HAL
open science

A bottom-up approach to energy management in a Smart Grid : adaptive and reactive model using a decentralized architecture for a generic user-centered system enabling large scale deployments

Victor Lequay

To cite this version:

Victor Lequay. A bottom-up approach to energy management in a Smart Grid : adaptive and reactive model using a decentralized architecture for a generic user-centered system enabling large scale deployments. Artificial Intelligence [cs.AI]. Université de Lyon, 2019. English. NNT : 2019LYSE1304 . tel-02519746v2

HAL Id: tel-02519746

https://theses.hal.science/tel-02519746v2

Submitted on 31 Mar 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THESE de DOCTORAT DE L'UNIVERSITE DE LYON

opérée au sein de

l'Université Claude Bernard Lyon 1

Ecole Doctorale Informatique et Mathématiques
ED512

Spécialité de doctorat : Intelligence artificielle

Discipline : Informatique

Soutenue publiquements le 11/12/2019, par :

Victor Lequay

Une approche ascendante pour la gestion énergétique d'une Smart-Grid : Modèle adaptatif et réactif fondé sur une architecture décentralisée pour un système générique centré sur l'utilisateur permettant un déploiement à grande échelle.

Devant le jury composé de :

Kheddouci, Hamamache	Professeur à l'Université Lyon 1	Président
Beal, Jacob	Chercheur à l'Université de l'Iowa	Rapporteur
Gleizes, Marie-Pierre	Professeure à l'Université Toulouse 3	Rapporteure
Diagonescu, Ada	Maître de conférences à Telecom Paris	Rapporteure
Hassas, Salima	Professeure à l'Université Lyon 1	Directrice de thèse
Lefort, Mathieu	Maître de conférences à l'Université Lyon 1	Co-encadrant
Mansour, Saber	Docteur, CTO à Ubiant	Encadrant entreprise
Schmeck, Hartmut	Professeur à l'Université de Karlsruhe	Invité

Thesis

A bottom-up approach to energy management in a Smart Grid : Adaptive and reactive model using a decentralized architecture for a generic user-centered system enabling large scale deployments.

Victor Lequay

Foreword

Abstract

The field of Energy Management Systems for Smart Grids has been extensively explored in recent years, with many different approaches being described in the literature. In collaboration with our industrial partner Ubiant, which deploys smart homes solutions, we identified a need for a highly robust and scalable system that would exploit the flexibility of residential consumption to optimize energy use in the smart grid. At the same time we observed that the majority of existing works focused on the management of production and storage only, and that none of the proposed architectures are fully decentralized. Our objective was then to design a dynamic and adaptive mechanism to leverage every existing flexibility while ensuring the user's comfort and a fair distribution of the load balancing effort ; but also to offer a modular and open platform with which a large variety of devices, constraints and even algorithms could be interfaced. In this thesis we realised (1) an evaluation of state of the art techniques in real-time individual load forecasting, whose results led us to follow (2) a bottom-up and decentralized approach to distributed residential load shedding system relying on a dynamic compensation mechanism to provide a stable curtailment. On this basis, we then built (3) a generic user-centered platform for energy management in smart grids allowing the easy integration of multiple devices, the quick adaptation to changing environment and constraints, and an efficient deployment.

Résumé

Le domaine de la gestion de l'énergie dans les smart grids a été largement exploré ces dernières années, de nombreuses approches différentes étant proposées dans la littérature. En collaboration avec notre partenaire industriel Ubiant, qui déploie des solutions d'optimisation énergétique du bâtiment, nous avons mis en évidence le besoin d'un système vraiment robuste et évolutif exploitant la flexibilité de la consommation résidentielle pour optimiser l'utilisation de l'énergie au sein d'une smart grid. Dans le même temps, nous avons observé que la majorité des travaux existants se concentraient sur la gestion de la production et du stockage et qu'aucune des architectures proposées n'étaient véritablement décentralisées. Notre objectif était alors de concevoir un mécanisme dynamique et adaptatif permettant de tirer parti de toute la flexibilité existante tout en garantissant le confort de l'utilisateur et une répartition équitable des efforts d'équilibrage ; mais aussi de proposer une plate-forme ouverte et modulaire avec laquelle une grande variété d'appareils, de contraintes et même d'algorithmes pourraient être interfacés. Dans cette thèse, nous avons réalisé (1) une évaluation des techniques actuelles de prévision de la consommation individuelle en temps réel, dont les résultats nous ont amenés à suivre (2) une approche ascendante et décentralisée pour l'ajustement diffus résidentiel reposant sur un mécanisme de compensation pour assurer un effacement stable. Sur cette base, nous avons ensuite construit (3) une plate-forme générique centrée sur l'utilisateur pour la gestion de l'énergie dans les réseaux intelligents, permettant une intégration aisée de plusieurs périphériques, une adaptation rapide à l'évolution de l'environnement et des contraintes, ainsi qu'un déploiement efficace.

Remerciements

D'abord je veux remercier Salima Hassas et Mathieu Lefort pour leur encadrement efficace et leur suivi continu, pour la confiance qu'ils m'ont accordée durant ces quatre années de travail. J'ai pu découvrir le monde de la recherche académique dans les meilleures conditions possibles et apprendre la rigueur nécessaire au travail scientifique grâce à leurs conseils avisés. Merci Salima pour ton guidage avisé à travers ce labyrinthe qu'est la thèse, pour ton optimisme et tes connaissance du monde des SMA. Merci également d'avoir su m'orienter vers le master IA, merci infiniment pour toutes ces opportunités. Merci Mathieu d'avoir passé tant de temps à corriger mes manuscrits malgré mes avancées laborieuses, d'avoir tenu ces longues discussions et d'avoir su faire face à mon entêtement quand il le fallait, j'ai énormément appris et j'espère que je saurai le transmettre comme tu l'as fait.

Je veux remercier également toute l'équipe d'Ubiant pour le cadre idéal dans lequel j'ai pu mener à terme ce travail, pour les connaissances techniques que j'ai pu acquérir et les relations humaines que j'ai pu construire. Merci en particulier à Emmanuel Olivier, Saber Mansour, Olivier Lefèvre, Sebastien Mazac, Cédric Morin, Sylvain Montessuit, Aurélien Langlois et Romain Gagou pour leur soutien et leur aide, tant morale que technique, sans laquel il aurait été difficile d'atteindre mes objectifs.

Je remercie bien sûr les membres du jury : les rapporteur.es Jacob Beal, Marie-Pierre Gleizes et Ada Diagonescu pour leurs questions, critiques et conseils constructifs qui ont considérablement nourri ma réflexion autour de mon travail et ont permis une prise de recul nécessaire. Merci beaucoup à Hamamache Kheddouci d'avoir accepté de présider mon jury mais également d'avoir été la porte d'entrée dans le monde de la recherche durant mes études d'ingénieur. Un grand merci encore à Hartmut Schmeck, professeur de l'Université de Karlsruhe, pour son aide tout au long de ma thèse à travers ses relectures, ressources et références ; en plus d'avoir fait partie de mon comité de thèse ainsi que de mon jury de thèse en tant qu'invité.

Je veux évidement remercier mes parents sans qui tout mon parcours scolaire n'aurait pas été possible, mon père qui m'a transmis sa passion de l'informatique et ma mère sans qui je n'aurais probablement même pas eu le bac... Merci à Alexis Ringot, mon ami et camarade d'études pendant ce long parcours à qui je dois la confiance sans laquelle je ne me serais pas lancé dans cette aventure. Un énorme merci enfin à tou.te.s celles et ceux qui ont su être là et m'aider pendant cette longue période, chacun.es à leur manière. Une phrase ne suffit pas, mais vous vous reconnaîtrez.

Contents

Acknowledgement **i**
Abstract i
Résumé ii
Remerciements iii
Contents v

1 Introduction **1**
Global context and outline of contributions.
1.1 Context 3
1.1.1 The energy transition 3
1.1.2 The meaning of "smart" 3
1.1.3 The Smart Grid 5
1.1.4 Ubiant and HEMIS 6
1.2 Problem Statement 6
1.2.1 Optimizing energy use 6
1.2.2 Putting the user back in the loop 7
1.2.3 Facilitating deployment 7
1.3 Contributions 8
1.3.1 Very Short Term Residential Load Forecasting 8
1.3.2 Decentralized Residential Load Shedding 8
1.3.3 Generic and Adaptive Energy Management System 9
1.4 Outline of this thesis 9

2 Background and State of the art **11**
Definition of key concepts and more detailed contextual information.
2.1 The challenge of energy management 13
2.1.1 Electricity 13
2.1.1.1 Main concepts 13
2.1.1.2 A bit of history 14
2.1.2 Conventional power grids 15
2.1.2.1 General composition 15
2.1.2.2 Power grid balancing 18

2.1.2.3 Issues of the current system 19
2.1.3 Smart-grids : the new paradigm 20
2.1.3.1 Impact of climate change 20
2.1.3.2 Opportunities of information technologies 21
2.1.3.3 A diversity of solutions 21
2.2 Leveraging residential consumption 24
2.2.1 The weight of residential consumption 24
2.2.1.1 Current situation and trends 24
2.2.1.2 Load consideration in research 25
2.2.2 Engaging users 26
2.2.2.1 Economical considerations 26
2.2.2.2 Social factors 27
2.2.2.3 Security and privacy 28
2.2.3 The need for a user-centered system 28
2.3 Enabling large scale deployments 29
2.3.1 Handling the variety of situations 29
2.3.1.1 Genericity 29
2.3.1.2 Scalability 29
2.3.2 Adapting to changes 30
2.3.2.1 Modularity 30
2.3.2.2 Adaptiveness 30
2.3.3 Limiting the cost 31
2.3.3.1 Simplicity 31
2.3.3.2 Robustness 31
2.4 Multi-Agent Systems 32
2.4.1 Concepts 32
2.4.2 Applications in smart grids 33
2.5 Summary and positionning 34

3 Short Term Residential Load Forecasting 37
Study of short term residential load forecasting techniques and opportunities for improvement.
3.1 Problem statement 39
3.1.1 Aggregation level 39
3.1.2 Time scales 40
3.1.3 Local variables 40
3.2 State of the art 41
3.2.1 Clustering 41
3.2.2 Linear Regression Models 42
3.2.3 Non-linear regression 42
3.2.3.1 Artificial Neural Networks 43
3.2.4 Exogenous Variables 44

3.3 Methodology . . . 45
3.3.1 Datasets . . . 45
3.3.1.1 Ubiant's office . . . 45
3.3.1.2 UCI's dataset . . . 46
3.3.2 Algorithm parameters . . . 47
3.4 Results . . . 48
3.4.1 Very Short Term Load Forecasting . . . 48
3.4.2 Influence of exogenous variables . . . 49
3.5 Summary . . . 51

4 Decentralized User-centered Residential Demand Response 53
Demonstration of the feasibility of a bottom-up approach to residential load management.
4.1 Problem Statement . . . 55
4.1.1 Demand response mechanism . . . 55
4.1.2 User constraints . . . 56
4.2 State of the art . . . 57
4.2.1 Indirect load control . . . 57
4.2.2 Direct load control . . . 58
4.3 Proposed approach . . . 60
4.3.1 The Smart Building Solution: Hemis . . . 61
4.3.2 The Push-Sum algorithm . . . 62
4.3.3 Coordination problem . . . 63
4.4 Decentralized Load Shedding Model . . . 64
4.4.1 General principle . . . 64
4.4.2 Agent self-evaluation . . . 65
4.4.3 Engaging the agent's flexibility . . . 66
4.4.4 Adjusting agent's participation . . . 67
4.5 Experiments . . . 68
4.5.1 Test protocol . . . 68
4.5.2 Results . . . 70
4.5.2.1 Reliability, robustness and scalability . . . 70
4.5.2.2 Self-evaluation . . . 73
4.5.2.3 Security margin . . . 75
4.6 Summary . . . 78

5 Generic Adaptive Energy Management for the Smart Grid 81
Decentralized platform for the real-time scheduling of connected appliances using a generic model.
5.1 Problem statement . . . 83
5.1.1 Energy use optimization . . . 83

5.1.2 User satisfaction . . . 85
5.1.3 Process simplification . . . 85
5.2 State of the art . . . 86
5.2.1 User acceptance . . . 86
5.2.2 Scalability, modularity and robustness . . . 87
5.2.3 Market-based models . . . 89
5.3 Proposed approach and agent model . . . 89
5.3.1 General process . . . 90
5.3.2 Local constraints satisfaction . . . 91
5.3.3 Energy optimization . . . 92
5.3.4 Effort distribution . . . 93
5.3.5 Communication framework . . . 94
5.4 Implementation . . . 95
5.4.1 Inertia-based devices . . . 96
5.4.2 Task-based devices . . . 96
5.4.3 Generators . . . 97
5.4.4 Energy Storage Systems . . . 97
5.4.5 Non controllable assets . . . 98
5.5 Formal properties of the model . . . 98
5.5.1 Nash equilibrium . . . 98
5.5.2 Pareto efficiency . . . 99
5.6 Simulations . . . 99
5.6.1 Reliability and robustness . . . 100
5.6.2 Scalability . . . 100
5.6.3 Fairness . . . 103
5.7 Summary . . . 104

6 Conclusion and perspectives **107**
Summary and discussion of future works.
6.1 Objectives . . . 109
6.1.1 Users . . . 109
6.1.2 Deployment . . . 109
6.2 Contributions . . . 110
6.2.1 Load Forecasting . . . 110
6.2.2 Decentralized Residential Demand Response . . . 110
6.2.3 Smart Grid Energy Management System . . . 111
6.3 Future work . . . 112
6.3.1 Real life testing . . . 112
6.3.2 Security . . . 112
6.3.3 Further studies on load forecasting . . . 113
6.3.4 User interface . . . 113

A Annexes **115**

A.1 Prosumption profiles used in simulations 117
A.2 Publications 120

Bibliography **121**

List of Figures

2.1 Basic structure of a large scale electric power system. 15
2.2 Evolution of fuel sources for electricity generation in France 16
2.3 Investment in renewable energies in the last decades 17
2.4 Evolution in total energy consumption by region[4]. 18
2.5 End uses of residential electricity consumption in France. Source : RTE, Bilan Electrique 2017. 25

3.1 (a) An artificial neuron. (b) A simple two-layer feed-forward neural network. This figure is provided by [Hippert01]. 43
3.2 The Self-Recurrent Wavelet Neural Network used by Chitsaz et al. [Chitsaz15] (figure from the original paper). 47
3.3 Correlation coefficient between the different variables of the ubiant's office's dataset and the instantaneous consumption. 49
3.4 The minute by minute dataset from the Ubiant office. Only instantaneous electricity consumption (top) and occupancy (down) of the whole office are shown. 50
3.5 Error of the forecasting algorithm over increasing horizons, with and without occupancy data as input. 51

4.1 Main phases of an agent process. 65
4.2 Target crossing detection process. Here a crossing is confirmed when the aggregate value is above $\hat{F} * (1 + H)$ (green areas) for a duration of t_{wait}. The agent then starts adjusting its share to the effort. 68
4.3 Probability distribution of failure with time depending on the total flexibility. 69
4.4 Dynamic adjustment during two consecutive events without the self-evaluation mechanism. 71
4.5 Dynamic adjustment by 2000 agents on 4 consecutive events. 72
4.6 Evolution of the MSE for 14 consecutive events with and without the full evaluation mechanism, averaged over 50 runs. Colored areas show the respective standard deviation. 74

4.7 Error evolution during six consecutive events. (a) the agents engage their full flexibility at once. (b) the agents take their mark into account and engage immediately (c) the agents consider their mark for both engagement and reaction time. 76
4.8 Error evolution during six consecutive events. (a) the agents engage their all the flexibility allowed by their mark. (b) the agents retain a margin to help recovering from future failure. 77

5.1 Overview of the agent process loop. 90
5.2 Time discretization. During each period, each agent iterates over its local process until the optimization converges. 91
5.3 Mean RMSE over 50 consecutive runs with the settings described in section 5.6.1. Almost half of the agents are added during the optimization period (the vertical line shows the exact time). 101
5.4 3200 task-based appliances coordinating to fit the production of a single power source (section 5.6.2). RMSE measured every 200ms of one optimization period, averaged over 60 consecutive runs. 102
5.5 Mean effort of the agents as they settle on a satisfying solution with and without the *fairness factor* (section 5.6.3). 104

A.1 Seven different load profiles for dishwashers used in section 5.6.1. Source : KIT, OSHv4 117
A.2 5 different load profiles for washing machines used in section 5.6.1. Source : KIT, OSHv4 118
A.3 Production curve of a solar panel array used in section 5.6.1. Source : Ubiant 119

1
Introduction

▷ *The objective of this introduction is to give the general context and motivation for our work, describe the issues we want to tackle and briefly outline our contributions.* ◁

Chapter content

1.1 Context **3**
1.1.1 The energy transition 3
1.1.2 The meaning of "smart" 3
1.1.3 The Smart Grid 5
1.1.4 Ubiant and HEMIS 6
1.2 Problem Statement **6**
1.2.1 Optimizing energy use 6
1.2.2 Putting the user back in the loop 7
1.2.3 Facilitating deployment 7
1.3 Contributions **8**
1.3.1 Very Short Term Residential Load Forecasting 8
1.3.2 Decentralized Residential Load Shedding 8
1.3.3 Generic and Adaptive Energy Management System 9
1.4 Outline of this thesis **9**

1.1 Context

1.1.1 The energy transition

The way we produce, transport, and use the energy we need in our everyday life is changing. Adding to the systemic inefficiencies of current power grid management systems (which will be explained in section 2.1.2.3), the pressure of the climate change fuels the search for new approaches to power grid management. Indeed, the rise of renewable energy sources bring in new constraints : their various size, location and power question the traditional way of managing energy production that took the availability and stability of power plants for granted. Distributed energy resources like solar panels or wind turbines, combined with a growing number of local storage solutions like home battery or electric vehicles, will create new forms of power grids which will require custom management systems for each configuration. Conveniently, the constant advances in information technologies offers new opportunities, allowing a more dynamic control of the various assets of the network.

To ease the burden on the production side, the idea of adjusting the demand has gained traction in the last decades and the democratization of connected equipment in households opens the door to residential demand response. Household consumption represents a significant share of the total energy demand, but the integration of residential buildings into a dynamic energy management process raises a number of issues, notably when considering the comfort of the inhabitants. There is a need for a system with the ability to handle the large variety of possible use cases and existing devices, as well as the inherent complexity of the human factor, in order to build a smart grid.

1.1.2 The meaning of "smart"

The word *smart* originally defines[1] something that is quick, sharp, vigorous, and derived from describing a stinging pain to become an adjective for cleverness, wit, quick thinking. But what interests us here is the more recent meaning of *smart*. Embedded in *smartphone* in its most famous application, the use of this adjective used as a prefix can be traced back to the Vietnam war in a far less fortunate example : "smart bombs". Already, it conveyed the ideas that are still referred to today in the terms *smart TV*, *smart homes*, *smart buildings*, *smart grids*, and even *smart car*, among others. From the smallest home appliance (*smart plug*) to the ubiquitous *smart grid*, what really does this term represent? One could argue that it is a mere marketing tool, which we will not deny. However, as this word continues to be used even in scientific papers, it seems interesting to analyse the set of ideas it conveys. If this term is so pervasive in today's

[1] source : Merriam-Webster 2019.

society, it is probably because its meanings are plenty, and cover a large part of what we envision as the future.

Smart as technologically advanced

The first idea that comes to mind when thinking of smart-something is one of improvement and optimization, that such an object or entity can do *more* that its previous iterations, provide more services or offer better solutions. This notion is very often directly linked to technological progresses, and more precisely, to the promises of information technologies. Computing power and connectivity are the two pillars of the *smart* era, opening a broad range of new possibilities in every field since the beginnings of modern computer sciences around the second world war. Machine learning makes the automation of more and more complex tasks possible, and plays a key role in what makes a system "smart" by allowing it to predict the evolution of its environment and the actions of its users, providing more relevant and customized services. Because more than just new services, what we call *smart* is the adaptiveness of the connected systems and their seamless integration into our everyday life.

Smart as sustainable

But *smart* does not always mean *more*. In a context of rising awareness of our impact on climate change, sustainability is a crucial point which is also implied here. It may seem paradoxical to think that adding more technology could be a way to limit the impact our growth has on climate. The cost of embodied energy is difficult to precisely evaluate and the only undisputed way to significantly reduce our environmental impact would be to reduce our need for energy and technology, which asks for concrete and much needed global behavioral changes. However, one can be optimistic regarding the future evolution of energy consumption in the digital field and consider that "smart" systems should be those which reduce our need for transportation and infrastructures (remote working) or help us adopt a more sustainable way of life (smart thermostats and light switches for example avoiding unnecessary consumption). Sustainability also substantially lies in the way the system is designed considering its life-cycle, its reusability and its openness to limit early replacements, costly maintenance or useless redundancies due to incompatibilities. Much like an intelligent species, a single-use system tailored for a precise use case only and unable to evolve through time can not pretend to be smart.

Smart as enabling

Among the number of advanced, adaptive services technological progress can offer, those we call smart are often those that tighten the link between humans and their tools. Using a combination of learning abilities and adaptiveness, services can become more intuitive over time, allowing an ever-growing share of the population to use them. The concept of *ambient intelligence* reflects this idea : a technology so seamlessly integrated in our everyday life and environment that its use becomes natural and users forget

its existence [Cook09]. This notion is key to the democratization of smart systems as they become more pervasive by making their use feel more natural and less constraining to the user. Real-time reactivity also plays a big part when it comes to offering an interaction that feels natural. Finally, interoperability and adaptiveness are crucial, removing the hassle of system reconfigurations and incompatibility, as well as ensuring an extended life-cycle.

1.1.3 The Smart Grid

Smart grids may be a solution to the quest for sustainability in our power grids, harnessing the potential of renewable energy and dynamic load adjustment to reduce consumption and the use of fossil fuels. As such, they are more of a general concept, a global trend regarding the evolution of power grids rather than a clearly defined entity, though the articles of Farhangi [Farhangi10] and Massoud Amin et al. [Massoud Amin05] provide a very good overview of the topic. We described above the different aspect of the adjective "smart", we will now see how they apply to the future of the power grid.
Firstly, the smart grid relies on information technologies to enable a more precise and timely management of the different elements constituting the grid, most notably the recently integrated equipment like distributed renewable sources, local storage systems and connected residential appliances. The growing number of sensors and connected devices brings new types of inputs to the management algorithms regarding the state of the grid, and the improving frequency and precision of the measurements allows for a more accurate and reactive control.
Secondly, reducing the carbon emissions of the power grid by allowing the integration of renewable energy sources is one of the main purposes of the smart grid. Consumption management along with eco-friendly behaviors prompted by more intuitive and connected interfaces are also very important in this regard. As we mentioned before, sustainability also lies in the diminution of deployment and maintenance resources needed. The abstraction of management algorithms would enables the implementation of open and scalable architectures in order to promote interoperability and to avoid unnecessary redundancies.
Thirdly, the ubiquity of connected devices in households and more importantly the deployment of smart meters enables the consumer to be more aware of its active role by interacting easily with the management systems. A big part of the smartness of a grid management system lies in its ability to engage the consumer into a proactive behavior via intuitive interfaces and appealing incentives. To this end, preserving the comfort and privacy of the user when controlling its equipment is crucial, and transparency and accountability are needed.

1.1.4 Ubiant and HEMIS

The work presented in this thesis has been achieved in collaboration with Ubiant, a company based in Lyon, France. Ubiant develops HEMIS (Home Energy Management Intelligent System), a software solution that allows users to control their connected homes in order to reduce their energy consumption while preserving their comfort. HEMIS is built around the idea of adaptiveness, starting from the observation that no situation is exactly the same : each household is built differently, sits in a different environment, and is inhabited by different persons with varying needs. The long term objective of Ubiant from the beginning has been to allow buildings to be autonomous by managing local production and storage assets [Mansour12]. This thesis is part of their ambition to propose a complete integration of the building into the smart grid by interfacing their building operating system with an upper layer of global energy management system.

1.2 Problem Statement

Our work was guided by three main objectives, which will be regularly recalled throughout this manuscript. They represent the requirements we identified as necessary to design a practical and efficient energy management system for the smart grid :

- ensuring the continuous supply of loads and optimizing energy use
- maintaining user comfort in residential building to improve acceptance
- adapting to the different use cases and situation while allowing a cheap deployment and maintenance of the system.

Here we want to provide a concise description of our goals regarding these key points, as they will be thoroughly detailed in chapter 2.

1.2.1 Optimizing energy use

The smart grid as we described it in Section 1.1.3 in an ensemble of technologies aiming at handling the profound paradigm shift underwent by the energy sector as production transitions to less predictable renewable sources. However, the fundamental issue associated with power grid management is the same regardless of the energy sources or the kind of grid : the balance between production and consumption must be maintained at all time. In Section 2.1.2.2 we will describe in detail the different ways this balance is kept, notably regarding how different mechanisms are used depending on the time scale. The role of an Energy Management System (EMS) is then to schedule the controllable assets to make sure demand is satisfied at all time, in a way that optimizes a

certain number of criteria such as cost and ecological footprint, for example. To achieve this, an EMS needs two critical types of information : an accurate forecast of both production and consumption plus a knowledge of the available means of action on both production and consumption. Ideally, these inputs would be available and perfectly accurate, allowing a "perfect" solution to exist at any given time that could be found using appropriate optimization techniques. This is not to say that this problem would be trivial, as even in perfect condition the large amount of possible solutions considering the number of variables asks for very efficient algorithms. However in this theoretical case, the field of multi-objective optimization already covers the issue extensively. In practice, the accuracy and availability of these inputs can vary greatly depending on the use case, complicating the optimization process and asking for new approaches to be found. In this thesis our primary objective is to design a system able to achieve this multi-objective optimization dynamically to handle the inaccuracy of forecasts (see Section 1.3.1) and the variable availability of controllable assets notably in a residential setting.

1.2.2 Putting the user back in the loop

As we put forward in section 1.1.3, the smart grid offers many opportunities for the user to actively participate in the energy management by taking control of its consumption to not simply represent a passive constraint for the EMS.
To this end, our system must adapt to the user and not the other way around. If residential load management is the keystone of the smart grid (see Section 2.2), it must not be achieved at the user's expense otherwise the appeal of the whole concept will be significantly diminished. Avoiding discomfort could be done in by either taking full control of the appliances without disturbing the user which means perfectly predicting their behavior, or by letting the users in charge, guaranteeing their comfort but potentially jeopardizing the optimization of energy use. We think that a mix of the two approaches could be achieved by allowing a smooth cooperation between the user and the system.

By that, we mean that the users must be able to precisely define their level of engagement in the process. With the growing number of connected appliances, the constant evolution of user interfaces, and the latent behavioral changes due to environmental awareness, it is very important that the way the users interact with the system should not be set in stone and instead be very flexible and modular to accommodate various use cases and user preferences.

1.2.3 Facilitating deployment

Our final objective is to develop a system that could be deployed on real settings in the years to come. It was then imperative to take into account the practical constraints and obstacles that are typically encountered when deploying a commercial solution of this

kind, using the experience of Ubiant in this field. Indeed, considering the yet unknown form of future grids, the future of connected appliances, and the varying degree on equipment of the connected buildings, we had to build our system in a way that would not rely on any of these parameters, or at least make as few assumptions as possible. Genericity, adaptiveness, scalability and robustness were the keywords in the work presented here in order to avoid being limited to only a few use cases. They will be thoroughly defined in Section 2.3.

1.3 Contributions

1.3.1 Very Short Term Residential Load Forecasting

As we mentioned in 1.2.1, an energy management system relies heavily on its ability to predict what the energy production and consumption will be, in order to plan adjustments if needed. If residential buildings are to be integrated in an energy management system to provide load adjustment services, being able to forecast their consumption is necessary. However, everyday use of household appliances rarely follows a fixed schedule, making residential consumption extremely variable. We studied various load forecasting methods to assess the achievable accuracy of very short term individual load forecasting, in order to evaluate the viability of a system that would heavily rely on such a forecast. If the lack of large enough data set prevented us to draw definitive conclusions, our preliminary results seem to show that the variability of residential consumption on the very short term is too high to be accurately predicted by existing techniques. However opportunities lie in the multiplication of sensors in the household providing contextual information, potentially allowing significant improvements in the forecast that still need to be proven using larger datasets.

1.3.2 Decentralized Residential Load Shedding

As a first step toward the objectives we listed in section 1.2, and considering the results of our exploratory work into load forecasting, we built a load shedding system based on Hemis, Ubiant's solution, which satisfies the requirements of the current french power grid operator regarding load shedding mechanisms. We chose a decentralized approach facilitating the scalability of the system, where each building estimates its adjustment capacity without infringing upon the user's control of their equipment. To compensate for unforeseen user's behavior, a coordination algorithm allows the connected building to compensate in real time the variations in shedding capacity, providing an adaptive and robust service. A decentralized reinforcement mechanism helps improve the service quality and further reduces the stress on the system and the users while ensuring equity among the latter. This model shows that scalability and reactivity can overcome the lack of reliability of the individual load forecasting and the stochas-

tic nature of user behavior. It provided both a practical and deployable system using existing systems, and a strong platform on which to build the principal contribution of this thesis.

1.3.3 Generic and Adaptive Energy Management System

The final contribution of this thesis is a multi-agent model designed as a modular backbone for energy management in a smart grid. The idea was to offer a generic framework on which a large variety of control algorithms, constraints and management policies could be implemented. Unlike the previous model which mainly focused on a specific use case (load shedding) and a precise spatial and temporal scale (buildings with hourly predictions), this system offers an abstracted representation of the objectives and constraints both local and global. With only limited implementation requirements, it provides a reliable coordination mechanism in which the various connected assets agree on an operational schedule satisfying the needs of the grid. We implemented our model and tested it on various situations using realistic data to demonstrate its reliability.

1.4 Outline of this thesis

This thesis is articulated around the three contributions presented in the previous section. Chapter 3 presents our exploratory work on short term load forecasting, the obstacles we encountered and the preliminary results we obtained. In Chapter 4, we describe the decentralized residential load-shedding system we designed, which fits the reliability requirement of today's utilities while handling the variability of household consumption thanks to our dynamic adjustment mechanism. Finally in Chapter 5 we explain how we used the load shedding system as a basis to develop a generic smart grid energy management system able to handle a large variety of use cases, mainly focusing on residential building integration. We conclude this thesis in Chapter 6 by a summary of our contributions, how we intend to continue to improve it in the coming months and what perspective it opens. Following this introduction, we first give in Chapter 2 an in-depth overview of the context of our work. Following the three main objectives stated in section 1.2, we explain the situation of the current power grids and the issues that smart grids aim to solve, then focus on the importance of residential load management and the necessity of a facilitated deployment. We also give an introduction to the multi-agent paradigm we use as a basis for our models.

2

Background and State of the art

▷ *This chapter's objective is to provide the reader with the necessary background to understand the rest of the thesis. We will introduce the main concepts and notions related to energy management, comfort preservation and practicality of the final system.* ◁

Chapter content

2.1 The challenge of energy management **13**
2.1.1 Electricity 13
2.1.1.1 Main concepts 13
2.1.1.2 A bit of history 14
2.1.2 Conventional power grids 15
2.1.2.1 General composition 15
2.1.2.2 Power grid balancing 18
2.1.2.3 Issues of the current system 19
2.1.3 Smart-grids : the new paradigm 20
2.1.3.1 Impact of climate change 20
2.1.3.2 Opportunities of information technologies 21
2.1.3.3 A diversity of solutions 21
2.2 Leveraging residential consumption **24**
2.2.1 The weight of residential consumption 24
2.2.1.1 Current situation and trends 24
2.2.1.2 Load consideration in research 25
2.2.2 Engaging users 26
2.2.2.1 Economical considerations 26
2.2.2.2 Social factors 27
2.2.2.3 Security and privacy 28
2.2.3 The need for a user-centered system 28
2.3 Enabling large scale deployments **29**
2.3.1 Handling the variety of situations 29
2.3.1.1 Genericity 29
2.3.1.2 Scalability 29
2.3.2 Adapting to changes 30
2.3.2.1 Modularity 30
2.3.2.2 Adaptiveness 30
2.3.3 Limiting the cost 31
2.3.3.1 Simplicity 31
2.3.3.2 Robustness 31
2.4 Multi-Agent Systems **32**
2.4.1 Concepts 32
2.4.2 Applications in smart grids 33
2.5 Summary and positionning **34**

2.1 The challenge of energy management

As we briefly explained in chapter 1, the first purpose of an Energy Management System is to optimize the energy use. In this section we will cover basic notions about electricity and power grids to better explain what optimizing means here and why it is necessary.

2.1.1 Electricity

Electricity is a physical phenomenon whose exploitation gave birth to tremendous progresses, mainly by allowing us to easily transmit energy from one point to another. Over the course of the last century, it has clearly become one of the backbones of our societies.

2.1.1.1 Main concepts

Electricity is the manifestation of the *electromagnetic force*, one of the four fundamental forces (with gravitation, weak interaction and strong interaction). This force creates an electric charge in certain particles, notably electrons with a negative charge and protons with a positive one. Being at the outer layer of atoms, electrons can move more or less freely in certain elements, called *conductors*. This movement of electrons, named electrical current and which *intensity* (I) is measured in Ampere, is triggered by a difference in electric potential between two points, called *voltage* (U) and measured in Volts.
In a simple electric circuit, charges flow in one direction only in what is called a direct current (DC). However, most commercial application use Alternating Current (AC) where the current periodically reverses direction, its voltage usually following a sine wave. The main advantage of alternating current is the possibility to modify its voltage using a transformer, which made its use more practical during the development of large scale power systems. Indeed, the *power* or amount of energy (in Watts (W)) conveyed by an electric current is equal to the intensity of the current multiplied by the voltage : $P = VI$. As the intensity is the main parameter responsible for losses due to heat in long power lines, being able to raise the voltage allows to lower the intensity to keep the same power output while limiting losses (see section 2.1.2.1).
It is a common misconception that charged particles carry this electrical energy along as it goes. In fact, electricity is a wave transmitting energy from one point to another in the same way that sounds travels through air, or a liquid transfers pressure in a piston. **As such, electricity in itself can not be stored.** It does not really matter on a small electronic circuit, but as soon as multiple power sources are present, ensuring that there is no excess or shortage of energy becomes an issue. The larger the grid, the more important the problem is, which is why power grids need management systems, as we will see in the next sections.

2.1.1.2 A bit of history

The existence of electricity, under one form or another, has been known since Ancient Egypt (around 3000 BCE). It is interesting to note that links between magnetism and electricity, as well as the notion of conductivity began to be theorized during Antiquity. At the time, the electricity was mostly observed in its static form, giving certain objects like pieces of amber the ability to attract other small objects. In fact the Latin word "electricus" comes from the Greek word "elektron" which literally means "amber". It was first used by English scientist William Gilbert in "De Magnete", the first comprehensive study of the phenomenon written almost five millenniums later in 1600. In the centuries following this pioneering work, numerous scientists of the time interested themselves in the matter, often giving their name to the properties their discovered. Benjamin Franklin famously showed that lightnings are electrical, Luigi Galvani studied the way by which neurons communicate via electric signals, Alessandro Volta invented the battery, André-Marie Ampère formalized the link between electricity and magnetism, Michael Faraday invented the electric motor in 1821 and Alexandre Becquerel discovered the photo-voltaic effect in 1839. The evolution of the various theories developed during this period has been notably summarized by E.T. Whittaker in 1910 [Edmund Taylor Whittaker10].

The end of the 19th century witnessed significant advances in the field of electrical engineering. The first distribution system is built in 1882 in the US, in Manhattan, by Thomas Edison. At its source, a direct current (DC) generator was able to power roughly 1200 light bulbs at the time and converted only 2.5% of the energy of the burning coal. Its profitability was made difficult by the limitation of DC current at the time, which made long distance transmission inefficient [Smil10]. Edison's competitor, George Westinghouse, installed the first commercial alternating current (AC) generator in the famous Niagara Falls power plant in 1896. It was hundreds of time more powerful than Edison's station and, thanks to the advantages of AC on long distance transmission (see 2.1.2.1), was able to power Buffalo City 40km away. It marked the beginning of AC domination and democratization over DC [Carlson13]. At the beginning of the 20th century, the invention of the induction motor (enabling to convert electricity back into mechanical motion) led to the invention of electric household appliances and the rapid growth in demand pushed the creation and expansion of electric utilities. Around 1915, nearly all states had their own regulated electric utility. Soon, larger companies controlled most of the transmission network and power plants over multiple states. In the last fifty years, various regulation policies were voted to consolidate the rules of the energy market and strengthen the transmission system in anticipation to the growing number of distributed renewable energy resources [1].

[1] "Public Utility Holding Company Act of 1935: 1935-1992," U.S. Energy Information Administration, January 1993

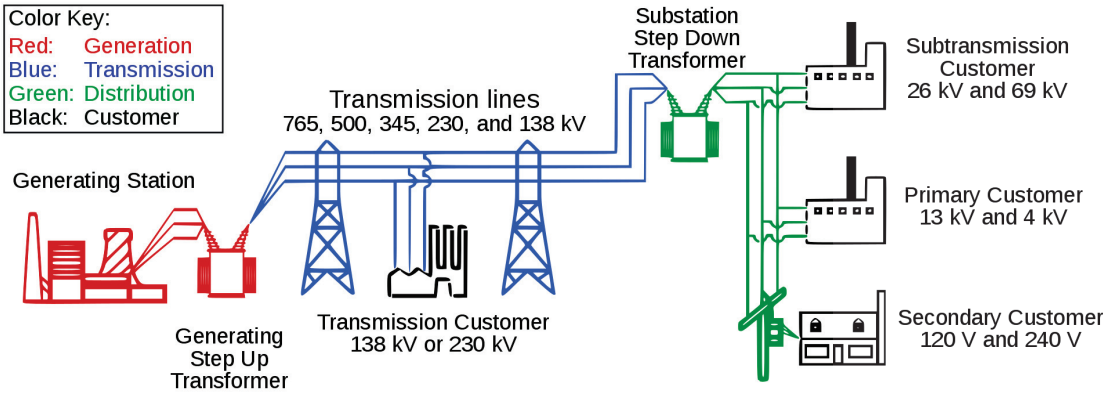

Figure 2.1: Basic structure of a large scale electric power system [2].

2.1.2 Conventional power grids

Today, the vast majority of developed countries has a well established power grid, often interconnected with one another, and all are very similar in their structure and management. Here we describe their composition, how they tackle the energy management problem and why new solutions are needed.

