DISS. ETH NO. 25404

CHARACTERIZATION AND USE OF FLEXIBLE ELECTRIC ENERGY UNITS

A thesis submitted to attain the degree of
DOCTOR OF SCIENCES OF ETH ZURICH
(DR. SC. ETH ZURICH)

presented by

FABIAN LUKAS MÜLLER
MSC ETH ZURICH
born on 01.08.1986
citizen of Baden, Switzerland

accepted on the recommendation of
Prof. Dr. John Lygeros, examiner
Prof. Dr. Antonio J. Conejo, co-examiner
Dr. Olle Sundström, co-examiner

2018

Fabian Lukas Müller: *Characterization and Use of Flexible Electric Energy Units,* © 2018

ABSTRACT

The main theme of this thesis is the characterization and use of the electric flexibility of populations of electric energy generation and consumption units. The goal of this work is to develop methods and tools that allow us to efficiently quantify and exploit the flexibility of different types of energy units. Electric flexibility is crucial for the safe and reliable operation of any electric power grid and is indispensable for the widespread integration of variable renewable energy resources.

The quantitative description of the flexibility available from populations of energy units is challenging. First, it requires knowledge about the dynamics and constraints of the units, which is not readily available in general. Second, the flexibility of energy units changes over time and can be affected by various sources of uncertainty. Finally, exploiting the flexibility of multiple energy units efficiently can be complex and computationally demanding.

The work in this thesis addresses these challenges by making the following main contributions:

First, we suggest a generic and quantitative description of the electric flexibility of an individual energy unit that is based on the set of all power trajectories the unit can implement over a given time horizon. Based on this concept of flexibility, we develop a procedure to estimate the flexibility of individual units. In contrast to traditional system identification methods, our procedure can be automated and requires only limited and readily available measurement data. This makes it an inexpensive and scalable tool to estimate the flexibility of large populations of energy units.

Second, we apply the flexibility estimation procedure in a real-life Demand-Response demonstration where the flexibility of more than 200 residential buildings with heat pump installations is estimated. We study the potential of the entire population of units to offer Demand-Response services to the power grid and show how the load-reduction potential and the corresponding rebound effect can be precisely quantified and predicted. Experimental results prove the effectiveness of

our approach and motivate the development of advanced methods for harvesting the flexibility of populations of energy units.

Third, it is shown that zonotopes can serve as inner-approximations of the set-based flexibility description of an energy unit. We formulate optimization problems to find optimal zonotopic approximations and evaluate their approximation quality for different types of units. Furthermore, we demonstrate that zonotopic approximations can be aggregated efficiently and that they are convenient for making market and control decisions on the aggregate level. In addition, we present a disaggregation algorithm to distribute aggregate-level control decisions among the individual units of a population in an economically optimal and computationally efficient way.

Finally, we investigate how the flexibility of various energy units can be exploited more efficiently with regard to providing Secondary Frequency Reserves. To this end, we consider multiple flexible energy units and control them in a coordinated manner. A common control approach based on optimal adjustable robust control policies is extended to properly model the power ramp-rate limits of individual units and to take into account the different timescales and lead times of energy and reserve markets. We study the problem of how to interact optimally with energy and reserve markets and investigate the synergy effect that can occur when aggregating energy units with complementary physical properties.

KURZFASSUNG

Das Hauptthema dieser Dissertation ist die Charakterisierung und Nutzung der elektrischen Flexibilität von Populationen elektrischer Erzeugungs- und Verbrauchseinheiten. Das Ziel dieser Arbeit ist die Entwicklung von Methoden und Hilfsmitteln, die uns erlauben die Flexibilität verschiedener Arten von Erzeugungs- und Verbrauchseinheiten zu quantifizieren und nutzbar zu machen. Elektrische Flexibilität ist wichtig für das sichere und zuverlässige Betreiben elektrischer Stromnetze und ist unverzichtbar für die grossflächige Einbindung erneuerbarer Energiequellen.

Die quantitative Beschreibung der Flexibilität von Populationen elektrischer Erzeugungs- und Verbrauchseinheiten ist anspruchsvoll. Erstens setzt sie die Kenntnis der Dynamik und der Beschränkungen elektrischer Erzeugungs- und Verbrauchseinheiten voraus. Entsprechende Informationen sind jedoch oft nicht leicht ermittelbar. Zweitens ändert sich die Flexibilität elektrischer Einheiten mit der Zeit und kann von mehreren Unsicherheitsfaktoren beeinflusst werden. Schliesslich kann die effiziente Nutzung der Flexibilität mehrerer elektrischer Einheiten komplex und rechenintensiv sein.

