *Consumer Electronics, IEEE Transactions on*, 53(2):653–660, 2007.
- [92] Robert H. Lande and John M. Connor. How High Do Cartels Raise Prices? Implications for Reform of the Antitrust Sentencing Guidelines. *Tulane Law Review*, pages 513–570, 2005.
- [93] C. Laughman, Kwangduk Lee, R. Cox, S. Shaw, S. Leeb, L. Norford, and P. Armstrong. Power signature analysis. *IEEE Power and Energy Magazine*, 1(2):56–63, 2003.

- [94] Thomas Léauté and Boi Faltings. E[DPOP]: Distributed Constraint Optimization under Stochastic Uncertainty using Collaborative Sampling. In *Proceedings of the IJCAI'09 Distributed Constraint Reasoning Workshop (DCR'09)*, pages 87–101, Pasadena, California, USA, July 2009.
- [95] Thomas Léauté and Boi Faltings. Distributed Constraint Optimization under Stochastic Uncertainty. In *Proceedings of the Twenty-Fifth Conference on Artificial Intelligence (AAAI'11)*, pages 68–73, San Francisco, USA, August 2011.
- [96] Thomas Leaute, Brammert Ottens, and Radoslaw Szymanek. Frodo 2.0: An open-source framework for distributed constraint optimization. In *Proceedings of the IJCAI09 Distributed Constraint Reasoning Workshop (DCR09)*, pages 160–164, 2009.
- [97] K.D. Lee, S.B. Leeb, L.K. Norford, P.R. Armstrong, J. Holloway, and S.R. Shaw. Estimation of variable-speed-drive power consumption from harmonic content. *IEEE Transactions on Energy Conversion*, 20(3):566–574, 2005.
- [98] W K Lee, G S K Fung, H Y Lam, F H Y Chan, and Mark Lucente. Exploration on load signatures. In *International Conference on Electrical Engineering (ICEE)*, volume 2, pages 690–694, 2004.
- [99] S.B. Leeb, S.R. Shaw, and Jr. Kirtley, J.L. Transient event detection in spectral envelope estimates for nonintrusive load monitoring. *IEEE Transactions on Power Delivery*, 10(3):1200–1210, 1995.
- [100] Hongyan Li and Leigh Tesfatsion. Development of open source software for power market research: The AMES test bed. *Journal of Energy Markets*, 2(2):111–128, 2009.
- [101] Marco Mamei and Franco Zambonelli. *Field-Based Coordination for Pervasive Multiagent Systems (Springer Series on Agent Technology)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2005.
- [102] S. Mariani and A. Omicini. Self-organising news management: the molecules of knowledge approach. In *Self-Adaptive and Self-Organizing Systems Workshops (SASOW), 2012 IEEE Sixth International Conference on*, pages 235–240.
- [103] F. Von Martial. *Coordinating Plans of Autonomous Agents*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1992.
- [104] S. D. J. McArthur, E. M. Davidson, V. M. Catterson, A. L. Dimeas, N. D. Hatziargyriou, F. Ponci, and T. Funabashi. Multi-agent systems for power engineering applications - Part 1: Concepts, approaches, and technical challenges. *IEEE Transactions on Power Systems*, 22(4):1743–1752, 2007.

-
- [105] P. McDaniel and S. McLaughlin. Security and privacy challenges in the smart grid. *Security Privacy, IEEE*, 7(3):75–77, may-june 2009.
- [106] Amnon Meisels. *Distributed Search by Constrained Agents: Algorithms, Performance, Communication*. Advanced Information and Knowledge Processing. Springer, Berlin, 2007.
- [107] W. Mert and W. Tritthart. Consumer acceptance of smart appliances. in: WP5 report from smart-a project. Technical report, 2008. http://www.smart-a.org/WP5_5_Consumer_acceptance_18_12_08.pdf.
- [108] R.-C. Mihailescu, M. Vasirani, and S. Ossowski. Towards agent-based virtual power stations via multi-level coalition formation. In *Proceedings of the International Workshop on Agent Technologies for Energy Systems (ATES@AAMAS)*, pages 107–109, 2010.