2.1.2.1 General composition

A power grid, like any electric circuit whatever its size, is composed of three parts: production, transmission/distribution and consumption (see figure 2.1). National power grids are supplied by large power plants, from which electricity is transmitted by a backbone of high voltage power lines to the loads where it is consumed.

Production

Alternators provide most of the world's electric power, using a principle discovered in 1832 by Faraday to convert mechanical energy to electrical energy. In such a generator, a turbine driven by wind, water, steam or gas rotates a magnet within a closed loop of copper wire, creating current. Major evolution is in power generation technologies has then for long mainly concerned the force used to rotate that magnet. The first power plants used water to move a turbine, then steam turbines were used, allowing the exploitation of various fuel such as coal, natural gas, oil, bio-fuels or wastes to burn and heat water into steam. Nuclear power plants appeared after the second world war as carbon-free and very efficient way of producing electricity and rapidly became a major source despite the significant risks linked to radioactive fuel waste. In France in particular, nuclear power currently amounts to 78% of the electricity production (fig. 2.2). The major and almost only source of clean, sustainable energy until then was hydroelectricity, using the potential energy of water to drive a turbine. Unfortunately its

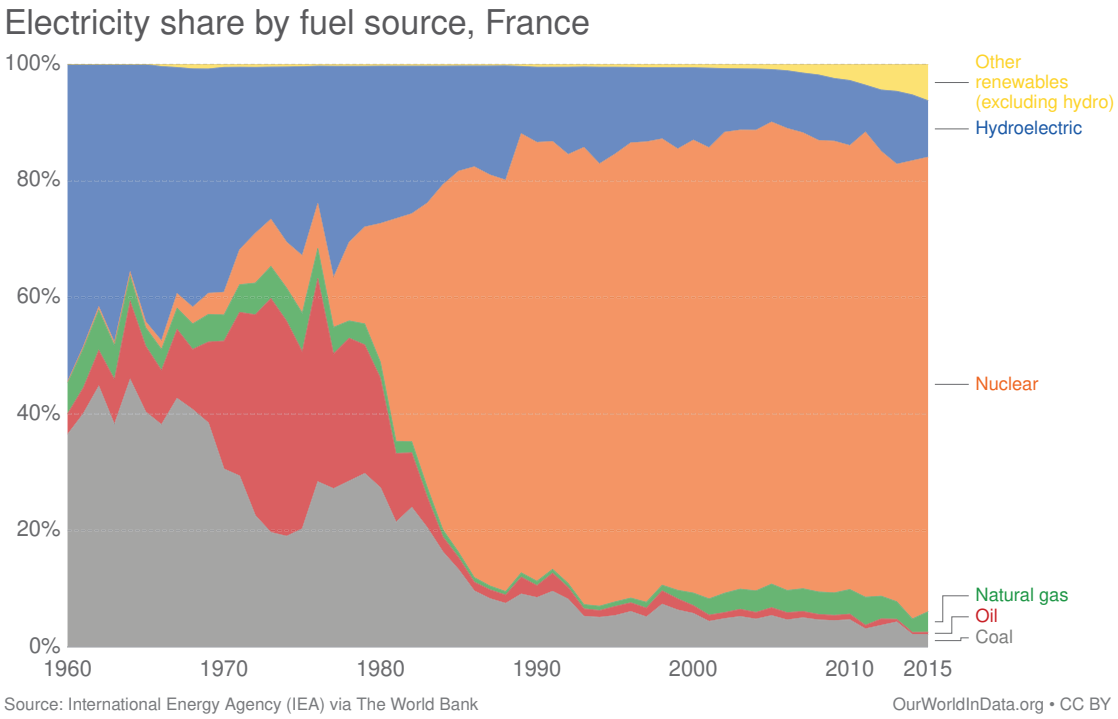

Figure 2.2: Evolution of fuel sources for electricity generation in France

limited availability due to geographical and environmental constraints prevents it from accommodating the growing demand, explaining its decreasing share in the energy mix over the years (fig. 2.2).

The last two decades however have seen the rise of many other renewable energy sources following massive worldwide investments in sustainable power generation (see fig. 2.3). In France for example, solar generation started growing steadily in 2008 to reach 8,5MW of installed capacity in 2018 [3]. In the same year, the cumulative capacity of installed wind turbines reached 15MW, and renewable energies (with hydroelectricity) supplied more than 20% of the national demand[3]. Sun, wind and water are not the only source of renewable energy. Geothermal energy which uses Earth internal heat to turn water into steam, as well as bio-fuel and waste burning are also amongst developing sources. A comprehensive survey of renewable energy sources can be found at [Ellabban14].

Transmission/Distribution

In large scale power grids, the centralization of power generation means that electricity must be transmitted over long distances. Unfortunately, power lines can not perfectly conduct electricity and tend to loose a portion of the energy transmitted as heat (Joule

[3] source : RTE, Panorama de l'électricité renouvelable 2018.

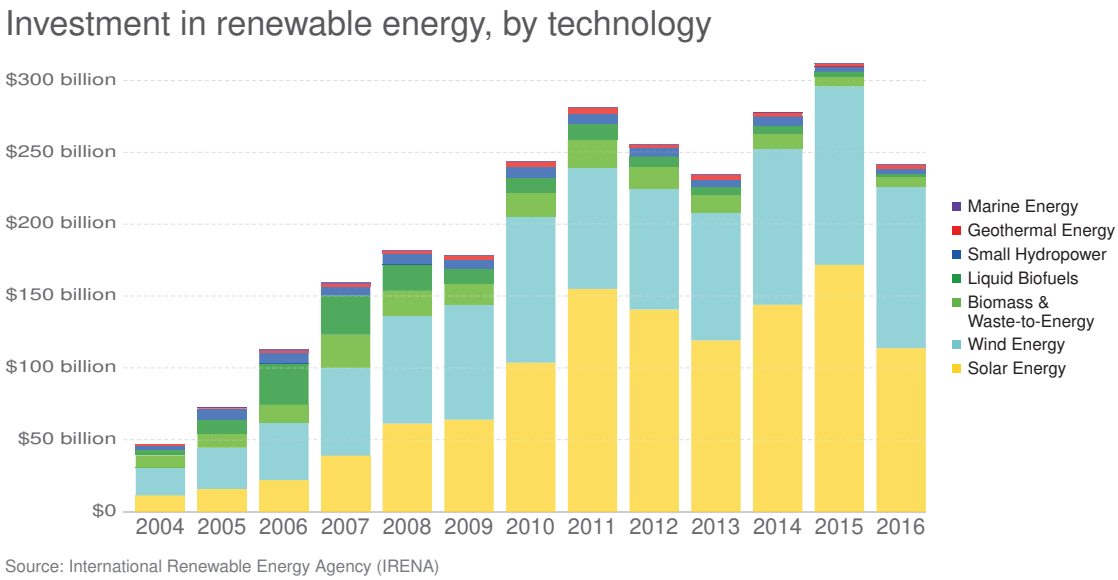

Figure 2.3: Investment in renewable energies in the last decades

heating). This loss being proportional to the intensity of the current passing through the wire, increasing its voltage allows to convey the same amount of energy more efficiently (see 2.1.1.1). Therefore, as depicted in figure 2.1, transmission grids are a network of transformers that increase the voltage coming from power plant up to 800kV (step up transformers) to send it through long distance power lines before reducing it again (step down transformers) to distribute it to end consumers. This results in a segmented network, with sections carrying very high voltage (138 to 765 kV) constituting the *transmission* network and sections connecting the end users at lower voltage (120V to 70kV) as part of the *distribution* network.

Consumption

Nowadays, electricity has a large variety of uses, powering everything from transport systems to farming equipment and of course our homes and industries. Due to the electrification of energy uses (transport, heating, household appliances), global consumption tends to generally increase over the years. If it stays relatively stable in developed countries, the trend is clear when considering Asia and the Middle-East as can be seen in figure 2.4. In 2017, Chinese electricity consumption grew by almost 6%, its fastest pace since 2014[4].

[4] source : Enerdata, Global Energy Statistical Yearbook 2018.

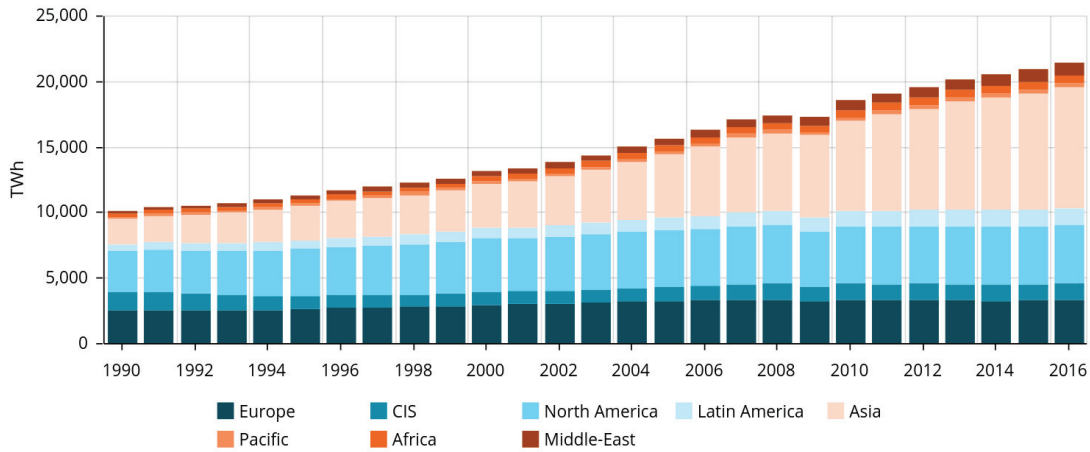

Figure 2.4: Evolution in total energy consumption by region[4].

2.1.2.2 Power grid balancing

The problem of storage

We explained in section 2.1.1.1 that electricity itself cannot be stored simply. It means that the electricity needed at any given time must be produced by the power plants at that same time, but also that all the power produced must be consumed. Power can however be stored in other forms. Until recently, the only efficient way to do so was to pump water up in the higher reservoir of a dam with surplus electric power, then releasing it through turbines when needed. This solution is highly flexible and very cheap, its only limitation being its total capacity as mentioned in 2.1.2.1. Significant research and development efforts have been undertaken in recent years, notably in chemical battery technologies [Aneke16], but it is still far from being a generally scalable, affordable and sustainable way of balancing a power grid. It is then necessary for a grid operator to constantly monitor and control generation and transmission system to maintain the desired quality of service to the end users.

Forecasting and scheduling

The first part of this balancing process is done ahead of time, using demand forecast over various horizons. On the long term, to design and build power plants and transmission equipment that are able to handle the load. On the medium to short term, to schedule power generation a day ahead to match the predicted consumption. This is done by stacking for each hour or half-hour the various energy sources in an order depending on their characteristics until the total production for the interval matches the forecast. Usually, energy sources are prioritized from the least flexible, like solar panels or wind turbines that cannot be controlled or nuclear reactors that are slow to ramp up of down, to the most reactive like hydraulic or gas turbines.

Error correction

This process however relies on the accuracy of the load forecast that cannot be perfect but also on the *controllability* of the majority of the generation panel. To compensate for unforeseen variations, a series of mechanisms is implemented in every national scale grids. A first layer of instantaneous corrections is automatically executed by generators and power electronics in a local reactive manner. They deal with small deviations in frequency and voltage that occur continuously and rely on simple electronics with almost no supervision. When the deviations are too important to be managed by this layer, coordinated operations are needed. Their names and details can differ from one operator to another but the following principles stay the same. When possible, power plants do not commit to their full capacity when scheduling their operations but instead keep what is often called a short-term operating reserve that allows them to produce more power if needed. If this reserve is not enough, complementary measures are executed that usually include the activation of peaking power plants that are able to ramp up almost instantaneously. Industrial customers often have contracts with power grid operators to provide demand response services, shutting down assembly lines at times if needed for example. In some cases however, these mechanisms can fail to compensate a peak in demand or a drop in production. As a last resort measure, parts of the distribution network can be switched of temporarily (brownout) to avoid a large scale failure of the power grid (blackout).

2.1.2.3 Issues of the current system

We saw in 2.1.2.2 that central power plants need to keep an operational margin in order to compensate for unforeseen peaks in demand or production failure. It means that today's generators cannot operate at their nominal range as they are dimensioned for power outputs that they do not reach most of the time. More importantly, as the energy sources transition from fossil fuels to intermittent renewable sources like wind and solar, the forecast uncertainty will increase while the available operational margin will decrease. Also, as a side effect of the centralization of power systems, high voltage power lines which allow the transmission of electricity over long distances represent a significant cost and generate between 5% and 10% losses in power globally. To keep the energy management system as it is today, the solution could be to increase the adjustment capacity of the network. This can be done either by building more peaking power plants or by developing large scale storage systems, which would be very expensive and polluting. Sustainable solutions, as we will explain in the next section, reside not only in more dynamic processes but mostly in more control over the consumption.

2.1.3 Smart-grids : the new paradigm

In section 1.1.3, we defined the general idea behind the term "Smart grid" as a natural improvement of the existing system thanks to new technologies. Here we will explain the motivations behind this concept and how it brings new solutions to the issues encountered by today's power grids.

2.1.3.1 Impact of climate change

The first ingredient to the rise of smart grids is climate change. Whether is was a trigger or only a catalyst, the ongoing environmental crisis has been instrumental in numerous changes in our societies.

On policies and people

Under the ecological pressure, various institutions around the globe took steps to reduce our impact on the environment. In the case of energy, this was notably done by setting emission reduction targets to limit the use of fossil fuel and promote renewable energy sources. The European Union for example took in 2010 a series of measures to reduce its greenhouse gases by 20%, increase the share of renewable energy to at least 20% of consumption, and achieve energy savings of 20% or more, all by 2020. More recently at the Paris Climate Conference (COP21) in 2015, nearly every countries in the world agreed on a global action plan to limit global warming below a 2°C. In May, European Parliament elections saw a net progression of environmentalists, with green political parties coming third and second in France and Germany, respectively.
As awareness about climate change grows, so does the general interest in ways to better manage and reduce individual energy consumption [Lee15], as we will see in the next section.

On the grid

Large scale power grids, as we explained in section 2.1.2.3, are not built to both integrate a large share of intermittent, distributed renewable energy sources and exploit the load flexibility of more energy-aware customers. Indeed, solutions under the current system would either be to increase the security margin of existing power plant and build more polluting fast-activating plants or to create large scale storage stations to compensate variations in production, both of which would be costly, but also inefficient as it would offer a temporary fix to a long term issue. In addition to climate-related concerns, many European countries in the last decades decided to progressively phase out nuclear power considering the potential risk for people and the environment, despite the important reliance on this kind of energy. It further boosted the development of renewable sources, particularly in Germany.

2.1.3.2 Opportunities of information technologies

At the same time climate change urges societies to take action and increases the pressure on the grid balancing mechanism, advances in information technologies and more importantly their integration in more and more devices brings new opportunities. As we explained in 2.1.2.2, the efficiency of an energy management system relies largely on its ability to predict the consumption to schedule production accordingly. To this end, electrical consumption is constantly monitored at various intervals, and the more frequent the measurements are, the more precise and useful the information can be. Until recently, small consumptions from residential or commercial buildings could not be automatically transmitted to the operator. It was measured at the distribution point, which lacked granularity, and the actual meter was difficult to access for the consumer himself. Smart meters bring a digital solution by allowing the grid operator to access the reading in real time and to make it available to the consumer. They are part of the constant progress of information technologies toward connectivity, which gave birth to the concept of *Internet of Things* : an ever-increasing number of connected sensors, controllers and interfaces integrated in everyday devices and appliances and offering a plethora of data, automation and remote control possibilities.

2.1.3.3 A diversity of solutions

To tackle the issue of the instability caused by intermittent renewable energy sources and thus to enable their integration, various approaches have been offered. The underlying principle is still the same : a number of assets must be coordinated in their operations so that the sum of their production and consumption matches a predefined value.

Supergrids and Virtual Power Plants

The need for better, more flexible energy management systems comes mostly from the difficulty to control the output of intermittent energy sources like wind turbines and solar panels. To tackle the problem at its source, two approaches offer to reduce this inherent variability by grouping renewable sources scattered over a large territory, as the variable power output is often directly linked to the meteorological conditions of the location of the wind turbine or solar panel. By aggregating multiple generators from various locations, the total production becomes less dependant on the weather and thus is more stable.

The notion of *Supergrids* relates to a range of enhancement and additions to the transmission system that would enable the interconnection of multiple national scale power grids and the integration of remote renewable energy sources, like offshore wind farms [Gordon06]. By facilitating the supply and exchange of renewable energy where it is needed, it would help secure and stabilise the grid while enabling the development

of more efficient and sustainable energy sources. The idea has gained significant traction in Europe where most inland countries are already interconnected [Van Hertem10].

While Supergrids are mainly a hardware approach to managing the variable availability of renewable energy, *Virtual Power Plants* take a more software-based path. Virtual Power Plants, as indicated by their name, offer to present a single entity to the power grid management system, hiding the variations of individual turbines or solar panels and behaving as a stable power plant, thus simplifying the integration of renewable energies. The concept, first formally presented in [Awerbuch97], could provide an interface to both the energy market and the grid balancing mechanism [Pudjianto07], and *can also be used to manage Demand Response* [Ruiz09].

Microgrids

The concept of Microgrids emerged in the 2000's [Lasseter01] as a way to handle the growing number of distributed generation units. These are dispatchable units using fossil fuels like fuel cells or CHP[5] units, or renewable energy sources like photo-voltaic panels, wind turbines, biomass or micro hydroelectricity [Lasseter02]. They are installed on low-voltage networks, close to customer load, and can provide power backup in case of blackouts. But the true potential of these new generation means lies in the way they are managed as all-day power supply. Indeed, these low-voltage, low-power units cannot be integrated as-is in the main power system [Piagi06]. The idea is then to create semi-autonomous sub-networks, "micro-grids", composed of low-voltage equipment such as distributed generation, storage and residential loads, that can be considered as one single electric entity by the main network. The difference with the Virtual Power Plant is that the different components of a microgrid are physically (not virtually) close to each other and connected to the same low-voltage network. Robert Lasseter [Lasseter01], one of the first researchers to study this concept back in 2001, describes a microgrid as *"a cluster of micro-sources, storage systems and loads which presents itself to the grid as a single entity that can respond to central control signals"*. This definition can be found in the majority of the papers related to microgrids [Dimeas05], but as the IEEE-PES Task Force on Microgrid Control puts it in a 2014 survey [Olivares14], *"a detailed definition of microgrids is still under discussion in technical forums"*. The size of a microgrid is also part of its definition, as many of the advantages of the concept disappear when the size of the system increase, in terms of capacity, number of appliances or topography (as line losses increase with the distance). Colson et al for example [Colson09] consider a power system of several MW or less, which can be illustrated by the isolated power system of Lemnos [European Commission96] which gather a total power of 14.84 MW using diesel generators, wind turbines and some solar panels.

If the microgrid concept was proposed as a way to better integrate renewable energy sources and distributed generation in general into the main power grid, consumers that are part of a microgrid would not just benefit from a more efficient, eco-friendly

[5]Combined Heat and Power

power system as a microgrid has other advantages. As Colson et al. [Colson09] put it : *"The primary goal of microgrid architectures is to significantly improve energy production and delivery for load customers, while facilitating a more stable electrical infrastructure with a measurable reduction in environmental emissions"*. First, its local generation capacity and single point of connection with the main grid means that it can disconnect from it in case of general blackout and still provide its consumers with a good quality current. Secondly, the general quality of service should be globally better as small disturbances in frequency and voltage that are usually observed in the main grid can be compensated by local sources. Finally, as a microgrid is considered as a single entity on the electricity market, it can trade energy with the main grid and sometimes make profit from the low cost of renewable energies.

Demand Response

The idea of controlling the demand to alleviate the burden of intermittent generation or reduce load peaks is not new [Sanghvi89] and is relatively straightforward. As we mentioned in 2.1.2.2, this method is regularly used with industrial consumers as their operations can be easily scheduled. A significant volume of the total consumption can be reduced by a handful of factories, which makes the process compatible with the current centralized control of power grids. Moreover, delayed operations can easily be compensated financially and have limited impact on people [Gils14]. However, industrial demand response is limited in terms of overall flexibility, as the volume of curtailed consumption is difficult to adjust precisely and in a reactive way : assembly lines cannot be partially stopped, and neither can they be switched on and off frequently and randomly. **To constitute a credible solution to the future fluctuations of production, demand response must be highly flexible and reactive.** We saw in 2.1.3.2 that the democratization of smart meter has accelerated in recent years. It opens the door to the management of residential loads which are potentially more flexible and could be more dynamically controlled. Residential load shedding [Haider16] raises two major challenges because if it solves the two issues of industrial demand response, it does not have its two key advantages : limited impact on comfort and large capacity.
Thus the first challenge comes from the impact of load shedding on the comfort of the inhabitants. Indeed, unlike industrial processes, domestic needs and associated electrical consumption can not systematically be delayed with a simple financial cost [Nguyen14]. Depending on the time of the day for example, a space heater could be switched off without being noticed or on the contrary be necessary for a sleeping baby. This complexity is addressed in depth in section 2.2.
The second challenge comes from the small volume of energy consumed by a residential building compared to factories. In order to be useful in a large scale power grid balancing effort, a residential load shedding system would need to aggregate a large number of buildings in order to reach a significant shedding capacity [Beal12]. As distributed demand side management could provide a substantial relief in handling the variability of renewable production without necessitating extensive modifications of

the existing system, our first contribution detailed in Chapter 4 focused on this approach. But demand side management can be applied at any scale : in residential microgrids for example, load management could be crucial to avoid relying on external sources in case of shortage, or to take advantage of excess production [Lim14].

2.2 Leveraging residential consumption

In the previous section we gave a broad overview of how power grids work and how new forms of energy managements could help solve the challenges faced by current systems. We saw that leveraging residential consumption could bring new opportunities in demand response and also in microgrid management. In this section we focus on the issue of residential consumption management to furthesr describe these opportunities, identify the associated challenges and the existing advances in the field.

2.2.1 The weight of residential consumption

2.2.1.1 Current situation and trends

In France, households and commercial buildings account for more than a third of the total electricity consumption[6]. The ongoing deployment of smart meter, which should soon cover the vast majority if not the integrality of consumers by 2021 as part of a European initiative, is the key to unlock the potential of residential load management as it enable its precise monitoring [McKenna12]. According to the French grid operator, air conditioning (both heating and cooling) and water heating accounted for roughly 40% of a household electricity consumption in 2017, as we see in Figure 2.5. With the democratization of smart thermostats, these devices represent a major tool to act on a building's consumption [Lu10]. Accounting for another quarter of electricity uses, connected household appliances like dishwashers, along with lighting, are beginning to spread and could also take part in demand response mechanisms in the near future.

The democratization of Plug-in Electric Vehicles (PEV) along with home storage batteries, micro-cogeneration and solar panels are bound to fundamentally change the way we consider the consumption of a residential building [Nguyen14]. As it becomes more and more possible for a home to generate electricity to satisfy its own needs and even to feed it to the grid, the word "prosumption" seems more fitting. This mix between production and consumption is very interesting from many point of views, regarding the improved autonomy from the grid in case of global failures, the reduction in transmission losses or the possibility of creating local energy market, and as such a prosuming building can be seen as a microgrid (see Section 2.1.3.3) of its own [Lopes07].

[6] source : RTE, Bilan Electrique 2017.

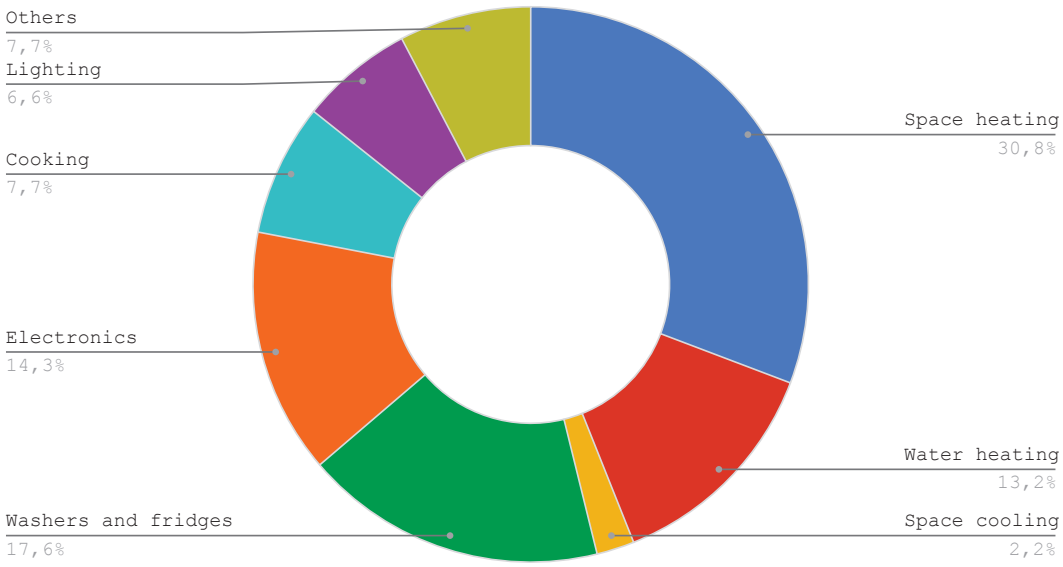

Figure 2.5: End uses of residential electricity consumption in France. Source : RTE, Bilan Electrique 2017.

2.2.1.2 Load consideration in research

Various ways of controlling residential consumption have been explored in the literature, with a clear focus on demand response applications [Nunna12, Albadi08], sometimes with both production and consumption [Kahrobaee13]. An extensive survey by Pierluigi Siano [Siano14] gives a broad overview of papers on the subject. Each time, the importance of high customer participation is highlighted, as it is the logical condition for a functioning demand response system. The state of the art of residential demand response will be detailed in section 4.2. In papers studying microgrid control however, loads are often neglected as researches focus mostly on production and storage management. Levron et al. [Levron13] for example do not take load into account and consider it constant over time. Others like [Basir Khan16, Dagdougui14, Karavas15] use time-varying loads, with various degree of realism in the load profiles used. Even when loads are integrated in the optimization process, they are often considered as either critical (that can not be switched off) or non-critical (that can be switched on or off at will) [Colson11b, Colson11a, Chaouachi13]. In [Fazal12], loads are more clearly categorized into "critical", "interruptible" and "deferrable". An in-depth review of load management in smart grids will be provided in section 5.1.2. We argue that such simplistic classification can be the source of inconveniences on the user's side, as loads do not systematically have the same priority depending on the time of the day or the user, for example a heating system is not always interruptible. Moreover, customer satisfac-

tion, comfort or proactive participation has been given less attention in existing works, as we will see next.

2.2.2 Engaging users

In 2010, The EU Commission Task Force for Smart Grids asserted that *"the acceptability of new services by the customers is a main concern"*.[7] The general idea is that in order to use the flexibility of residential loads for grid balancing, involving the user in a positive and active way is needed. There are many aspects to this issue that we will cover here.

2.2.2.1 Economical considerations

The economics of energy management are complex and multifaceted and we do not pretend to provide a deep hindsight into this field, instead focusing on the most relevant aspects. The general principle is simple : on one end of the electrical grid, suppliers face a challenge as profitability goes generally against environmental considerations. On the other, consumers have it easier as reducing their environmental footprint also means reducing their electricity bill. For this reason, a significant share of researches in residential demand response try to leverage electricity prices to influence consumption. Price incentives are indeed an effective way to cheaply and rapidly enroll a large number of consumer into a load shedding program, which is why time-based pricing contract have been used by utilities since the early days of large scale power systems. These include "time-of-use pricing" with predetermined time periods with different prices, or "critical peak pricing" where certain days have significantly higher electricity prices [Newsham10]. The main drawback of this approach is the reliance on human reactivity which can lead to inconsistent results and can also affect the customer negatively by increasing their bill if they do not adjust their behavior correctly [Eid16]. In order to increase the reliability and potentially reduce the impact on the customer's bill, well designed interfaces can go a long way [Becker12] but human reactivity is still a bottleneck. To bypass this obstacle, the idea of controlling residential consumption directly via connected equipment has been extensively explored, where the dynamic electricity price is used as a signal to switch off appliances or delay their use [Kahrobaee13]. The advantage here is that users no longer have to worry about consuming when price is too high which can help convincing them to enroll in such programs. However, more work regarding the satisfaction of their preferences could be done, as pointed out in [Callaway11]. In this regard, the work of Beal et al. [Ranade10, Beal12] proposes a demand response system where the user can indicate the flexibility of each connected device by choosing between four different levels, the first allowing it to be switched off at any time and the last forbidding any interruption.

[7] Final deliverable, expert Group 1, functionalities of smart grids and smart meters. EU Commission Task Force for Smart Grids.

2.2.2.2 Social factors

Economical considerations aside, awareness of climate change and its probable human causes has increased the consumers willingness to reduce their environmental footprint, and a significant part of the population is ready to do so by reducing energy usage at home [Semenza08]. To encourage this kind of behaviors, social incentives can be leveraged.

Gamification and communities

It is a well known fact that peer pressure is a very powerful tool to influence someone into changing their behavior. And indeed, as the concepts of residential demand response and microgrid grow in popularity, the idea of connecting users participating in the same community or load shedding system has emerged, using the heuristics and dynamics of games to design interfaces destined to stimulate user engagement. Gnauk et al. [Gnauk12] for example use an explicit scoring and competition system to encourage the participant in indicating flexible periods where the system can delay the starting of appliances. Similarly, [AlSkaif18] implements a rewarding system promoting energy saving and peak shaving behaviors. In another example, [Kashani17] went as far as creating carbon emission challenges accessible via the social network Facebook. Our partner Ubiant also takes this path, using physical interfaces such as the *Luminion*[8] which shows the user how its current energy consumption compares to both its own average and the community's.

The importance of fairness

If gamification can be an effective way to incentivise participation to grid balancing programs, the notion of fairness in itself is crucial. Various works, like [Koutitas12], acknowledge the lack of research towards residential demand response algorithms that would consider fairness in their optimization criteria. [Vuppala11] argue that a fair pricing system can lead to more user engagement in demand response scheme. With the same observation, [Baharlouei13] propose a mechanism that fairly rewards participants considering their level of commitment without penalizing too much those who do not participate much. It seems indeed straightforward that users expect such a system to treat every participant in the same manner and a failure to do so would inevitably lead to users dropping out of the program. However, it is important to note that although the fairness of a system could be measured to some extent, by counting how many times each consumer takes part in a load shedding event for example or by comparing the capacity it consents to curtail, it should also take into account more subjective factors like the inconvenience caused to the user which is not necessarily proportional to the capacity.

[8] ubiant.com

2.2.2.3 Security and privacy

The rise of the Internet of Things is accompanied by growing concerns about security and privacy. The plethora of connected devices offering remote access to indoor microphones, cameras and other sensors are indeed prime targets for cyber attacks. [Sicari15] provides an in-depth survey of the technical challenges of IOT security and give pointers for solutions. A more recent paper by [Conti18] gives an up-to-date summary of the current state of the art.

Regarding the smart grids in particular, the deployments of smart meters have been the topic of debate and even faced strong opposition, notably in France and also in Austria. The discussions revolve mostly around the grid operator monitoring in real-time the electrical consumption of the building, possibly revealing private information like occupancy or used devices for example [Revuelta Herrero18, Lisovich10]. Privacy concerns in demand response architecture has been the focus of numerous papers, [Lisovich10] for example provide an in-depth analysis of the issue and points to the lack of regulatory safeguards in this regard. [Efthymiou10] also acknowledge the importance of security and privacy for the acceptance of smart grid and smart metering networks and offer a method to anonymize electrical consumption data. Going further, [Wicker11] demonstrates the need for a privacy-preserving demand response system and advocates for a minimization or even an absence of centralized data collection. Indeed, a lot of privacy and security concerns can be avoided when sensible information simply does not leave the household. The use of blockchains to secure individual data in this context has also been studied, notably in [Knirsch18, Dorri17, Pop18], bringing in the drawbacks of the blockchain, mostly in terms of resources.

2.2.3 The need for a user-centered system

In this section we saw that residential loads have a major role to play in future energy management systems, and that their controllability depends largely on the involvement of end users. We described the various methods used to engage users in demand response program, from purely economic arguments to more subtle social engineering techniques. However we find that the vast majority of approaches on smart grid energy management follow a top-down approach, where the need and constraints of the grid prevail on those of the users. When load is prioritized, it is not really integrated in the optimization process which is not desirable either from our point of view. We argue that the satisfaction of the users needs should be highly prioritized, as residential consumers are bound to become the primary users and beneficiaries of smart grids. Also, we claim that user preferences, constraints and satisfaction are not static parameters and should not considered as such, and that no two users are alike on these grounds. On this topic we align with the position of Jin et al. [Jin17] who designed a user-centered system that learns preferences user the SMARTER method [Edwards94]. Others like [Monyei18] and [Fan10] also follow this path and prioritize users well being.

2.3 Enabling large scale deployments

In section 2.1 we explained how a smart energy management system could take advantage of real time information and control in order to handle the variability of intermittent generation and leverage flexibility on the demand side. In section 2.2, we showed the importance of managing residential consumption and argued that the well-being of the user should be placed at the core of the system. Here, we focus on the challenges brought by the variety and number of situations a smart grid energy management system would have to handle. Indeed, in order to be viable, the solution will have to minimize the costs related to deployment and maintenance but also be able to adapt to changes in the grid composition.

2.3.1 Handling the variety of situations

2.3.1.1 Genericity

We saw in 2.1.3 that smart grids can take multiple forms. From a handful of households equipped with solar panels and storage, to an offshore wind farm generating megawatts of electricity, or a large demand response scheme involving thousands of consumers, the use cases are extremely diverse. In the literature, many papers focus on a particular setting : [Zamora10, Levron13, Nguyen12] offer a control algorithm from microgrids with storage for example, [Fan12, Jansen10, Paetz13] researched ways to integrate electric vehicles and [Hernandez-Aramburo05] looked into microgrids powered by fuel cells. While such research into the specificities of each kind of configuration is essential, a deployable platform must allow these various algorithms and controllers to coexist inside the smart grid management model [Kuznetsova14]. In addition to a large variety of use cases, the heterogeneity of the actual devices on site must also be dealt with. In practice, for the same kind of device each brand will likely have its own characteristics and controls that the system will need to be interfaced with, meaning that the less requirement our system has regarding this interface, the better [Mauser16].

Our goal in this thesis is not to design an energy management system that would pretend to fit each and every possible use cases. However, even with a clear focus on residential load management and microgrids, the number of possible configurations calls for an open model that could not only be applied to a large variety of settings but also that could work with third-party appliances and controllers with limited requirements [Mauser15].

2.3.1.2 Scalability

From optimizing energy use in a single autonomous household to managing a whole smart city, the main principles are the same but the complexity is not. The combination

of a large number of devices, each with its own characteristics and constraints, makes a hard optimization problem [Logenthiran11]. This issue needs to be considered as the development of smart grids is likely to be incremental, typically going from small scale experimentations of a dozen of households to a single district, then maybe to an entire city depending on the situation [Frey13]. If a different system must be deployed each time the scale changes significantly, necessitating the replacement of equipment, adding infrastructure and modifying interfaces, this development will be costly and unpractical. It is important then that an energy management system should be **scalable, i.e that its performances do not depend on the number of controlled assets** [Veit13]. In this regard, decentralized approaches have a clear advantage over centralized ones where the cost in resources and especially in complexity is directly linked to the number of controlled assets [Dagdougui14].

2.3.2 Adapting to changes

Power grids are not static structures, whatever their type or size. Over time, almost every aspect of it will change: its components' properties change with wear until they are replaced, some are added by new users, settings and preferences vary, and global constraints can also be modified by new regulations or other external reasons. For an Energy Management System to last and provide a steady quality of service in the long term, it is then crucial that such changes do not require profound and expensive updates.

2.3.2.1 Modularity

It is maybe trivial to acknowledge that an energy management system should be able to accept addition or removal of assets without major disruption. However, the frequency of such changes in some cases could prove challenging depending on the model. It is important for example that a homeowner can replace its appliances or install new equipment with only minor intervention to the management system. The interest of such a "plug-and-play" feature is advocated in [Colson10], with [Mauser15] following a modular approach in the same objective. In this regard, we consider an energy management system as **modular** if new physical equipment can be added dynamically without modifying the system in place.

2.3.2.2 Adaptiveness

We explained why a viable energy management system must be generic enough to accommodate various settings but also modular to allow changes to be made to the grid composition. In a shorter time scale, parameters like user preferences or components properties can be subject to real time changes that also need to be accounted

for. Thermal comfort is a good example, with user-defined set-points for indoor temperature that need to be reached at the correct time. A plethora of parameters can be learned : battery charging time, room occupancy, ideal indoor temperature, preferred lighting, user's priorities regarding economical and ecological concerns, etc... For example, Jin et al. [Jin17] use machine learning to estimate usage patterns, and Mazac et al. [Mazac14] propose a generic algorithm to learn relevant patterns (room temperature increase, occupancy, etc...) from raw sensor data, enabling the dynamic adaptation to real environment rather than relying on theoretical models. Learning user preferences and behavioral patterns is a complex task and out of the scope of this thesis. Using a similar approach, Nigon [Nigon17] uses a multi-agent system to learn the behavior of complex systems. However, regardless of the achievable accuracy on this field, the ability to handle the variability of these parameters is an important factor of resilience and robustness.

2.3.3 Limiting the cost

2.3.3.1 Simplicity

Cost is often a decisive criteria in the choice of a system considered for large scale deployments. As computing power and ICT in general gets cheaper and cheaper with technological progress, we argue that the final price of an Energy Management System will largely be determined by its simplicity and its robustness which impact not only its deployment but also its maintenance in the long run. Without advocating for a simplistic approach to energy management, it seems quite straightforward that ease of understanding translates into a facilitated user engagement and above all a quicker development of new features and adaptations.

2.3.3.2 Robustness

The robustness of an energy management system is defined by its fault tolerance and relies on its ability to react and adapt to changes without breaking, but also on the structural strength of its architecture. This resilience largely depends on the amount of critical nodes on which depend the operations. In this regard the degree of centralization of the system is a good indicator of its resilience. If many approaches in the literature are centralized, like in [Koutitas12, Chakraborty15, Conti10], a clear trend toward decentralization can be observed [Olivares14]. However, it is interesting to note that among the works using a decentralized approach most of the proposed designs use some sort of hierarchical structures like [Pournaras14] and only a few offer a strictly flat organization [Karavas15]. In another example, Basir Khan et al. [Basir Khan16] use "facilitators" to manage communication between the different entities.