Die Arbeit dieser Dissertation liefert zu diesen Herausforderungen folgende Hauptbeiträge:

Erstens schlagen wir eine generische und quantitative Beschreibung der elektrischen Flexibilität einzelner elektrischer Erzeugungs- und Verbrauchseinheiten vor, die auf der Menge aller Leistungstrajektorien basiert, welche eine Einheit über einen gegebenen Zeithorizont umsetzen kann. Auf diesem Konzept von Flexibilität aufbauend entwickeln wir ein Verfahren zur Identifikation der Flexibilität einzelner Erzeugungs- und Verbrauchseinheiten. Im Unterschied zu herkömmlichen Methoden der Systemidentifikation kann unser Verfahren automatisiert werden und erfordert nur eine beschränkte Anzahl einfach erhältlicher Messdaten. Dies macht unser Verfahren zu einem kostengünstigen und skalierbaren Hilfsmittel, mit dem die Flexibilität von grossen Populationen elektrischer Einheiten identifiziert werden kann.

Zweitens wenden wir das Verfahren zur Identifikation der Flexibilität in einem realen Demand-Response-Projekt an, in dem die Flexibilität von über 200 Wohnhäusern mit Wärmepumpen identifiziert wird. Wir untersuchen das Potential der gesamten Population zur Erbringung von Demand-Response-Dienstleistungen und zeigen, dass die Lastreduktion und der zugehörige Erholungseffekt genau quantifiziert und vorhergesagt werden können. Wir illustrieren die Wirksamkeit unseres Ansatzes anhand mehrerer Demand-Response-Experimente. Die Resultate motivieren die Entwicklung verfeinerter Methoden zur Nutzung der Flexibilität von Populationen elektrischer Einheiten.

Drittens zeigen wir, dass Zonotopen als Innenapproximation der oben eingeführten, Mengen-basierten Beschreibung von Flexibilität dienen können. Wir formulieren das Finden optimaler zonotopischer Approximationen als Optimierungsproblem und evaluieren die Approximationsqualität von Zonotopen für verschiedene Arten von elektrischen Einheiten. Ausserdem zeigen wir, dass zonotopische Approximationen effizient aggregiert werden können und sich gut eignen, um Markt- und Regelentscheidungen auf aggregierter Ebene zu treffen. Zudem stellen wir einen Disaggregations-Algorithmus vor, welcher Regelentscheidungen ökonomisch optimal und rechnerisch effizient unter den Einheiten einer Population verteilt.

Schliesslich untersuchen wir, wie die Flexibilität verschiedener Erzeugungs- und Verbrauchseinheiten zur Bereitstellung von Sekundärregelleistung effizienter genutzt werden kann. Dazu betrachten wir mehrere Einheiten und regeln sie auf koordinierte Weise. Wir ergänzen einen bekannten Regelansatz, der auf optimalen, anpassbaren und robusten Regelstrategien beruht, um die genaue Modellierung der Beschränkungen von Leistungsänderungsraten einzelner Einheiten. Zudem berücksichtigen wir die verschiedenen Zeitskalen und Vorlaufzeiten von Energie- und Reservemärkten. Wir untersuchen, wie man mit Energie- und Reservemärkten optimal interagiert und studieren den Synergieeffekt, der entstehen kann, wenn Einheiten mit komplementären physikalischen Eigenschaften aggregiert werden.

ACKNOWLEDGMENTS

First and foremost, I would like to thank my supervisor, Prof. John Lygeros, for advising me during my PhD studies. Through his support, guidance, and the freedom he gave me to carry out my own research, John contributed greatly to this work. He has always been supportive and I benefited a lot from the discussion with him and the comments and feedback he provided.

My sincere thanks also go to Dr. Olle Sundström for supervising my studies at IBM Research–Zurich and for agreeing to serve as a co-examiner of this thesis. I am deeply grateful for the numerous discussions with Olle, his helpful comments and inspiring ideas. It was a great pleasure working with Olle in the Smart Grid team.

I would also like to thank Prof. Antonio J. Conejo for his kind willingness to co-examine this thesis.