- [109] R.-C. Mihailescu, M. Vasirani, and S. Ossowski. Dynamic coalition adaptation for efficient agent-based virtual power plants. In *Proceedings of the Ninth German Conference on Multi-Agent System Technologies (MATES'11)*, pages 101–112, 2011.
- [110] R.-C. Mihailescu, M. Vasirani, and S. Ossowski. Dynamic coalition formation and adaptation for virtual power stations in smart grids. In *Proceedings of the 2nd Int. Workshop on Agent Technologies for Energy Systems (ATES@AAMAS)*, pages 85–88, 2011.
- [111] R.-C. Mihailescu, M. Vasirani, and S. Ossowski. A game-theoretic coordination framework for energy-efficient demand-side management in the smart grid. In *Proceedings of the Ninth European Workshop on Multi-agent Systems (EU-MAS)*, 2011.
- [112] Pragnesh Jay Modi, Wei-Min Shen, Milind Tambe, and Makoto Yokoo. Adopt: asynchronous distributed constraint optimization with quality guarantees. *Artificial Intelligence*, 161(12):149–180, 2005.
- [113] A. Molderink, V. Bakker, M.G.C. Bosman, J.L. Hurink, and G. J M Smit. Domestic energy management methodology for optimizing efficiency in smart grids. In *PowerTech, 2009 IEEE Bucharest*, pages 1–7, 2009.
- [114] Roger Myerson. *Game theory: Analysis of conflict*. Harvard University Press, 1991.
- [115] T. Nagata and H. Sasaki. A multi-agent approach to power system restoration. *Power Systems, IEEE Transactions on*, 17(2):457–462, 2002.

- [116] Elena Nardini, Andrea Omicini, Mirko Viroli, and Michael Ignaz Schumacher. Coordinating e-health systems with TuCSon semantic tuple centres. *Applied Computing Review*, 11(2):43–52, 2011.
- [117] Guy R. Newsham and Brent G. Bowker. The effect of utility time-varying pricing and load control strategies on residential summer peak electricity use: A review. *Energy Policy*, 38(7):3289 – 3296, 2010.
- [118] Kah-Hoe Ng and G.B. Sheble. Direct load control - A profit-based load management using linear programming. *Power Systems, IEEE Transactions on*, 13(2):688–694, 1998.
- [119] Leslie K. Norford and Steven B. Leeb. Non-intrusive electrical load monitoring in commercial buildings based on steady-state and transient load-detection algorithms. *Energy and Buildings*, 24(1):51 – 64, 1996.
- [120] Martin Nowak and Karl Sigmund. A strategy of win-stay, lose-shift that outperforms tit-for-tat in the prisoner’s dilemma game. *Nature*, 364:56–58, 1993.
- [121] K.E. Nygard, S.B. Ghosn, M.M. Chowdhury, D. Loegering, R. McCulloch, and P. Ranganathan. Optimization models for energy reallocation in a smart grid. In *Computer Communications Workshops (INFOCOM WKSHPS), 2011 IEEE Conference on*, pages 186–190, 2011.
- [122] Department of Energy and Climate Change. Smarter grids: The opportunity. <http://www.decc.gov.uk/assets/decc/smartergridsopportunity.pdf>, 2009.
- [123] U.S. Department of Energy. Grid 2030: A national vision for electricity’s second 100 years. Technical report, 2003. http://energy.gov/sites/prod/files/oeprod/DocumentsandMedia/Electric_Vision_Document.pdf.
- [124] The Association of Home Appliance Manufacturers. The home appliance industry’s principles & requirements for achieving a widely accepted smart grid. Technical report, AHAM, 2009.
- [125] Andrea Omicini and Sascha Ossowski. Objective versus subjective coordination in the engineering of agent systems. In Matthias Klusch, Sonia Bergamaschi, Pete Edwards, and Paolo Petta, editors, *Intelligent Information Agents*, volume 2586 of *Lecture Notes in Computer Science*, pages 179–202. Springer Berlin Heidelberg, 2003.