2.4 Multi-Agent Systems

Multi-Agent Systems (MAS) seem to be a promising approach to provide the scalability, adaptability and robustness needed for a reliable Energy Management System, judging by the significant amount of works using this paradigm in the literature [Kantamneni15]. In this section we will describe in more details what defines a MAS with a special focus on their application to EMS.

2.4.1 Concepts

In computer sciences, a Multi-Agent System or MAS is typically a software representation of multiple entities that collaborate to solve a given problem. It emerged as research field of its own as a mix of artificial intelligence, distributed system and generally software engineering, but also inherited from sociology, artificial life, game theory and economy [Ferber97]. It can be tempting to call any system composed of several distinct entities a Multi-Agent System, however as Rohbogner et al. point out in an thorough analysis of MAS applications in Smart Grids : "It seems that engineers use the term 'agent' without a common understanding of what it actually embodies." [Rohbogner13].

The term "agent" comes from the Latin verb *agere* meaning "to conduct, manage, perform, do". If various definitions of the concept of intelligent agent exist [Franklin97], Wooldridge and Jennings [Wooldridge02] describe a software artifact exhibiting reactivity, pro-activeness and social ability. An agent is then defined by its ability to respond (react) to changes in its environment, initiate actions on its own and communicate with other agents in order to satisfy its objectives. The properties that interest us in our work are the inherent robustness and scalability of distributed architecture, along with the modularity and adaptability that are among the main focuses of the research in MAS.

Designing a MAS consist in defining the local process that will allow a satisfying global behavior of the population of agents to achieve the given objectives. The complexity of this local process is what separates the two main approaches in the field : reactive or cognitive models. In reactive models, agents are designed with limited proactiveness and only display simple behaviors are described as "reactive". Their actions are then comparable to reflexes, as rules dictate the immediate reaction that should follow a given external signal. Such agents often have no clear notion of goal, but they can be designed so that a complex behavior *emerges* from their collective actions. The term *self-organization* defines the spontaneous creation of order in a population of agents without the existence of a predetermined authority or control system. Designing a system achieving a given task following this approach is generally not a straightforward process as it can be difficult to design the local processes and behaviors which, when put together, will produce the desired result. This is the reason why *bio-inspiration* is a trademark of Multi-Agent System research. Indeed, nature is not short of interesting

examples of self-organization, one of the most used and cited being ants colony and particularly their foraging technique in which they use pheromone trails to reinforce shortest paths [Kakas11]. This indirect communication using markers on the environment that influence the actions of others is called *stigmergy* and is a key concept in MAS. On the other end of the spectrum, cognitive models give the agent a knowledge of the purpose of the whole system, its situation, and their respective roles in it. This approach is more straightforward to design as the final objectives can be more explicitely translated than in a more reactive model, but it also generally requires more computing resources for each agent. More complex agents allows for more complex organisations and roles, with hierarchies, coalitions and negotiations being very present in the literature, which has been comprehensively surveyed in [Horling04].

2.4.2 Applications in smart grids

The distributed nature of a power grid and the need for robustness and reactivity in energy management led many researchers toward multi-agent approaches, as noted by [Olivares14]. The literature on this field has been recently reviewed by [Harmouch16] and [Kantamneni15], who note the advantages of MAS approaches considering scalability and resiliency. These same considerations led us to also follow a multi-agent approach in our designs.

As pointed out by [Rohbogner13], a majority of MAS-oriented approaches to smart grid energy management are local market models where agents representing the different assets at varying scales trade energy, for example in [Chung13, Dimeas05, Foo. Eddy15] or [Linnenberg11]. In these examples as in the majority of models following this approach, a hierarchical organisation is used where coordinator agents centralize and match bids at one layer then trade similarly at a potential upper layer. These systems are designed for real-time negotiations and not for scheduling energy use on a given interval as they are primarily based on a price-following competitive system. Their inherently centralized architecture can be a major drawback regarding fault-tolerance, as the absence of horizontal communication mechanism makes the whole system dependant on the central coordinator [Dagdougui14].

Market-based MAS are not the only centralized approaches. In many models in the literature, agents have fixed roles in the system with varying responsibilities. In [Colson11a] for example, an "observer agent" is in charge of communicating the state of the grid to the others. In [Pipattanasomporn09], a "control agent" has a similar role of oversight, while a "database agent" works as a communication manager. [Massimo Cossentino11] designed their MAS for microgrid control as a "society" of agents where a "supervisor society" is in charge of the decision making process with a "broker agent" and a "policy manager". While these models are sometimes not completely centralized, they do present critical agents on which a large part of the system's working relies, which renders them more vulnerable to failures.

Other drawbacks of centralization are the lack of modularity [Foo. Eddy15, Karavas15] but also the potential security and privacy issue inherent to a mechanism where informations of the different stakeholders (agents) are necessarily explicitly shared [Klaimi18]. In the distributed MAS proposed by [Jiang06] for example, there is no central agent but a common shared repository where agents register their capabilities, allowing others to request services and place bids. While this solution might be more fault-tolerant as the decision-making process is decentralized, sensible information about usage patterns and needs are shared openly. Occurrences of fully decentralized models, like the one proposed in [Dagdougui14] which uses a peer-to-peer coordination mechanism sharing only necessary information, are rare in the literature. Some argue that "in a system where the presence of strong coupling between various operating units requires a minimum level of coordination, a fully decentralized control is unable to achieve stable operation based on local information alone" [Foo. Eddy15]. As a compromise, holonic architectures are an interesting form of hierarchy that stems from the concept of *holons*, where an each element is composed of one or more subordinate elements, and can itself be a member of a super-ordinate one. These nested structure can be found in many instances in biological or social systems and are particularly fit to represent power grids, as a load can be part of a larger entity (the house for example) which itself can be a element of a sub-network (a microgrid) which is connected to the wider powergrid, etc... [Frey13]. This form of organisation, also used by [Mauser16] in their smart building energy management system, offers the needed modularity of frequently changing configurations and limits the optimization process of each agent to a manageable scope.

2.5 Summary and positionning

The objective of this chapter was to give a comprehensive overview of the motives and context of this thesis. In section 2.1 we explained the principles behind current large scale power systems and their drawbacks, and how a panel of solutions could enable a paradigm change in the way energy is managed and enable the integration of renewable energies. In this thesis, we focus on residential consumption management through the concepts of demand response and residential microgrid, which are addressed in Chapter 4 and Chapter 5 respectively. In this regard, we explained section 2.2 the importance of user acceptance and even user engagement. While a diversity of solutions for user consideration in specific use cases are proposed in the literature, one of our main concern in this work was to propose an open platform enabling various kinds of interactions with the users, making sure that the interfaces can evolve and be adapted to different needs without changing the system itself. Thus we consider the user as an external constraint that the system has to deal with, raising the question of how to ensure a fair treatment by the system while still letting them in control of their appliances. Specific mechanisms such as gamification (see section 2.2.2.2) or financial incentive (see

section 2.2.2.1) are out of the scope this thesis and are complementary of your model, e.g. by including them in the user constraints measure (as discussed in section 5.3.2) For the same reason, in section 2.3 we defended the need for an easily deployable system, which implies a high robustness, adaptiveness and reliability in order to handle the diversity of use cases when dealing with residential buildings and smart-grids in general. These requirements, which were not previously clearly acknowledged in the literature, make the multi-agent paradigm particularly suited for the design of energy management systems, so in section 2.4 we gave an overview of the field with a focus on their architectures and how it affects their robustness. In the literature we find a clear argumentation in favor of decentralization regarding robustness and scalability, however all the approaches we reviewed featured a centralization of some sort, either for control purposes or in the communication protocol. Considering this, we aimed at showing that a fully decentralized energy management system was possible, i.e with no dependency between any agent in order to provide the fault tolerance needed for a robust energy management system. Before describing our core contributions in Chapters 4 and 5, Chapter 3 describes our exploratory work on short-term individual load forecasting.

3

Short Term Residential Load Forecasting

▷ In this chapter we identify the challenges related to the problem of individual short-term residential load forecasting and give an overview of the existing works. Then we describe our approach using exogenous variables, the obstacles we met, and the preliminary results we obtained. ◁

Chapter content

3.1 Problem statement **39**
- 3.1.1 Aggregation level 39
- 3.1.2 Time scales 40
- 3.1.3 Local variables 40

3.2 State of the art **41**
- 3.2.1 Clustering 41
- 3.2.2 Linear Regression Models 42
- 3.2.3 Non-linear regression 42
  - 3.2.3.1 Artificial Neural Networks 43
- 3.2.4 Exogenous Variables 44

3.3 Methodology **45**
- 3.3.1 Datasets 45
  - 3.3.1.1 Ubiant's office 45
  - 3.3.1.2 UCI's dataset 46
- 3.3.2 Algorithm parameters 47

3.4 Results **48**
- 3.4.1 Very Short Term Load Forecasting 48
- 3.4.2 Influence of exogenous variables 49

3.5 Summary **51**

3.1 Problem statement

The prediction of future consumption is crucial for energy management as it tells the operating system what the power grid, whatever its size, will need to provide and when. In chapter 2 we explained the challenges encountered by traditional power grid management, how it relied on a day-ahead load forecasting on the global scale to schedule its production, and why being able to act on the demand side is a crucial step toward smarter grids as the energy source are increasingly difficult to schedule.
We also explained that residential loads account for a significant share of the demand, and argued that managing them requires the inhabitants to be proactive or at least compliant with the process, implying that the load management system must be able to deal with the preferences and needs of every building individually. To this end, *an accurate prediction of the building consumption is an important, if not crucial tool*.

In this chapter we study the state of the art in short term residential load forecasting techniques, discuss and do preliminary test of their accuracy and identify potential opportunities for improvement. In this section we define precisely the kind of settings we focus on in terms of aggregation level and time scales, and we bring up the possibility of local variables exploitation. In section 3.2 we review the state of the art in load forecasting with a clear focus on the use cases that interest us. In section 3.3 we describe how we tested the algorithm we selected, and on which datasets, and we analyse the preliminary results we obtained in section 3.4 before concluding on this exploratory work in section 3.5.

3.1.1 Aggregation level

In the traditional way of managing a power grid, generation is scheduled on a day to day basis to match the global load forecast. Generation being centralised, this global forecast accounts for a large number of consumers : the french grid for example powers more than thirty millions households[1]. At this scale, variations of individual consumption are hidden by the global trend due to the averaging effect. As intermittent generation becomes more and more pervasive in the production landscape, the scale at which energy is to be managed tends to be considerably reduced [Hernandez23]. Typically, the concept of microgrid described in 2.1.3.3 gathers local generation and loads in sub-networks of small sizes, often less than a dozen of residential buildings. In such a situation, the typical variation of one building's consumption represents a far more significant share of the network's consumption that on a national scale. The question is then : *can a disaggregated consumption be accurately predicted, ideally at the level of a single household, with state of the art methods ?*

[1]CRE, Les marchés de détail de l'électricité et du gaz naturel, 31/03/2018

3.1.2 Time scales

Multiple interdependent parameters related to time can be considered in load forecasting. In the same way the aggregation level is significantly reduced by the profound changes in energy management, the time scales used for microgrid energy management are different from the one used in global power grids.

The key parameter is the time steps or interval used, meaning the duration between two consecutive predicted points. As load forecasting relies mostly on past data, the duration of this interval is directly linked to the granularity of the available data : predicting the load hour by hour can only be done if it has been measured at least at this rate. In the next sections, *real-time* will refer to an interval of one minute or less between two measures.

Another important parameter is the horizon, corresponding to the duration over which the load needs to be predicted. Today, production is scheduled a forecast of the next 24 hours and, on the larger scale, investments in grid infrastructure and new power plants are planned using yearly projections. Theoretically, the further the horizon, the better it is for the energy management system regarding operational scheduling, as long as the forecast is accurate. In the case of microgrids however, the volatility of production and demand makes long-term forecasting difficult and ask for reactivity rather than long-term planning.

The last parameter is the amount of past measurement fed to the load forecasting algorithm. It can range from many years of past consumption data to the last few minutes only. This depends obviously on the total amount of data available but it is potentially also limited by the specifications of the device where the computation is made, for example in decentralized energy management models where it would be processed locally with limited resources. We did not have any objective or constraints for this parameter as its limitations are mostly technical and not theoretical. Obtaining good results with a limited amount of past data is definitely interesting in many regards but it was not our focus here.

The focus of this thesis is the management of energy in the context of future smart grids and in terms of time scales, the most demanding configurations are microgrids. Indeed, as the demand in such disaggregated environment is highly volatile [Chitsaz15] and the production and storage asset's capacity is limited, reactivity is key in order to maintain the grid balance. In such cases, a minute by minute interval, which is often called "Very Short-Term Load Forecasting (VSTLF)" [Hernandez23], would be a reasonable objective.

3.1.3 Local variables

The consumption of a household depends on multiple factors. Some of them, like time of the day, of the year, or the weather, are already used in large scale load forecasting.

However these variables only provide information on large time and spatial scales, they do not vary by the minute (weather changes are mostly relevant on an hourly basis) and do not change between two neighboring houses. It is then unlikely that they could be used to predict changes in the consumption on smaller scales. More local information however, like the presence of the occupant in the house or the temperature inside a room are very likely to be correlated to changes in the building's consumption. The developments of microgrids often goes hand in hand with the integration of connected equipment in buildings, so measurements allowing us to monitor these variables along with consumption are becoming more and more available. Ubiant, through the deployment of its Home Energy Management solution, had begun gathering anonymized records from various sensors in inhabited households. One of our objectives in this work was then to determine whether these exogenous variables could be used to improve the accuracy of the forecast.

3.2 State of the art

A very extensive survey on load forecasting techniques for smart-grids has been written by Hernandez et al. in 2014 [Hernandez23] and is a recommended read for a comprehensive view of the topic, listing the various objectives and approaches and comparing the performances of the studied works. It covers a broad range of time scales but focuses mostly on short term (hours to weeks) and medium to long term load forecasting (months to years), with less than ten percent of the reviewed papers relating to very short term, i.e under an hour. It highlights the lack of works focusing on disaggregated environment, citing the growing need for accurate forecasts on small sized networks such as microgrids and smart buildings. The authors also acknowledge the potential of using relevant input data other than historical load to improve the accuracy of the prediction, however they found a small minority of works in this direction, mostly exploiting only weather data. More recently, Deb et al. [Deb17] published a similar survey with detailed explanation of a large variety of methods, sharing similar observations. Here, as stated in the previous section, we want to focus on the very short term load forecasting in disaggregated environment, with a particular interest on the use of exogenous variables. In this section we review the main approaches to load forecasting, i.e classification, linear regression and non-linear regression, the latter representing the large majority of the contemporary literature on the topic.

3.2.1 Clustering

The principle of classification methods for load forecasting is to choose among a number of past load patterns the one that is the most likely to appear in the next steps. The most used method is the k-means algorithm which is mostly combined with other algorithms, like Self-Organizing Maps [Chicco04, Zhang12, Tsekouras07] or genetic al-

gorithms [Yang05] to predict the consumption for the next 24 hours. It is notably used by Valgaev et al. [Valgaev16], using calendar informations in conjunction with the past consumption period to find the most likely pattern for the next period. However this method, which is also currently used by Ubiant to forecast hourly, weekly and monthly consumption, is not well suited for smaller time scale as many patterns might not reproduce exactly at the same time from one day to another.

3.2.2 Linear Regression Models

Linear models appeared in 40s to forecast demand, and among them appeared the very popular ARMA models [Abou-Hussien81]. ARMA or ARIMA for Autoregressive (Integrated) Moving Average is an efficient tool which capture the seasonal patterns of time series, making it well suited to forecast the total load of large power systems on the long term [Edwards12, Weron07]. Indeed, the averaging effect and the relative similarity of load profiles present smooth aggregated load curves. On a small scale, it has been used to predict peak loads [Rajurkar85] or to predict electricity prices [Conejo05, Contreras03]. The general principle is that the forecast value is a weighted sum of the past values. In practice, this approach requires an extensive knowledge of the problem to manually tune the parameters of the model. As said in [Hernandez23] : "Demand prediction is a complex problem due to its non-linearity, and it is therefore a big challenge to translate it into linear models". This observation is shared by [Chitsaz15] whose authors join the consensus among recent papers that non-linear regression methods easily outperform linear regression thanks to the recent advances in neural networks.

3.2.3 Non-linear regression

Non-linear models and notably Artificial Neural Networks (ANN) in general appeared in the 80s as an efficient way to represent both the regularities and stochasticity of electrical consumption time series. They have been extensively studied and used in load forecasting context. A highly cited review on neural networks for load forecasting written by Hippert et al. [Hippert01] makes a good starting point on the matter. It confirmed the interest of the approach at the time, recommending more rigorous testings. Among the more recent papers on the subject, we want to highlight an interesting benchmark of non-linear models on hourly consumption prediction published by [Edwards12] and a similar work by [Mocanu14] providing a comparison of different deep learning methods, but targeting aggregated environments.

3.2.3.1 Artificial Neural Networks

Artificial Neural Networks are a statistical learning method inspired by biological neurons and constitute the basis of what are called *deep learning* algorithms. They are composed of formal neurons which are mathematical functions taking a number of inputs , combining them linearly (in the majority of models) before using a non-linear function called the *activation function* to determine its final input (see figure 3.1a). The kind of neurons used and the way they connect with each other define the kind of network being used. In one of the most common model called a Multi-Layer Perceptron (MLP), these neurons are grouped in *layers*, where the output from each neuron of one layer is used as input for each neuron of the next layer. In its simplest form, only two layers are used : the first, called the *hidden layer*, is directly connected to the input variables and sends its outputs to the second, called the *output layer*, which then expresses the final result of the network for a given set of input variables. The term *deep learning* is used when more than one hidden layer is used. Figure 3.1b depicts such an architecture.

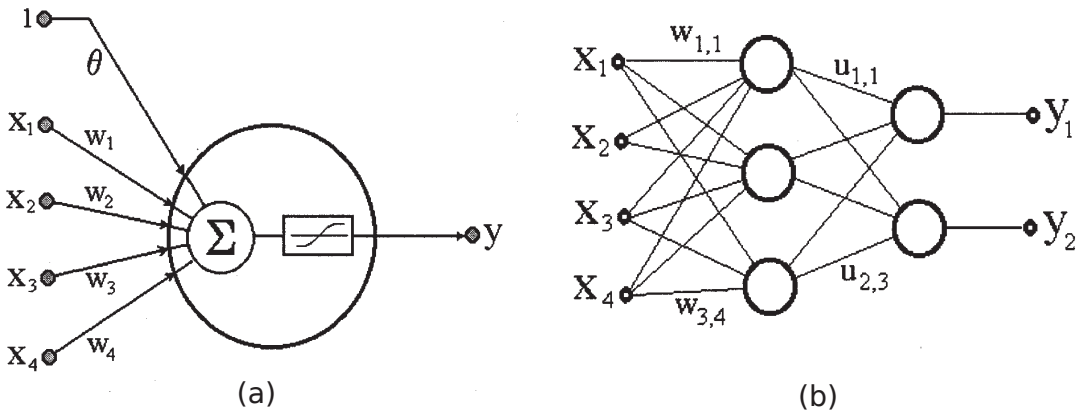

Figure 3.1: (a) An artificial neuron. (b) A simple two-layer feed-forward neural network. This figure is provided by [Hippert01].

In order to produce satisfying results the network must then be *trained* by tuning the weight of each neuron. This learning process is the most important step in the design of a neural network and various approaches exist, usually using a descent-based method to minimize a manually defined error function. The performances of the network, as well as the learning speed are related to its architecture and its learning algorithm but also to the choice of the error function and of the input variables. Theses variables can prove difficult to select in some cases, which is why *hybrid approaches* have been studied, where an optimization algorithm is used to automate the learning process. This can be done for example using meta-heuristics such as Particle Swarm Optimization [Liu14], evolutionary algorithms [Amjady10] or even a Fruit Fly Optimization algorithm [Li13]. Another interesting approach are *Ensemble models* which have been the focus of a survey

and benchmark by Wang et al.[Wang17]. Unlike hybrid models which aim at perfecting the predictions of a single algorithm by using different methods for the different steps of the forecasting process, ensemble model bet on quantity and run multiple learning processes in parallel. In a heterogenous ensemble model for example, different algorithms are trained on the same set of data in hope that they complement each others and combine their strong points to offer a satisfying performance [Jetcheva14]. In the same objective, a homogeneous model will use only one algorithm but will run multiple instances on different subsets of the input data. Gaillard et al. [Gaillard16] used such an approach to predict the evolution of electricity prices. In the conclusion of their survey, Wang et al. [Wang17] noted the lack of works considering disaggregated environment and short term forecasting but also state that "The study of incorporating occupancy information into prediction model has a greater potential to improve the prediction performance."

Among the approaches using a single non-linear model, the Self-Recurrent Wavelet Neural Network (SRWNN) developed by Chitsaz et al. [Chitsaz15] has been designed with single residential loads as a target. In their model, the inputs are chosen using a custom feature selection algorithm explained in [Amjady11]. Their neural network uses the Morlet wavelet function as the activation function in the hidden layer, which result is then fed back in with the next input (recurrent Elman network architecture). The model is then fitted using the Levenberg-Marquardt algorithm which uses a gradient descent.

3.2.4 Exogenous Variables

In the survey by Hernandez et al. [Hernandez23], the need for new variables in the forecasting process "such as data which anticipate the operation of the main loads" is mentioned but not focused on. In large scale demand forecasting (hourly, daily and monthly load forecast on large geographical areas), the weather and calendar data have long been used [Abou-Hussien81, Drezga98, Valgaev16].

However, these variables carry less information on smaller scale : the weather does not discriminate between neighboring homes, and appliance usage is not regular to the hour, let alone to the minute. The interest of more local variables, notably occupancy, has been brought up by several studies [Nguyen13] ; although this information is harder to obtain. As Wand and Srinivasan explain it : "Most of the studies did not adopt occupancy information because it is hard to acquire quality data; for example, lack of occupancy sensors and other privacy concerns have led to data unavailability." [Wang17]. Jetcheva et al. [Jetcheva14] share the same insight when describing their approach to the hourly forecast of a industrial building consumption via an ensemble model, where they lacked such occupancy data. Kwok et al. [Kwok11] successfully use occupancy data to improve the predictive accuracy of the cooling load of a large commercial building (hourly). Similarly, works presented in [Massana16, Shi16, Kamaev12]

and [Li15] all leverage occupancy and sometimes device-specific consumption to improve their forecast accuracy. All these works however focus on non-residential building and aim at predicting hourly consumption only, with the exception of [Mena14] which predicts the next 60 minutes. The model presented by Aung et al. [Aung12] consider a residential building but uses only past temperature data and above all predicts only the peak load for the next time unit, not the actual consumption.
To the best of our knowledge, despite a number of ongoing work on real-time activity recognition in residential buildings [Nguyen13], there is no published work on the use of occupancy data to predict the consumption of a single home in the very short term, i.e hourly or less.

3.3 Methodology

Our survey of existing approaches to load forecasting distinctly showed the lack of study on disaggregated environment (typically a single household), on the very short term (less than an hour) and therefore no studies on the potential of local variables to increase the accuracy of such a forecast. Our primary objective here was to assess the performance of a state of the art algorithm on a significantly more demanding dataset, i.e on a minute by minute forecast of a single residential building energy consumption. Considering the high variability of real-time consumption, we expect that the result would be significantly worse than on traditional aggregated hourly data. A second objective was then to determine whether the use of local variables could improve these performances. We chose the model presented by Chitsaz et al. [Chitsaz15] for our tests as it was a recent model designed for disaggregated environment with promising results according to the authors. The code was also accessible which facilitated our work. In this section we describe the datasets we used for this exploratory study, and the hyper parameters we used for the algorithm.

3.3.1 Datasets

3.3.1.1 Ubiant's office

To our knowledge there is no public open dataset for minute by minute measurement on multiple variables for residential buildings, we turned to the data coming from Ubiant's deployed systems. Unfortunately, the number of available real-time measurements was very limited due to the early stage of deployments.

We set up a platform for automatic gathering, cleaning and qualification of real-time data, which would allow us to work on an increasing amount of relevant datasets in the future. Energy consumption is given by the smart meter with a frequency varying between 2 and 30 seconds depending on the device. Other sensors also either do not have a regular output (for example a movement sensor only transmits when move-

ments are detected) or have a different rythm depending on the brand. To convert this irregular data stream to an exploitable dataset, we made the time scale constant either by aggregating or by interpolating the measures depending on the original granularity. The final data sets are then composed of minute by minute data from a single building, with at least the instantaneous energy consumption (in kW) of the building.

The only dataset containing enough data at this time comes from Ubiant's own offices and contains a week of real-time (one minute interval) monitoring from the smart meter and more than a hundred connected sensors measuring movements, doors and windows opening, indoor temperature and air quality (humidity, CO2, VOC) and brightness. Among these sensor data are also recorded switch pushes which, with many other events like movements, doors and windows opening and CO2 changes, are human activity indicators that are processed inside HEMIS (Ubiant's solution) to create a virtual occupancy sensor that is also recorded in real time. Ultimately, we had a week long record of minute by minute measurement from the Ubiant's office in Lyon, France. It consisted of 14678 data points going from 09/06/17 at 09:29 to 19/06/17 at 14:06. Table 3.1 describes the 9 different variables that are recorded on multiple rooms and aggregated at the building level (totalling 55 measurements each minutes).

Variable	Description	Unit	Min	Max
IPOW	Instantaneous consumption	kW	0.92	4.81
OCC	Occupancy	%	0.00	100
OPE_CLO	Last open/close sensor value	0/1	0	1
TMP	Room temperature	°C	20.68	27.76
BRI	Brightness in the room	lux	0.57	1020
CO2	CO2 concentration	ppm	326	951
VOC	VOC[2] concentration	ppm	1.00	1661.5
HUM	Relative humidity	%	30.20	57.20

Table 3.1: Description of the variables recorded by Hemis in the dataset from the Ubiant office

3.3.1.2 UCI's dataset

The algorithm we wanted to test had been designed to predict hourly electricity use, so it could be expected that its performances on our minute by minute dataset would be much poorer. Also, this dataset coming from Ubiant's office was very small in regard to the amount of data usually needed to correctly train a neural network, and aggregating the data time-wise to increase the time scale would have reduced even more the amount of data available. To provide grounds for comparison, we chose to assess the performances of the algorithm on an open dataset from the UCI machine learning repository[3]. This dataset contains the electricity consumption of individual

[3] https://archive.ics.uci.edu/ml/datasets/ElectricityLoadDiagrams20112014

buildings (in kW) recorded every 15min for four years from 2011 to 2014, which sums up to 140256 data points.

3.3.2 Algorithm parameters

The algorithm used by Chitsaz et al. [Chitsaz15] is a Self-Recurrent Wavelet Neural Network (SRWNN) which architecture is slightly more complex than that of the typical Feed-Forward Neural Network described in 3.2.3.1, and can be visualized in figure 3.2.

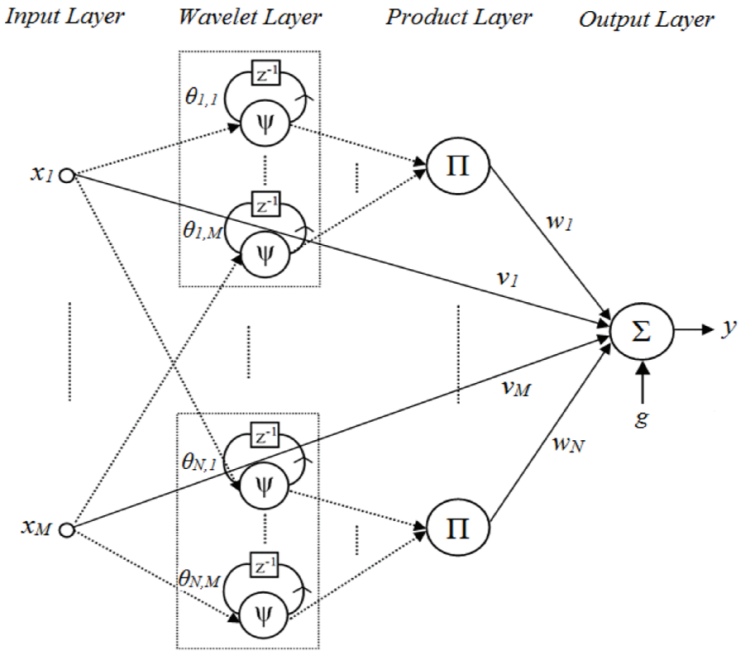

Figure 3.2: The Self-Recurrent Wavelet Neural Network used by Chitsaz et al. [Chitsaz15] (figure from the original paper).

The input variables $x_1 ... x_M$ are first processed by the *wavelet layer* which consists of NxM neurons using a feedback loop (hence the Self-Recurrent part of the name) that stores the previous state of the network and a Morlet wavelet function [Grossmann84] as the activation function $\psi_{i,j}$:

$$\psi_{i,j}(x_j) = \psi\left(\frac{x_j + \psi_{i,j}z^{-1} \cdot \theta_{i,j} - b_i}{a_i}\right), \qquad \psi(x) = e^{-0.5x^2}cos(5x) \tag{3.1}$$

With $i \in N$, $j \in M$, a_i and b_i a scaling and shifting parameter respectively, $\psi_{i,j}z^{-1}$ the value of the previous iteration and $\theta_{i,j}$ the weight of this feedback. The M inputs

are processed in parallel in N groups and the results sent to a *product layer* where the product of the M wavelet function are combined as follows :

$$\Psi_i = \prod_{j=1}^{M} \psi_{i,j}$$

These products are finally combined with the direct inputs in a weighted sum so that the actual output of the SRWNN is as below :

$$y = \sum_{i=1}^{N} w_i \cdot \Psi_i + \sum_{j=1}^{M} v_j \cdot x_j + g$$

w_i and v_j are respectively the weight of the product neurons and of the direct inputs in the final sum and g is the bias of the output neuron. In total, the network has $M + (M+3)N + 1$ free parameters that need to be tuned by the training process, which we will not describe in details here. Instead we refer the interested reader to the original paper [Chitsaz15] where the learning algorithm is described. We followed the implementation of the authors and used the recommended hyper parameters for the learning algorithm. For both datasets, we used 2/3 of the data for training on 1/3 as test. We asked the algorithm to predict the electricity over an horizon H of 1, 5 and 15 next periods for each dataset. The inputs where chosen empirically to be the last 96 measures of electricity consumption. As an evaluation criteria, we used the classical Root Mean Square Errors (RMSE) defined by the following formula, where L_t^{obs} and L_t^{pred} being the observed and predicted load for the predicted period t.

$$RMSE = \sqrt{\frac{1}{H} \sum_{t=1}^{H} (L_t^{obs} - L_t^{pred})^2} \tag{3.2}$$

3.4 Results

3.4.1 Very Short Term Load Forecasting

Table 3.2 shows the forecast error on different horizons using the SRWNN on the two datasets presented in the previous section. As expected, the forecast error increases with the forecast horizon, and the accuracy of the prediction on minute by minute data is significantly worse than on 15min interval. However, the minute by minute dataset was far too small to allow for an efficient training, and these tests do not allow us to truly compare the performances of the algorithm on the two time scales as they are not performed on similar datasets.

Dataset	Scale	Interval	Method	Error on prediction horizon 1	5	15
UCI	Agregated	15min	SWRNN	0.0636	0.16	0.225
Ubiant	Individual	1min	SWRNN	0.217	0.411	0.496

Table 3.2: Prediction error of tested methods for different forecast horizon on the two datasets

3.4.2 Influence of exogenous variables

As detailed in section 3.3.1.1, the dataset coming from Ubiant's office contains a number variables coming from various sensors in real-time. Figure 3.3 shows the correlation coefficient found between the different variables and the instantaneous electricity consumption. It appears that the most correlated variable is the occupancy, with CO2 levels and brightness following closely behind.

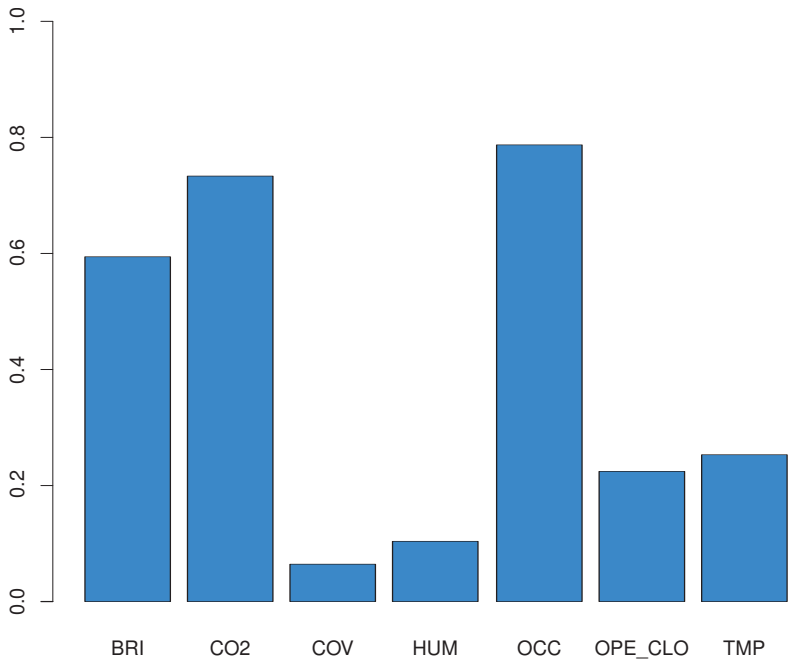

Figure 3.3: Correlation coefficient between the different variables of the ubiant's office's dataset and the instantaneous consumption.

A graph of both electricity consumption and occupancy can be seen on figure 3.4. It illustrates the correlation between occupancy and instantaneous electricity consumption, but also that occupancy is clearly not the only factor influencing consumption. Indeed, if the large periods of occupancy and higher consumption undoubtedly match, variations in occupancy do not systematically translate to changes in energy consumption.

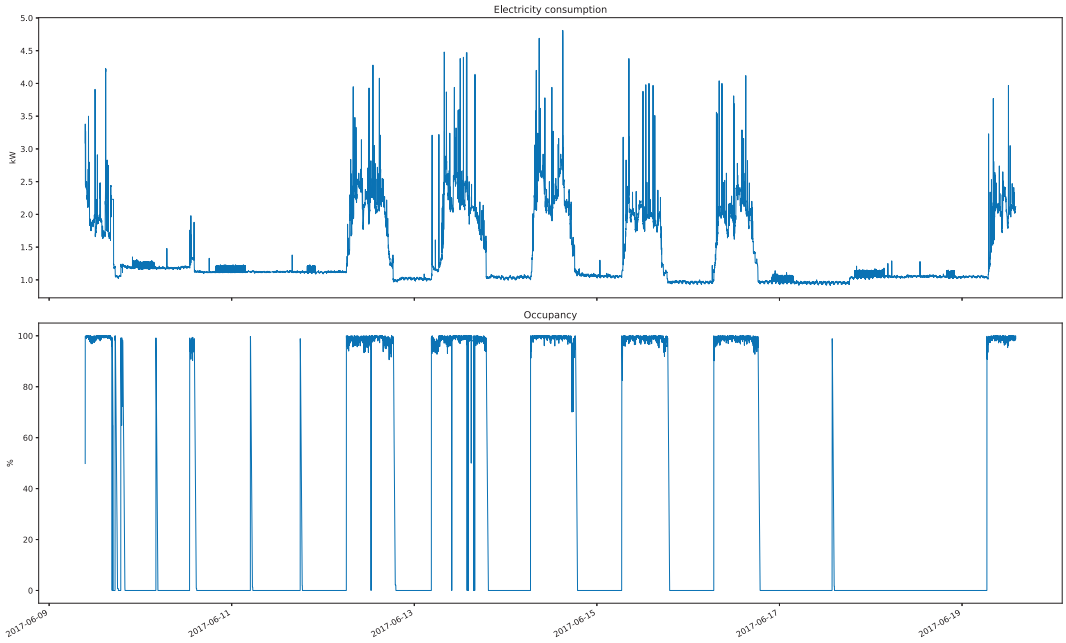

Figure 3.4: The minute by minute dataset from the Ubiant office. Only instantaneous electricity consumption (top) and occupancy (down) of the whole office are shown.

To confirm the potential of this additional information, we fed the SRWNN with past occupancy data in addition to past consumption. More precisely, the inputs where the last 96 occupancy measures along with the last 96 electricity consumption values, and the rest of the training was done as explained in section 3.3.2. In figure 3.5 which shows the error of the algorithm with and without occupancy data as input, we can see an improvement on the forecast over short horizons, the error being more noticeably smaller from horizon 1 to 5. It seems to indicate that exogenous data can indeed help the prediction, even without having to develop dedicated new architectures. These results have to be confirmed on other and larger datasets.

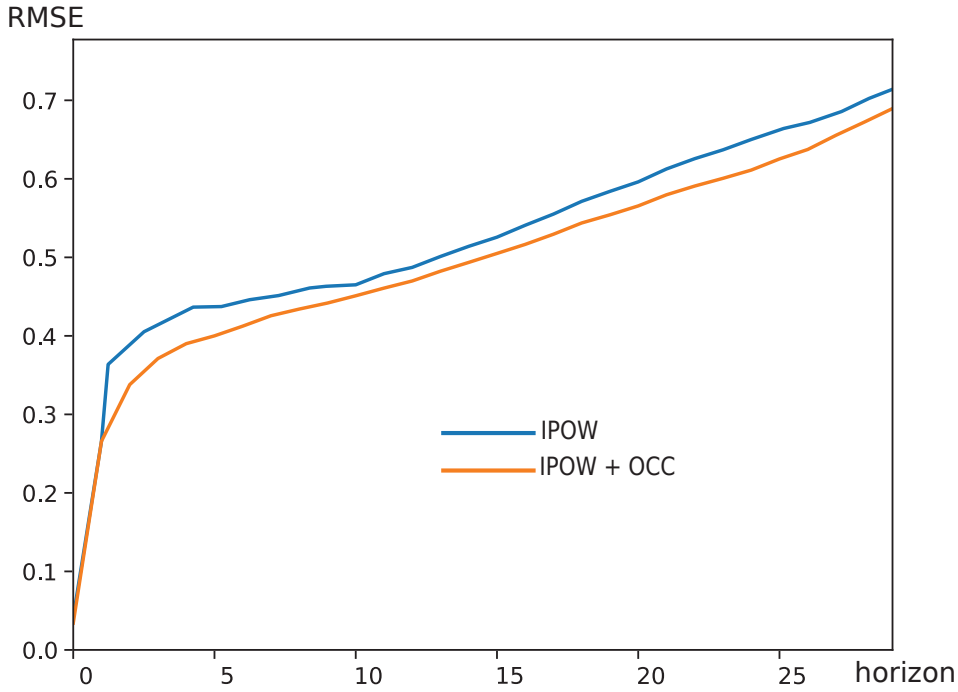

Figure 3.5: Error of the forecasting algorithm over increasing horizons, with and without occupancy data as input.