Many thanks go to Bernhard Jansen for the countless discussions, Bernhard's time and patience to answer myriads of questions, his dedication to our work, and our fruitful collaboration. The beneficial influence that Bernhard's constant supply of the finest Italian espresso had on the successful completion of this thesis should not be underestimated.

I am deeply grateful to Stefan Woerner for our valuable and instructive collaboration. I greatly enjoyed working with Stefan from whom I learned interesting concepts at the intersection of optimization and control.

Special thanks go to Jácint Szabó who showed me how mathematicians think. I am deeply grateful for all the discussions with Jácint; they always revealed new ways forward and enlightened me repeatedly.

I would like to thank Axel Tanner and Metin Feridun for our common lunch breaks, the Binding field walks, and the refreshing and inspiring exchange of ideas.

Also, I would like to thank Marco Laumanns, Ulrich Schimpel and Douglas Dykeman for their time and constant support.

Furthermore, I enjoyed interacting with Tim Nonner, Carl Binding, Dieter Gantenbein, Mike Nidd, and Michael Baentsch.

Fabian Müller
June 2018

CONTENTS

1 INTRODUCTION 1
- 1.1 Background & Motivation 1
- 1.2 Contributions 7
- 1.3 Thesis outline 9
- 1.4 Publications 10

2 DESCRIPTION OF FLEXIBILITY 13
- 2.1 Definitions of flexibility 13
- 2.2 Model of an energy unit 15
- 2.3 Description of flexibility 17

3 ESTIMATION OF FLEXIBILITY 21
- 3.1 Introduction 21
- 3.2 System modeling 22
- 3.3 System identification 24
- 3.4 Estimation of flexibility 33
- 3.5 Conclusion & Future work 36

4 DEMONSTRATION OF DEMAND RESPONSE 37
- 4.1 Introduction 37
- 4.2 Experimental setup 40
- 4.3 Population modeling 44
- 4.4 Experimental results 48
- 4.5 Conclusion & Future work 54

5 AGGREGATION OF FLEXIBILITY 55
- 5.1 Introduction 55
- 5.2 Problem formulation 57
- 5.3 Feasible sets and their aggregation 58
- 5.4 Computing zonotopic approximations 61
- 5.5 Conclusion & Future work 69
- 5.6 Appendix: Facets of PE-zonotopes 70

6 DISAGGREGATION OF FLEXIBILITY 73
6.1 Introduction 73
6.2 Problem formulation 74
6.3 Cost of flexibility 76
6.4 Disaggregation algorithm 80
6.5 Use-cases for zonotopic flexibility 86
6.6 Conclusion & Future work 91

7 ENERGY AND RESERVE BIDDING 93
7.1 Introduction 93
7.2 Problem formulation 95
7.3 Energy baseline and power reference 99
7.4 Description of flexibility 104
7.5 Solving the bidding problem 109
7.6 Conclusion & Future work 114
7.7 Appendix: Bounds on discrete-time state dynamics . . . 115

8 SYNERGY EFFECT OF AGGREGATION 117
8.1 Introduction 117
8.2 Adjustable power reference 119
8.3 The bidding problem 121
8.4 Synergy I: Battery and freezer 124
8.5 Synergy II: Battery and turbine 133
8.6 Conclusion & Future work 135

9 DISCUSSION 137
9.1 Description and estimation of flexibility 137
9.2 Demonstration of Demand Response 140
9.3 Aggregation and disaggregation of flexibility 143
9.4 Using flexibility for regulation services 146