- [126] M. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

- [127] Sascha Ossowski and Ana Garca-Serrano. Social structure in artificial agent societies: Implications for autonomous problem-solving agents. In Jörg Müller, Anand Rao, and Munindar Singh, editors, *Intelligent Agents V: Agents Theories, Architectures, and Languages*, volume 1555 of *Lecture Notes in Computer Science*, pages 133–148. Springer Berlin Heidelberg, 1999.
- [128] Sascha Ossowski and Ronaldo Menezes. On coordination and its significance to distributed and multi-agent systems: Research articles. *Concurr. Comput. : Pract. Exper.*, 18(4):359–370, 2006.
- [129] Pascal Paillier and David Pointcheval. Efficient public-key cryptosystems provably secure against active adversaries. In *ASIACRYPT*, pages 165–179, 1999.
- [130] Girish Keshav Palshikar and Manoj M. Apte. Collusion set detection using graph clustering. *Data Min. Knowl. Discov.*, 16:135–164, April 2008.
- [131] Oliver Parson, Siddhartha Ghosh, Mark Weal, and Alex Rogers. Non-intrusive load monitoring using prior models of general appliance types. In *Twenty-Sixth Conference on Artificial Intelligence (AAAI-12)*, pages 356–362, July 2012.
- [132] Shwetak Patel, Thomas Robertson, Julie Kientz, Matthew Reynolds, and Gregory Abowd. At the flick of a switch: Detecting and classifying unique electrical events on the residential power line. pages 271–288. 2007.
- [133] B. Peleg and P. Sudholter. *Introduction to the Theory of Cooperative Games*. Springer Verlag, 2007.
- [134] Adrian Petcu and Boi Faltings. Approximations in distributed optimization. In Peter Beek, editor, *Principles and Practice of Constraint Programming - CP 2005*, volume 3709 of *Lecture Notes in Computer Science*, pages 802–806. Springer Berlin Heidelberg, 2005.
- [135] Adrian Petcu and Boi Faltings. MB-DPOP: A new memory-bounded algorithm for distributed optimization. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence, IJCAI'07*, pages 1452–1457, San Francisco, CA, USA, 2007. Morgan Kaufmann Publishers Inc.
- [136] Adrian Petcu, Boi Faltings, and David C. Parkes. M-dpop: Faithful distributed implementation of efficient social choice problems. *J. Artif. Int. Res.*, 32(1):705–755, 2008.
- [137] M. Pielke, M. Troschel, M. Kurrat, and H.-J. Appelrath. Operation strategies to integrate CHP micro units in domestic appliances into the public power supply. In *Proceedings of the VDE-Kongress*, 2008.

- [138] M. Pipattanasomporn, H. Feroze, and S. Rahman. Multi-agent systems in a distributed smart grid: Design and implementation. In *Power Systems Conference and Exposition, 2009. PSCE '09. IEEE/PES*, pages 1–8, 2009.
- [139] M. Pipattanasomporn and S. Rahman. Intelligent distributed autonomous power systems (IDAPS) and their impact on critical electrical loads. In *Critical Infrastructure Protection, First IEEE International Workshop on*, pages 101–109, 2005.
- [140] European SmartGrids Technology Platform. Vision and strategy for european electricity networks of the future. Technical report, European Union, 2006. http://ec.europa.eu/research/energy/pdf/smartgrids_en.pdf.
- [141] J. Ponstein. Seven kinds of convexity. *SIAM Review*, 9(1):115–119, January 1967.
- [142] Evangelos Pournaras, Martijn Warnier, and Frances M. T. Brazier. Local agent-based self-stabilisation in global resource utilisation. *Int. J. Autonomic Comput.*, 1:350–373, December 2010.
- [143] D Pudjianto, C Ramsay, and G Strbac. Virtual power plant and system integration of distributed energy resources. *IET Renewable Power Generation*, 1(1):10–16, 2007.
- [144] S. Rahman, M. Pipattanasomporn, and Y. Teklu. Intelligent distributed autonomous power systems (idaps). In *Power Engineering Society General Meeting, 2007. IEEE*, pages 1–8, 2007.