3.5 Summary

Our goal in this exploratory work was to determine the feasibility of individual very short-term load forecasting. Indeed, the emergence of small scale network such as microgrid comes with a need for a precise load forecasting which is crucial to manage production accordingly. At the scale of a small microgrid or even a single residential building however, the demand is far more volatile than that of an aggregated environment, as demonstrated by [Chitsaz15]. The reduction of network size also brings a reduction in time scales, as the small capacity of distributed generation and storage assets require reactivity in order to maintain the grid stability. As the variations in residential consumption is naturally linked to the activities of the inhabitants, we also wanted to evaluate the improvement in accuracy that could be achieved when using data from sensors in the household. We reviewed the literature on the topic focusing on approaches considering disaggregated environment and very short term load forecasting, and we studied their use of exogenous data such as occupancy. We observed, as did other surveys [Hernandez23, Wang14, Deb17], that a very small share of works is interested in these aspects. We identified a neural network model by Chitsaz et al. [Chitsaz15] that seemed promising in its results on hourly forecast of a single building, the authors having designed the system to be particularly fitted to disaggregated envi-

ronments. We gathered a minute-by-minute dataset from measurements in the Ubiant's office in order to test its accuracy on very short term forecasting, along with an existing dataset from the UCI with measurements every 15 minutes. As the Ubiant's datasets also contains exogenous data, notably occupancy, we confirmed the significant correlation of this variable with the instantaneous energy consumption and observed the influence of using it as an additional input variable on the neural network. Both tests gave expected results : the accuracy is significantly lower when the time scale reduces, and occupancy data seems to improve it. Unfortunately, the very limited size of the minute-by-minute dataset prevent us from drawing more definitive conclusion than mere trends, as it does not allow the neural network to learn from enough past data. As future research we want to validate our preliminary results on larger datasets that are currently collected thanks to our new deployments. However, as human behavior and in our case appliance usage is not perfectly regular it is likely that the accuracy of forecasting algorithms relying only on past consumption data will be limited under a certain granularity. The use of occupancy data and other sensor can certainly reduce the prediction error by adding context, and the democratization of connected equipments and commercially available sensors will allow a growing amount of data to be exploited to this end. However, it seems clear for now that the energy management systems for small scale networks can not fully rely on the accuracy of load forecasts to plan their operations, and must instead be able to cope with a highly volatile demand, which will be one of our main objective in the design of our energy management systems, described in the next chapters.

4

Decentralized User-centered Residential Demand Response

▷ In this chapter we show how a bottom up approach can be used to build a reliable Demand Response system thanks to a decentralized coordination algorithm. Instead of the users having to adapt to the grid balancing constraints, we put them at the core of the system which adapts in real-time to unforeseen variations. This reactive user-centered mechanism is built on top of a scalable and robust communication layer to offer a reliable service. ◁

Chapter content

4.1 Problem Statement **55**
4.1.1 Demand response mechanism 55
4.1.2 User constraints 56
4.2 State of the art **57**
4.2.1 Indirect load control 57
4.2.2 Direct load control 58
4.3 Proposed approach **60**
4.3.1 The Smart Building Solution: Hemis 61
4.3.2 The Push-Sum algorithm 62
4.3.3 Coordination problem 63
4.4 Decentralized Load Shedding Model **64**
4.4.1 General principle 64
4.4.2 Agent self-evaluation 65
4.4.3 Engaging the agent's flexibility 66
4.4.4 Adjusting agent's participation 67
4.5 Experiments **68**
4.5.1 Test protocol 68
4.5.2 Results 70
4.5.2.1 Reliability, robustness and scalability 70
4.5.2.2 Self-evaluation 73
4.5.2.3 Security margin 75
4.6 Summary **78**

4.1 Problem Statement

We saw in chapter 2 that power grids are currently under pressure due to significant changes on the production side. Indeed, increased irregularities in electricity generation ask for new paradigms to be found. A highly studied solution is Demand Response, in which consumers actively modify their consumption in response to grid needs. Various demand response mechanisms already exist, and they have been in use for a long time. However, they only exploit a fraction of the potential of demand response, as they usually exclude residential consumers (Section 2.1.3.3). Indeed, leveraging residential consumption into demand response adds two major constraints.

First, in order to obtain a load shedding capacity significant enough to matter on the scale of a national power grid, a large number of residential loads have to be aggregated and mobilized.
Secondly we saw in section 2.2.2 that the user's comfort should not be compromised on, requiring the system to handle the variability of residential consumption, which is very difficult to forecast (chapter 3). The objective of this work is to show that a bottom up approach can be used to leverage residential consumption to provide reliable demand response while satisfying these constraints. In this section we further detail the problem we want to tackle and the properties we want to obtain. In section 4.2 we review the literature related to residential load shedding and position our contribution. We explain our approach in section 4.3 and the model we designed in section 4.4. The results of our tests are presented in section 4.5 before a summary of this chapter in section 4.6.

4.1.1 Demand response mechanism

Our goal here is to design what is called a Demand Response Operator (DRO) : a system or entity that manages a set of connected buildings to adjust their consumption at the utility's request. The designation and the precise specifications can change depending on the country and the utility, but the general principle is the same regardless. This adjustment is typically done by reducing the consumption, which is called load shedding or curtailment, but demand response could also compensate a surplus in production by consuming more. In this work we chose to only consider load shedding because cases of overproduction are still rare even if they are beginning to occur in some parts of the world with a significant share of renewable energy production. Even then, excess of energy is less of a major issue than the opposite and can be dealt with relatively easily. For this reason current utilities mostly need demand response for load shedding purposes.

In the following sections, we assume that the process unrolls as follows:

1. The utility sends a load shedding offer to the DRO with a requested capacity in Watts, due to be reached at a given starting time and kept constant until a given

ending time.

2. The DRO answer whether or not they accept the offer
3. At the given starting time, the loads managed by the DRO reduce their consumption to reach a total curtailment matching the request. The stability of the reduced consumption is verified by the utility.
4. At the ending time the loads resume their normal consumption and the utility either rewards or penalizes the DRO regarding the quality of the load shedding.

Currently, in centralized nation-wide power grids, load shedding is relevant only if the available capacity is significant at the scale of the power grid. In France for example, the order of magnitude used to negotiate load shedding offers is the megawatt (1 MW = 1000 kW), while the average household instantaneous consumption rarely exceeds a few kilowatts. Numbers vary depending on the exact situation, but in order to gather enough shedding capacity from residential consumption the number of connected buildings must at least be in the thousands, and the more the better. This means that to provide a relevant service to the power grid, the residential load shedding system must be highly **scalable** (see Section 2.3.1.2). It also need to be **reliable**, i.e it must maintain a stable curtailment for the whole duration of the event. Finally, the large number of distributed components in the system combined with the high stakes of the grid stability requires the system to be **robust**, meaning that weakness points should be avoided as much as possible.

It follows from the process described previously that the Demand Response Operator must be able to *precisely know the available flexibility* of the managed loads (i.e the total shedding capacity of the system) at all time. When using industrial loads to provide demand response, knowing the available flexibility in advance and maintaining a stable demand during a requested period is straightforward as the consumption comes from managed and scheduled processes. Reaching the same level of control in a residential setting brings in additional constraints that we will described in the next section.

4.1.2 User constraints

The available flexibility of a residential load logically varies with the current consumption, but also depends largely on how the load shedding system can act on the various appliances. In section 4.2 we give an overview of the approaches used to control the building's consumption, and most importantly how they take the user's comfort into account, as the impact of demand side management on the users comfort and way of life is often neglected or under-emphasized [Giebel07]. In this thesis we argue that as residential loads exist solely to provide comfort to the users, a load shedding system that would impact negatively this comfort goes against the very purpose of the power grid by trying to maintain it. In order for the users to accept the idea of load shedding

and to actively engage in a Demand Response process, it must be integrated in such a way that the user can clearly and simply specify how the system can operate while maintaining its comfort. The system's interface with the users must then be easy to set up and to adjust in real time. Most importantly for us, it must **never interrupt or prevent the user from turning on or off an appliance**, even during a load shedding. These constraints mean that in order to be reliable and offer the necessary curtailment stability the system must be able to cope with unpredictable changes in the user's behavior, thus being highly **adaptive** and **reactive**, i.e to be able to compensate sudden variations in the available flexibility.

In addition to a non-intrusive integration into the life of the users, the system must also ensure that the burden of the load shedding is spread among the participants in a fair way. Depending on the situation, it is possible that the totality of the participants in the system might not be required for a given load shedding event, and in any case the amount of participation of each building will vary. The participation to a load shedding effort can be burdening (decreased comfort) or rewarding (financial incentive), sometimes both, depending on the situation. The fairness of the system then lies in ensuring that any preference regarding the participation of a building over another is justified by the needs of the system and that every participant is treated in the same way.

It should be noted that we choose to not consider any economic model in our work regarding the way the DRO is rewarded for a load shedding event, or the way the users can obtain a financial compensation for the reduction of their consumption. Various commercial schemes could be used regardless of the way the Demand Response System is implemented. Indeed, assuming that precise consumption is monitored and logged at all time as it is a prerequisite for any load shedding to be possible, all accounting can be done a posteriori. We left these questions out of the scope of this thesis, however price incentive-based systems will be discussed in section 4.2.

4.2 State of the art

We covered the inception and the main trends of load shedding in section 2.1.3.3. Here we review the different ways residential load shedding has been implemented in the literature focusing on the criteria and properties we find relevant in section 4.1, namely scalability, reliability and the respect of user comfort.

4.2.1 Indirect load control

The different approaches present in the literature can be described as exerting either a direct control on the loads or an indirect one. In the later case, which is the most widespread, the power consuming appliances are controlled by the end user only. The load management system then consists in incentives aimed at encouraging the desired

behavior from the user, typically using the electricity price as a signal (time-of-use pricing or real-time pricing). This technique is often used by utilities who offer contracts in which the cost of electricity varies depending on the needs of the grid, thus pushing the users to avoid consuming when the price is high and deferring the use of power-hungry appliances to periods where the price is low [Samadi12]. Multiple variations of this method have been proposed, some with fixed rates depending on the time and day, others with a more dynamic pricing communicated in advance to the user in different ways (flashing indicator, smartphone apps, etc...) [Mohsenian-Rad10]. The incentive can also be unrelated to the electricity price and rely mostly on social pressure and rewards. Some works have been done on interfaces allowing the user to compare its consumption to its neighbor's, or to know the amount of stress the grid is under at any time, encouraging changes in energy use (see section 2.2.2.2).
The main advantage of these indirect approaches is their *scalability* [Weller88]. Indeed, they generally do not require any complex installation on the user side, and can then be deployed at very large scale with very little resources. However, what these approaches gain on volume, they loose it on control. The principal and blocking issue with incentive-based load management control is that the actual result is hardly predictable [Callaway11]. Users will react to changes in the signal in a time and fashion that is not known by the utility, making precise scheduling of load shedding difficult by relying on statistics only. Consequently, those mechanisms are *not reliable enough* to be used alone as a Demand Response system as defined in section 4.1.1. Their great simplicity and scalability are still very interesting properties that should be aimed at.
Regarding the user's comfort one could argue that incentive-based approaches, by simply encouraging the desired behavior, does not get into the way of the every day life of the users. However, if they do not directly constrain the user, not following the incentive can have significant consequences, notably with price incentives where unwanted behavior (consuming when the demand is high for example) can be penalized by considerably higher electricity bills. In order to not be bothered by the economics of the system, the user must constantly make sure that their consumption matches the constraints of the grid, which does not fit our requirements for a user-centered system.

4.2.2 Direct load control

Approaches that fall under the direct control category do not rely on the user alone to act on its consumption and use a direct connection to the actual building and its appliances in order to modify their operation schedule. Contrary to the indirect control described previously, *the reliability of the system is significantly higher* depending on the level of control the user maintains on its appliances, and corresponds to our needs in this work. However, other major drawbacks can be found depending on the control architecture used. Regardless of the technical implementation, these approaches can be further classified regarding the way decisions are taken.

In most cases the system is centralized, meaning that a single algorithm chooses the next action of each controlled appliance, in a top-down fashion. In France, the biggest (and until a few years ago the only) energy provider uses a "centralized remote control by musical frequency" which turns on and off water heater in households via power-line communication [Tavernier46]. This method allows the quick mobilisation of a large adjustment capacity but lacks in precision as the orders are blindly broadcast, even if they can be limited to selected sub-network. The impact on the user's comfort is limited as the only affected appliances electric water heater that benefit from a good thermal inertia. The same approach has been followed and enhanced by others in commercial systems that were able to control more appliances in households, but with limited success as reviewed by [Eid15]. The most obvious issue with this kind of system is the *lack of robustness* due to the single point of failure [Lu11]. Security measures like redundancy can not completely compensate the *inherent fragility of a centralized model*. *Scalability* is also an issue as the complexity of the decision process is generally in direct relation to the number of connected appliances or buildings [Ramchurn11]. The principal advantage of a centralized approach is its straightforwardness, as all the information needed for the decision making process is available for system to use classical multi-objectives optimization algorithms.
Instead of a purely centralized system, distributed approaches have been explored where the decision making process is shared between the participants in the system. However, many works in this field use hierarchical architectures which, while indeed not being totally centralized, still present critical elements. The work presented in [Karfopoulos15] uses such an architecture, modeling "coordinator agents" aggregating information at various levels. Similarly, Lim et al. [Lim14] present an agent-based load shedding system where a control agent receives bids from production agents and load agents before dispatching orders. In another example, The work presented in [Pournaras14] use a hierarchical architecture where information goes up and decisions go down. Regarding user satisfaction, their system uses two different polls to evaluate the user's willingness to participate in load shedding effort and configure their system accordingly. While this can improve the reliability of the load shedding by lowering the share of uncooperative participants, it does not prevent the system from impacting the user's comfort at time, and no solution is provided to cope with real-time changes in the user's behavior.

A mix between direct and indirect control, sometimes referred to as "dynamic pricing with enabling technologies", offers interesting properties. It consists of in-house systems controlling the appliances by reacting to a broadcast signal, often price-like. It is similar in a way to the legacy system of the french utility mentioned above but provides more precision as the signal is not reduced to a on/off order. The work of Beal et al. [Beal12] is a representative example. Each appliance is equipped with a smart plug which will switch on or off randomly. The probability of switching is determined by two factors : a broadcast incentive depending on the needs of the grid, and the user preferences regarding the priority of the appliance in question. The broadcast signal

allows a considerable scalability as it does not depend on the number of connected devices, however the stochastic nature of the system means that its reliability depends on the number of connected devices, possibly preventing a progressive deployment which would likely be required as we explained in Section 2.3.1.2. Indeed, the most recent paper on the model states that "the controller scales well on a range of at least 10 000 to 10 000 000 devices" [Papalexopoulos13]. The user preferences are represented by four "colors" indicating for each smart plug when the device can be controlled by the system : "Anytime", "Peak power" (for peak shaving adjustments), "Emergency only" and finally "Never". Moreover, a one-hour override button allows the user to manually disable the system at any time. This approach provides a simple way for the user to set its preferences, however the randomness of the switching can be inconvenient and the interface does not allow precise settings such as preferred time intervals, optimal indoor temperature, etc... This approach is similar to the one used in the hierarchical system by Karfopoulos et al. [Karfopoulos15] where the user's input is limited to a priority list indicating the level of disruption that is tolerated for each appliance.

Few works offer a fully decentralized model. Among them, Mohsenian-Rad et al. [Mohsenian-Rad10] use a smart pricing system on which agents adjust their consumption, taking into account precise user preferences. Each agent tries to maximize its benefits in a game-theoretic setting, relying on the external price and direct exchanges with other agents to evaluate its options. This system gathers many interesting properties such as robustness and user consideration, although the scalability is not mentioned in the paper. However it lacks the control and precision needed in our settings, as the price signal is only an incentive. The reliance on inter-agent communication can also reduce the robustness of the system.

4.3 Proposed approach

To build a Demand Response system satisfying the scalability and the reliability constraints stated in section 4.1 while solving the issues relative to the user's comfort noted in section 4.2.2, we chose to follow a bottom-up approach using a multi-agent system. Instead of a centralized model where each building consumption is modified following a load shedding request, our system allow a global load shedding capacity to emerge from the individual capacities of the connected buildings. To this end we built a dynamic coordination algorithm, described in section 4.4, which relies on two essential elements. The first is a building's energy management system called HEMIS, developed by Ubiant, which is able to compute the flexibility of the building at all time and act on its consumption while maintaining the user's comfort. The second is a gossip-based decentralized communication protocol named Push-Sum [Kempe03] allowing a fast and reliable propagation and aggregation of information between the participants. In this section we describe HEMIS and the Push-Sum algorithm before formalizing the precise problem our system needed to solve following this approach.

4.3.1 The Smart Building Solution: Hemis

Developed by Ubiant in collaboration with the LIRIS laboratory, Hemis is a smart building energy management system [Mansour12]. It provides an intuitive interface enabling the user to set its preferences regarding indoor temperature and lightning by defining an optimal set point but also a comfort margin for the system to make adjustments if the need to act on the consumption arises. For example, the target temperature for one room could be set at 22°C with a margin of +/- 2°C, allowing the system to reduce the energy consumption while staying within the tolerance margins of the user. This process differs from the usual energy saving processes which maintain an optimal comfort at all time, and is specifically designed to offer load shedding services to the grid. It allows the users to easily inform the system of both their needs and their willingness to participate in load shedding at the same time. Hemis also collects the data coming from various sensors around the connected house that monitor movements, door openings, light switches and air quality among other variables, and stores it locally, along with the real-time electricity consumption. This historical data is then combined with contextual information such as the time period and the weather to be processed by classification algorithms, in order to identify common patterns and try to predict the occupancy of the building. This forecast is then used in combination with the user's preferences to estimate the load shedding capacity of the building, that we call flexibility in the following sections. This flexibility is the estimated share of the predicted electricity consumption that could be reduced if needed. The accuracy of this information can vary widely depending on the situation, basically whether the users behave predictably or not. Indeed, it is important to notice that the users stay in control of their equipment at all time. For example if an appliance is switched on and contradicts the predictions, it will not be forced off again by the system. Instead the flexibility is simply updated with this new piece of information.

Hemis is designed with interoperability and genericity in mind, in order to accommodate for the various building configurations it is presented with. Based on a multi-agent system solving a multi-constraints optimisation problem, this solution is able to maintain a desired level of comfort (indoor temperature, lighting, air quality, etc... with respect to the user's habits and preferences) while automatically lowering the overall energy consumption. Its underlying principle is a virtualization of both abstract objectives and physical entities of the building. Connected appliances (heaters, lights, sensors, electrical rolling shutter, etc...) are represented in the system by reactive agents, each linked to an area of the building and to one or several environmental factors. For example, an electric heater will obviously act on the room's temperature, but also on the electricity consumption. A rolling shutter would be linked to the brightness of the room, but will also impact its temperature in addition to the power consumption. The user simply interacts with the system by setting objectives on each of the environmental factors, thanks to an intuitive interface. The agents then collaborate using shared marking spaces to find the best trade-off between their different and sometimes antagonistic

objectives.

In our approach, each building connected to the system is equipped with Hemis, and as such is able to communicate its flexibility for any time period. This information is an estimation, not a guarantee of shedding capacity, as the users maintain the full control of their devices at all time. However it is the necessary information for the agents to coordinate their action and adjust their consumption to reach the desired shedding capacity. As the only requirement for our system is the ability to predict a flexibility and to control the loads when needed, a system different from Hemis could be used as long as it fullfils them.

4.3.2 The Push-Sum algorithm

The Push-Sum algorithm, presented by Kempe et al. [Kempe03], allows a group of n nodes to collectively process the value of an aggregate information (sum, mean, product, maximum or minimum value) in a completely decentralized way using what is called an epidemic or gossip propagation model. A node then holds a value x and the objective is to know the mean value $\frac{1}{n}\sum x$ after a sufficient number of rounds. It is supposed that the nodes are able to communicate with each others, forming a connected graph. This algorithm works for any graph topology, as long as the network is a connected graph, meaning that there exists a path between each pair of nodes. Time is divided in rounds during which the nodes all follow the same process. The node uses two variables : a sum s initialized to x and a weight w initialized to 1. The algorithm followed by each node i is then as follows. At the beginning of each round, a set of J nodes *including* i are selected at random and sent a pair $(\alpha s, \alpha w)$ with $\alpha = 1/J$.
At the end of the round, the node updates the value of s and w by summing all pair $(\hat{s}_r, \hat{w}_r)$ it received during the round : $s = \sum \hat{s}_r$ and $w = \sum \hat{w}_r$.
The new estimate value of the mean is then $\frac{s}{w}$.

This algorithm allows the computation of various aggregates other than a mean. In order to compute a sum instead of a mean, only the weight w needs to be initialized differently. In this case, w will be set to 0 except for one agent whose w will be set to 1. Computing minima or maxima can be done by keeping the minimum or the maximum of the received values and propagating it with $\alpha = 1$.

Kempe et al. show that, "given a γ and a δ, the relative error in the approximation of the real value *agg* is at most γ with probability at least $1 - \delta$ in at most $O(\log n + \log \frac{1}{\gamma} + \log \frac{1}{\delta})$ rounds" [Kempe03]. However the convergence speed can vary depending on the topology of the network and the potential connectivity constraints. To illustrate the performances of the protocol in a representative use case, we simulated a network of 2000 nodes connected via *websockets* to represent a large district of connected buildings using an internet connection. We initialized x for each node to a random value between 0 and 3000 (the order of magnitude of a residential instantaneous consumption in Watts), and let the node communicate with only one other ran-

dom node at each round. This limitation constitutes a worst case scenario on connectivity constraints, considering today's networks capacities. We measured that in this case the maximum number of rounds needed until each node converges to the exact mean $\hat{x}$ is 25. Using our implementation in Java and standard hardware, with one thread per node dispatched on multiple computers with two different IP networks, these 25 rounds represent approximately one second. It shows that the logarithmic complexity of the Push-Sum protocol, combined with the speed of internet communications and the relative small number of nodes needed for our use case in practice, allows for a comfortable scalability.

In addition to a fast convergence, the nature of this algorithm also makes it possible to add and remove any number of nodes at any time. This *plug-and-play* feature provides a great robustness and facilitates an incremental deployment of a system using on this protocol. It is also relevant to notice that this algorithm allows for a straightforward implementation of monitoring features, for example by using a decentralized sum to know the size of the population or to count the occurrence of a property among the nodes, by setting $x = 1$. All the information and shared variables can be accessed by communicating with any of the nodes, thus eliminating the need for a parallel and potentially centralized monitoring system. This versatility, adding to the aforementioned scalability, modularity and robustness of the Push-Sum algorithm makes it a highly reliable and efficient protocol to use in our decentralized system.

4.3.3 Coordination problem

In the following of this chapter, each building taking part in the Demand Response program is considered as an agent a in the system. Each of these agents is equipped with the system described in section 4.3.1 and is able to know at all time t the total flexibility of the building $f_a^{tot}(t)$ that is the power (in Watts) that can be removed from the normal consumption of the building. This flexibility is an estimate and can change depending on the user's behavior. For a population of N agents the estimated total flexibility, or load shedding capacity, of the system is then

$$F^{tot}(t) = \sum_{a=0}^{N} f_a^{tot}(t) \tag{4.1}$$

As defined in section 4.1.1, a load shedding request (LSR) is a set $\{\hat{F}, t_d, t_f\}$ with $\hat{F}$ the expected power consumption reduction which is constant between t_d and t_f. Here we assume that the LSR is received by a subset of agents which transmit it to the whole population, either by broadcast or in an epidemic fashion, eliminating the need for any central coordination mechanism. It is also supposed that the the estimated total flexibility $F_{tot}(t)$ is higher than the target $\hat{F}$ by a margin decided by the operator, otherwise the request is declined. This margin mostly depend on the predicted reliability of the system which can be based on past load sheddings, but this question is out of the scope

of this thesis as is mainly a practical consideration depending on the situation and does not change the problem addressed here.

To reach the target capacity $\hat{F}$ during the whole event, each agent a must engage a share $\hat{f}_a(t)$ of its total flexibility $f_a^{tot}(t)$. As the system does not prevent the user from using their appliances as they wish, potentially reducing the available flexibility $f_a^{tot}(t)$ of the agent (see section 4.3.1), it is likely that the actual flexibility $f_a(t)$ of some agents will deviate from $\hat{f}_a(t)$. During the load shedding event the actual total flexibility $F(t)$ is then

$$F(t) = \sum_{a=0}^{N} f_a(t) \tag{4.2}$$

To ensure the load shedding stability the system must be able to dynamically react to these changes so that the agents adjust their respective flexibility $f_a(t)$ to maintain $F(t) = \hat{F}$ at all time t. In the next section, we propose an architecture that provides this stability, in a fair, robust and scalable fashion.

4.4 Decentralized Load Shedding Model

4.4.1 General principle

In order to solve the problem defined in section 4.3.3, we chose a decentralized approach based on a multi-agent system using an epidemic communication algorithm. An agent is a building able to communicate its flexibility thanks to the ambient intelligence solution described in 4.3.1, which **maintains the user's comfort at all time**. Thanks to the aggregate computation algorithm described in section 4.3.2, each agent also knows at all time the value of various variables concerning the whole population which allows it to take decisions locally without the need for a centralized or hierarchical control, thus enabling a **scalable and robust** system. The aim of this decentralized system is then to determine for each load shedding event the amount of participation of each agent. For this purpose each agent locally evaluates itself on its ability to fullfil its commitment and assigns itself a mark, representing its level of performance, following a process described in section 4.4.2 When the agent receives a LSR, it will engage a certain amount of its total flexibility depending of course on the capacity requested but also on its mark and its total flexibility in order to **fairly spread the requested adjustment**. The engagement process is described in section 4.4.3. During the load shedding, each agent knows at all time a decentralized estimate of the total capacity engaged by the population and **reactively** adjust its own engagement if it does not match the objective, as explained in section 4.4.4, in order to ensure the **reliability** of the system. Following the event, the agent will update its mark according to its performance.

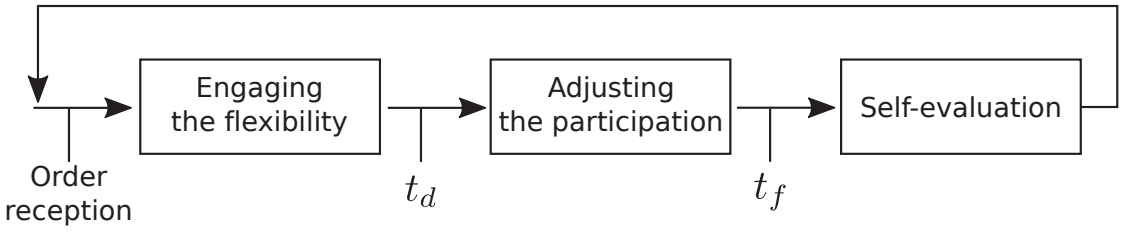

Figure 4.1: Main phases of an agent process.

4.4.2 Agent self-evaluation

The mark g_a of the agent a assesses its reliability, meaning its ability to execute the shedding it committed to. Indeed, an agent unable to correctly estimate its flexibility negatively impacts the whole system and is the sign that the user is reluctant to participate in load shedding: its participation to future events must then be reduced. To this end, the agent computes during the whole event the difference between the engaged flexibility $\hat{f}_a(t)$ (see section 4.4.3) and the actual shedding $f_a(t)$ it achieved. To sanction even more the big deviations compared to the small differences, the agent measures the quadratic error $(\hat{f}_a(t) - f_a(t))^2$.

To obtain an error relative to the shedding capacity of the agent, the sum of the differences during the event is weighted by the inverse of the squared initial flexibility. That way, if the error represents a significant fraction of the initial flexibility, the agent is penalized more. Thus an error $e_a > 0$ computed as follows:

$$e_a = \sum_{t=t_d}^{t_f} \frac{(\hat{f}_a(t) - f_a(t))^2}{\hat{f}_a(t)^2} \tag{4.3}$$

If $e_a = 0$ the agent did not deviate from its initial pledge. On the contrary, $e_a = 1$ means that it was not able to leverage any flexibility at all. If the agent did not participate to the last event, it assigns itself an average error $e_a = 0.5$ to avoid being influenced by an event it did not took part. In order to evaluate the agent on its global reliability and not only on the last event, an exponential moving average $\bar{e}_a$ is used following the formula $\bar{e}_a = \tau e_a + (1 - \tau)\bar{e}_a$, τ being the degree of weighting decrease (or smoothing constant). Such a moving average provides a good approximation of the real average while eliminating the need to store a complete history of past values.

Eventually, the final mark g_a of the agent is obtained by normalizing the error by the minimal and maximal average errors of the population:

$$g_a = \frac{\bar{e}_a - \bar{e}_{max}}{\bar{e}_{min} - \bar{e}_{max}} \in [0, 1] \tag{4.4}$$

Using this formula, the most reliable agent will have $g_a = 1$ and the least reliable $g_a = 0$. $\bar{e}_{min}$ and $\bar{e}_{max}$ are obtained thanks to the epidemic communication algorithm described in section 4.3.2.

4.4.3 Engaging the agent's flexibility

When receiving a LSR the agent computes its engagement considering its current total flexibility $f_a^{tot}(t)$ and its reliability mark g_a. The principle is straightforward, with the engaged flexibility $\hat{f}_a(t)$ being proportional to the reliability. However two important elements are added to the equation.
First, an agent must be able to take part to the event, even with a reliability mark of 0. Otherwise, it would not be able to improve its mark and so would never be able to participate again. In order to prevent agents with such low reliability from taking a large share of the capacity, the minimum engagement f_a^{base} depends on how the importance of the agent's flexibility $f_a^{tot}(t)$ compares to the minimum and maximum flexibility among the population, respectively f_{min} and f_{max} (known via decentralized aggregation) :

$$f_a^{base} = f_{min} * \frac{f_{max} - f_a^{tot}(t)}{f_{max} - f_{min}} \tag{4.5}$$

Second, in order to preserve some capacity to compensate for other agents potential failure (section 4.4.4), the agents do not engage their full flexibility regardless of their mark and retain a fraction $m \in [0, 1]$ of their flexibility.
The agent's engagement is then given by the following equation.

$$\hat{f}_a(t) = (f_a^{tot}(t)g^a + f_a^{base}(t)(1 - g_a))(1 - m) \tag{4.6}$$

As this engagement aims at reaching the requested capacity $\hat{F}$, it is bounded by the difference between $\hat{F}$ and $F(t)$ the total flexibility of the population (equation 4.2):

$$\hat{f}_a(t) = \min(\hat{f}_a(t), (\hat{F} - F(t))) \tag{4.7}$$

Furthermore, as the agents all get the LSR almost at the same time (depending on the speed at which the information is propagated), it is possible that less reliable agents engage themselves first, leaving only a small fraction of the requested capacity for the most reliable to participate in. To prevent this, the agents take advantage of the delay before the beginning of the event to spread their reaction in time. If t_r is the moment the LSR is received by the agent and $\delta \in [0, 1[$ a security margin before the beginning of the event t_d, each agent chooses the time t_e when it will compute and communicate its engagement regarding its reliability as follows :

$$t_e = t_r + (t_d - t_r) * (1 - g^a) * \delta \tag{4.8}$$

This way, the moment the agent will engage depends linearly on his reliability : the most reliable agent reacts immediately when receiving the LSR while the less reliable waits until the last moment.

4.4.4 Adjusting agent's participation

As soon as a LSR is received and until the end of the event, agents must adjust their participation when necessary to keep the total load shedding $F(t)$ stable at the objective $\hat{F}$. To this end, we designed a mechanism that allows each agent to adjust its share on a collective effort to quickly reach a global objective by summing their contributions. The first step is to detect a gap between the estimated sum of the contributions $F(t)$ and the global objective $\hat{F}$. It may seem like an easy task, but here individual contributions are constantly changing, adding to the small variations created by the decentralized computation of the aggregates. To avoid an oscillatory behaviour when $F(t) \simeq \hat{F}$, we used two complementary mechanisms that are often found in control systems.

The first one is the use of a ratio $0 \leq H \leq 1$ implementing an hysteresis around the target $\hat{F}$, damping small variations at low frequency due to the continuous variations of the agent's flexibility (see figure 4.2). A threshold crossing is then detected if $\left|\frac{F(t)}{\hat{F}} - 1\right| \geq H$. This ratio must be defined considering the stability requirements of the utility and the order of magnitude of the capacity measured.

The second mechanism allowing a reliable detection is a kind of low-pass filter erasing quick variations. It is implemented with a timer t_{wait}. To be taken into account, a crossing must be observed for at least t_{wait} time steps. This delay has to be set so that the new value of $F(t)$ is known by the whole population before any new adjustment, thus avoiding any unwanted oscillations of $F(t)$. This could be done in a dynamic way by the agents themselves but we let this issue out of the scope of this work, as it can easily be tuned empirically. These two systems are combined to lead to an efficient threshold crossing detection by each agent despite the frequent variations due to the convergence of the collectively computed aggregate (see figure 4.2).
Once a crossing is detected the agent adjusts its engagement $\hat{f}_a(t)$ using equation 4.9 which spreads the required adjustment over the population, proportionally to the current participation $\hat{f}_a(t)$ of each agent.

$$\hat{f}_a(t+1) = \hat{f}_a(t) + \hat{f}_a(t)\frac{\hat{F} - F(t)}{F(t)} \tag{4.9}$$

As $\hat{f}_a(t) \leq f_a^{tot}(t)$, it is possible that the population becomes unable to reach the target if the total available flexibility drops too much. Situation will not prevent the system

from working normally but could be damageable for the grid in practice depending on the case, which is why we stated in section 4.3.3 that a LSR should not be accepted if the initial estimated total flexibility is not significantly higher than the target.

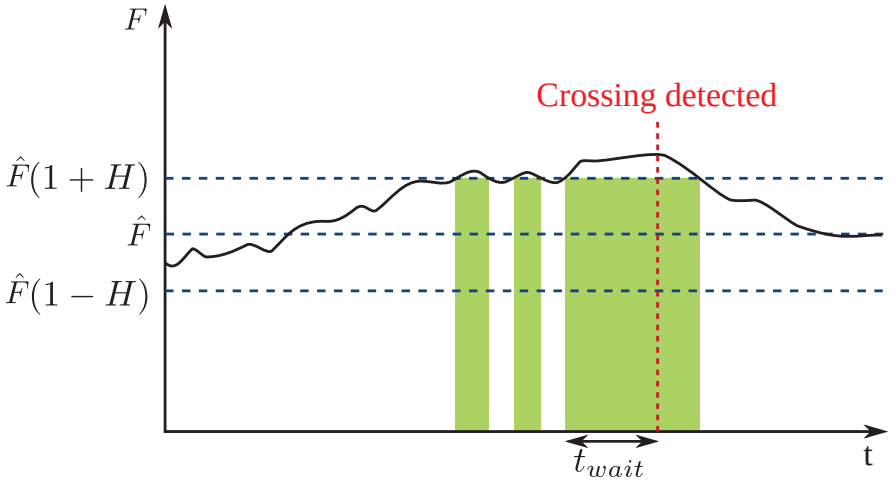

Figure 4.2: Target crossing detection process. Here a crossing is confirmed when the aggregate value is above $\hat{F} * (1+H)$ (green areas) for a duration of t_{wait}. The agent then starts adjusting its share to the effort.

4.5 Experiments

4.5.1 Test protocol

To evaluate our model we simulated n agents, each with a total flexibility $f_a^{tot}(t)$ (in Watts) randomly chosen in the set $\{100, 200, 400, 500, 2000, 3000\}$ corresponding to a range of flexibility going from switching off or dimming lightning to switching off or delaying space heating. The total flexibility engaged by an agent at any instant t is available at once. However, it can vary afterward depending on the user's actions (turning appliances on or off for example) or on external conditions (weather, price of electricity, etc...). To simulate a realistic user behavior, we acted on two factors : the probability of a failure to maintain the engaged flexibility, and the moment this failure indeed happens during the load shedding event. Each agent a is then attributed a failure probability p_a following an exponential distribution $p_a = 0.9e^{-0.3x_a}$ with x_a being the index of the agent a in the population, in order to obtain a small number of agents with a high failure probability and vice-versa. During the load shedding event, if an agent fails, its flexibility drops to zero, simulating a refusal from the user to take part in the load shedding. In a real situation, the flexibility would not systematically drop to zero, but this worst case scenario has the advantage of showing both the reliability and

the robustness of our system, on top of facilitating the reading of the results by making the drops more obvious. Indeed, it not only demonstrates the behavior of our system when the total flexibility drops abruptly, it also simulates the sudden deconnection of a number of agents and shows the reaction of our system in such case.

We wanted to verify on one hand that our system was able to resist to a significant drop in the total flexibility (as long as it stays above the requested capacity $\hat{F}$, see section 4.4.4), and on the other hand that the stability of the load shedding is maintained even when the system has to adjust very frequently (the drop in total flexibility is spread out in time). To this end, we combined two different types of failure which randomly determine the precise time when an agent could fail depending on its total flexibility.
The first type is meant to simulate the case of a household with a small load shedding capacity, which could represent a simple dimming of the lights. A reduction of luminosity being easily noticeable by the user, they are more likely to either react quickly or not react at all. For agents with a small flexibility, i.e $f_a^{tot}(t) \leq 1000$, the moment t_{fail} is then chosen as follows : $t_{fail} = t_d + (t_f - t_d) * (1 - \sqrt[3]{\alpha})$ with $\alpha \in [0,1]$ drawn randomly, t_d and t_f being the starting and ending time of the load shedding (see section 4.1.1). The agent is in this case very likely to fail right at the beginning of the event, and the probability diminishes with time, as illustrated by figure 4.3.
If the total flexibility if bigger ($f_a^{tot}(t) > 1000$ in our tests), $t_{fail} = t_d + (t_f - t_d) * \sqrt[3]{\alpha}$ which corresponds to the contrary of the previous failure type. As time passes, the agents is more and more likely to fail (see figure 4.3). This scenario simulates the cases of buildings with a higher flexibility which would correspond to turning off heating equipment, or delaying home appliance activation. In those cases, the effect of the shedding is not noticed immediately by the users, who are then more likely to react as time passes.