BIBLIOGRAPHY 149

NOTATION 169

ACRONYMS 177

LIST OF FIGURES

Figure 2.1 Typical feasible set of a PE-unit 19
Figure 3.1 Energy dynamics of a TCL 23
Figure 3.2 Aggregate power consumption of a building . . 25
Figure 3.3 Histogram of power changes 26
Figure 3.4 Estimation of exact HP switching time 28
Figure 3.5 Estimates of charge rate, loss rate, and duty cycle 31
Figure 3.6 Visualization of flexible energy 35
Figure 4.1 Market-based direct-control DR setup 41
Figure 4.2 Building-level communication and control setup 43
Figure 4.3 Expected maximum aggregate load reduction . 45
Figure 4.4 Expected maximum response power 46
Figure 4.5 Expected maximum rebound power 47
Figure 4.6 Results from load-reduction experiment 49
Figure 4.7 Results from rebound-damping experiment . . . 51
Figure 4.8 Achieved aggregate load-reduction values . . . 52
Figure 5.1 Aggregation setup 57
Figure 5.2 Polytopic feasible set and approx. zonotope . . 60
Figure 5.3 Feasible set constraints approximation 63
Figure 5.4 Evaluation of offset differences 66
Figure 5.5 PE-polytope and inscribed zonotope 67
Figure 6.1 Aggregation and disaggregation setup 75
Figure 6.2 Example of flexibility cost function 77
Figure 6.3 Aggregation of cost functions 79
Figure 6.4 The cost of offering regulation power 89
Figure 7.1 Piece-wise affine power reference 100
Figure 7.2 Structures of policy matrices 103
Figure 7.3 Ramp-rates of target power trajectory 106
Figure 7.4 Auxiliary state dynamics 108
Figure 7.5 Max. SFR capacity for different ramp-rates . . . 112
Figure 8.1 Unit response to SFR activation 130

Figure 8.2 Synergy factors for freezer and batteries 132

LIST OF TABLES

Table 4.1 Load-reduction and rebound prediction error statistics 53
Table 5.1 Parameters of PE-units 64
Table 5.2 PEV parameters 68
Table 6.1 Solving time of LP-based disaggregation 86
Table 6.2 Solving time of SG-based disaggregation 86
Table 6.3 Baseline costs 91
Table 7.1 SFR capacities and power ramp-rates 113
Table 8.1 Freezer warehouse parameters 127
Table 8.2 Results of battery-freezer aggregation 133
Table 8.3 Results of battery-steam turbine aggregation . . 134

LIST OF ALGORITHMS

Algorithm 6.1 Projected subgradient algorithm 83
Algorithm 6.2 Distribution of $\beta^{(\mathcal{J})*}$ 84

1 INTRODUCTION

1.1 BACKGROUND & MOTIVATION

This section provides an overview of recent and ongoing developments in electric power systems that motivate the work presented in this thesis.

1.1.1 *Decarbonization of the energy sector*

Electric power systems are currently undergoing radical changes driven by political decisions, economic considerations, and technical innovations. A key driver is the desire to significantly reduce global greenhouse gas (GHG) emissions, as expressed by most United Nations member states in the *Kyoto Protocol* [1] and the *Paris Agreement* [2]. For example, the countries of the European Union (EU) target a reduction of GHG emissions by 20% by 2020, and by 80–95% by 2050 compared with the emission levels of 1990 [3].

Because the energy sector currently produces the lion's share of man-made GHG emissions, meeting these reduction targets requires a fundamental transition from an energy generation heavily reliant on fossil fuels to a highly efficient low-carbon energy system [4]. However, uncertainty remains about how this transition could be achieved in a timely and cost-efficient manner, and various transitions pathways have been proposed [3].

Recently, two key measures have emerged to drive this transition, namely, to improve energy efficiency, and to increase the share of energy from renewable energy sources (RES) like wind, solar, hydro and biomass. Energy efficiency targets are now in place in at least 149 countries [5]. The EU, for instance, is aiming for a 20% improvement of energy efficiency by 2020 [6]. Similarly, various countries have set goals regarding the deployment of RES: 145 parties that submitted contributions under the *Paris Agreement* referred to renewable energy action to mitigate and adapt to climate change [5]. The EU is committed

to a 20% share of renewable energy in the overall energy consumption by 2020 [7]. The EU *Energy Roadmap 2050* is even more ambitious and involves a share of RES of up to 75% in gross final energy consumption and 97% in electricity consumption by 2050 depending on the scenario considered [3].

1.1.2 *Characteristics and impacts of renewable energy*

RES have experienced a remarkable evolution over the past decade involving significant technological advances and sharp cost reductions. Supported by favourable regulations and financial incentives, renewable generation has grown rapidly, and accounted for almost two-thirds of all new generating capacity installed globally in 2016 [5]. In 2015, RES provided about 23.5% of all electricity generated, the bulk of which came from hydro, followed by wind, biomass and solar photovoltaic (PV) [5]. In particular, the installed capacity of wind and solar PV, collectively referred to as variable renewable energy sources (VRES), has increased in many countries, exceeding a share of 45% of the total electricity generation in Denmark, and 20% in Lithuania, Ireland, Spain and Germany [5], for example. The potential of RES is immense compared with current global energy use [8, 9], and RES are expected to play a crucial role in the decarbonization of the energy sector. However, the integration of RES into existing power systems is far from trivial.