- [145] Talal Rahwan, Sarvapali D. Ramchurn, Viet Dung Dang, Andrea Giovannucci, and Nicholas R. Jennings. Anytime optimal coalition structure generation. In *Proceedings of the 22nd National Conference on Artificial Intelligence*, pages 1184–1190, 2007.
- [146] Sarvapali Ramchurn, Perukrishnen Vytelingum, Alex Rogers, and Nicholas R. Jennings. Putting the "Smarts" into the Smart Grid: A grand challenge for artificial intelligence. *Communications of the ACM*, 55(4):86–97, 2012.
- [147] Sarvapali Ramchurn, Perukrishnen Vytelingum, Alex Rogers, and Nick Jennings. Agent-based homeostatic control for green energy in the smart grid. *ACM Transactions on Intelligent Systems and Technology*, 2(4):1–35, May 2011.
- [148] Sarvapali D. Ramchurn, Perukrishnen Vytelingum, Alex Rogers, and Nicholas R. Jennings. Agent-based control for decentralised demand side management in the smart grid. In *Proc. of The 10th International Conference on Autonomous Agents and Multiagent Systems AAMAS*, pages 5–12, 2011.

- [149] H.F. Rashvand, K. Salah, J.M.A. Calero, and L. Harn. Distributed security for multi-agent systems - review and applications. *Information Security, IET*, 4(4):188–201, december 2010.
- [150] Leandro Chaves Rêgo and Joseph Y. Halpern. Generalized solution concepts in games with possibly unaware players. *International Journal of Game Theory*, 41(1):131–155, 2012.
- [151] Valentin Robu, Ramachandra Kota, Georgios Chalkiadakis, Alex Rogers, and Nicholas Jennings. Cooperative virtual power plant formation using scoring rules. In *AAAI-12: Twenty-Sixth Conference on Artificial Intelligence*, pages 370–376, 2012.
- [152] Davide Rossi, Giacomo Cabri, and Enrico Denti. Tuple-based technologies for coordination. In Andrea Omicini, Franco Zambonelli, Matthias Klusch, and Robert Tolksdorf, editors, *Coordination of Internet Agents*, pages 83–109. Springer Berlin Heidelberg, 2001.
- [153] M. H. Rothkopf, A. Pekec, and R. M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1995.
- [154] H. Salehfar, P. J. Noll, B.J. LaMeres, M.H. Nehrir, and V. Gerez. Fuzzy logic-based direct load control of residential electric water heaters and air conditioners recognizing customer preferences in a deregulated environment. In *Power Engineering Society Summer Meeting, 1999. IEEE*, volume 2, pages 1055–1060 vol.2, 1999.
- [155] T. W. Sandholm, K. Larson, M. Andersson, O. Shehory, and F. Tohme. Coalition structure generation with worst case guarantees. *Artificial Intelligence*, 111(1-1):209–238, 1999.
- [156] F. C. Santos, J. M. Pacheco, and T. Lenaerts. Evolutionary dynamics of social dilemmas in structured heterogeneous populations. In *Proc. Natl. Acad. Sci.*, volume 103, pages 3490–3494, USA, 2006.
- [157] L. J. Savage. The theory of statistical decision. *Journal of the American Statistical Association*, 46(253):55–67, 1951.
- [158] C. Schulz, G. Roder, and M. Kurrat. Virtual power plants with combined heat and power micro-units. In *Future Power Systems, 2005 International Conference on*, pages 1–5, 2005.
- [159] Michael Schumacher. *Objective Coordination in Multi-agent System Engineering: Design and Implementation*. Springer-Verlag, Berlin, Heidelberg, 2001.

- [160] S. Sen and P. Dutta. Searching for optimal coalition structures. *In Proceedings of the Fourth International Conference on Multiagent Systems*, pages, pages 286–292, 2000.
- [161] L. S. Shapley. A value for n-person games. *Contributions to the theory of games*, 2:307–317, 1953.
- [162] Onn Shehory and Sarit Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(12):165 – 200, 1998.
- [163] V. Silva, V. Stanojevic, D. Pudjianto, and G. Strbac. Value of smart domestic appliances in stressed electricity networks. in: WP4 report from smart-a project. Technical report, Imperial College, London, 2009. http://www.smart-a.org/W_P_4_D_4_4_Energy_Networks_Report_final.pdf.