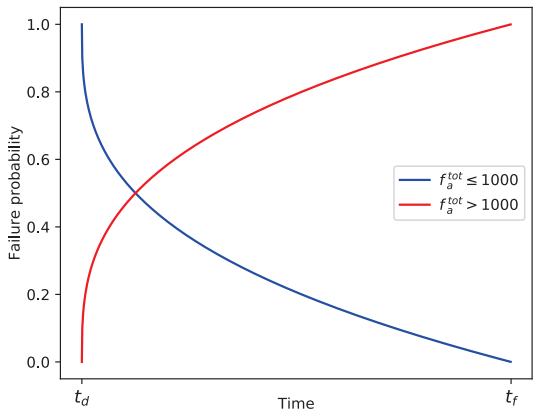

Figure 4.3: Probability distribution of failure with time depending on the total flexibility.

The visualize the results, time is indexed using the iterations of the Push-Sum algorithm used by the agents to aggregate the shared variables (see section 4.3.2). In order to reduce the simulation time, both the length of the load shedding events and the interval separating them last only 200 iterations, which is extremely short (less than a few seconds in real settings) compared to the duration of a real-life load shedding event (several minutes or several hours, with a preparation time of at least half an hour). The hysteresis is fixed at $H = 0.01$ and the delay t_{wait} at 10 iterations. See section 4.4.4 for details on how these parameters should be set.

4.5.2 Results

4.5.2.1 Reliability, robustness and scalability

In order to test the efficiency of the multi-agent system in maintaining the curtailment stability, we analysed the performances of our model without the self evaluation mechanism and without the security margin, which will both be evaluated in the following tests. When an agent receives a LSR it engages its full flexibility, considering its mark as being the highest possible ($g_a = 1$) as if it was fully reliable regardless of its actual performances. This way, the less reliable agents are not filtered out after the first event, maintaining the pressure on the adjustment mechanism. It also helps observing the effect of only the adjustment mechanism on the quality of the shedding. To increase the readability of the results and to reduce the simulation time, the size of the population is limited to $n = 20$ agents in this first setup. We show in the next experiment that the size of the population does not affect the performances of our system.
Every other parts of the model are working as described in section 4.4.

Figure 4.4 shows the 20 agents performing consecutive load shedding events. As the evaluation mechanism is disabled, each event is independant in regards to the performance of the system, which is why we chose to show only two to allow a more detailed view of the process.

The continuous green line represents the total engaged flexibility $F(t)$. At the beginning of the simulation, there is none. When the LSR is received (t = 30), the engagement phase starts. We observe a peak in engaged flexibility that reaches the total available flexibility F_{tot} (dashed line) as every agent engages its total flexibility at once (the self-evaluation is deactivated), then adjusts back to the requested capacity $\hat{F} = 10000$ W after a time corresponding to the delay t_{wait}, confirming the correct behavior of the agents during the coordination phase. The actual load shedding begins at t = 230, indicated by the dash-dotted line of the target $\hat{F}$ rising to the desired value (10000). Following the scenarios described in section 4.5.1 the total F_{tot} drops gradually during the event, with the failure of part of the agents. As expected a corresponding drop can be seen in the engaged flexibility $F(t)$, but it is quickly compensated by other agents after a delay of at least t_{wait}. This delay is the shortest time allowed to detect a variation, as described

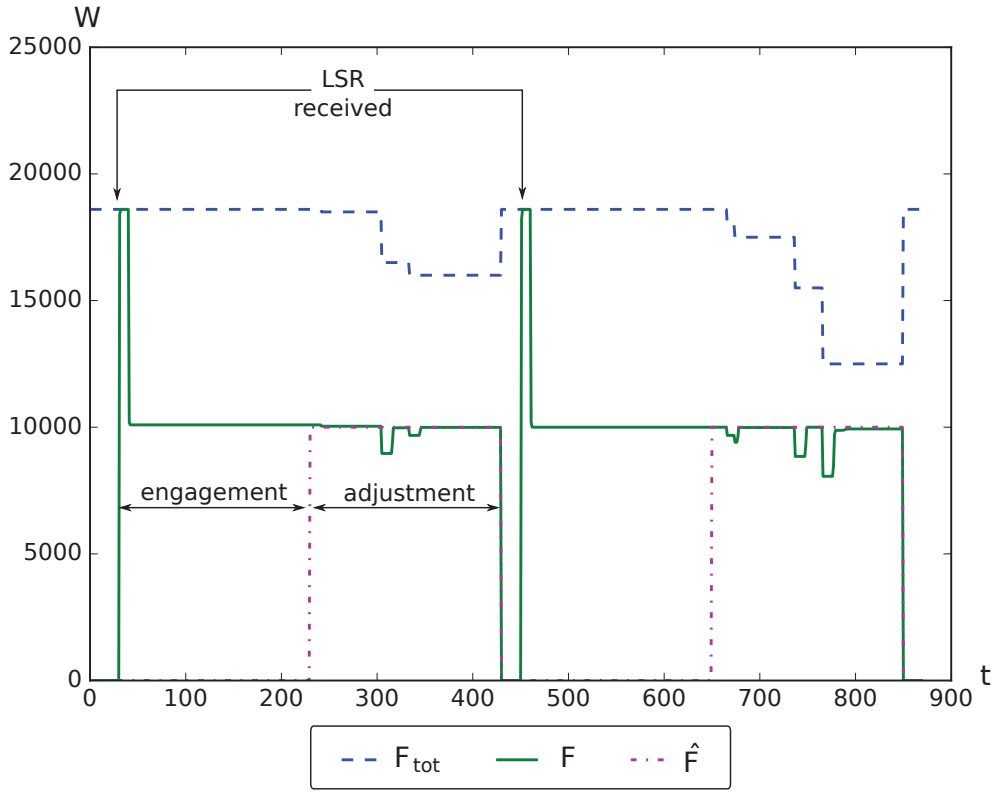

Figure 4.4: Dynamic adjustment during two consecutive events without the self-evaluation mechanism.

in section 4.4.4, as it dampens the potential swings due to the epidemic propagation of information. This dynamic recovery from local failure without any over-compensation, which can be observed each time the engaged flexibility deviates from the objective, confirms the efficiency of our adjustment mechanism to provide a reliable load shedding. As we mentioned in section 4.5.1, this also shows the robustness of our system in case of deconnexion from a number of agents, as it corresponds to the complete flexibility loss experienced here.

Scalability

To validate that our system is able to handle large populations, we ran simulations with different population sizes using otherwise the protocol described in section 4.5.1. The requested capacity $\hat{F}$ is set to 40% of the total flexibility. Figure 4.5 shows that the behavior of the system is similar to the above experiment with a smaller population.

In order to assess the reliability of our system independently from the population size, we measured for each event the Mean Square Error normalized (nMSE) by the

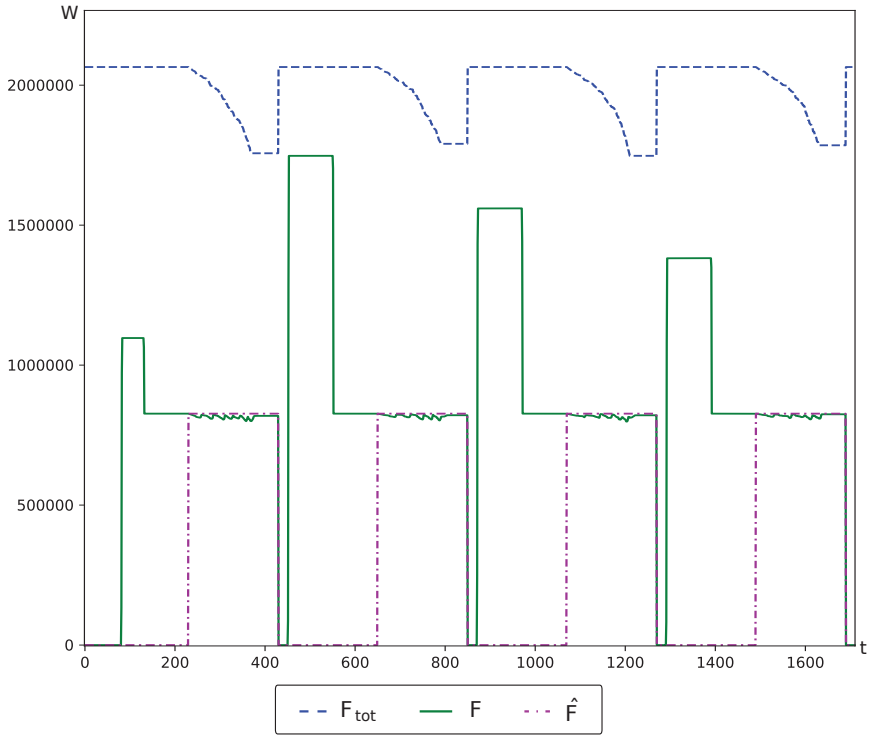

Figure 4.5: Dynamic adjustment by 2000 agents on 4 consecutive events.

requested capacity.

$$nMSE = \frac{\sqrt{\sum_{t=t_d}^{t_f}(F(t) - \hat{F})^2}}{(t_f - t_d).\hat{F}}$$

It allows us to compare the performances of the system when the scale changes. Table 4.1 shows the mean nRMSE computed over several repetitions for each population size, as well as the standard deviation. We see that the error does not feature any clear trend correlated with an increase in population size and neither does the standard deviation, proving the scalability of the model.

Pop. size	Nb. runs	nMSE	Std. dev.
10	100	0.2	0.01
200	60	0.24	0.03
2000	40	0.22	0.03
200000	20	0.26	0.05

Table 4.1: Average nRMSE on the population size (section 4.1).

4.5.2.2 Self-evaluation

The role of the self-evaluation mechanism, that we disabled in our previous tests, is to improve the reliability of the load shedding event after event by prioritizing the most reliable agents during the engagement phase.

To demonstrate the effect of this mechanism we simulated 14 consecutive events using the protocol described in section 4.5.1. For each event, we calculate the mean square error as

$$MSE = \frac{\sum_{t=t_d}^{t_f}(F(t) - \hat{F})^2}{t_f - t_d}$$

Figure 4.6 shows the average MSE over the 50 runs, along with the standard deviation, in two situations. First, the blue line and area correspond to the results obtained without the evaluation mechanism. As expected the error shows noticeable variations due to the random failing probabilities of the agents, but no apparent trend as the least reliable agents continue to participate in the same proportion from one event to the other. Second, the black line and gray area show the results obtained with the full evaluation mechanism. The curve follows a clear downward trend, proving the ability of the system to reduce the error over time. We notice that the error reaches a plateau from the 11th event onward, as the evaluation mechanism cannot prevent even the best agents from failing.

The self-evaluation mechanism acts on two separate parts of the model, namely the agent's engaged flexibility and reaction time. Figure 4.7 shows the detail of the first six consecutive load shedding events of a single run in three different situations which will be described in the next paragraphs, allowing us to better illustrate the behavior of the system. This time we measured at each iteration the error $|F(t) - \hat{F}|$, i.e the absolute difference between the requested curtailment and the actual load shedding. To allow for a relevant comparison of these situations, the random generator seed is the same.

The first situation, in Figure 4.7a, serves as a reference point by showing the performance of the system without the evaluation mechanism, as in section 4.5.2.1. The agents *immediately* engage their *entire* flexibility for each load shedding event. On the graph we distinguish the six event by the blue areas. Without the evaluation mechanism, the magnitude of the peaks does not show any particular trend, as expected. Indeed, the total flexibility of the agents is the same at each event and is totally engaged each time. The variations in error is then only due to random changes in information propagation through the gossip algorithm.

In the second case, we tested the first part of the evaluation mechanism, meaning that the agents this time engage a *share* of their total flexibility considering their mark, which is updated after each event (see section 4.4.2). The agents still engage their flexibility *immediately* as they receive the LSR. Figure 4.7b shows the errors in this case. We clearly see that a majority of peaks has significantly reduced in amplitude. We notice

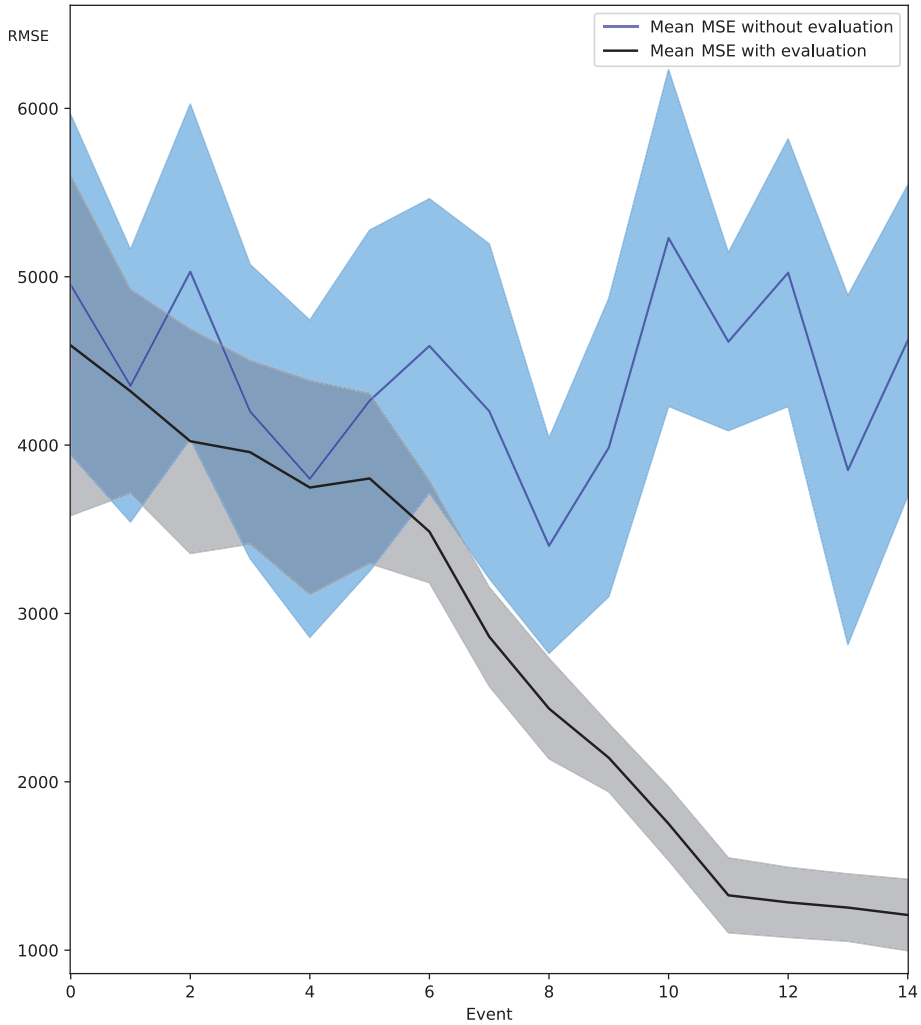

Figure 4.6: Evolution of the MSE for 14 consecutive events with and without the full evaluation mechanism, averaged over 50 runs. Colored areas show the respective standard deviation.

that some peaks are still present as they correspond to agents failing for the first time, which cannot be avoided by the evaluation mechanism. This is logically the case for the first event as no agent has had the opportunity to evaluate itself yet. As soon as the second event however, two average peaks in the first case are considerably reduced in the second, meaning that agents that failed in the first event had a far less negative impact on the second. The effect is clearly noticeable in the following events. This proves the relevance of our evaluation mechanism and its capacity to improve the quality of the load shedding through time.

In a third case, which performances are plotted on figure 4.7c, the evaluation mechanism is in full effect. The agents consider their mark to moderate not only the flexibility they engage but also the speed at which they react to the LSR, letting higher ranking agents weight in first but most importantly allowing the information to spread correctly during the engagement phase. For the same reason as in the second case, the first peak and some of the highest peaks stay identical in these simulations. However every other peaks has almost disappeared compared to the two previous cases, meaning that delaying the introduction of the agents regarding their mark further improved the reliability of the whole system.

4.5.2.3 Security margin

In section 4.4.3 we explain how the agents retain a security margin m during the engagement phase. This constitute a reserve to compensate the potential failure of less reliable agents during the load shedding. To demonstrate the interest of this mechanism, we ran the same simulation as in section 4.5.2.2 with 6 consecutive events and the evaluation mechanism fully activated, but this time we represented a worst case scenario where the most reliable agents ($g_a > 0.99$) systematically fail during the last three events. Results are shown in figure 4.8. If $m = 0$ (figure 4.8a), the agents do not keep a margin at all, they engage their entire flexibility considering their current mark. The global effort is then spread amongst the most reliable agents only. When they fail, the recovery is difficult and slow as a number of agents are not engaged yet and take time to react. If $m = 0.5$ (figure 4.8b), the agents only engage half of their flexibility at each event, spreading the effort on a larger share of the population as the requested capacity is not reached by the most reliable agents immediately. When these agents fail, the recovery is significantly better, as showed by shorter peaks in error and significantly fewer steps.

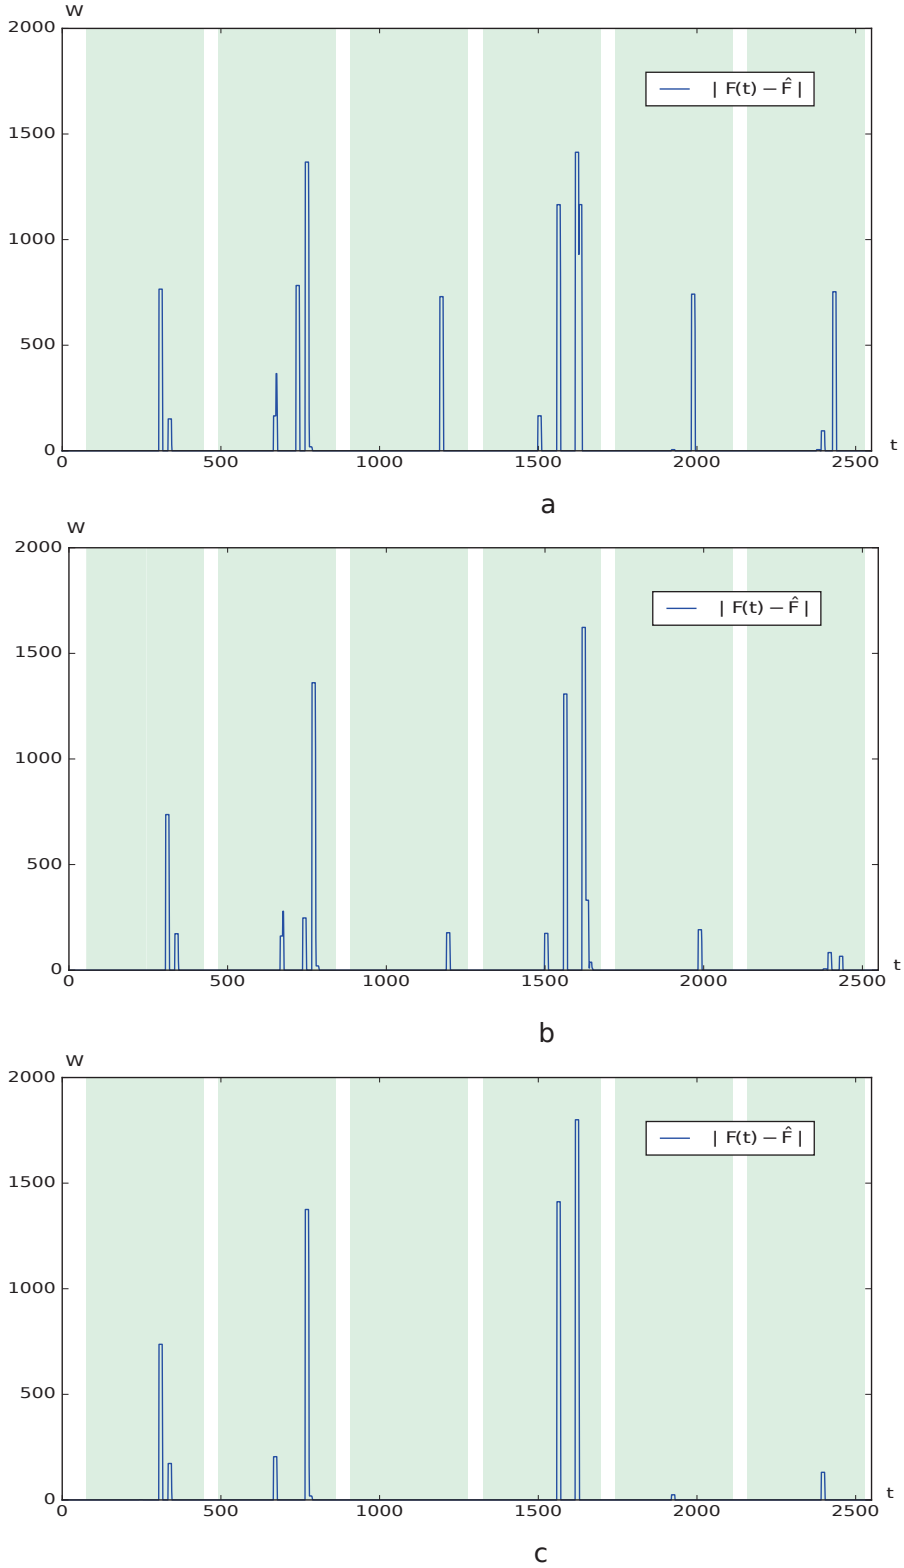

Figure 4.7: Error evolution during six consecutive events. (a) the agents engage their full flexibility at once. (b) the agents take their mark into account and engage immediately (c) the agents consider their mark for both engagement and reaction time.

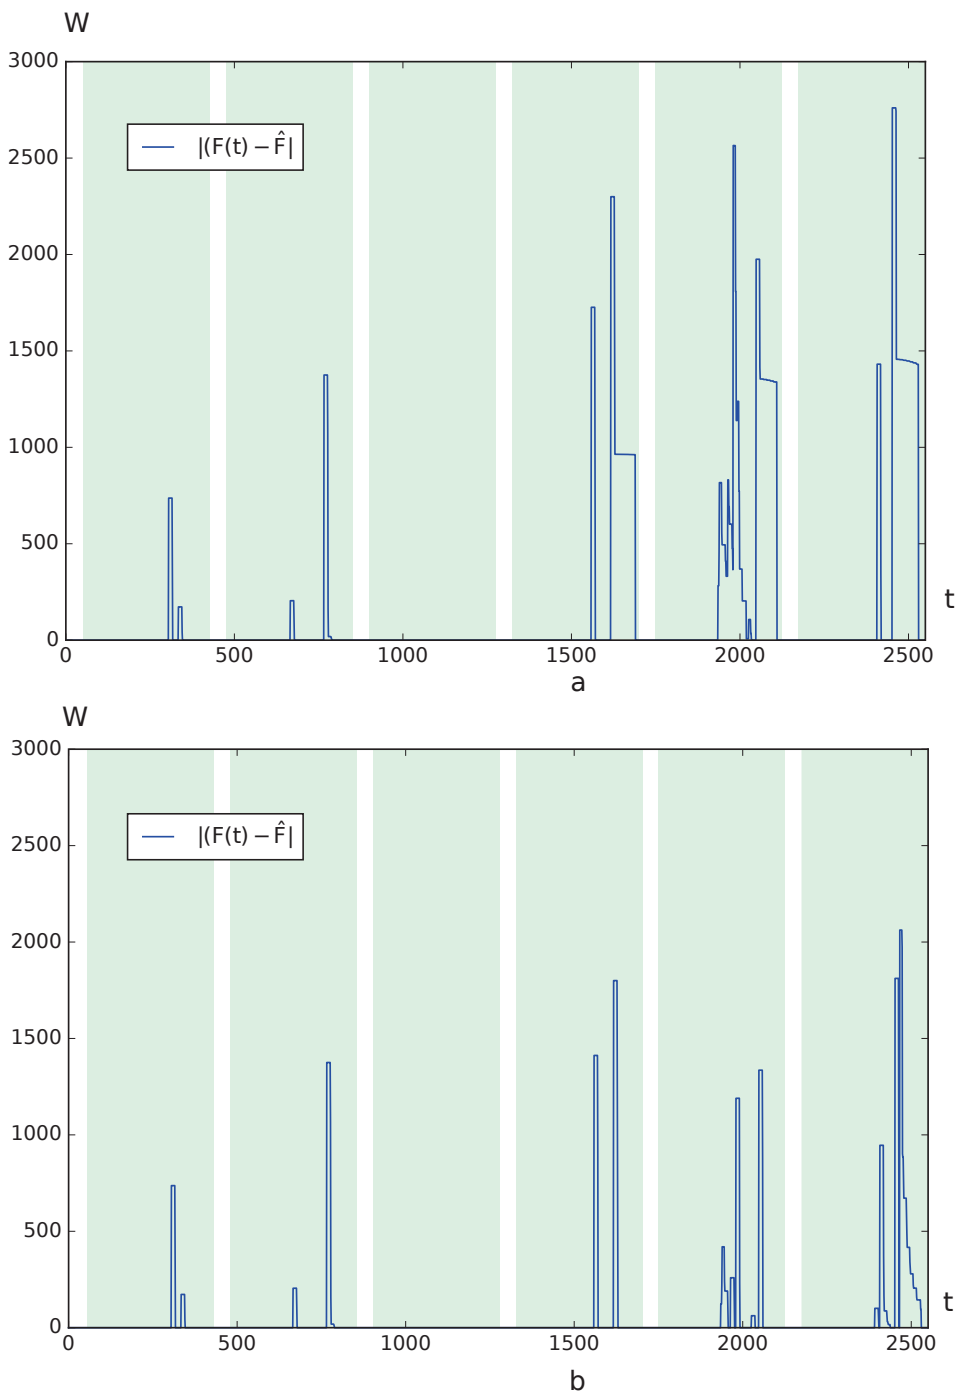

Figure 4.8: Error evolution during six consecutive events. (a) the agents engage their all the flexibility allowed by their mark. (b) the agents retain a margin to help recovering from future failure.

4.6 Summary

Distributed load shedding is prone to quickly develop in the upcoming years. As a straightforward solution to the biggest current problems of the energy sector, it offers an ecological way to keep the network running. The biggest obstacle faced by existing systems and ongoing work is to be able to deploy such a system while not affecting the user's comfort with a model light enough to be reliable and cost effective on a large scale. Yet, with a simple centralized approach, taking the user's comfort into account implies a constant and heavy two-way communication between the server and the nodes, limiting the system's scalability and flexibility. In this chapter we showed that a fully decentralized architecture, able to manage a large number of consumers simultaneously following a bottom-up approach, is a possible solution to these reliability and scalability issues. By delegating the comfort management to the building itself and by letting the agents independently decide how to take part in a load shedding effort we obtain a highly resilient system which can scale up easily.

We presented a coordination mechanism associated with a self-evaluation process for the agents. It uses an existing gossip-like algorithm allowing each agent to know the state of the system at all time without the need for a central coordinator nor a high connectivity, while being powerful enough to allow real-life use even when the number of connected building is high. Of all the shared variables, the most important one is the total engaged flexibility of the whole population. It allows each agent to adjust dynamically its own share of the effort to maintain the stability of the load shedding. For this adjustment to be efficient, a hysteresis mechanism combined with a low-pass filter is used to ignore small variations due to the decentralized communication and also to avoid oscillatory behaviors.

On top of this communication layer, the self-evaluation process enables the most reliable agents to take the most part of the load shedding effort. To this end, each agent measures its performance during the event. If its actual consumption reduction does not match the one it committed to, its mark drops. In subsequent events, it will engage a smaller part of its available flexibility in order to lower its impact on the stability of the shedding. This potential inability to fulfill its engagement comes directly from the decisive choice to not interfere with the user at any time, which means that they are free to turn on or off their devices even during a load shedding event. The dynamic adjustment mechanism handles these unforeseen variations in real-time, while the self evaluation process improves the whole system's resilience and stability in the long term.

To validate our approach, we measured the performances of our system by realistically simulating the agents failures. We varied both to failure probability of the agent and the moment the failure happens to put the system under various types of stress. We first confirmed the efficiency of the dynamic adjustment mechanism and showed that each variation for the target capacity is compensated as soon as it is detected. We

then proved the relevance of the self-evaluation process by showing the evolution of the load shedding stability on consecutive events. As self-evaluation reduces the impact of the less reliable agents on the shedding quality, the error decreases with time as the number of failures diminishes. Comparatively, the number of failures without the self-evaluation remains constant.

We showed here that a decentralized multi-agent system could bring the scalability and the reactivity needed to compensate in real time the unpredictability of user behaviors for a reliable load shedding service. However, this approach still relies heavily on the ability of the building's energy management system to predict its flexibility. In the next chapter, we present a model for decentralized load management in a microgrid, based on the bottom-up approach described here, where we provide a per-appliance control enabling a finer dynamic consumption scheduling.

5

Generic Adaptive Energy Management for the Smart Grid

▷ In this chapter we present a generic platform providing a coordination algorithm for the dynamic scheduling of connected assets in a smart-grid. The decentralized architecture using a gossip-based communication protocol provides both scalability and robustness, while the modularity of the model allows the quick integration of various constraints and equipment. It is designed with reactivity and adaptability in mind to handle high variability in production and demand, leveraging the flexibility of residential buildings without imparing the users comfort. ◁

Chapter content

5.1 Problem statement **83**
5.1.1 Energy use optimization 83
5.1.2 User satisfaction 85
5.1.3 Process simplification 85
5.2 State of the art **86**
5.2.1 User acceptance 86
5.2.2 Scalability, modularity and robustness 87
5.2.3 Market-based models 89
5.3 Proposed approach and agent model **89**
5.3.1 General process 90
5.3.2 Local constraints satisfaction 91
5.3.3 Energy optimization 92
5.3.4 Effort distribution 93
5.3.5 Communication framework 94
5.4 Implementation **95**
5.4.1 Inertia-based devices 96
5.4.2 Task-based devices 96
5.4.3 Generators 97
5.4.4 Energy Storage Systems 97
5.4.5 Non controllable assets 98
5.5 Formal properties of the model **98**
5.5.1 Nash equilibrium 98
5.5.2 Pareto efficiency 99
5.6 Simulations **99**
5.6.1 Reliability and robustness 100
5.6.2 Scalability 100
5.6.3 Fairness 103
5.7 Summary **104**

5.1 Problem statement

The load shedding system presented in chapter 4 was a first step toward our goal of developing a smart grid management system. We showed that is was possible to provide residential demand response while maintaining the user's comfort thanks to an adaptive and dynamic decentralized architecture. Load shedding is a very powerful solution to the issues encountered by current power grids regarding the increased pressure on the grid stability caused by intermittent power sources. However, as the share of renewable increases in the production panel, the traditional way the energy is managed is challenged as a whole as new solutions emerge (see section 2.1).

In its broadest sense, a smart grid can be defined as a group of electrical assets which production or consumption can be controlled and/or predicted to some extent, by an energy management system. Energy management in a smart grid can then be defined as a dynamic multi-objective optimization problem where [1] the sum of their energy output must be as close as possible to a given objective to ensure the stability of the grid (section 5.1.1) and [2] the local operational constraints must be satisfied (notably considering the user's comfort, see section 5.1.2).

Our objective is to design an Energy Management System (EMS) to control the different assets of any given smart grid to solve the aforementioned multi-objective optimization problem in real-time with as few limitations as possible regarding the use cases. It means that our system must *handle the considerable diversity* of both possible situations and devices, ranging from a Virtual Power Plant composed of a mix of fuel cells and solar panel to a wind turbine-powered microgrid comprising a few homes with connected appliances. To facilitate a quick and cheap deployment of the solution on new configurations, it must also be *highly modular and generic*. As we stated in section 1.2, our work revolves around three main axes which are the optimization of energy use, the respect of user comfort and the ease of deployment. In this section we will define for each of these axes the problems we want to solve with our approach.

5.1.1 Energy use optimization

Energy management in smart grids covers a wide range of domains from power electronics to machine learning. Part of the research on energy management focuses on real-time generation regulation, sometimes referred to as the primary control level [Olivares14]. This primary control aims at maintaining the grid quality of service regarding frequency and voltage on the very short term, typically under a minute. The response time required by such systems is very short as reactivity is crucial to the grid operations and the solutions come mostly from electrical engineering [Guerrero13], thus we will not focus on this level of control here. Relying on this primary layer, a secondary control layer operates on the medium to long term and is responsible for the reliable operations of the grid regarding more abstract objectives, the most im-

portant being the balance between consumption and production. As we detailed in section 2.1.3.2, three general types of energy management have emerged over the last decade.

- Demand Side Management (DSM) eases the pressure on the supply side by offering intelligent load control, for example by spreading water heating device activation to minimize consumption peaks [Ranade10].
- Virtual Power Plants (VPP) make scattered, intermittent generation compatible with traditional power grid by virtually aggregating into a single, more controllable entity [Pudjianto07].
- Microgrids (MG) are sub-networks that consist in local energy production and consumption, behaving as a single entity for the main grid and facilitating the integration of intermittent energy sources, mixing the concepts of VPP and DSM [Lasseter01].

As a growing amount of research is being done on smart grids, a number of systems and approaches are proposed for each of these use cases and their variants. However, we argue that the underlying optimization problem is the same regardless of the approach, as we explain next.

At any given time, an electrical device will consume or produce electricity. For simplicity this energy output, which can be negative (if the device is a consumer), positive (if it produces electricity) or null, is often called *prosumption* and the device itself a *prosumer* (see section 2.2.1.1). As a smart grid is a group of such assets, it becomes a prosumer itself. The first goal of an EMS is then to control the energy output, or prosumption, of the smart grid, which is the sum of the individual prosumptions of its components. This target energy output can be defined as a vector $\boldsymbol{O}$ which corresponds to the desired prosumption for each of the next H discrete time periods. The values of this vector, the horizon H and the duration Δp of each time periods, but also the way they are communicated to the system, all vary depending on the use case.
For example, in the case of Demand Response (DR), the target output is usually given by the utility (see Chapter 4) which will request temporary changes in the expected consumption to compensate variations in production. A load shedding (reduction in consumption) corresponds to an increased prosumption for a defined period of time.
For a Virtual Power Plant (VPP), the desired prosumption depends among other things on the current price for energy sale. In order to sell energy when the prices are high, the VPP will want to maximise its prosumption. On the contrary, it could try to store energy or ramp down its production to minimize the prosumption when selling prices are low.
In the case of an islanded Microgrid, high demand might not be met by the local production which can not rely on the main grid to fill in. Also, excess of production can not be absorbed by the few loads. To avoid damaging equipments or causing a blackout,

the energy output has to be null at all time.
Finally, a connected microgrid could behave as a DR provider by adapting its prosumption to the needs of the grid, as a VPP by following the selling price of electricity, or it could choose to maximize the use of local production and behave as an islanded microgrid.

The primary objective of an EMS is then to schedule the grid's operations so that *the actual power output matches the target power output at all time*, or more precisely to minimize the distance between the two vectors.

5.1.2 User satisfaction

As we already stated previously (see section 2.2), our thesis is particularly focused on smart grid comprising connected households. When dealing with residential load management, the end users are fundamental actors of the system. Space and water heating, which directly impact the user's comfort, account for nearly two thirds of home energy use in the US[1] for example, making their control an essential part of a smart grid system. The efficiency and large scale deployment of such systems thus depends on the user's acceptance.

In chapter 4 we explained the necessity of residential user's acceptance and gave the rules an EMS must follow to increase its chances to obtain it :

- *Energy management should not prevent or force any action.* Users have to know that they always have the option of controlling their appliances while allowing EMS to take control as long as their preferences are respected.
- *Equity must be guaranteed among grid participants.* The global cost of grid balancing must be shared equally between participants, considering individual preferences regarding the willingness to participate.

It is relevant to note that these considerations also apply to industrial or commercial users, although their constraints are generally less subject to frequent changes. We also purposely chose to not include any form of economic consideration, regarding for example a financial incentive to participate in the grid balancing as a user, as these questions are out of the scope of this thesis.

5.1.3 Process simplification

A smart energy management system must be designed to accommodate a large number of different use cases. Even when considering microgrid settings only, the amount of

[1] source : www.eia.gov/consumption/residential/data/2015

possible combinations of connected appliances in smart-homes is considerable and be subject to change. Yet the cost of deploying such a system at a large scale is directly related to the amount of expert knowledge needed to install and maintain it. The real-life scalability of an EMS then depends on the *straightforwardness* of its working, which is related to the amount of work any installation or modification requires. The individual components of a smart grid such as energy sources, storage equipment or household appliances all have their own constraints, working and often proprietary protocols. As we aim to integrate as many different kinds of assets and not design each and every part of a smart grid, the core of the system should be both highly *generic*, enabling the configuration and implementation of any use cases and devices. It should also be *scalable* and enable the incremental addition of many entities into the grid without impairing its performances. Finally, it is essential that a modification to one part of the system, like the constraints of an asset or the topology of the network for example, does not necessitate an intervention on any other parts. Such modularity would prevent costly maintenance and significantly improve the robustness of the system.

5.2 State of the art

Instead of focusing on a particular use case, we looked at the general picture and evaluated the existing approaches on the key points developed in section 5.1 : user acceptance, scalability, modularity, and robustness.

5.2.1 User acceptance

A significant share of the research on energy management systems integrating residential loads relies on strong assumptions regarding load control and user acceptance. Levron et al. [Levron13] for example do not take load into account and consider it constant over time, which is unrealistic. Colson et al. [Colson10] work with a time-varying load but do not describe the origin of the information. Similarly, [Umeozor16] assumes a perfectly predicted load which is also unrealistic as we saw in Chapter 3, and do not present a way to manage it, only mentioning that it could be controlled.

When the load is considered as part of the optimization process, it is often simply split between critical (that can not be powered off) and non-critical (that can be powered off at will) over which the user has no control [Colson11a, Parisio11, Logenthiran12]. This approach is too simplistic as the non-critical loads like dishwashers or HVAC[2] equipment, while indeed being less sensible to changes in schedules, still significantly impact the user's comfort depending on how they are managed.