In traditional power systems, electric energy is produced centrally by large-scale generation units located in the vicinity of areas with high electricity demand, such as urban and industrial areas. Generation units are connected to the high-voltage transmission grid and provide electricity to the consumers attached to the low-voltage distribution grid. This clear separation of electricity generation and consumption results in predominantly unidirectional power flows from high- to low-voltage levels of a power grid [10]. The electricity generated by conventional generation units, such as thermal and hydro-power plants, can be scheduled precisely and reliably to meet the total demand, which is assumed to be uncontrollable and stochastic, but can be forecast accurately on the level of the transmission grid. These traditional power systems must undergo fundamental structural and operational changes if large shares of RES are to be integrated.

First, RES come in different shapes and sizes. Large-scale installations, in particular solar PV, wind, and hydro-power plants, are only profitable if installed at locations where the corresponding primary energy is available abundantly. Such locations do not, in general, coincide with areas of high demand. Thus, linking resource-rich areas with demand may require extending the high-voltage infrastructre of a power grid [5]. In contrast, small-scale RES, in particular residential solar PV and wind turbines, can be widely distributed and are connected directly to the low-voltage distribution grid. The collective generation of small-scale RES can put a high stress on the low-voltage power-grid infrastructure, and can even lead to reverse power flows. To avoid line congestions and component overloads, it can be necessary to reinforce distribution grids and install more advanced monitoring and control equipment [11].

Second, the electricity generation of VRES, such as wind and solar PV and, at a slower timescale, hydro power, is determined by weather conditions. As a consequence, it is inherently uncertain and fluctuates according to the instantaneous availability of the primary source of energy. In addition, the power output of VRES is not fully controllable but can only be curtailed. These key characteristics of VRES – uncertainty, intermittency and limited controllability – pose serious challenges to System Operators (SOs) who are responsible for keeping the electricity supply and demand in a power grid balanced at all times. Increasing the share of VRES in a power grid can lead to larger and more rapid variations of the net load (defined as the difference between forecast load and expected generation) whose compensation by traditional power plants requires fast power ramp-rates and can invlove frequent start-up and run-down cycles of generation units [12, 13].

Third, the integration of RES can have significant impacts on electricity markets [14]. Most importantly, the variable operating cost of RES, and in particular of solar PV, wind, and hydro power, is lower than the operating costs of thermal generation units [15, 16]. Thus, as their penetration increases, VRES displace more expensive thermal generation and decrease electricity prices during hours in which VRES operate [14, 16]. Lower wholesale electricity prices affect the profitability of conventional generation units, such as coal- and gas-fired power plants, which have traditionally been providing the power reserves necessary to balance mismatches between the supply and demand of electricity.

If such power plants cease operation, large volumes of highly reliable power reserves become unavailable. In addition, the total rotational inertia of the power system decreases as large turbines are disconnected from the grid. This affects the power system stability and makes frequency control more challenging because the frequency dynamics of a power system become faster as the system's rorational inertia decreases [17].

1.1.3 *The need for flexibility*

The flexibility of a power system describes the extent to which it can adjust the generation and consumption of electric energy in response to variability, expected or otherwise [18]. With growing shares of VRES, both the quantity and the quality of the flexibility required for the safe and reliable operation of a power grid are changing: the additional variability and intermittency introduced by VRES must be balanced continuously, and the discontinuation of conventional sources of flexibility must be compensated. Characterizing sufficient flexibility requirements of a power system for different levels of VRES penetration has been the focus of many studies, see [19, 20] for overviews. It is expected that the amount of flexibility required is increasing with growing shares of VRES [21–23], and that this flexibility must be accessible at faster timescales, with shorter lead-times [17, 21, 24, 25], and at different levels and topological regions of a power grid [26, 27].

The key to the successful integration of VRES is the availability of adequate sources of flexibility [18, 28]. The total flexibility of a power system consists of a number of factors. Traditionally, large-scale and quickly dispatachable generation units, such as coal-, gas-, and pumped hydro-power plants, have been providing the flexibility required to balance variable consumption, compensate for power-plant outages, and deal with other contingencies [29]. Other sources of flexibility are interconnections among adjacent power grids, and energy storage systems [29].