- [164] Reid G. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. *Computers, IEEE Transactions on*, C-29(12):1104–1113, 1980.
- [165] Leen-Kiat Soh and Costas Tsatsoulis. A real-time negotiation model and a multi-agent sensor network implementation. *Autonomous Agents and Multi-Agent Systems*, 11(3):215–271, 2005.
- [166] S. Srivastava, J.R.P. Gupta, and M. Gupta. PSO & neural-network based signature recognition for harmonic source identification. In *TENCON 2009 - 2009 IEEE Region 10 Conference*, pages 1–5, 2009.
- [167] G. Strbac. Demand side management: Benefits and challenges. *Energy Policy*, 36(12):4419–4426, 2008.
- [168] Lukas G. Swan and V. Ismet Ugursal. Modeling of end-use energy consumption in the residential sector: A review of modeling techniques. *Renewable and Sustainable Energy Reviews*, 13(8):1819 – 1835, 2009.
- [169] K. Sycara, A. Pannu, M. Williamson, Dajun Zeng, and K. Decker. Distributed intelligent agents. *IEEE Expert*, 11(6):36–46, 1996.
- [170] K. Sycara and Dajun Zeng. Coordination of multiple intelligent software agents. *International Journal of Cooperative Information Systems*, 05(02):181–211, 1996.
- [171] G. Tel. *Introduction To Distributed Algorithms, 2nd ed.* Cambridge Univ Press, 2000.

- [172] Sylvie Thiebaux, Carleton Coffrin, Hassan Hijazi, and John Slaney. Planning with MIP for supply restoration in power distribution systems. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2900–2907, Beijing, August 2013.
- [173] Spyridon L. Tompros, Nikolaos P. Mouratidis, M. Draaijer, Andreas Foglar, and Halid Hrasnica. Enabling applicability of energy saving applications on the appliances of the home environment. *IEEE Network*, 23(6):8–16, 2009.
- [174] D. Trudnowski, M. Donnelly, and E. Lightner. Power-system frequency and stability control using decentralized intelligent loads. *2005/2006 PES TD Conference and Exposition*, pages 1453–1459, 2005.
- [175] Akshay Uttama-Nambi S. N., Thanasis G. Papaioannou, Dipanjan Chakraborty, and Karl Aberer. Sustainable Energy Consumption Monitoring in Residential Settings. In *2nd IEEE INFOCOM Workshop on Communications and Control for Smart Energy Systems (CCSES 2013)*, pages 1–6, 2013.
- [176] Adriaan van der Welle, Carlos Madina, Christos Kolokathis, Jaap Jansen, and Angel Diaz. Financial and socio-economic impacts of embracing the Fenix concept. Technical report, ECN, 2009. <http://fenix.iwes.fraunhofer.de/html/documents.htm>.
- [177] Stijn Vandael, Nelis Boucké, Tom Holvoet, Klaas De Craemer, and Geert Deconinck. Decentralized coordination of plug-in hybrid vehicles for imbalance reduction in a smart grid. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2, AAMAS '11*, pages 803–810. International Foundation for Autonomous Agents and Multiagent Systems, 2011.
- [178] Stijn Vandael, Sachin Kamboj, Tom Holvoet, Geert Deconinck, and Willett Kempton. A comparison of two GIV mechanisms for providing ancillary services at the university of delaware. In *Proceedings of the 4th IEEE International Conference on Smart Grid Communications (SmartGridComm 2013)*, pages 211–216, 2013.
- [179] Matteo Vasirani and Sascha Ossowski. A proportional share allocation mechanism for coordination of plug-in electric vehiclecharging. *Eng. Appl. of AI*, 26(3):1185–1197, 2013.
- [180] Matteo Vasirani, Sascha Ossowski, Ramachandra Kota, Renato L. G. Cavalcante, and Nicholas R. Jennings. Using coalitions of wind generators and electric vehicles for effective energy market participation. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 3, AAMAS '11*, pages 1099–1100. International Foundation for Autonomous Agents and Multiagent Systems, 2011.