Among the few works that clearly integrate user feedback or preferences into the process, Pournaras et al. [Pournaras14] use answers from one survey to adjust the flex-

[2]Heating, Ventilation and Air Conditioning

ibility of agents and evaluate the resulting discomfort on a second survey. This method is not practical enough as the surveys do not allow the users to precisely define how they want their appliances to be managed, they only give two vague options regarding whether the user is favorable to a higher amount of control or not. Beal et al. [Beal12] tested a system where each connected device was given a load shedding priority chosen by the user on a scale of one to four. These approach eventually take a real user setting into account, however it does not fit the need for an intuitive interface between the user and the system or for a real-time adjustment capacity as we detailed in section 4.2.2. [Pipattanasomporn09] also mentions user-defined priorities but do not give many details regarding how they are implemented and use a critical/non-critical distinction in their simulations.

It is worth noting that price incentive-based DSM systems exist [Mohsenian-Rad10, Veit13, Fan10], providing a fair user-centered way to leverage residential loads. However, it does not provide the reactivity and controllability needed for a dynamic grid balancing because it relies exclusively on the users to act on the consumption, as we explained in section 4.2. In a variation of this approach, [Fan10] integrate the user's preferences in their optimization via a generic parameter allowing a variety of settings to be implemented.

The work presented by [Ramchurn11] is a very interesting approach to load management considering the user's preferences. For each kind of load they associate a *marginal comfort cost* to a given solution (deferring or reducing the load) which is to be minimized by the system. [De Oliveira11] also focuses on user's comfort by modeling different kinds of services and how to evaluate the satisfaction of the user. However they offer a static optimization algorithm which ability to handle real-time variations in user's behavior is not discussed. Indeed, while these approaches do provide promising ways to handle the user's preferences, they do not satisfy our need for an adaptive and scalable system, as we will see next.

5.2.2 Scalability, modularity and robustness

"Modern grids include significant numbers of prosumers, making the scaling of control algorithms a pending challenge" (Frey et al. [Frey13]). Indeed, the complexity of the optimization problem grows with the number of connected homes and appliances.

A significant part of the existing works relies on centralized models that gather all the parameters of the considered network into a single fitness function subject to various constraints. Model Predictive Control is often used in this case, and consists in formulating an optimal schedule for every asset by minimizing the global fitness function over a given time horizon. The first step of the schedule is then implemented and the resulting changes integrated in the fitness function, before the process is repeated. As representative examples, Parisio et al. [Parisio11] used this method to manage a microgrid and Ouammi et al. went beyond that to control a "cooperative network of

smart microgrids" [Ouammi15]. Power-flow equations can also be used to formulate the fitness function, as in [Levron13]. In cases were the number of solutions is limited it is possible to find the absolute best schedule, as in [Hernandez-Aramburo05] where all possible solutions are computed in advances to be accessed in real-time. In this case however only the production from gas engine, which is time-independent, is optimized. When considering energy storage systems and thermal loads, which are time-dependent, this kind of approach is unsuitable. To find the optimal solution at each time step, various algorithms can be used from gradient-based linear optimization to dynamic programming. These approach have three main drawbacks. In terms of **scalability**, these optimization methods reach their limits when facing high-dimensional state variables and multiple dynamic and conflicting objective to be optimized [Li16]. Their **modularity** and adaptability is also very limited as a modification of the grid (adding or removing an asset, updating a constraint) systematically requires changes in the fitness function. Most importantly, regarding **robustness**, if alternatives to standard optimization exist such as distributed computation using bio-inspired heuristics [Colson10], too often they are implemented in a centralized way, offering a single point of failure [Basir Khan16] which can only be mitigated partially by using redundancy for example.

As previously said, there are many different use cases labelled a smart grid, and almost as many proposed systems, each with their specific configurations and constraints. For example, some focus on fuel consumption minimization [Hernandez-Aramburo05], some on the integration of electric vehicles [López13], others specifically on storage system management [Xu11]. Some works [Frey13, Pipattanasomporn09] advocate the idea of a generic framework enabling the integration of any device and accommodating for different management policies. However the former only provides guidelines and a survey focused on generic architectures while the latter provides a role-based model, lacking robustness.

On these observations, a number of works propose a multi-agent paradigm using different approaches and a comprehensive survey can be found in [Kantamneni15]. Role-based models are frequent but often lack the fault-tolerance capacity as they may include an "optimizing agent" [Ramachandran13], a "schedule agent" [Zhao15], a "facilitator" [Basir Khan16] or even a "central coordinator agent" [Anvari-Moghaddam17], which would bring the drawbacks of centralized systems. Moreover, Rohbogner et al. [Rohbogner13] find that some approaches, notably market-oriented designs such as the one presented in [Linnenberg11] or [Mengelkamp17], lack the proactive and self-adaptive capabilities that are part of the commonly accepted definition of multi-agent systems. Fulfilling the requirement for decentralization, Colson et al.[Colson11a] offer a simple model for microgrid power management using three different agent types to demonstrate that decentralized multi-agent systems are "a viable alternative and deserve investigation, especially considering the daunting microgrid control problem." However, their approach lacks genericity and relies on direct communication between

the different types of agents which implies complex exchanges and can limit the scalability and modularity of the system.

5.2.3 Market-based models

A significant amount of work has been focused on market-based system for energy management in recent years, and a comprehensive survey can be found in [Abdella18]. If use cases vary, from large scale demand response to small microgrid settings, the underlying logic stay the same : the idea is that producers and consumers trade energy on a competitive market in order to reach an equilibrium where offers satisfy the demands. The concept has been used in various approaches, with recent examples in [Mengelkamp18, Zhang16, Tushar18]. In the literature we had access to, every proposed approach used a form of centralization [Jogunola17] either to simply share the energy bids or to directly determine the best match. Even if we set apart this structural issue, we chose not to follow a market-based approach not only because of its competitive nature which could enable greedy behaviors that do not fit with the explicit objectives of fairness and cooperation we wanted to achieve, but also because of the nature of the market model where the emergence of an equilibrium is reached indirectly through the market's laws of offer and demand and not via a transparent and easily explainable process that we deem necessary (see Section 2.3.3.1).

5.3 Proposed approach and agent model

In section 5.1 we listed three main objectives. Our system must be able to schedule the operations of the different assets of a smart grid so that its energy output matches the desired value at any time. As our focus is on integrating residential loads without impairing user comfort, this scheduling process must be highly reactive and adaptive to handle the variations in residential consumption. It must also ensure that the burden of load adjustments is spread among the users considering their willingness to participate. Finally, the architecture should allow an easy, modular and scalable deployment while staying robust and reliable.

To achieve these requirements, we present a framework that allows different energy optimization algorithms, and even different architectures, to coexist. We showed in chapter 4 that a decentralized multi-agent system can be used to build a reactive, scalable and robust coordination mechanism, so we exploited the same bases in a similar approach. The load shedding system was designed to be an efficient solution using an existing smart building management system (HEMIS) to offer a service fitting the current requirement of traditionnal utilities. Here we go beyond this limited use case and do not rely on an existing system to manage the devices in the building. Instead, in this model each device affecting the stability of the grid by consuming and/or produc-

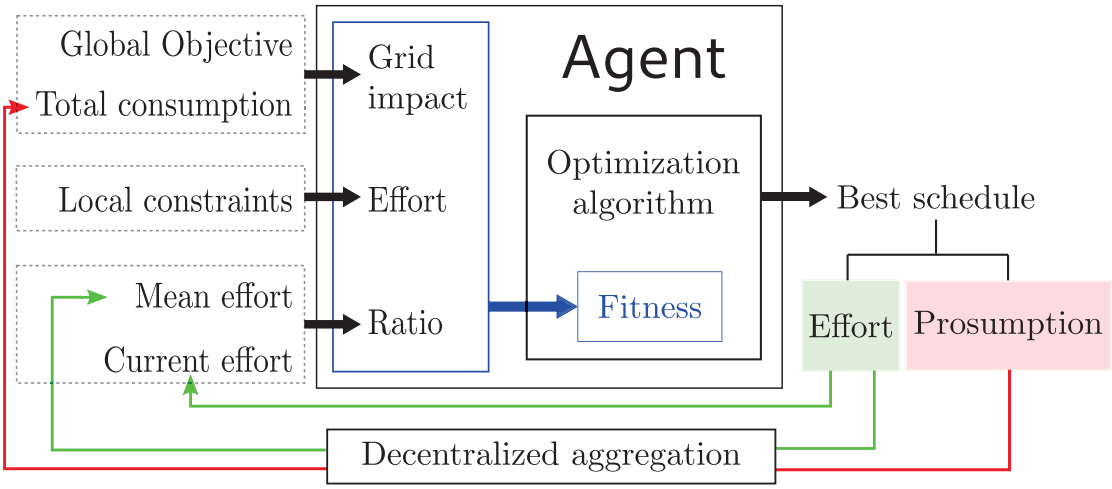

Figure 5.1: Overview of the agent process loop.

ing (prosuming) electricity is represented by an agent. It can be a home heating system, a solar panel, an electric vehicle, a washing machine, a fuel cell, etc...

The smart grid EMS is then constituted by a population of agents, pursuing the goal of achieving the dynamic multi-objective optimization described in section 5.1 : match the energy output to a given objective $\boldsymbol{O}$ over H discrete time steps while satisfying the local constraints. As we explained in 5.1.1, this target vector $\boldsymbol{O}$, its size H, the duration of the time steps Δp and the way these parameters are set all depend on the situation. As opposed to role-based models presented in section 5.2.2, each agent in our model follows the same general process presented in the next section.

5.3.1 General process

Each agent has the ability to optimize its future actions considering the energy optimization objective (section 5.3.3), and to estimate the amount of effort these actions require regarding its local constraints (section 5.3.2). Its prosumption forecast and estimated effort are shared with the other agents of the network so that the total estimated total consumption and the average effort are known by the whole population (section 5.3.5). Agents then repeat the process by adjusting their actions consequently, depending on how the global prosumption objective is met and how their effort compares to the average. figure 5.1 gives an overview of the local agent loop.

As shown in figure 5.2, the agents are synchronised with a global timeline which is divided into discrete periods p, during which they will iterate over a local feedback loop to schedule their operations over a certain horizon H, aiming to converge toward a satisfying solution, where the sum of their scheduled prosumption equals $\boldsymbol{O}$ at all time. If the agents are synchronised with regards to the global timeline using their internal

clock, their local timeline can differ from one agent to another, i.e the iterations i can have different durations and thus vary in number during one period p.

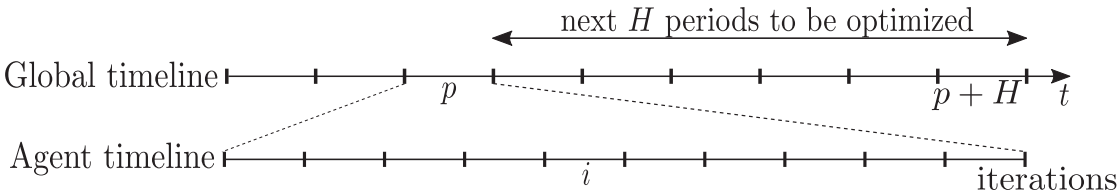

Figure 5.2: Time discretization. During each period, each agent iterates over its local process until the optimization converges.

More precisely, at each iteration i, the agent chooses an operational schedule s using an optimization algorithm that minimizes a fitness function $f(s_i)$, which is a weighted sum of the impact of the schedule on the grid output $f^c(s_i)$ (section 5.3.3) and the effort it represents $f^e(s_i)$ (section 5.3.2). These two antagonistic components are prioritized by a dynamic ratio $r(i)$ (section 5.3.4) in the following way :

$$f(s_i) = (1 - r(i)) \cdot f^e(s_i) + r(i) \cdot f^c(s_i) \tag{5.1}$$

An operational schedule s can take many forms depending on the actual device the agent represents. For a dishwasher, it could consist in a list of tasks with their starting time. A space heater could use a time series of temperature set points. This representation is specific to each agent and is not communicated to the others. When a schedule is selected, only its corresponding prosumption and effort are shared in a decentralised manner with the entire population of the smart grid (see section 5.3.5).

5.3.2 Local constraints satisfaction

The first component of a schedule's fitness $f(s_i)$ is the satisfaction of local constraints $f^e(s_i)$. In most state of the art approaches (see section 5.2.1), the local constraints only consider the operational financial cost to be minimized. As explained in section 5.1.2 however, the integration of residential loads and their constraints associated to user's satisfaction add new rules that can not be reduced to a financial cost only. Comfort is difficult to measure in an absolute way, and the needed real-time adaptiveness prevents any form of fixed scales to measure it, regardless of the use case. Moreover, we need to make sure the burden of compromising between the global objective and their local constraints is spread among the agents (see section 5.3.4), particularly when these local constraints reflect the impact of the grid balancing on the user's comfort. In this objective, the degree of satisfaction of the local constraints must be expressed in the same way for all agents.

In our model, we propose the notion of **effort** as a unified measure to evaluate an agent's operational schedule s_i regarding its own constraints. The effort $f^e(s_i)$ is a normalized value that allows the implementation of a wide range of constraints depending on the use case. It represents the quality of a given operational schedule s with $f^e(s_i) = 0$ corresponding to the best possible schedule at the current iteration i regarding the user's tolerance. For example, in the case of a heating system where the user can define optimal temperatures and tolerance margins, $f^e(s_i) = 0$ when the schedule s perfectly matches optimal temperatures and $f^e(s_i) = 1$ when it reaches the tolerance margin. The implementation could allow $f^e(s_i) > 1$ in cases where the balance of the grid is an absolute priority, in an islanded microgrid for example. This definition is flexible and can be adapted to match different objectives and policies. The only strong requirements are that $f^e(s_i) \geq 0$ and that its definition is consistent for all agents in the system. We provide the detailed implementations used for our simulations in section 5.4.

5.3.3 Energy optimization

The primary objective of the EMS is to optimize its energy use (see section 5.1.1). The second component of a schedule's fitness $f(s_i)$ (eq. 5.1) is then its impact $f^c(s_i)$ on the grid prosumption. This impact depends on two H-dimensional vectors:

- $\boldsymbol{C}(s_i)$ is the prosumption schedule corresponding to the operational schedule s for iteration i.
- $\boldsymbol{C}^{tot}(i-1)$ is the sum of the agents' last proposed schedules prosumption at the beginning of the iteration i. It represents the predicted total prosumption of the smart-grid, including the prosumption of this agent's previously selected schedule $\hat{s}_{i-1}$: $\boldsymbol{C}(\hat{s}_{i-1})$.

For the EMS to satisfy its primary objective, $\boldsymbol{C}^{tot}(i)$ must be as close to $\boldsymbol{O}$ as possible. Minimizing the impact of a schedule s on the total energy output of the grid will be obtained by minimizing

$$f^c(s_i) = \frac{||(\boldsymbol{C}^{tot}(i) - \boldsymbol{C}(\hat{s}_{i-1}) + \boldsymbol{C}(s_i) - \boldsymbol{O}||^2}{H} \cdot \frac{1}{f^c(s_i^{loc})} \tag{5.2}$$

which returns the squared mean of the current total consumption taking into account the consumption of the new schedule $\boldsymbol{C}(s_i)$. $f^c(s_i^{loc})$ is used to normalize $f^c(s_i)$. It corresponds to the impact of a schedule s^{loc} that would only satisfy local constraints, meaning that $f^e(s_i^{loc}) = 0$ (the evaluation of local constraints satisfaction is described in section 5.3.2). Indeed, the schedule that minimize $f^e(s, i)$ has always the worst prosumption schedule, otherwise the agent would consume or produce energy uselessly.

This means that $f^c(s,i) \in [0,1]$ where $f^c(s_i) = 0$ when s allows the total prosumption to perfectly match the objective of the grid and $f^c(s_i) = 1$ when it corresponds to the worst consumption profile for the grid.

5.3.4 Effort distribution

Third component of the fitness function (equation 5.1), the ratio $r(i) \in [0,1]$ determines the current priority of the agent. At the beginning of an optimization period, $r(i) = 0$ meaning that only the effort $f^e(s,i)$ is minimized. As the end of the optimization period comes closer, the ratio increases with the pressure to balance the grid $f^c(s,i)$ until a satisfying solution is found (i.e the total prosumption of the grid matches the objective). $r(i) = 1$ means that the selected schedule only tries to minimize the impact on the grid.

The ratio is computed as follows :

$$r(i) = \left(\frac{t_i - t_p}{\Delta t}\right)^{q(i)} \qquad q(i) = \frac{f^e(\hat{s}_{i-1})}{\overline{f^e(i)}} \tag{5.3}$$

The time stamp t_p marks the beginning of the current period p, t_i corresponds to the time at the beginning of iteration i, Δt is the duration of a period, and $\overline{f^e(i)}$ is the average effort consented by the community.

The interest of this ratio is double. As detailed in section 5.1.1, the grid balance is the primary objective of the system. This ratio ensures that while the local constraints like the user's comfort are satisfied at first, they are progressively outweighed by the pressure to maintain grid stability. At the end of the optimization period, it is the only criteria considered by the agents when choosing a schedule. This allows an adaptive and dynamic compromise to be made between the secondary and primary objectives defined in section 5.1.

This ratio also pushes towards an equal distribution between participants of the amount of effort they consent. As stated in 5.1.2, a fair sharing of the cost of grid balancing between the participants is an important step toward the acceptance of such systems by the users which is an objective of our model. As the notion of effort is relative, this fair sharing of effort corresponds to a minimization of the effort variance among the participants. The rate at which the ratio $r(i)$ increases thus depends on the effort currently consented by the agent compared to the mean effort of the participants of the smart grid. Indeed, thanks to the component $q(i)$ in equation (5.3) which can be refered to as a *fairness factor*, an agent making comparatively less effort than the average will more quickly increase its ratio. On the contrary, the ratio will increase at a slower rate for agents requiring more effort from the users. This enables the progression towards a fair repartition of the effort among the agents.

5.3.5 Communication framework

Once the agent selected its best schedule $\hat{s}_i$ with respect to the fitness function 5.1 at the end of the iteration i, its prosumption $\boldsymbol{C}(\hat{s}, i)$ is communicated to the population along with the associated effort $f^e(\hat{s}, i)$. The total current prosumption forecast $\boldsymbol{C}^{tot}(i)$ of the smart-grid and the mean effort $\overline{f^e(i)}$ are then updated to allow the other agents to react (see section 5.3.1).

As the only communication requirements are the aggregation of these two variables, this model allows us to use the *Push-Sum* protocol [Kempe03], as we did for our load shedding system presented in chapter 4. This gossip-like algorithm allows a large number of agents to compute aggregates — notably means and sums — in a fully decentralized manner while only requiring the synchronisation of the participants on *rounds*, which can be done by using the local clocks of the agents. It is inherently robust toward agent's failure as it allows the dynamic addition or removal of participant in real time. This "plug-and-play" capability of the system associated with its inherent scalability makes it particularly well fitted for large scale real-life deployments where configurations and use cases vary frequently. Furthermore, as detailed in section 4.3.2, its convergence speed allows the agents to obtain the correct value in a matter of seconds which is negligible, even compared to the demanding time scales considered in 5.3. In place of Push-Sum, one could use a central node gathering all the individual variables then broadcasting the aggregated results, or another kind of distributed aggregation protocol.

However the use of this algorithm satisfies the criteria of robustness and scalability, but also offers a fundamental layer of privacy and security. In section 2.2.2.3, we explained the need for privacy-preserving architectures, with [Wicker11] advocating for the absence of centralized data collection. Here not only does our model require no centralization or any sustained peer-to-peer communication, with the Push-Sum algorithm the agents never even share the real value of their variables directly, but rather an updated share of the aggregate (see section 4.3.2 for details). Preserving the privacy of users and securing the communications are not our main focus in this work, however we argue that these considerations should be part of the design of the system, thus our communication requirement limited to the strict minimum.

Another advantage brought by the simplicity of our framework is its modularity regarding the network topology. Our model enables a completely flat network, where each device can communicate with all the population without intermediaries. However this single-layered communication is not required, and the creation of subnetworks is entirely possible if the need arises. For example, in some cases a centralized aggregation could be used in parts of the network like inside a smart home, where existing solutions mostly consist in a box controlling the connected equipments. Implementing such a configuration is straightforward, by using an agent representing the home as a gateway, aggregating the prosumption and effort of the agents inside the home before sharing in

it a decentralized way with the rest of the grid.

5.4 Implementation

In this section we want to show how the different types of devices that are usually part of smart grids can be integrated by using the generic effort function to implement their respective constraints. The objective here is not to give an exhaustive list of all the possible devices or constraints, as this list is constantly growing and changing. Instead, the following implementations are rather representative and realistic examples of how the constraints of the most common smart grid components can be implemented. We aim to demonstrate the modularity of our model, and how it can be adapted to the different use cases by fitting a wide range of constraints.

In this work we focused on residential smart grids, so in this section we mainly detailed how household devices can be integrated. Generally, a smart grid is not composed of controllable assets only, so the EMS must also consider non controllable ones in order to take informed decisions. In the controllable category we usually find :

- Inertia-based devices : notably heating and cooling system (HVAC[3], fridges, etc.) they account for the biggest share of adjustable loads in residential buildings [Nguyen13, Wright07]. Consequently, smart thermostats, space heaters and fridges are already deployed in a growing number of households[4].
- Task-based devices : mainly including household appliances that draw a significant amount of power in a short period of time such as dishwashers. This category could also includes industrial equipments and is commonly referred to as "deferrable loads".
- Generators : this category includes all the power sources that can be programmed, ranging from small fuel cells or micro hydro generators to larger traditional power plants.
- Energy storage systems : they are essential assets for facilitate the integration of renewable sources.

In the non controllable category we put :

- Renewable energy sources like wind turbines and solar panels that depend directly on the weather.
- Global household consumption : a large part of a household's power consumption is far from being entirely controllable and can only be forecast to a certain extent (see chapter 3).

[3]Heating, Ventilation and Air-Conditioning
[4]Source : full report at www.iea.org/efficiency2017.

It is important to note that regardless of this classification the different assets are integrated in the same way using the single effort function described in section 5.3.2 which implementation varies, as we will see in the next sections.

5.4.1 Inertia-based devices

Common heating or cooling devices are thermostatic loads that are activated or not depending on the current temperature relative to a target. An operational schedule s for such a load can then be a time series giving the state of the device for each of the next steps. This state $q \in \mathbb{Z}$ would correspond to a given energy consumption (given by the device's specifications) and a known influence on the room temperature. This influence can be computed using various method, either by observation or by using thermal models, most often a combination of both [Lu10], but this is out of the scope of our work. Using this information, the temperature schedule $\boldsymbol{T}(s_i)$ corresponding to an operational schedule can be computed iteratively.

User preferences for such an agent would typically include a target schedule $\hat{\boldsymbol{T}}$ of the desired temperature for each step and a tolerance margin λ in degree. The comfort criteria is then defined as the mean square error (MSE) between the estimated and target temperature divided by the square of the margin:

$$f^e(s_i) = \frac{1}{H\lambda^2} \cdot ||\hat{\boldsymbol{T}} - \boldsymbol{T}(s,i)||^2$$

.

This way, as defined in section 5.3, $f^e(s,i) = 0$ when the target temperature is exactly matched and $f^e(s_i) \geq 1$ when the temperature is predicted to reach or exceed the tolerance margin of the user in average. To optimize the operational schedule in our simulations (see section 5.6), we chose to use a genetic algorithm. However, it is important to note that a different optimization algorithm could be used here depending on the performance needed.

5.4.2 Task-based devices

Task-based devices are typically household appliances : dishwashers, washing machines, dryers, or even rice cookers. They can achieve different tasks, each task k consisting of a duration $d(k)$ and a load profile $c(k)$, which are fixed and given by their specifications. At a time t_s the user would select a task (a washing cycle for example) and define the time t_u when the task must be finished (ex : 6pm). An operational schedule s for such a load is then a tuple (k, t_k) with k being the user-selected task and t_k the candidate starting time of the task.

We defined the effort consented by the user as the interval between the time $(t_k + d(k))$ the task is actually completed by the appliance and the desired completion time t_u

set by the user. The worst case scenario ($f^e(s_i) = 1$) being the task starting immediately ($t_k = t_s$), the effort associated to a given schedule is evaluated as follows :

$$f^e(s_i) = \left| \frac{t_u - (t_k + d(k))}{t_u - (t_s + d(k))} \right|$$

.

When the tasks have a fixed duration and consumption as supposed in this case, the number of possibilities for the optimization of the starting time t_k is quite limited (equal to the number of steps of the optimization horizon H), allowing to simply compute the fitness of each possibility to select the best solution. This is the method we use in our simulation (section 5.6).

5.4.3 Generators

The operation schedule of a generator such as a fuel cell or a diesel-powered unit would be an array containing its power output for each of the next M steps. Optimizing it can then be done by a genetic algorithm, as for inertia-based devices (see section 5.4.1). Eventually, the goal of the agent could be to minimize its running time to reduce the fuel expenses notably. The effort of a schedule could then be defined relatively to a worst case scenario, maybe considering a price forecast, with the best case scenario being not turning on at all (or when electricity price exceeds fuel cost). Depending on the use case, periods of time when the user would rather have the cell not activated (during the night for example) could also be added to the effort criteria.

5.4.4 Energy Storage Systems

Energy Storage Systems (ESS) are a key actor in the smart grid paradigm [Ibrahim08]. Indeed, when deferring or shedding the demand is not enough, storage is the only way to maintain the stability of a grid with intermittent energy production. Moreover, Plug-In Electric Vehicles (PEV) are predicted to become an important part of the residential energy consumption in the next years [Callaway11]. If the charging process is open to third-party control a plugged vehicle could become a proper storage asset or at least a controllable load, and managing it would require high-level constraints like the *user-specified charge completion time*.

Regardless of the method used, the goal of an energy storage system (besides providing grid balancing or economic optimization services) is to maximize its lifetime [Basir Khan16]. Thus in our simulation we defined the effort associated with a charge/discharge schedule as the amount of energy exchanged by it, relative to its maximal instantaneous power :

$$f^e(s_i) = \sum_{p}^{H} \left(\frac{|C(s_i)_p|}{max} \right) \cdot \frac{1}{H}$$

This maximal power *max* is the same for charging and discharging in our implementation, but one could easily consider different values or a time-dependant variable. Also the number of switches between charge and discharge caused by the schedule coud be taken into account [Levron13].

5.4.5 Non controllable assets

Non-controllable loads by definition do not need an optimization mechanism. They simply communicate their planned consumption (or production) to the system like any other agent, and *their effort is set to the average effort of the population*. This way, it does not affect the pressure on the controllable assets.

5.5 Formal properties of the model

Our objective in this section is to demonstrate the ability of our coordination mechanism to systematically reach an optimum garanteeing a stable grid while maximizing the satisfaction of local constraints. In our model, an equilibrium is reached when the sum of all prosumptions is equal to the objective, i.e $\boldsymbol{C}^{tot} = \boldsymbol{O}$. Here we want to prove that such equilibrium is a Nash equilibrium and is also pareto efficient, the actions of each agent are the proposed schedules and its utility being the effort that has to be minimized. To this end, we make the following hypothesis.

1. The time of convergence of the communication algorithm is negligible so that all shared values are considered as immediately available to all agents. As we explained in section 5.3.5, our implementation using the Push-Sum algorithm makes this assumption reasonnable.
2. We consider that only one agent changes its schedule at any time. In practice overlaps are rare and absorbed by subsequent changes.
3. A correct solution always exists, ensuring that an equilibrium can be achieved. It reflects the reality as the worst scenario could always be to turn of the loads to avoid the damages of a blackout.
4. For any agent, the effort value is different for two different schedules.

5.5.1 Nash equilibrium

Once the equilibrium reached, if any agent unilaterally changes its schedule, it will lead to an unbalanced grid which is not acceptable. Thus any equilibrium is systematically a Nash equilibrium.

5.5.2 Pareto efficiency

To prove the Pareto efficiency we will consider, ad absurdum, that the reached equilibrium E is not Pareto efficient which means that at least one agent A can reduce its effort by changing its schedule. This necessarily leads to an unbalanced grid by creating periods with excess production and periods with excess consumption, called slots in the following demonstration. Consequently, another agent F (or group of agents) will have to fill the slots FS freed by A, and one agent G (or group of agent) will have to get out the slots NS filled by A. In the following, we will considered that there is only one agent G and one agent F as the reasoning is exactly the same with multiple agents applying the argumentation iteratively on each agent of the group. We distinguish three possible cases:

- $F \neq G$ and NS is not the preferred solution of G. G has to move to other slots were its effort is smaller. Another unbalance is created, which has to be compensated, so that the whole reasoning has to be applied recursively.
- $F \neq G$ and NS is the preferred solution of G. In order to move, G has to decrease its effort so that this new equilibrium cannot dominate E in the sens of Pareto. Since in the algorithm, all the agents initially proposes their preferred solution, this case will happen at some point.
- $F = G$, meaning that the agents can compensate the imbalance caused by A. For the new equilibrium to Pareto dominates E, FS has to provide strictly smaller effort (hyp. 5). Considering that all agents always have updated information (hyp. 2), that only one agent acts at a time (hyp. 3) and that in the algorithm each agent proposes the schedule that fits the most its local constraints and grid pressure (equation 5.1), there is 2 possible sub-cases:
  - if FS were free at some point, F should have proposed this schedule and was forced to move (because this is not E). So that to get back these slot it has to move another agent and the same reasoning applies recursively.
  - if NS were free, the same applies to A.

5.6 Simulations

In order to validate our approach we simulated different use cases using the implementation presented in section 5.4. Our goal was to confirm the following key behaviors:

- reliability : the system converges to an optimal solution ($\boldsymbol{C}_{tot} = \boldsymbol{O}$) consistently
- robustness : agents can be seamlessly added during the process
- scalability : the system can handle a large number of agent

- fairness : the effort is fairly spread amongst the agents

5.6.1 Reliability and robustness

To show the ability of the multi-agent system to quickly reach a solution to optimize the energy use, we simulated an islanded residential microgrid. The only energy sources are a solar panel which production data comes from a test array deployed by Ubiant in the south of France, and a generator. The loads are 7 connected dishwashers, 5 connected washing machines and 3 electric heaters. The effort function of the dishwashers and washing machines is computed following the implementation proposed in section 5.4.2, and the heater's using the one detailed in section 5.4.1. The prosumption profiles of the solar panel, washing machine and dishwasher can be found in annexe A. The target energy output $\boldsymbol{O}$ is zero at all time (islanded microgrid), the optimization period were set to a duration Δp of 1 minute, and the horizon H to the next 24 hours (1440minutes). To demonstrate the robustness of the system, the 7 dishwashers were added to the system during the optimization period.

We run the same simulation for one optimization period 50 times, sampling the total energy output $\boldsymbol{C}^{tot}$ every 2ms (sampling steps on the figures). At each time we measured the RMSE between $\boldsymbol{C}^{tot}$ and the target energy output $\boldsymbol{O} = \boldsymbol{0}$ using the following formula :

$$RMSE = \sqrt{\sum_{p=1}^{H} \boldsymbol{C}_p^{tot} \cdot \frac{1}{H}}$$

Figure 5.3 represents the mean RMSE of the consecutive runs as well as its standard deviation. At the beginning of the optimization periods, the agents converge toward a first solution in less than a tenth of a second (50 sampling steps), demonstrating the efficiency and reliability the model. At the time marked by the vertical line, the 7 dishwashers are added. The curve shows that with the new constraints induced by the added agents, the system takes some time to converge again toward $RMSE = 0$ as expected. This shows the robustness of our adaptive mechanism as well as the consistency of the system's behavior in converging rapidly toward a satisfying solution.

5.6.2 Scalability

Here we want to demonstrate the ability of our system to find the best solution independently from the size of the population. We defined an isolated microgrid setting where an hypothetical external source (a solar farm for example) steadily provides 400kW for 16 hours during the day, from 2am to 6pm. As a non-controllable source, its effort simply follows the mean and its operational schedule is fixed (see section 5.4.5). We simulated 3200 identical task-based appliances with only one task consuming 1kW of electricity during 2 consecutive hours. Their effort function $f^e(s, i)$ is implemented as

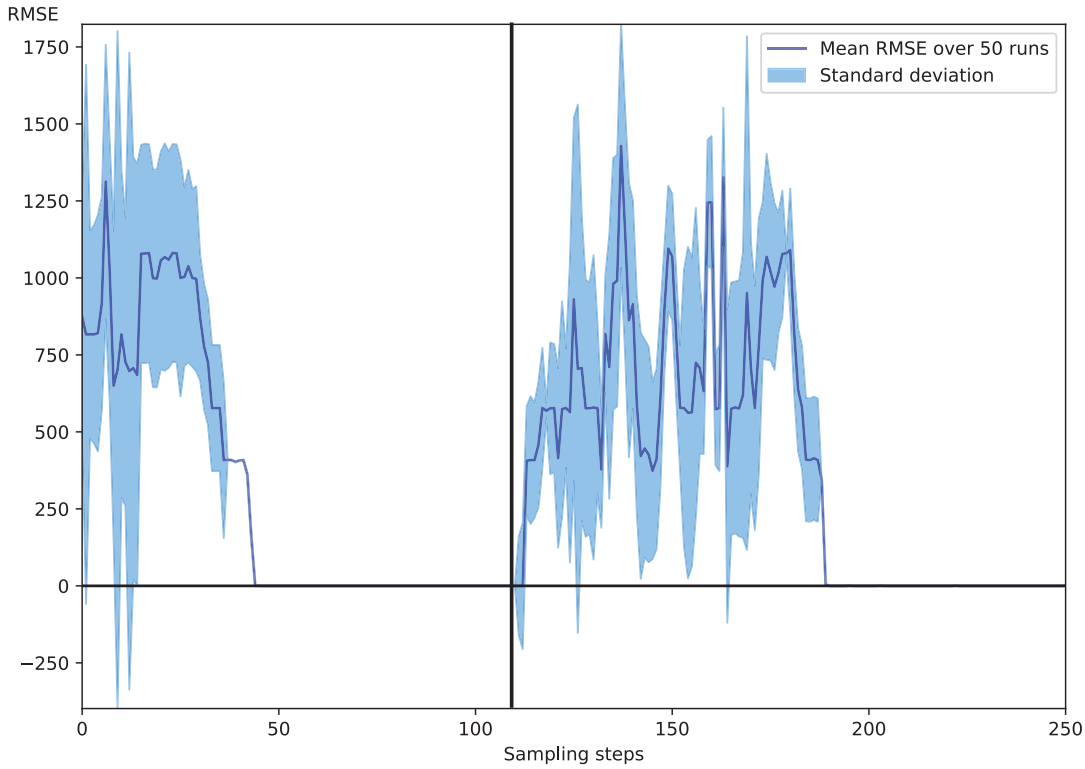

Figure 5.3: Mean RMSE over 50 consecutive runs with the settings described in section 5.6.1. Almost half of the agents are added during the optimization period (the vertical line shows the exact time).

described in section 5.4.2 and their target completion time t_u is 12am, meaning that they all have the same latitude as to when to start their task. The objective is identical to the previous experiment (section 5.6.1) : reaching $\boldsymbol{O} = \mathbf{0}$ for the next 24h, i.e $H = 24$, the only change being the duration of each period which is set to $\Delta p = 1h$. The sampling rate is also modified to one measure every 200ms.

The simulated situation represents an extreme case as the only way for the system to run all the tasks is to evenly spread them during the interval where energy is provided. As the tasks last 2 hours each, the agents could create 1 hour gaps that could not be used by another agent, preventing the population from converging to a satisfying solution. Figure 5.4 shows the mean and standard deviation of the evolution of the RMSE during a single optimisation period, repeated 60 times with varying random seeds. We see that the system never fails to converge toward a correct solution where the RMSE = 0.

We also ran the same experiment multiple times while varying the population size to evaluate its influence on the convergence time. The production of the power source was adapted each time to fit exactly the demand. Table 5.1 shows the results where

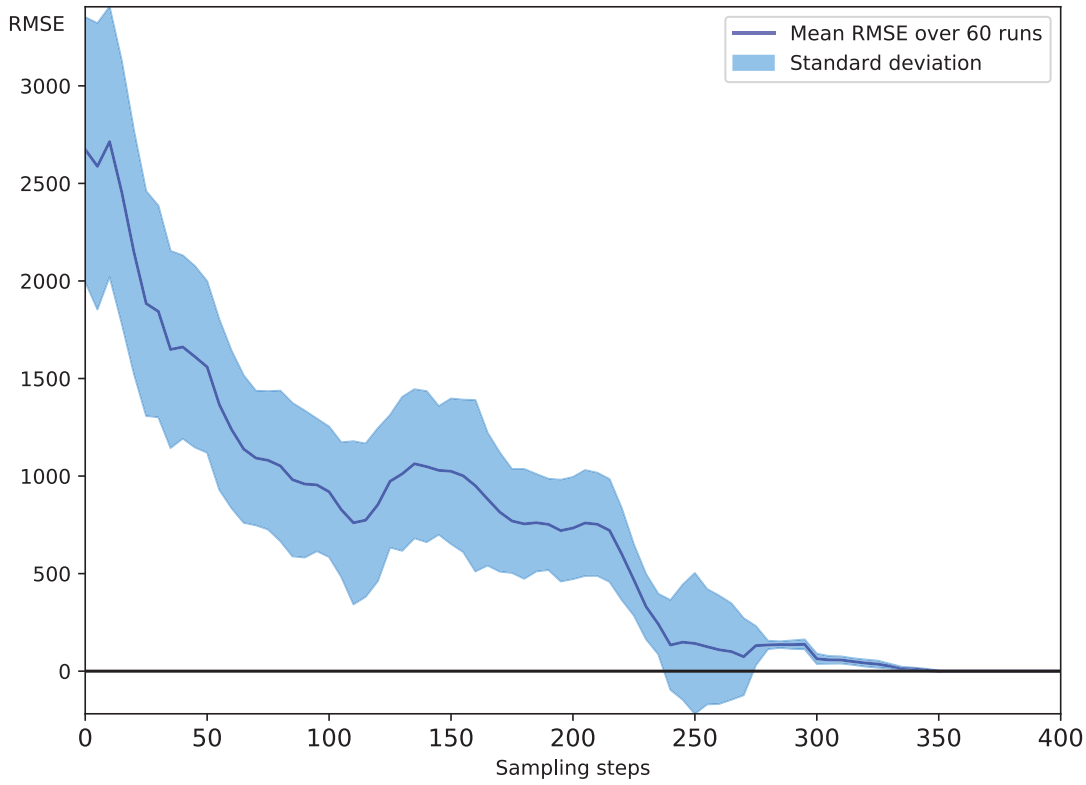

Figure 5.4: 3200 task-based appliances coordinating to fit the production of a single power source (section 5.6.2). RMSE measured every 200ms of one optimization period, averaged over 60 consecutive runs.

we notice that the convergence time does not increase linearly with the population but rather logarithmically, allowing the system to reach a satisfying solution in a few minutes even with a very large number of agents.