However, a significant, yet under-exploited, source of flexibility consists of various types of distributed, small-scale energy units, on both the demand and the supply side. Different concepts and definitions of the flexibility of individual energy units have been proposed in the

literature. Here, we restrict ourselves to the *electric* flexibility, which we define as the ability of an electric energy unit to adjust its electric power in- and/or outflow from/to the power grid according to an external signal. Sources of small-scale electric flexibility include heating, ventilation, and cooling systems, plug-in electric vehicles (PEVs), stationary batteries, various household appliances, such as dishwashers and washing machines, as well as micro-generation units and flywheels. The flexibility of these Distributed Energy Units (DEUs) can be used to complement, or even substitute, the flexibility that has traditionally been provided by large-scale supply-side units.

The management of DEUs for the purpose of providing services to transmission or distribution grid operators is referred to as Demand Response (DR). Historically, DR has included only consumption units. However, we suggest to extend this term to include any flexible energy unit independent of whether it can be assigned uniquely to either the supply or the demand side. This broader definition of DR is reasonable because several types of flexible energy units, such as batteries, capacitors, and flywheels, are well suited for DR and can act as both generation and consumption units. DR is considered a key enabler of high VRES penetration because it can be used for different purposes such as peak load shaving [30], valley filling [31], or the provision of other dedicated flexibility services, such as frequency regulation [32–34], for example.

DR schemes can be sorted into two main categories, namely, direct control and indirect control [35]. Indirect control refers to a setup in which an incentive signal, such as a forecast of the electricity price, is broadcast to the participating units which are expected to adjust their electricity consumption or generation accordingly. The fact that only a single incentive signal needs to be broadcast and that the control of the electric devices is the responsibility of their operators makes indirect control a simple, privacy-preserving and scalable DR approach. However, the response sensitivity of the units with regard to the incentive signal is unknown and must be estimated.

Direct control, on the other hand, refers to the case where energy units are controlled directly via a control signal that is applied to individual units. Direct control requires two-way communication between the controller and each unit to send control commands and receive feedback.

The main advantage of direct control is that it allows the controller to coordinate the DR of a group of units. This makes direct control an accurate and versatile method to provide different services [35, 36].

1.1.4 *Smart Grids and Aggregators*

The exploitation of the flexibility of small-scale DEUs can be challenging for several reasons. First, to use the flexibility of a unit, the electricity consumption or generation of the unit must be controllable. This usually requires retrofitting the unit with communication and control infrastructure, which can be time-consuming and costly. Second, flexible units, such as heating and cooling systems, generally serve a primary purpose other than providing flexibility to the power grid. Thus, adjusting the electricity consumption or generation of such units must never interfere with their primary purpose, such as keeping the indoor air temperature of a building within comfortable limits, for example. Finally, all physical and operational constraints of a unit must be respected when predicting, trading and activating its flexibility. That is, to safely and efficiently exploit the flexibility of a unit, its dynamics and constraints must be known or estimated.

The difficulties related to the use of the flexibility of DEUs are mainly rooted in the lack of monitoring, communication and control infrastructure. To alleviate these difficulties, the concept of *Smart Grids* has emerged, referring to a vision of novel electricity systems that leverage the potential of digital information and communication technology to implement advanced monitoring, communication, and control techniques with the aim of improving the reliability, security, and efficiency of the electricity system [37]. Smart Grids are expected to provide the measurement data and control capabilities necessary to identify and utilize the flexibility of DEUs.

To achieve a significant impact on the grid level, gain access to energy and reserve markets, benefit from synergy effects, and reduce the complexity of planning, trading, and control, it is crucial to aggregate the flexibility of large numbers of DEUs, and to manage their electricity generation and consumption in a coordinated way [38–41]. In direct-control schemes, so-called Aggregators (AGs) [42, 43] perform the tasks of identifying, aggregating and controlling the flexibility of a

group of DEUs, sometimes referred to as *Virtual Power Plants* [44]. An AG concludes contracts with individual DEUs, defining the way of communication and control between the DEU and the AG, the specifics of the flexibility offered, as well as the details on how the DEU is reimbursed for the flexibility provided [42, 45]. Thus, an AG acts as an intermediary between a group of DEUs, electricity markets, and the SO.