- [181] Meritxell Vinyals, A Bistaffa, Alessandro Farinelli, and Alex Rogers. Stable coalition formation among energy consumers in the smart grid. In *Proceedings of the 3th International Workshop on Agent Technologies for Energy Systems (ATES@AAMAS)*, pages 73–80, 2012.
- [182] Mirko Viroli and Matteo Casadei. Biochemical tuple spaces for self-organising coordination. In John Field and VascoT. Vasconcelos, editors, *Coordination Models and Languages*, volume 5521 of *Lecture Notes in Computer Science*, pages 143–162. Springer Berlin Heidelberg, 2009.
- [183] Mirko Viroli, Matteo Casadei, Elena Nardini, and Andrea Omicini. Towards a pervasive infrastructure for chemical-inspired self-organising services. In Danny Weyns, Sam Malek, Rogrio Lemos, and Jesper Andersson, editors, *Self-Organizing Architectures*, volume 6090 of *Lecture Notes in Computer Science*, pages 152–176. Springer Berlin Heidelberg, 2010.
- [184] P. Vytelingum, T. D. Voice, S. D. Ramchurn, A. Rogers, and N. R. Jennings. Agent-based micro-storage management for the smart grid. *Proc. of the 8th Conf. Autonomous Agents And MultiAgent Systems, AAMAS*, pages 39–46, 2010.
- [185] Lingfeng Wang, Zhu Wang, and Rui Yang. Intelligent multiagent control system for energy and comfort management in smart and sustainable buildings. *IEEE Transactions on Smart Grid*, 3(2):605–617, 2012.
- [186] Yodai Watanabe, Junji Shikata, and Hideki Imai. Equivalence between semantic security and indistinguishability against chosen ciphertext attacks. In *Proceeding of Public Key Cryptography (PKC '03)*, volume 2567 of *Lecture Notes in Computer Science*, pages 71–84. Springer, 2003.
- [187] J. Watts, D. and H. Strogatz, S. Collective dynamics of small-world networks. *Nature*, 393(6684):440–442, 1998.
- [188] B. M. Weedy and B. J. Cory. *Electric Power Systems*. John Wiley Sons, 2004. 4th edition.
- [189] Deh-Chang Wei and Nanming Chen. Air conditioner direct load control by multi-pass dynamic programming. *Power Systems, IEEE Transactions on*, 10(1):307–313, 1995.
- [190] W. Wichakool, A. T Avestruz, R.W. Cox, and S.B. Leeb. Modeling and estimating current harmonics of variable electronic loads. *Power Electronics, IEEE Transactions on*, 24(12):2803–2811, 2009.
- [191] M. Wooldridge. *An Introduction to MultiAgent Systems, 2nd ed.* John Wiley and Sons, 2009.

-
- [192] H.-T. Yang and K.-Y. Huang. Direct load control using fuzzy dynamic programming. *Generation, Transmission and Distribution, IEE Proceedings-*, 146(3):294–300, 1999.
- [193] William Yeoh, Ariel Felner, and Sven Koenig. BnB-ADOPT: An asynchronous branch-and-bound DCOP algorithm. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 2, AAMAS '08*, pages 591–598, Richland, SC, 2008. International Foundation for Autonomous Agents and Multiagent Systems.
- [194] M. Yokoo, T. Ishida, E.H. Durfee, and K. Kuwabara. Distributed constraint satisfaction for formalizing distributed problem solving. In *Distributed Computing Systems, 1992., Proceedings of the 12th International Conference on*, pages 614–621, 1992.
- [195] G. Zhang, J. Jiang, N. Xia, and Z. Su. Partical swarms cooperative optimization for coalition generation problem. In *Simulated Evolution and Learning, Springer Berlin/Heidelberg*, pages 166–173, 2006.
- [196] M. G. Zimmermann and V. M. Eguiluz. Cooperation, social networks, and the emergence of leadership in a prisoners dilemma with adaptive local interactions. In *Physical Review E (Statistical, Nonlinear, and Soft Matter Physics)*, volume 72(5), pages 56–118. APS, 2005.