Pop. size	Nb. runs	Avg. conv. time	Std. dev.
10	1000	2s	0.2s
100	100	11s	1s
3200	60	70s	5s
10000	50	121s	13s
100000	40	155s	31s

Table 5.1: Average convergence time depending on the population size (section 5.6.2).

5.6.3 Fairness

We explained in section 5.3.1 how our system minimizes the effort of each agent while spreading it fairly amongst the agents thanks mostly to the ratio $r(i)$ which allows each agent to adjust its effort by comparing it to the mean effort of the population. Computed by equation 5.3, this ratio prioritizes either the local constraint satisfaction or the energy use optimisation depending on two factors : the time left before the end of the optimization period and the effort consented by the agent compared to the mean effort of the population. The objective of this last element, which we refered to as the *fairness factor* in section 5.3.4, is to push the agent to converge to a solution where the effort is fairly distributed by reducing the deviation from the mean.

To demonstrate this property, we used the same setup as scalability experiment described previously (section 5.6.2), with the following modifications. To the 3200 task-based devices we add 16 fully charged battery packs able to provide 400kWh each, and we consider that the solar panel array can be disconnected if useless. The battery packs' effort function is computed using the equation provided in section 5.4.4, i.e they try to minimize their own prosumption. If the appliances do not consider the mean effort of the population in their optimization, they would all schedule their operation at their prefered time, between 10pm and 12am. In such situation, their effort would be 0 but the battery's effort would be maximal. A better effort distribution would be to exploit the production from the solar panel, leading to a smaller deviation to the mean effort.

In this experiment we ran the same coordination period with and without the *fairness factor* ($q(i) = 1$) and measured the mean effort and the mean deviation from this mean among the population every 200ms. We then averaged the measures over 60 runs with variable random seeds to validate the observations statistically. The results displayed on figure 5.5 show that without the fairness factor, the mean effort stays almost constant, as does the deviation from this mean. As expected, with the *fairness factor* the mean effort is higher but the deviation is significantly smaller, as agents tend toward a fair solution where the effort is spread evenly among them.

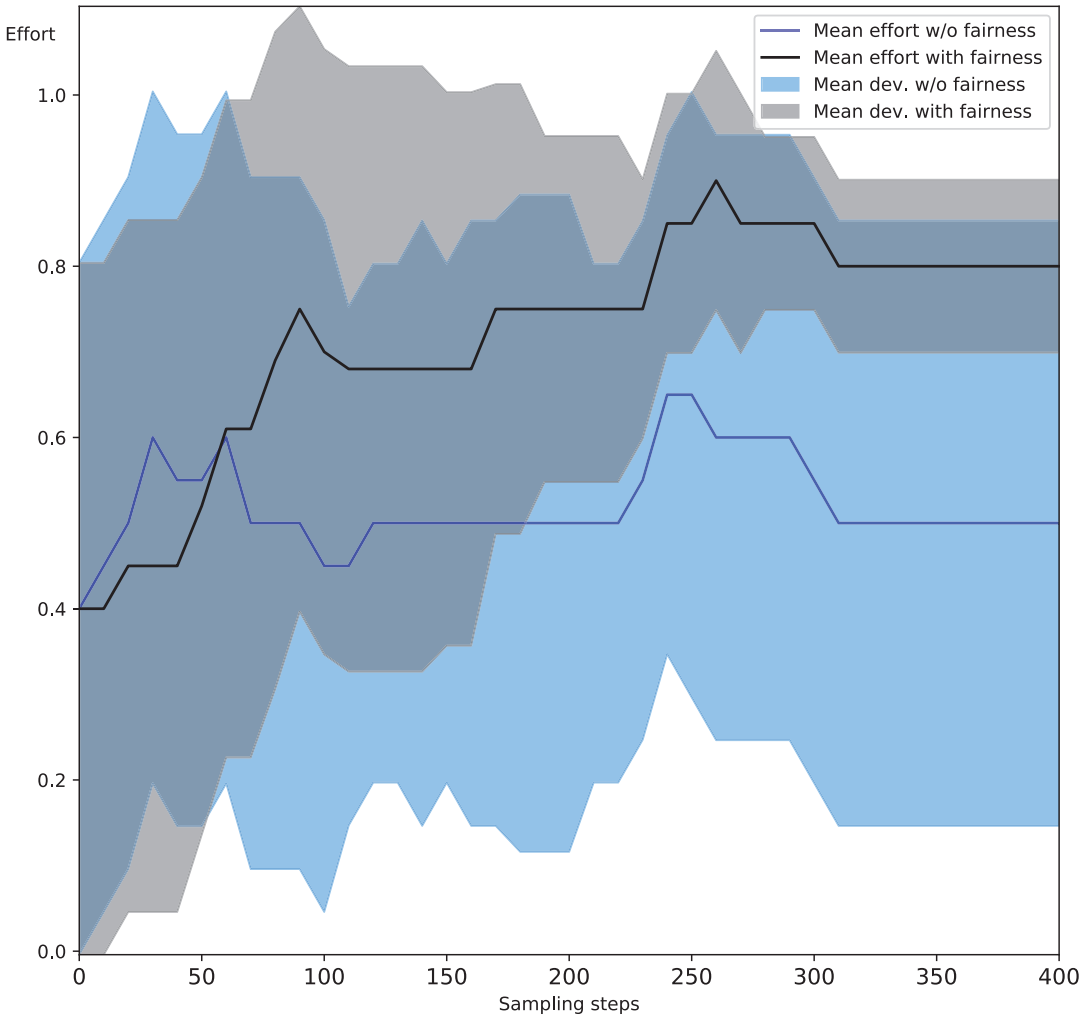

Figure 5.5: Mean effort of the agents as they settle on a satisfying solution with and without the *fairness factor* (section 5.6.3).

5.7 Summary

To sustain and accelerate the current energy transition from large scale fossil power plant to renewable but intermittent power sources, smart grids integrating residential load management are a promising solution as we explained in section 2.2. Particularly, microgrids can combine the flexibility of residential load adjustement with the efficiency of local renewable production while at the same time relieving the main grid balance system by behaving as a stable autonomous unit. However, leveraging residential consumption to balance a microgrid comes with two major challenges. As we

stated repeatedly throughout this thesis, the user must be the priority of a smart grid management system. As such, the impact of grid balancing on their comfort should be minimized. In order to handle the changes in the user preferences and behavior, the management system must then be highly adaptive and reactive. At the same time, the diversity in household appliances and grid assets and the need for a cheap and efficient deployment calls for a modular, reliable and scalable system.

In this chapter we described an approach to smart grid dynamic load balancing enabling a user-centered residential load adjustment. By adopting a bottom-up approach, we offer a way for the users to take control over their commitment to participate in demand response mechanisms, as we did in our previous model (see chapter 4). Instead of considering the user as a fixed constraint and making strong assumption regarding the flexibility of household appliances like a majority of existing works (section 5.2.1), we designed an adaptive system that handles the diversity of use cases and user constraints. The core of our model is built around the compromise between the user's comfort and the grid stability, which is dynamically adjusted in real time. To build a complete solution that could be deployed to the mass market, we chose a *decentralized infrastructure* with a *modular* design in order to considerably ease the process. The *generic model* proposed here facilitates the integration of a *large variety of devices* by simply defining the way they evaluate their own running cost and their impact on the user's comfort via the associated criterion. The iterative optimization process allows for a *self-adaptive dynamic reaction to changes in the user's behavior* and most notably to changes in the grid's configuration, as we demonstrated by simulating the real-time addition of new agents during an optimization period, in a realistic microgrid setting. We proved the pareto efficiency of our coordination algorithm, which shows that our model allows the agents to converge to a solution where the grid stability is ensured while equity regarding the consented effort is maintained. In addition to the generic device representation mentioned before, our *light communication requirements* allowed us to use a gossip aggregation algorithm as a backbone to our smart grid management system, which enables a truly decentralized system.

6
Conclusion and perspectives

▷ *Here we will briefly recall the challenges we addressed in this thesis and review the solutions we offered. We then discuss the results we obtained and draw the perspectives for futur works opened by our approach.* ◁

Chapter content

6.1 Objectives **109**
6.1.1 Users 109
6.1.2 Deployment 109
6.2 Contributions **110**
6.2.1 Load Forecasting 110
6.2.2 Decentralized Residential Demand Response 110
6.2.3 Smart Grid Energy Management System 111
6.3 Future work **112**
6.3.1 Real life testing 112
6.3.2 Security 112
6.3.3 Further studies on load forecasting 113
6.3.4 User interface 113

6.1 Objectives

Smart grids have been a trending topic in the last decade and will without a doubt continue to be for the years to come, as they encompass all the emerging use cases, approaches and technologies enabling the transition toward a greener grid. To deal with the plethora of issues caused by the increasing intermittency of energy sources, a mix of various systems managing different parts and aspects of the future grid will be needed, with new generation of power electronics guaranteeing the quality of service in the short term and higher-level energy management systems operating on the medium to long term. In this thesis we primarily focused on the latter, with the *general objective to develop an EMS for residential smart grids*. In this field, plenty of approaches have already been explored, each addressing specific concerns from battery management to failure recovery. However, adding to the objective of optimizing energy use in smart grids, we identified two important aspects that we decided to focus on.

6.1.1 Users

Demand response will play a major role in tomorrow's management systems. Instead of always building more production and storage assets to compensate for imbalances in the grid, controlling the consumption is more likely to be cheaper, straightforward and efficient. A large part of this consumption is residential, the energy being used to power air conditionning, household appliances and the various electronics of our everyday life. The combined growth of environmental concerns and connected equipments encourages users to be an active part in the energy transition, primarily by reducing their consumption. Residential Demand Response and Microgrid management systems could push the idea a step further by directly controlling a part of the user's equipment in order to finely adjust the consumption to the needs of the grid. This idea can raise concerns regarding the user's comfort that could be seriously reduced if their appliances are not properly managed. A major goal of this thesis was then to *facilitate the democratization of residential demand response by ensuring the user's comfort at all time*.

6.1.2 Deployment

Another foreseeable aspect of the future smart grids is the cohabitation of heterogenous systems with different algorithms, requirements and communication protocols; managing networks of unique topology, configuration, location and constraints; composed of assets from competing brands presenting various characteristics and interfaces; used by inhabitants with their own involvement, preferences and tech-savviness. In this extremely diverse environment, experienced first-hand by our industrial partner Ubiant in the deployment of smart-home solutions, we identified the *vital need for a unifying open model* which would enable a sane cohabitation of proprietary components and an

easy adaptation to a large range of use cases. Also, because these very diverse configurations and situations tend to change with time (adding or removing, modifying management policies, etc...), this *genericity* should be combined with a *high modularity* to offer a practical platform for actual deployments with limited maintenance costs.

Challenges in moving agent-based technology to the real world Deployed applications of agent-based systems Emerging applications of agent-based systems Integrated applications of agent-based and other technologies User studies of deployed agent-based systems

6.2 Contributions

6.2.1 Load Forecasting

As a first step toward the design of a user-friendly EMS, we evaluated the capabilities of existing load forecasting algorithms to determine if anticipating the user's consumption and behavior on the short term would be possible (chapter 3). We implemented the most promising algorithm to measure its performances on very short term individual load forecasting, using available data from Ubiant deployed solutions. The limited datasets used prevent us from drawing definitive conclusions but our preliminary results tend to show that existing methods struggle to handle the important variability of small scale consumption on the very short term. We showed that the use of exogenous variables like occupancy could improve the accuracy of the forecast but this hypothesis would require further testing with an extensive dataset, which is currently in the making at Ubiant. Regarding our objectives, we decided that the currently achievable accuracy was far from enough to credibly rely on predicting consumption, even though the use of exogenous variable such as occupancy seems promising. Thus, to handle the stochasticity of user behaviors while optimizing the energy use, *our system would need to be highly reactive and robust*.

6.2.2 Decentralized Residential Demand Response

Our goal was then to design a user-friendly, robust and reactive system that would still be simple and generic enough to provide the modularity and genericity needed for large scale deployment. This challenge, added to the primary objective of energy use optimization, called for a very adaptive and resilient solution, which directed us toward a multi-agent approach. As a first step, we focused on demand response only and designed a system to fit the current framework of the french grid operator (chapter 4). Building on top of Ubiant's smart home energy management system, we designed a decentralized coordination mechanism allowing the agents to engage their flexibility to answer to a load shedding request. To maintain a stable curtailment, the system does not prevent the users to use their appliances but rather relies on the reactivity and

size of the population of connected buildings to compensate for unpredicted changes in the user's behavior. On top of this real-time adjustement algorithm, a self-evaluation mechanism allows the overall performances to improve over time by limiting the impact of the less reliable elements. The advantages of this model are twofolds. Firstly, contrary to existing works where the user either loose temorarily the control of its appliances or is constrained to follow dynamic prices, our bottom-up approach allows the user to maintain a full control of their appliances at any time while still offering any flexibility to be exploited. Secondly, while a majority of the literature offers centralized systems with either a direct control or a broadcast signal, the fully decentralized horizontal architecture we designed using a fast gossip communication algorithm enables a precise monitoring and control of the system without a single point of failure, each component (building) following the same process and possessing the same amount of information. The scalability and robustness brought by this flat design are crucial for a cheap and reliable large scale deployment which then increases the stability of the demand response by involving a large number of participants.

6.2.3 Smart Grid Energy Management System

Our residential demand response system was dedicated to distributed load shedding and relied on the prerequisite that the building energy management solution was able to provide a relatively accurate estimation of the flexibility of the building at anytime. Our goal was then to go further that simple demand response by providing an architecture for the management of smart grids in general, including production and storage control, the typical use case being microgrids. To handle the considerable diversity of use cases (see above section 6.1.2), we chose to abstract the coordination mechanism from the local characteristics of each component by defining the notion of effort which allows each device to handle its own constraints and objectives and express them in an unified way to the system. With this approach, integrating a new asset only requires to implement the fitness function converting these constraints into this notion of effort. Like our first model (chapter 5), the coordination mechanism is designed to allow a fully decentralized architecture using the same gossip communication protocol, this time to reach a consensus over their respective operational schedules. This is done by iterating over a process identical for all agents where they start by proposing the schedule that fits their local interest best (lowest effort), then making progressive modifications to it until the global prosumption schedule of the grid matches the objective or until no agent can make further efforts. As each agent is aware of the average effort consented by the population, it can adjust its own effort accordingly so that the burden of load balancing is spread evenly over the population. We showed that this simple algorithm allows the agents to converge quickly to a satisfying solution, even with a high number of conflicting constraints (all the appliances want to be scheduled at the same time). We proved that this solution is systematically a Nash equilibrium and is also Pareto efficient. This system shares the scalability and robustness of the demand

response system thanks again to the absence of any critical component, and has even fewer communication requirements as only two variables (the total schedule and the mean effort) need to be aggregated via the Push-Sum protocol.

6.3 Future work

6.3.1 Real life testing

Using data from various sources, we were able to simulate a realistic use case to test the performances and validate the properties of our energy management system. However, more testing needs to be carried out on real-life settings in order to consolidate our confidence in the model and better analyse its potential limitations. In the coming months, the actual deployment of the solution on an actual connected building equipped with production and storage assets is planned. We will then be able to work on more detailed local optimization algorithm and to study the behavior of the system depending on parameters such as connectivity, time scales and accuracy requirements on real equipment. The genericity of our approach will be put to the test with the integration of a variety of actual appliances, and we plan to stress the system to evaluate its robustness and adaptiveness. A large part of this incoming work relates to engineering, but it will allow us to collect crucial data in order to fully evaluate the performances of our system. A broader question relates to the notion of fair effort distribution. Indeed, if our system guarantees a theoretical convergence toward a solution when agents tend to consent to the same effort, this value is relative and depends exclusively on the implementation of the fitness function. For a part of the assets which are not directly linked to the user's well being, this function is usually well defined and boils down to an operating cost. However, for household appliances and HVAC systems which directly influence the user's comfort, evaluating the quality of the solution can be more complex and will need a comprehensive evaluation on the long run.

6.3.2 Security

Althought this relates more to engineering, further work could be carried on the securization of communication between the agents as the reliability of the system in this regard might and should be a major concern for the users. Our design is inherently secure, on one hand because only anonymized information necessary for the global scheduling are transmitted which do not include sensitive or detailed information about the users or their behavior ; on the other hand because the Push-Sum protocol used to convey these data makes it extremely difficult to trace a given information back to its original sender. However, the network as a whole could benefit from existing cryptography methods to provide an additionnal layer of protection. Using standard

web communication protocols would allow this, but many different methods can be employed depending on the actual link between the different assets.

In the same category, the use of a blockchain parallel to the coordination mechanism has been mentioned. It could be used to keep track of the energy transactions between the agents [Mengelkamp17], but also of the amount of effort consented by them at each step, to later apply potential rewards and tariffs. A blockchain is far from being the only way of implementing such a ledger, nor is it the cheapest or the simplest. However, it would be relevant with our affirmed position on decentralization and security [Knirsch18]. Before that, several obstacles inherent to current blockchain technologies such as energy consumption must be overcome.

6.3.3 Further studies on load forecasting

Our preliminary and exploratory work on short term individual load forecasting presented in chapter 3 yielded few significant results. The first tests we were able to perform on state of the arts algorithms for short term load forecasting were enough to confirm the initial intuition that the attainable accuracy on individual load is very low. However, if the dataset used allowed us to identify a possibility for improvement by exploiting exogenous variable such as occupancy, its size was insufficient to draw definitive conclusion. Datasets containing frequent monitoring (one minute interval at least) of multiple variables related to electric consumption are rare, mostly because there is only a small amount of inhabited building equipped with monitoring equipment, and because gather and cleaning the data is difficult. Ubiant, through the deployment of smart home solutions, has the potential to gather extensive datasets of this kind and has begun to do so using the tools we developed for our work, obviously with the informed consent of the homeowners (even if data is systematically anonymized). With more data at hand, a better and deeper analysis of the capabilities and shortcomings of very short-term individual load forecasting can be envisionned, possibly opening new doors regarding residential load management by anticipating behaviors and needs.

6.3.4 User interface

The developpement of advanced interfaces in line with our user-centered approach was not our focus in this thesis, but the ability to design intuitive interfaces was one of the motives behind the idea of a simple and unifying model. When presented with rudimentary interfaces which lack precision or are bothersome to adjust, it is likely that the settings will only fit loosely the actual preferences of the users. The consequences are a bigger uncertainty for the management system, as the users could frequently want to override or change the preferences in real time. On the contrary, if the users are able to easily and precisely communicate their preferences to the system, they will be encouraged to set them in a way that actually fits their needs, naturally increasing the

reliability of the management system and potentially reducing the final operating cost of the grid. Smooth communication with the users could also be of prime importance when dealing with emergency situations regarding the stability of the grid. In some cases, notably in small isolated microgrid, the risk of black-out could happen despite an efficient EMS, as production can depend heavily on external factors such as the weather. In such cases, a quick and clear understanding of the situation by the users could help reduce the inconvience but could also mitigate it, for example if manual intervention is needed to curtail the demand (some devices can not yet be managed by the EMS , as we saw in section 5.4). Also, as we mentioned in section 2.2.2.2, gamification and social incentives can be important tools for the active involvement of users. Thanks to the straightforward logic used by our coordination algorithm, it could be possible to design interfaces and that allow the users to understand how the balancing mechanism works, which would help them to better engage in the system but most importantly to better welcome the adjustments made by the management system and maybe even actively adapt their energy use themselves by seeing this as a social game. By leveraging the intuitive notion of effort which is at the core of our model, we could inform the users of their relative level of participation in the collective balancing effort in real time, which is both a good indicator of the impact of their energy usage but also a good way to show the interest of the adjustments made by the system.

As we mentionned previously on our perspectives regarding load forecasting, big opportunities lie in the ability to learn and forecast the users behavior, as it comes in pair with the idea of automatically learning and updating their preferences. The possibilities are numerous, from simply learning the occupancy patterns of the building to the precise temperature preferences for each room depending on the time of day and day of the year, but also potentially knowing the lighting habits and needs to be able to exploit these small but abundant sources of flexibility. Eventually, these considerations lead back to the idea of ambient intelligence (see section 1.1.2), with the combination of smart interfaces and advanced prediction capabilities allowing for a smooth user experience, which improves the performances of the system in a virtuous circle.

A

Annexes

▷

Annexes

◁

Chapter content

A.1 Prosumption profiles used in simulations **117**
A.2 Publications **120**

A.1 Prosumption profiles used in simulations

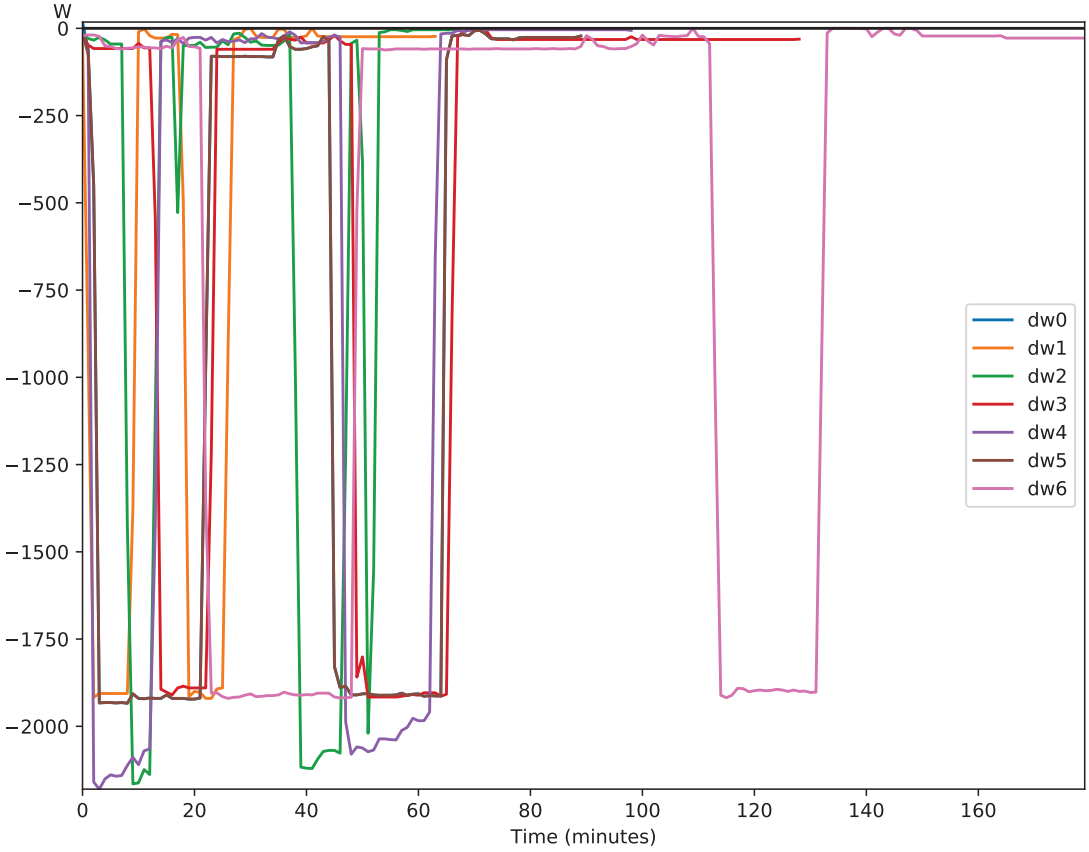

Figure A.1: Seven different load profiles for dishwashers used in section 5.6.1. Source : KIT, OSHv4

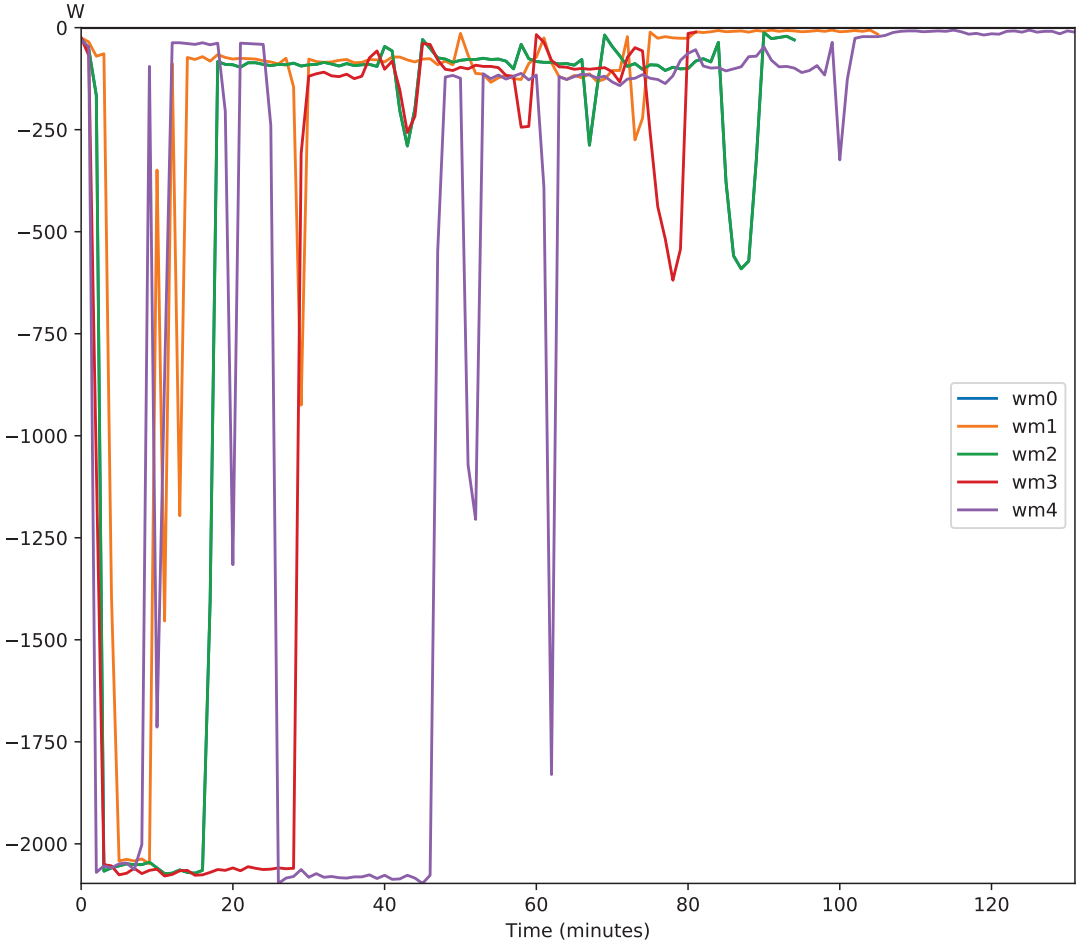

Figure A.2: 5 different load profiles for washing machines used in section 5.6.1. Source : KIT, OSHv4

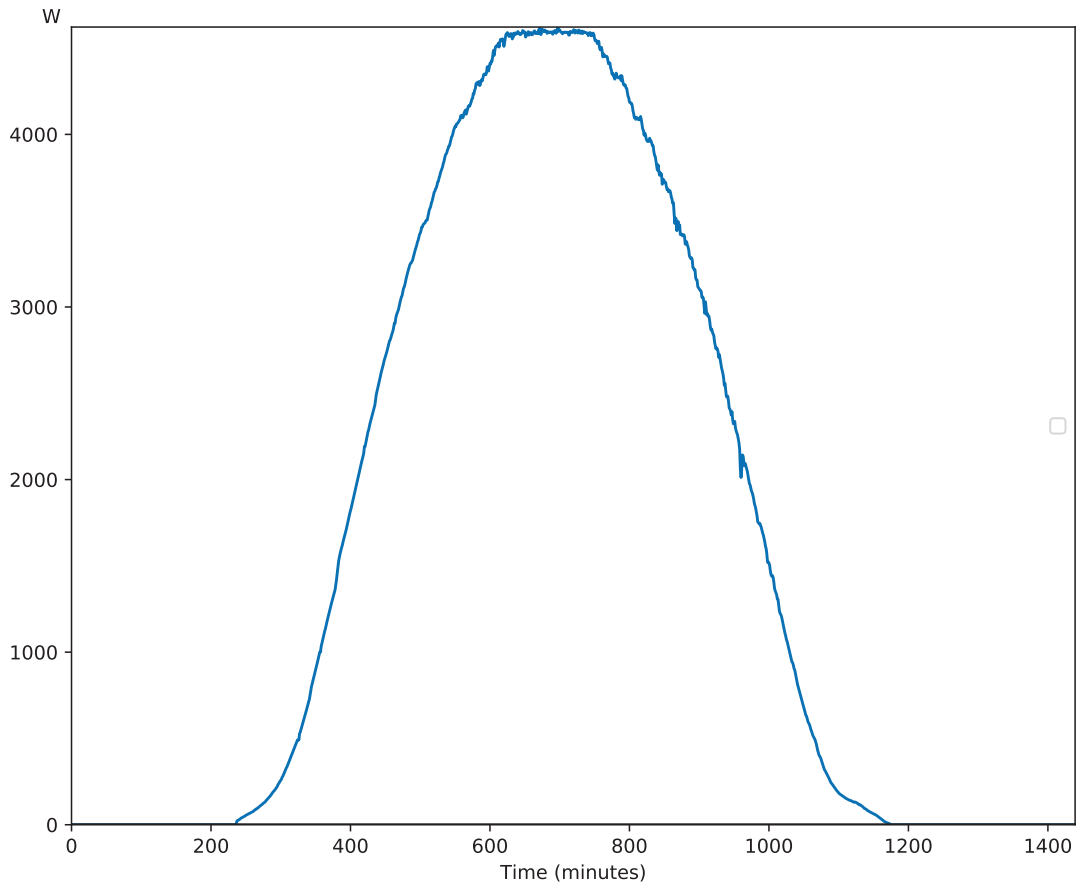

Figure A.3: Production curve of a solar panel array used in section 5.6.1. Source : Ubiant

A.2 Publications

[1] V. Lequay, M. Lefort, S. Hassas, and S. Mansour. Ajustement Diffus et Adaptatif de la Consommation Électrique Résidentielle par un Système Multi-Agents. *In Journées Francophones Sur Les Systèmes Multi-Agents*, 2016.

[2] V. Lequay, M. Lefort, S. Mansour, and S. Hassas. Flexible Load Shedding using Gossip Communication in a Multi-Agents System. *In IEEE 10th International Conference on Self-Adaptive and Self-Organizing Systems*, 2016.

[3] V. Lequay, M. Lefort, S. Mansour, and S. Hassas. Ajustement Diffus et Adaptatif de la Consommation Électrique Résidentielle par un Système Multi-Agents. *Revue d'Intelligence Artificielle*, 2017.

Bibliography

[Abdella18] Juhar Abdella & Khaled Shuaib. *Peer to Peer Distributed Energy Trading in Smart Grids: A Survey*. Energies, vol. 11, no. 6, page 1560, June 2018.

[Abou-Hussien81] M S Abou-Hussien, M S Kand, M A Tantawy & S A Farghal. *An Accurate Model for Short-Term Load Forecasting*. page 8, 1981.

[Albadi08] M.H. Albadi & E.F. El-Saadany. *A Summary of Demand Response in Electricity Markets*. Electric Power Systems Research, vol. 78, no. 11, pages 1989–1996, November 2008.

[AlSkaif18] Tarek AlSkaif, Ioannis Lampropoulos, Machteld van den Broek & Wilfried van Sark. *Gamification-Based Framework for Engagement of Residential Customers in Energy Applications*. Energy Research & Social Science, vol. 44, pages 187–195, October 2018.

[Amjady10] Nima Amjady, Farshid Keynia & Hamidreza Zareipour. *Short-Term Load Forecast of Microgrids by a New Bilevel Prediction Strategy*. IEEE Transactions on Smart Grid, vol. 1, no. 3, pages 286–294, December 2010. 00181.

[Amjady11] N Amjady, F Keynia & H Zareipour. *Wind Power Prediction by a New Forecast Engine Composed of Modified Hybrid Neural Network and Enhanced Particle Swarm Optimization*. IEEE Transactions on Sustainable Energy, vol. 2, no. 3, pages 265–276, July 2011. 00137.

[Aneke16] Mathew Aneke & Meihong Wang. *Energy Storage Technologies and Real Life Applications – A State of the Art Review*. Applied Energy, vol. 179, pages 350–377, October 2016.

[Anvari-Moghaddam17] Amjad Anvari-Moghaddam, Ashkan Rahimi-Kian, Maryam S. Mirian & Josep M. Guerrero. *A Multi-Agent*

Based Energy Management Solution for Integrated Buildings and Microgrid System. Applied Energy, vol. 203, pages 41–56, October 2017.

[Aung12] Zeyar Aung, Mohamed Toukhy, John R Williams, Abel Sanchez & Sergio Herrero. *Towards Accurate Electricity Load Forecasting in Smart Grids*. page 7, 2012.

[Awerbuch97] Shimon Awerbuch & Alistair Preston. The Virtual Utility: Accounting, Technology & Competitive Aspects of the Emerging Industry. Springer US, Boston, MA, 1997. OCLC: 851741542.

[Baharlouei13] Zahra Baharlouei, Massoud Hashemi, Hamed Narimani & Hamed Mohsenian-Rad. *Achieving Optimality and Fairness in Autonomous Demand Response: Benchmarks and Billing Mechanisms*. IEEE Transactions on Smart Grid, vol. 4, no. 2, pages 968–975, June 2013.

[Basir Khan16] M. Reyasudin Basir Khan, Razali Jidin & Jagadeesh Pasupuleti. *Multi-Agent Based Distributed Control Architecture for Microgrid Energy Management and Optimization*. Energy Conversion and Management, vol. 112, pages 288–307, March 2016.

[Beal12] Jacob Beal, Jeffrey Berliner & Kevin Hunter. *Fast Precise Distributed Control for Energy Demand Management*. In Self-Adaptive and Self-Organizing Systems (SASO), 2012 IEEE Sixth International Conference, pages 187–192. IEEE, 2012.

[Becker12] Birger Becker, Anna Kellerer & Hartmut Schmeck. *User Interaction Interface for Energy Management in Smart Homes*. In 2012 IEEE PES Innovative Smart Grid Technologies (ISGT), pages 1–8, Washington, DC, USA, January 2012. IEEE. 00026.

[Callaway11] D S Callaway & I A Hiskens. *Achieving Controllability of Electric Loads*. Proceedings of the IEEE, vol. 99, no. 1, pages 184–199, 2011.

[Carlson13] W. Bernard Carlson. Tesla: Inventor of the electrical age. Princeton Univ. Press, Princeton, NJ, 2013. OCLC: 847108358.

[Chakraborty15] Shantanu Chakraborty, Shin Nakamura & Toshiya Okabe. *Real-Time Energy Exchange Strategy of Optimally Cooperative Microgrids for Scale-Flexible Distribution System*. Expert Systems with Applications, vol. 42, no. 10, pages 4643–4652, June 2015. 00032.

[Chaouachi13] A. Chaouachi, R. M. Kamel, R. Andoulsi & K. Nagasaka. *Multiobjective Intelligent Energy Management for a Microgrid*. IEEE Transactions on Industrial Electronics, vol. 60, no. 4, pages 1688–1699, April 2013.

[Chicco04] G. Chicco, R. Napoli, F. Piglione, P. Postolache, M. Scutariu & C. Toader. *Load Pattern-Based Classification of Electricity Customers*. IEEE Transactions on Power Systems, vol. 19, no. 2, pages 1232–1239, May 2004.

[Chitsaz15] Hamed Chitsaz, Hamid Shaker, Hamidreza Zareipour, David Wood & Nima Amjady. *Short-Term Electricity Load Forecasting of Buildings in Microgrids*. Energy and Buildings, vol. 99, pages 50–60, 2015. 00049.

[Chung13] Il-Yop Chung. *Distributed Intelligent Microgrid Control Using Multi-Agent Systems*. Engineering, vol. 05, no. 01, pages 1–6, 2013. 00015.

[Colson09] C. M. Colson, M. H. Nehrir & C. Wang. *Ant Colony Optimization for Microgrid Multi-Objective Power Management*. In Power Systems Conference and Exposition, 2009. PSCE'09. IEEE/PES, Colson, pages 1–7. IEEE, 2009.

[Colson10] C. M. Colson, M. H. Nehrir & S. A. Pourmousavi. *Towards Real-Time Microgrid Power Management Using Computational Intelligence Methods*. In IEEE PES General Meeting, Colson, pages 1–8. IEEE, 2010.

[Colson11a] C. M. Colson & M. H. Nehrir. *Algorithms for Distributed Decision-Making for Multi-Agent Microgrid Power Management*. In 2011 IEEE Power and Energy Society General Meeting, Colson, pages 1–8. IEEE, 2011.

[Colson11b] C.M. Colson, M.H. Nehrir & R.W. Gunderson. *Multi-Agent Microgrid Power Management*. IFAC Proceedings Volumes, vol. 44, no. 1, pages 3678–3683, January 2011.

[Conejo05] A.J. Conejo, M.A. Plazas, R. Espinola & A.B. Molina. *Day-Ahead Electricity Price Forecasting Using the Wavelet Transform and ARIMA Models*. IEEE Transactions on Power Systems, vol. 20, no. 2, pages 1035–1042, May 2005.

[Conti10] S. Conti, R. Nicolosi & S. A. Rizzo. *Optimal Dispatching of Distributed Generators in an MV Autonomous Micro-Grid to Minimize Operating Costs and Emissions*. In 2010 IEEE International Symposium on Industrial Electronics, pages 2542–2547. IEEE, 2010. 00010.

[Conti18] Mauro Conti, Ali Dehghantanha, Katrin Franke & Steve Watson. *Internet of Things Security and Forensics: Challenges and Opportunities*. Future Generation Computer Systems, vol. 78, pages 544–546, January 2018.

[Contreras03] J. Contreras, R. Espinola, F.J. Nogales & A.J. Conejo. *ARIMA Models to Predict Next-Day Electricity Prices*. IEEE Transactions on Power Systems, vol. 18, no. 3, pages 1014–1020, August 2003.