In direct-control DR schemes, each DEU is controlled individually. Consequently, the control decisions made by the AG on the aggregate level must be distributed among the individual units. The disaggregation task can be challenging because an AG in general will try to find disaggregation policies that *i)* are optimal with regard to particular objectives, such as minimizing the operating costs of the DEUs or the reimbursements paid for activating their flexibility, or *ii)* activate the DEUs according to certain priorities, such as energy needs or prevention of short-cycling. Simultaneously, the disaggregation policies must respect the physical and operational constraints of each individual unit. Thus, the problem of disaggregating control decisions among individual units can be complex and computationally demanding.

1.2 CONTRIBUTIONS

This thesis focuses on the direct control of multiple flexible DEUs. We consider the four key tasks of an AG, namely, to first estimate the flexibility of individual DEUs, to subsequently compute the flexibility available from a group of units, to optimally exploit this aggregate flexibility when making decisions on energy and flexibility markets, and finally to distribute the aggregate-level control decisions in an optimal and feasible way. More specifically, the main contributions of this thesis are:

1. *We develop a procedure to estimate the flexibility of a population of thermostatically controlled DEUs (Chapter 3, Paper 1).* The procedure identifies the dynamics and constraints of individual energy units based on which their flexibility can be determined. The identification procedure can be automated and requires only limited, readily available measurement data. This makes it an inexpensive and scalable tool for AGs. The procedure is applied in practice to esti-

mate the flexibility of residential buildings with heat pump (HP) installations.

2. *We show that the coordinated direct control of populations of flexible DEUs can provide precise DR services to the power grid (Chapter 4, Paper 2).* Based on the flexibility estimated for more than 200 residential buildings with HPs, we show how their aggregate flexibility can be predicted accurately and used for providing DR services to the power grid. Experiments prove the effectiveness of our approach and motivate the development of advanced methods for harvesting the flexibility of populations of energy units.

3. *We derive inner-approximations of the flexibility of individual DEUs that are well suited for aggregation and disaggregation purposes (Chapter 5, Paper 3).* The approximations are based on zonotopes, a subclass of polytopes, that can reproduce important features of the original flexibility of an energy unit. We formulate optimization problems to find optimal zonotopic approximations and evaluate their approximation quality for different types of approximation quality measures and energy units. In addition, it is shown that the zonotopic flexibility can be aggregated efficiently for large numbers of units and that it can be used conveniently by AGs to make control and market decisions.

4. *We present a disaggregation procedure that distributes aggregate-level control decisions among the individual units of a population (Chapter 6, Paper 4).* The procedure relies on zonotopic descriptions of the flexibility and flexibility cost functions when distributing aggregate-level control decisions in an economically optimal and computationally efficient way. The procedure is applied to the problem of distributing energy reference schedules among DEUs.

5. *We study the problem of an AG to maximize its economic profit from trading on energy markets while simultaneously offering Secondary Frequency Regulation (SFR) (Chapter 7, Paper 5).* We extend an existing problem formulation, which is based on robust optimization and affinely adjustable control policies, to take power ramp-rate constraints into account explicitly and to model all the different timescales of the energy and reserve markets involved. This

makes our approach applicable to various types of flexible DEUs that operate in real market settings.

6. *We investigate the synergy effect that can occur when multiple flexible DEUs with complementary physical properties are controlled in a coordinated manner (Chapter 8, Paper 6).* Based on the problem formulation derived in Chapter 7, we demonstrate that aggregations of select DEUs can offer significantly larger amounts of SFR than the units could individually. This synergy effect is a strong motivation for the direct control of aggregations of flexible DEUs.

1.3 THESIS OUTLINE

This thesis comprises theoretical and practical work. In the theoretical part (Chapter 2–Chapter 6), we build the foundation for describing, estimating, aggregating, and disaggregating the flexibility of DEUs. The practical part (Chapter 4, Chapter 7, and Chapter 8) illustrates two use-cases of flexibility, namely, to offer load-reduction services, and to exploit flexibility when trading energy and SFR services simultaneously.

More precisely, Chapter 2 introduces the concept, the definition, and the quantitative description of the electric flexibility of an individual energy unit that underlies this work. A flexibility estimation procedure is presented in Chapter 3 and used to estimate the flexibility of residential buildings with HP installations. Based on the flexibility estimated, Chapter 4 demonstrates the potential of a population of flexible DEUs to provide DR services in a real-life setting and motivates the development of advanced methods for harvesting this potential. In Chapter 5, inner-approximations of the flexibility of individual DEUs are introduced that can be aggregated efficiently and that lie at the core of the disaggregation algorithm derived in Chapter 6. Chapter 7 deals with the economically optimal use of flexibility in energy and SFR markets. In Chapter 8, we study the synergy effect that can occur when the flexibility of multiple DEUs is aggregated and controlled in a coordinated manner with the aim of offering SFR services. We conclude and provide suggestions for future work in Chapter 9.