[Cook09] Diane J. Cook, Juan C. Augusto & Vikramaditya R. Jakkula. *Ambient Intelligence: Technologies, Applications, and Opportunities*. Pervasive and Mobile Computing, vol. 5, no. 4, pages 277–298, August 2009.

[Dagdougui14] Hanane Dagdougui & Roberto Sacile. *Decentralized Control of the Power Flows in a Network of Smart Microgrids Modeled as a Team of Cooperative Agents*. IEEE Transactions on Control Systems Technology, vol. 22, no. 2, pages 510–519, March 2014. 00043.

[De Oliveira11] Grégory De Oliveira, Mireille Jacomino, Duy Long Ha & Stéphane Ploix. *Optimal Power Control for Smart Homes*. IFAC Proceedings Volumes, vol. 44, no. 1, pages 9579–9586, January 2011. 00023.

[Deb17] Chirag Deb, Fan Zhang, Junjing Yang, Siew Eang Lee & Kwok Wei Shah. *A Review on Time Series Forecasting Techniques for Building Energy Consumption*. Renewable and Sustainable Energy Reviews, vol. 74, pages 902–924, July 2017. 00046.

[Dimeas05] A.L. Dimeas & N.D. Hatziargyriou. *Operation of a Multi-agent System for Microgrid Control*. IEEE Transactions on

Power Systems, vol. 20, no. 3, pages 1447–1455, August 2005.

[Dorri17] Ali Dorri, Salil S. Kanhere, Raja Jurdak & Praveen Gauravaram. *Blockchain for IoT Security and Privacy: The Case Study of a Smart Home*. In 2017 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops), pages 618–623, Kona, HI, March 2017. IEEE.

[Drezga98] I. Drezga & S. Rahman. *Input Variable Selection for ANN-Based Short-Term Load Forecasting*. IEEE Transactions on Power Systems, vol. 13, no. 4, pages 1238–1244, Nov./1998.

[Edmund Taylor Whittaker10] Edmund Taylor Whittaker. A History of the Theories of Aether and Electricity from the Age of Descartes to the Close of the Nineteenth Century. Longmans, Green and Company, 1910.

[Edwards94] Ward Edwards & F.Hutton Barron. *SMARTS and SMARTER: Improved Simple Methods for Multiattribute Utility Measurement*. Organizational Behavior and Human Decision Processes, vol. 60, no. 3, pages 306–325, December 1994.

[Edwards12] Richard E. Edwards, Joshua New & Lynne E. Parker. *Predicting Future Hourly Residential Electrical Consumption: A Machine Learning Case Study*. Energy and Buildings, vol. 49, pages 591–603, June 2012.

[Efthymiou10] Costas Efthymiou & Georgios Kalogridis. *Smart Grid Privacy via Anonymization of Smart Metering Data*. In 2010 First IEEE International Conference on Smart Grid Communications, pages 238–243, Gaithersburg, MD, USA, October 2010. IEEE.

[Eid15] Cherrelle Eid, Paul Codani, Yurong Chen, Yannick Perez & Rudi Hakvoort. *Aggregation of Demand Side Flexibility in a Smart Grid: A Review for European Market Design*. In 2015 12th International Conference on the European Energy Market (EEM), pages 1–5, Lisbon, Portugal, May 2015. IEEE.

[Eid16] Cherrelle Eid, Elta Koliou, Mercedes Valles, Javier Reneses & Rudi Hakvoort. *Time-Based Pricing and Electricity Demand Response: Existing Barriers and next Steps*. Utilities Policy, vol. 40, pages 15–25, June 2016.

[Ellabban14] Omar Ellabban, Haitham Abu-Rub & Frede Blaabjerg. *Renewable Energy Resources: Current Status, Future Prospects and Their Enabling Technology*. Renewable and Sustainable Energy Reviews, vol. 39, pages 748–764, November 2014.

[European Commission96] European Commission & and Development Science Directorate-General XII Research. Final report: Project JOULE II : JOU-CT92-0053 : Development and implementation of an advanced control system for the optimal operation and management of medium-sized power systems with a large penetration from renewable power sources JOULE II: JOU2-CT92-0053. EUROP, Luxembourg, 1996.

[Fan10] Zhong Fan. *Distributed Demand Response and User Adaptation in Smart Grids*. arXiv:1007.5425 [cs], July 2010. 00060.

[Fan12] Zhong Fan. *A Distributed Demand Response Algorithm and Its Application to PHEV Charging in Smart Grids*. IEEE Transactions on Smart Grid, vol. 3, no. 3, pages 1280–1290, September 2012.

[Farhangi10] H. Farhangi. *The Path of the Smart Grid*. IEEE Power and Energy Magazine, vol. 8, no. 1, pages 18–28, January 2010.

[Fazal12] Rehan Fazal, Jignesh Solanki & Sarika Khushalani Solanki. *Demand Response Using Multi-Agent System*. In 2012 North American Power Symposium (NAPS), pages 1–6, Champaign, IL, USA, September 2012. IEEE.

[Ferber97] Jacques Ferber. Les systèmes multi-agents: Vers une intelligence collective. Informatique, Intelligence Artificielle. InterÉd, Paris, 2. tirage edition, 1997. OCLC: 247784735.

[Foo. Eddy15] Y. S. Foo. Eddy, H. B. Gooi & S. X. Chen. *Multi-Agent System for Distributed Management of Microgrids*. IEEE

Transactions on Power Systems, vol. 30, no. 1, pages 24–34, January 2015. 00113.

[Franklin97] Stan Franklin & Art Graesser. *Is It an Agent, or Just a Program?: A Taxonomy for Autonomous Agents*. In J. G. Carbonell, J. Siekmann, G. Goos, J. Hartmanis, J. van Leeuwen, Jörg P. Müller, Michael J. Wooldridge & Nicholas R. Jennings, editeurs, Intelligent Agents III Agent Theories, Architectures, and Languages, volume 1193, pages 21–35. Springer Berlin Heidelberg, Berlin, Heidelberg, 1997.

[Frey13] Sylvain Frey, Ada Diaconescu, David Menga & Isabelle Demeure. *A Holonic Control Architecture for a Heterogeneous Multi-Objective Smart Micro-Grid*. pages 21–30. IEEE, 2013.

[Gaillard16] Pierre Gaillard, Yannig Goude & Raphaël Nedellec. *Additive Models and Robust Aggregation for GEFCom2014 Probabilistic Electric Load and Electricity Price Forecasting*. International Journal of Forecasting, vol. 32, no. 3, pages 1038–1050, July 2016. 00052.

[Giebel07] Gregor Giebel & Georges Kariniotakis. *Best Practice in Short-Term Forecasting. A Users Guide*. In European Wind Energy Conference and Exhibition 2007, pages 5–pages, 2007. 00025.

[Gils14] Hans Christian Gils. *Assessment of the Theoretical Demand Response Potential in Europe*. Energy, vol. 67, pages 1–18, April 2014.

[Gnauk12] Benjamin Gnauk, Lars Dannecker & Martin Hahmann. *Leveraging Gamification in Demand Dispatch Systems*. In Proceedings of the 2012 Joint EDBT/ICDT Workshops on - EDBT-ICDT '12, page 103, Berlin, Germany, 2012. ACM Press.

[Gordon06] S. Gordon. *Supergrid to the Rescue [Electricity Supply Security]*. Power Engineer, vol. 20, no. 5, page 30, 2006.

[Grossmann84] A. Grossmann & J. Morlet. *Decomposition of Hardy Functions into Square Integrable Wavelets of Constant Shape*. SIAM Journal on Mathematical Analysis, vol. 15, no. 4, pages 723–736, July 1984.

[Guerrero13] Josep M. Guerrero, Mukul Chandorkar, Tzung-Lin Lee & Poh Chiang Loh. *Advanced Control Architectures for Intelligent Microgrids, Part I: Decentralized and Hierarchical Control*. IEEE Transactions on Industrial Electronics, vol. 60, no. 4, pages 1254–1262, 2013.

[Haider16] Haider Tarish Haider, Ong Hang See & Wilfried Elmenreich. *A Review of Residential Demand Response of Smart Grid*. Renewable and Sustainable Energy Reviews, vol. 59, pages 166–178, June 2016.

[Harmouch16] Fatima Zahra Harmouch, Nissrine Krami, Driss Benhaddou, Nabil Hmina, Elmajid Zayer & El Hassane Margoum. *Survey of Multiagents Systems Application in Microgrids*. In 2016 International Conference on Electrical and Information Technologies (ICEIT), pages 270–275. IEEE, 2016. 00005.

[Hernandez-Aramburo05] C.A. Hernandez-Aramburo, T.C. Green & N. Mugniot. *Fuel Consumption Minimization of a Microgrid*. IEEE Transactions on Industry Applications, vol. 41, no. 3, pages 673–681, May 2005.

[Hernandez23] Luis Hernandez, Carlos Baladron, Javier M. Aguiar, Belen Carro, Antonio J. Sanchez-Esguevillas, Jaime Lloret & Joaquim Massana. *A Survey on Electric Power Demand Forecasting: Future Trends in Smart Grids, Microgrids and Smart Buildings*. IEEE Communications Surveys & Tutorials, vol. 16, no. 3, pages 1460–1495, 23.

[Hippert01] Henrique Steinherz Hippert, Carlos Eduardo Pedreira & Reinaldo Castro Souza. *Neural Networks for Short-Term Load Forecasting: A Review and Evaluation*. IEEE Transactions on power systems, vol. 16, no. 1, pages 44–55, 2001.

[Horling04] Bryan Horling & Victor Lesser. *A Survey of Multi-Agent Organizational Paradigms*. The Knowledge Engineering Review, vol. 19, no. 4, pages 281–316, December 2004.

[Ibrahim08] H Ibrahim, A Ilinca & J Perron. *Energy Storage Systems—Characteristics and Comparisons*. Renewable and Sustainable Energy Reviews, vol. 12, no. 5, pages 1221–1250, June 2008.

[Jansen10] Bernhard Jansen, Carl Binding, Olle Sundstrom & Dieter Gantenbein. *Architecture and Communication of an Electric Vehicle Virtual Power Plant*. pages 149–154. IEEE, October 2010. 00081.

[Jetcheva14] Jorjeta G. Jetcheva, Mostafa Majidpour & Wei-Peng Chen. *Neural Network Model Ensembles for Building-Level Electricity Load Forecasts*. Energy and Buildings, vol. 84, pages 214–223, December 2014.

[Jiang06] Zhenhua Jiang. *Agent-Based Control Framework for Distributed Energy Resources Microgrids*. In 2006 IEEE/WIC/ACM International Conference on Intelligent Agent Technology, pages 646–652, Hong Kong, China, 2006. IEEE.

[Jin17] Xin Jin, Kyri Baker, Dane Christensen & Steven Isley. *Foresee: A User-Centric Home Energy Management System for Energy Efficiency and Demand Response*. Applied Energy, vol. 205, pages 1583–1595, November 2017.

[Jogunola17] Olamide Jogunola, Augustine Ikpehai, Kelvin Anoh, Bamidele Adebisi, Mohammad Hammoudeh, Haris Gacanin & Georgina Harris. *Comparative Analysis of P2P Architectures for Energy Trading and Sharing*. Energies, vol. 11, no. 1, page 62, December 2017.

[Kahrobaee13] Salman Kahrobaee, Rasheed A. Rajabzadeh, Leen-Kiat Soh & Sohrab Asgarpoor. *A Multiagent Modeling and Investigation of Smart Homes With Power Generation, Storage, and Trading Features*. IEEE Transactions on Smart Grid, vol. 4, no. 2, pages 659–668, June 2013.

[Kakas11] Antonis C. Kakas, David Cohn, Sanjoy Dasgupta, Andrew G. Barto, Gail A. Carpenter, Stephen Grossberg, Geoffrey I. Webb, Marco Dorigo, Mauro Birattari, Hannu Toivonen, Jon Timmis, Jürgen Branke, Hannu Toivonen, Alexander L. Strehl, Chris Drummond, Adam Coates, Pieter Abbeel, Andrew Y. Ng, Fei Zheng, Geoffrey I. Webb & Prasad Tadepalli. *Ant Colony Optimization*. In Claude Sammut & Geoffrey I. Webb, editeurs, Encyclopedia of Machine Learning, pages 36–39. Springer US, Boston, MA, 2011.

[Kamaev12] V. A. Kamaev, M. V. Shcherbakov, D. P. Panchenko, N. L. Shcherbakova & A. Brebels. *Using Connectionist Systems*

for Electric Energy Consumption Forecasting in Shopping Centers. Automation and Remote Control, vol. 73, no. 6, pages 1075–1084, June 2012. 00013.

[Kantamneni15] Abhilash Kantamneni, Laura E. Brown, Gordon Parker & Wayne W. Weaver. *Survey of Multi-Agent Systems for Microgrid Control*. Engineering Applications of Artificial Intelligence, vol. 45, pages 192–203, October 2015.

[Karavas15] Christos-Spyridon Karavas, George Kyriakarakos, Konstantinos G. Arvanitis & George Papadakis. *A Multi-Agent Decentralized Energy Management System Based on Distributed Intelligence for the Design and Control of Autonomous Polygeneration Microgrids*. Energy Conversion and Management, vol. 103, pages 166–179, October 2015. 00088.

[Karfopoulos15] E. Karfopoulos, L. Tena, A. Torres, Pep Salas, Joan Gil Jorda, A. Dimeas & N. Hatziargyriou. *A Multi-Agent System Providing Demand Response Services from Residential Consumers*. Electric Power Systems Research, vol. 120, pages 163–176, March 2015. 00048.

[Kashani17] Atieh Kashani & Yusuf Ozturk. *Residential Energy Consumer Behavior Modification via Gamification*. In 2017 IEEE 6th International Conference on Renewable Energy Research and Applications (ICRERA), pages 1221–1225, San Diego, CA, November 2017. IEEE.

[Kempe03] David Kempe, Alin Dobra & Johannes Gehrke. *Gossip-Based Computation of Aggregate Information*. In Proceedings. 44th Annual IEEE Symposium on Foundations of Computer Science, 2003, pages 482–491. IEEE, 2003.

[Klaimi18] Joelle Klaimi & Meritxell Vinyals. *Decentralised District Multi-Vector Energy Management: A Multi-Agent Approach*. In Luis M. Camarinha-Matos, Hamideh Afsarmanesh & Yacine Rezgui, editeurs, Collaborative Networks of Cognitive Systems, volume 534, pages 551–559. Springer International Publishing, Cham, 2018.

[Knirsch18] Fabian Knirsch, Andreas Unterweger, Günther Eibl & Dominik Engel. *Privacy-Preserving Smart Grid Tariff Decisions with Blockchain-Based Smart Contracts*. In Wilson Rivera, editeur, Sustainable Cloud and Energy Services,

pages 85–116. Springer International Publishing, Cham, 2018.

[Koutitas12] George Koutitas. *Control of Flexible Smart Devices in the Smart Grid*. IEEE Transactions on Smart Grid, vol. 3, no. 3, pages 1333–1343, September 2012.

[Kuznetsova14] Elizaveta Kuznetsova, Yan-Fu Li, Carlos Ruiz & Enrico Zio. *An Integrated Framework of Agent-Based Modelling and Robust Optimization for Microgrid Energy Management*. Applied Energy, vol. 129, pages 70–88, September 2014. 00102.

[Kwok11] Simon S.K. Kwok, Richard K.K. Yuen & Eric W.M. Lee. *An Intelligent Approach to Assessing the Effect of Building Occupancy on Building Cooling Load Prediction*. Building and Environment, vol. 46, no. 8, pages 1681–1690, August 2011.

[Lasseter01] B. Lasseter. *Microgrids [Distributed Power Generation]*. In 2001 IEEE Power Engineering Society Winter Meeting. Conference Proceedings (Cat. No.01CH37194), volume 1, pages 146–149, Columbus, OH, USA, 2001. IEEE.

[Lasseter02] Robert Lasseter, Abbas Akhil, Chris Marnay, John Stephens, Jeff Dagle, Ross Guttromsom, A. Sakis Meliopoulous, Robert Yinger & Joe Eto. *Integration of Distributed Energy Resources. The CERTS Microgrid Concept*. Rapport technique LBNL–50829, 799644, April 2002.

[Lee15] Tien Ming Lee, Ezra M. Markowitz, Peter D. Howe, Chia-Ying Ko & Anthony A. Leiserowitz. *Predictors of Public Climate Change Awareness and Risk Perception around the World*. Nature Climate Change, vol. 5, page 1014, July 2015.

[Levron13] Yoash Levron, Josep M. Guerrero & Yuval Beck. *Optimal Power Flow in Microgrids With Energy Storage*. IEEE Transactions on Power Systems, vol. 28, no. 3, pages 3226–3234, August 2013.

[Li13] Hong-ze Li, Sen Guo, Chun-jie Li & Jing-qi Sun. *A Hybrid Annual Power Load Forecasting Model Based on Generalized Regression Neural Network with Fruit Fly Optimization Algorithm*. Knowledge-Based Systems, vol. 37, pages 378–387, January 2013.

[Li15] Kangji Li, Chenglei Hu, Guohai Liu & Wenping Xue. *Building's Electricity Consumption Prediction Using Optimized Artificial Neural Networks and Principal Component Analysis*. Energy and Buildings, vol. 108, pages 106–113, December 2015.

[Li16] Ting Li, Bo Yang & Dong Liu. *Applications of Multiobjective Particle Swarm Optimization Algorithms in Smart Grid: A Comprehensive Survey*. TelKOMNIKA (Telecommunication Computing Electronics and Control), vol. 14, no. 3A, 2016.

[Lim14] Yujin Lim, Hak-Man Kim & Tetsuo Kinoshita. *Distributed Load-Shedding System for Agent-Based Autonomous Microgrid Operations*. Energies, vol. 7, no. 1, pages 385–401, January 2014. 00025.

[Linnenberg11] Tobias Linnenberg, Ireneus Wior, Sebastian Schreiber & Alexander Fay. *A Market-Based Multi-Agent-System for Decentralized Power and Grid Control*. In Emerging Technologies & Factory Automation (ETFA), 2011 IEEE 16th Conference On, pages 1–8. IEEE, 2011.

[Lisovich10] Mikhail A. Lisovich, Deirdre K. Mulligan & Stephen B. Wicker. *Inferring Personal Information from Demand-Response Systems*. IEEE Security & Privacy Magazine, vol. 8, no. 1, pages 11–20, January 2010.

[Liu14] Nian Liu, Qingfeng Tang, Jianhua Zhang, Wei Fan & Jie Liu. *A Hybrid Forecasting Model with Parameter Optimization for Short-Term Load Forecasting of Micro-Grids*. Applied Energy, vol. 129, pages 336–345, September 2014.

[Logenthiran11] T. Logenthiran, Dipti Srinivasan & Ashwin M. Khambadkone. *Multi-Agent System for Energy Resource Scheduling of Integrated Microgrids in a Distributed System*. Electric Power Systems Research, vol. 81, no. 1, pages 138–148, January 2011. 00258.

[Logenthiran12] Thillainathan Logenthiran, Dipti Srinivasan, Ashwin M. Khambadkone & Htay Nwe Aung. *Multiagent System for Real-Time Operation of a Microgrid in Real-Time Digital Simulator*. IEEE Transactions on Smart Grid, vol. 3, no. 2, pages 925–933, June 2012. 00241.

[Lopes07] J.A. Peças Lopes, N. Hatziargyriou, J. Mutale, P. Djapic & N. Jenkins. *Integrating Distributed Generation into Electric Power Systems: A Review of Drivers, Challenges and Opportunities*. Electric Power Systems Research, vol. 77, no. 9, pages 1189–1203, July 2007.

[López13] M.A. López, S. Martín, J.A. Aguado & S. de la Torre. *V2G Strategies for Congestion Management in Microgrids with High Penetration of Electric Vehicles*. Electric Power Systems Research, vol. 104, pages 28–34, November 2013.

[Lu10] Jiakang Lu, Tamim Sookoor, Vijay Srinivasan, Ge Gao, Brian Holben, John Stankovic, Eric Field & Kamin Whitehouse. *The Smart Thermostat: Using Occupancy Sensors to Save Energy in Homes*. In Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems - SenSys '10, page 211, Zürich, Switzerland, 2010. ACM Press.

[Lu11] Shuai Lu, Nader Samaan, Ruisheng Diao, Marcelo Elizondo, Chunlian Jin, Ebony Mayhorn, Yu Zhang & Harold Kirkham. *Centralized and Decentralized Control for Demand Response*. In Innovative Smart Grid Technologies (ISGT), 2011 IEEE PES, pages 1–8. IEEE, 2011. 00090.

[Mansour12] Saber Mansour, Nicolas Wiest, Olivier Lefevre & Sebastien Mazac. *Hemis: Hybrid Multi-Agent Architecture for Energy Management and Home Automation*. In 2012 IEEE Sixth International Conference on Self-Adaptive and Self-Organizing Systems, pages 229–230, Lyon, France, September 2012. IEEE.

[Massana16] Joaquim Massana, Carles Pous, Llorenç Burgas, Joaquim Melendez & Joan Colomer. *Short-Term Load Forecasting for Non-Residential Buildings Contrasting Artificial Occupancy Attributes*. Energy and Buildings, vol. 130, pages 519–531, October 2016.

[Massimo Cossentino11] Massimo Cossentino, Carmelo Lodato, Salvatore Lopes, Marcello Pucci, Gianpaolo Vitale & Maurizio Cirrincione. *A Multi-Agent Architecture for Simulating and Managing Microgrids*. Proceedings of the Federated

Conference on Computer Science and Information Systems, pages 619–622, 2011. 00013 OCLC: 839426912.

[Massoud Amin05] S. Massoud Amin & B.F. Wollenberg. *Toward a Smart Grid: Power Delivery for the 21st Century*. IEEE Power and Energy Magazine, vol. 3, no. 5, pages 34–41, September 2005.

[Mauser15] Ingo Mauser, Julian Feder, Jan Müller & Hartmut Schmeck. *Evolutionary Optimization of Smart Buildings with Interdependent Devices*. In Antonio M. Mora & Giovanni Squillero, editeurs, Applications of Evolutionary Computation, volume 9028, pages 239–251. Springer International Publishing, Cham, 2015.

[Mauser16] Ingo Mauser, Jan Müller, Florian Allerding & Hartmut Schmeck. *Adaptive Building Energy Management with Multiple Commodities and Flexible Evolutionary Optimization*. Renewable Energy, vol. 87, pages 911–921, March 2016. 00033.

[Mazac14] Sébastien Mazac, Frédéric Armetta & Salima Hassas. *On Bootstrapping Sensori-Motor Patterns for a Constructivist Learning System in Continuous Environments*. In Alife 14: Fourteenth International Conference on the Synthesis and Simulation of Living Systems, 2014.

[McKenna12] Eoghan McKenna, Ian Richardson & Murray Thomson. *Smart Meter Data: Balancing Consumer Privacy Concerns with Legitimate Applications*. Energy Policy, vol. 41, pages 807–814, February 2012.

[Mena14] R. Mena, F. Rodríguez, M. Castilla & M.R. Arahal. *A Prediction Model Based on Neural Networks for the Energy Consumption of a Bioclimatic Building*. Energy and Buildings, vol. 82, pages 142–155, October 2014.

[Mengelkamp17] Esther Mengelkamp, Benedikt Notheisen, Carolin Beer, David Dauer & Christof Weinhardt. *A Blockchain-Based Smart Grid: Towards Sustainable Local Energy Markets*. Computer Science - Research and Development, August 2017.

[Mengelkamp18] Esther Mengelkamp, Johannes Gärttner, Kerstin Rock, Scott Kessler, Lawrence Orsini & Christof Weinhardt.

Designing Microgrid Energy Markets. Applied Energy, vol. 210, pages 870–880, January 2018.

[Mocanu14] Elena Mocanu, Phuong H. Nguyen, Madeleine Gibescu & Wil L. Kling. *Comparison of Machine Learning Methods for Estimating Energy Consumption in Buildings*. In 2014 International Conference on Probabilistic Methods Applied to Power Systems (PMAPS), pages 1–6, Durham, United Kingdom, July 2014. IEEE.

[Mohsenian-Rad10] Amir-Hamed Mohsenian-Rad, Vincent W. S. Wong, Juri Jatskevich, Robert Schober & Alberto Leon-Garcia. *Autonomous Demand-Side Management Based on Game-Theoretic Energy Consumption Scheduling for the Future Smart Grid*. IEEE Transactions on Smart Grid, vol. 1, no. 3, pages 320–331, 2010.

[Monyei18] Chukwuka G. Monyei & Serestina Viriri. *An Improved Comfort Biased Smart Home Load Manager for Grid Connected Homes Under Direct Load Control*. In Francisco Javier de Cos Juez, José Ramón Villar, Enrique A. de la Cal, Álvaro Herrero, Héctor Quintián, José António Sáez & Emilio Corchado, editeurs, Hybrid Artificial Intelligent Systems, volume 10870, pages 526–536. Springer International Publishing, Cham, 2018. 00000.

[Newsham10] Guy R. Newsham & Brent G. Bowker. *The Effect of Utility Time-Varying Pricing and Load Control Strategies on Residential Summer Peak Electricity Use: A Review*. Energy Policy, vol. 38, no. 7, pages 3289–3296, July 2010.

[Nguyen12] Cuong P. Nguyen & Alexander J. Flueck. *Agent Based Restoration With Distributed Energy Storage Support in Smart Grids*. IEEE Transactions on Smart Grid, vol. 3, no. 2, pages 1029–1038, June 2012.

[Nguyen13] Tuan Anh Nguyen & Marco Aiello. *Energy Intelligent Buildings Based on User Activity: A Survey*. Energy and Buildings, vol. 56, pages 244–257, January 2013.

[Nguyen14] Duong Tung Nguyen & Long Bao Le. *Joint Optimization of Electric Vehicle and Home Energy Scheduling Considering User Comfort Preference*. IEEE Transactions on Smart Grid, vol. 5, no. 1, pages 188–199, January 2014.

[Nigon17] Julien Nigon. *Apprentissage artificiel adapté aux systèmes complexes par auto-organisation coopérative de systèmes multi-agents*. Thèse de doctorat, Université de Toulouse, Toulouse, France, November 2017.

[Nunna12] H. S. V. S. Kumar Nunna & Suryanarayana Doolla. *Demand Response in Smart Distribution System With Multiple Microgrids*. IEEE Transactions on Smart Grid, vol. 3, no. 4, pages 1641–1649, December 2012.

[Olivares14] Daniel E. Olivares, Ali Mehrizi-Sani, Amir H. Etemadi, Claudio A. Canizares, Reza Iravani, Mehrdad Kazerani, Amir H. Hajimiragha, Oriol Gomis-Bellmunt, Maryam Saeedifard, Rodrigo Palma-Behnke, Guillermo A. Jimenez-Estevez & Nikos D. Hatziargyriou. *Trends in Microgrid Control*. IEEE Transactions on Smart Grid, vol. 5, no. 4, pages 1905–1919, July 2014.

[Ouammi15] Ahmed Ouammi, Hanane Dagdougui, Louis Dessaint & Roberto Sacile. *Coordinated Model Predictive-Based Power Flows Control in a Cooperative Network of Smart Microgrids*. IEEE Transactions on Smart Grid, vol. 6, no. 5, pages 2233–2244, September 2015.

[Paetz13] Alexandra-Gwyn Paetz, Thomas Kaschub, Patrick Jochem & Wolf Fichtner. *Load-Shifting Potentials in Households Including Electric Mobility - A Comparison of User Behaviour with Modelling Results*. In 2013 10th International Conference on the European Energy Market (EEM), pages 1–7, Stockholm, Sweden, May 2013. IEEE.

[Papalexopoulos13] Alex Papalexopoulos, Jacob Beal & Steven Florek. *Precise Mass-Market Energy Demand Management Through Stochastic Distributed Computing*. IEEE Transactions on Smart Grid, vol. 4, no. 4, pages 2017–2027, 2013. 00017.

[Parisio11] Alessandra Parisio & Luigi Glielmo. *Energy Efficient Microgrid Management Using Model Predictive Control*. In 2011 50th IEEE Conference on Decision and Control and European Control Conference, pages 5449–5454. IEEE, 2011.

[Piagi06] Paolo Piagi & Robert H. Lasseter. *Autonomous Control of Microgrids*. In Power Engineering Society General Meeting, 2006. IEEE, Lasseter, pages 8–pp. IEEE, 2006.

[Pipattanasomporn09] Manisa Pipattanasomporn, Hassan Feroze & 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. IEEE, 2009.

[Pop18] Claudia Pop, Tudor Cioara, Marcel Antal, Ionut Anghel, Ioan Salomie & Massimo Bertoncini. *Blockchain Based Decentralized Management of Demand Response Programs in Smart Energy Grids*. Sensors, vol. 18, no. 2, page 162, January 2018.

[Pournaras14] Evangelos Pournaras, Matteo Vasirani, Robert E. Kooij & Karl Aberer. *Decentralized Planning of Energy Demand for the Management of Robustness and Discomfort*. IEEE Transactions on Industrial Informatics, vol. 10, no. 4, pages 2280–2289, November 2014.

[Pudjianto07] D. Pudjianto, C. Ramsay & G. Strbac. *Virtual Power Plant and System Integration of Distributed Energy Resources*. IET Renewable Power Generation, vol. 1, no. 1, page 10, 2007.

[Rajurkar85] K. P. Rajurkar & J. L. Nissen. *Data-Dependent Systems Approach to Short-Term Load Forecasting*. IEEE Transactions on Systems, Man, and Cybernetics, vol. SMC-15, no. 4, pages 532–536, July 1985.

[Ramachandran13] Bhuvaneswari Ramachandran, Sanjeev K. Srivastava & David A. Cartes. *Intelligent Power Management in Micro Grids with EV Penetration*. Expert Systems with Applications, vol. 40, no. 16, pages 6631–6640, November 2013.

[Ramchurn11] Sarvapali D. Ramchurn, Perukrishnen Vytelingum, Alex Rogers & Nicholas R. Jennings. *Agent-Based Homeostatic Control for Green Energy in the Smart Grid*. ACM Transactions on Intelligent Systems and Technology, vol. 2, no. 4, pages 1–28, July 2011. 00094.

[Ranade10] Vinayak V. Ranade. *Model and Control for Cooperative Energy Management*. Raytheon BBN Technologies, 2010.

[Revuelta Herrero18] Jorge Revuelta Herrero, Álvaro Lozano Murciego, Alberto López Barriuso, Daniel Hernández de la Iglesia, Gabriel Villarrubia González, Juan Manuel Corchado Rodríguez & Rita Carreira. *Non Intrusive Load*

Monitoring (NILM): A State of the Art. In Fernando De la Prieta, Zita Vale, Luis Antunes, Tiago Pinto, Andrew T. Campbell, Vicente Julián, Antonio J.R. Neves & María N. Moreno, editeurs, Trends in Cyber-Physical Multi-Agent Systems. The PAAMS Collection - 15th International Conference, PAAMS 2017, volume 619, pages 125–138. Springer International Publishing, Cham, 2018. 00000.

[Rohbogner13] Gregor Rohbogner, Ulf Hahnel, Pascal Benoit & Simon Fey. *Multi-Agent Systems' Asset for Smart Grid Applications*. Computer Science and Information Systems, vol. 10, no. 4, pages 1799–1822, 2013.

[Ruiz09] N. Ruiz, I. Cobelo & J. Oyarzabal. *A Direct Load Control Model for Virtual Power Plant Management*. IEEE Transactions on Power Systems, vol. 24, no. 2, pages 959–966, May 2009. 00433.

[Samadi12] Pedram Samadi, Hamed Mohsenian-Rad, Robert Schober & Vincent W. S. Wong. *Advanced Demand Side Management for the Future Smart Grid Using Mechanism Design*. IEEE Transactions on Smart Grid, vol. 3, no. 3, pages 1170–1180, September 2012.

[Sanghvi89] A.P. Sanghvi. *Flexible Strategies for Load/Demand Management Using Dynamic Pricing*. IEEE Transactions on Power Systems, vol. 4, no. 1, pages 83–93, Feb./1989.

[Semenza08] Jan C. Semenza, David E. Hall, Daniel J. Wilson, Brian D. Bontempo, David J. Sailor & Linda A. George. *Public Perception of Climate Change*. American Journal of Preventive Medicine, vol. 35, no. 5, pages 479–487, November 2008.

[Shi16] Guang Shi, Derong Liu & Qinglai Wei. *Energy Consumption Prediction of Office Buildings Based on Echo State Networks*. Neurocomputing, vol. 216, pages 478–488, December 2016.

[Siano14] Pierluigi Siano. *Demand Response and Smart Grids—A Survey*. Renewable and Sustainable Energy Reviews, vol. 30, pages 461–478, February 2014.

[Sicari15] S. Sicari, A. Rizzardi, L.A. Grieco & A. Coen-Porisini. *Security, Privacy and Trust in Internet of Things: The Road*

Ahead. Computer Networks, vol. 76, pages 146–164, January 2015.

[Smil10] Vaclav Smil. Energy transitions: History, requirements, prospects. Praeger, Santa Barbara, Calif, 2010. OCLC: ocn505017554.

[Tavernier46] Édouard Tavernier. *Contrôle à Distance et Automatisé Des Installations Électriques*. La Houille Blanche, no. 2, pages 103–116, March 1946.

[Tsekouras07] George J. Tsekouras, Nikos D. Hatziargyriou & Evangelos N. Dialynas. *Two-Stage Pattern Recognition of Load Curves for Classification of Electricity Customers*. IEEE Transactions on Power Systems, vol. 22, no. 3, pages 1120–1128, August 2007.

[Tushar18] Wayes Tushar, Chau Yuen, Hamed Mohsenian-Rad, Tapan Saha, H. Vincent Poor & Kristin L. Wood. *Transforming Energy Networks via Peer-to-Peer Energy Trading: The Potential of Game-Theoretic Approaches*. IEEE Signal Processing Magazine, vol. 35, no. 4, pages 90–111, July 2018.

[Umeozor16] Evar Chinedu Umeozor & Milana Trifkovic. *Operational Scheduling of Microgrids via Parametric Programming*. Applied Energy, vol. 180, pages 672–681, October 2016. 00013.

[Valgaev16] Oleg Valgaev, Friedrich Kupzog & Harmut Schmeck. *Low-Voltage Power Demand Forecasting Using K-Nearest Neighbors Approach*. In Innovative Smart Grid Technologies-Asia (ISGT-Asia), 2016 IEEE, pages 1019–1024. IEEE, 2016. 00005.

[Van Hertem10] Dirk Van Hertem, Mehrdad Ghandhari & Marko Delimar. *Technical Limitations towards a SuperGrid : A European Prospective*. In 2010 IEEE International Energy Conference, pages 302–309, Manama, Bahrain, December 2010. IEEE.

[Veit13] Andreas Veit, Ying Xu, Ronghuo Zheng, Nilanjan Chakraborty & Katia Sycara. *Multiagent Coordination for Energy Consumption Scheduling in Consumer Cooperatives*. page 8, 2013.

[Vuppala11] Sunil K. Vuppala, Kumar Padmanabh, Sumit Kumar Bose & Sanjoy Paul. *Incorporating Fairness within Demand Response Programs in Smart Grid*. In ISGT 2011, pages 1–9, Anaheim, CA, USA, January 2011. IEEE.

[Wang14] Yu Wang, Shiwen Mao & R.M. Nelms. *Optimal Hierarchical Power Scheduling for Cooperative Microgrids*. In 2014 IEEE 11th International Conference on Mobile Ad Hoc and Sensor Systems, pages 497–498, Philadelphia, PA, USA, October 2014. IEEE.

[Wang17] Zeyu Wang & Ravi S. Srinivasan. *A Review of Artificial Intelligence Based Building Energy Use Prediction: Contrasting the Capabilities of Single and Ensemble Prediction Models*. Renewable and Sustainable Energy Reviews, vol. 75, pages 796–808, August 2017.

[Weller88] H.G. Weller. *Managing the Instantaneous Load Shape Impacts Caused by the Operation of a Large-Scale Direct Load Control System*. IEEE Transactions on Power Systems, vol. 3, no. 1, pages 197–199, Feb./1988.

[Weron07] Rafał Weron. Modeling and forecasting electricity loads and prices a statistical approach. John Wiley & Sons, Ltd., Chichester, 2007. OCLC: 938938870.

[Wicker11] Stephen Wicker & Robert Thomas. *A Privacy-Aware Architecture for Demand Response Systems*. In 2011 44th Hawaii International Conference on System Sciences, pages 1–9, Kauai, HI, January 2011. IEEE.

[Wooldridge02] Michael J. Wooldridge. An introduction to multiagent systems. J. Wiley, New York, 2002.

[Wright07] Andrew Wright & Steven Firth. *The Nature of Domestic Electricity-Loads and Effects of Time Averaging on Statistics and on-Site Generation Calculations*. Applied Energy, vol. 84, no. 4, pages 389–403, 2007.

[Xu11] Lie Xu & Dong Chen. *Control and Operation of a DC Microgrid With Variable Generation and Energy Storage*. IEEE Transactions on Power Delivery, vol. 26, no. 4, pages 2513–2522, October 2011.

[Yang05] H.T. Yang, S.C. Chen & P.C. Peng. *Genetic K-Means-Algorithm-Based Classification of Direct Load-Control*

Curves. IEE Proceedings - Generation, Transmission and Distribution, vol. 152, no. 4, page 489, 2005.

[Zamora10] Ramon Zamora & Anurag K. Srivastava. *Controls for Microgrids with Storage: Review, Challenges, and Research Needs*. Renewable and Sustainable Energy Reviews, vol. 14, no. 7, pages 2009–2018, September 2010. 00306.

[Zhang12] Tiefeng Zhang, Guangquan Zhang, Jie Lu, Xiaopu Feng & Wanchun Yang. *A New Index and Classification Approach for Load Pattern Analysis of Large Electricity Customers*. IEEE Transactions on Power Systems, vol. 27, no. 1, pages 153–160, February 2012.

[Zhang16] Ke Zhang, Yuming Mao, Supeng Leng, Sabita Maharjan, Yan Zhang, Alexey Vinel & Magnus Jonsson. *Incentive-Driven Energy Trading in the Smart Grid*. IEEE Access, vol. 4, pages 1243–1257, 2016.

[Zhao15] Bo Zhao, Meidong Xue, Xuesong Zhang, Caisheng Wang & Junhui Zhao. *An MAS Based Energy Management System for a Stand-Alone Microgrid at High Altitude*. Applied Energy, vol. 143, pages 251–261, April 2015.