1.4 PUBLICATIONS

The following publications are included in parts or in an extended version in this thesis:

1. F. L. Müller, B. Jansen, and O. Sundström. "Autonomous estimation of the energetic flexibility of buildings." In: *Amer. Control Conf.* IEEE, 2017, pp. 2713–2718.

2. F. L. Müller and B. Jansen. "Large-scale demonstration of precise demand response provided by residential heat pumps." In: *Submitt. to Energy* (2018), pp. 1–8. URL: http://arxiv.org/abs/1806.07670.

3. F. L. Müller, O. Sundström, J. Szabo, and J. Lygeros. "Aggregation of energetic flexibility using zonotopes." In: *54th IEEE Conf. Decis. Control*. IEEE, 2015, pp. 6564–6569.

4. F. L. Müller, O. Sundström, J. Szabo, and J. Lygeros. "Aggregation and disaggregation of energetic flexibility from distributed energy resources." In: *IEEE Trans. Smart Grid* (2017), pp. 1–10. arXiv: arXiv:1705.02815.

5. F. L. Müller, S. Woerner, and J. Lygeros. "Ramp-rate-constrained bidding of energy and frequency reserves in real market settings." In: (2018), pp. 1–8. URL: https://arxiv.org/abs/1804.03892.

6. F. L. Müller, S. Woerner, and J. Lygeros. "Unlocking the potential of flexible energy resources to help balance the power grid." In: *Submitt. to IEEE Trans. Smart Grids* (2018), pp. 1–8. URL: http://arxiv.org/abs/1806.08237.

The following publications were part of my PhD research, are however not covered in this thesis.

7. A. Caracas, F. L. Müller, O. Sundström, C. Binding, and B. Jansen. "VillaSmart : Wireless sensors for system identification in domestic buildings." In: *IEEE PES Innov. Smart Grid Technol. Eur.* IEEE, 2013, pp. 1–5.

8. F. L. Müller, O. Sundström, C. Binding, and D. Dykeman. "Power reference tracking of a large-scale industrial freezer system for ancillary service delivery." In: *IEEE PES Innov. Smart Grid Technol. Eur.* IEEE, 2013, pp. 1–5.

9. F. L. Müller, C. Binding, O. Sundström, and M. Bengsch. "Minimum-cost charging of electrical storage heaters." In: *IEEE Int. Conf. Smart Grid Comm.* IEEE, 2014, pp. 740–745.

2 DESCRIPTION OF FLEXIBILITY

This chapter reviews different concepts of electric flexibility and introduces the definition of flexibility that underlies the work presented in this thesis. A quantitative description of the flexibility of individual energy units is proposed that is applicable to different types of units and that is well suited for control and optimization purposes. The description also serves as the starting point for the aggregation and disaggregation of flexibility presented in Chapter 5 and Chapter 6, respectively.

2.1 DEFINITIONS OF FLEXIBILITY

In the context of electric power grids, the term *flexibility* has been used ambiguously to describe certain properties of electric power systems on the one hand, and of individual electric energy units on the other hand. The flexibility of a power system refers to the ability of the system to respond to unforeseen fluctuations in electric energy demand and supply, as well as to power plant outages and other contingencies [5, 29]. This *system flexibility* has been studied extensively both for existing power systems [19, 54–56] and for future systems with higher shares of variable renewable generation [20, 24, 57, 58]. The flexibility of a power system is made up of the cumulative flexibility available from all the electric energy units connected to the power grid. The characterization and use of the flexibility of individual and of populations of these energy units is the focus of this thesis. This *unit flexibility* refers to the ability of an electric energy unit to adjust the volume or timing of its electric energy intake from or of its energy output to the power grid according to an external signal.

A wide range of measures have been proposed to quantify the flexibility of energy units. Among them are flexibility metrics that are designed for particular applications, such as specific DR products. In [59], the flexibility of a unit is determined by its load-shifting potential,

i.e., by the amount of electric energy consumption or generation that can be deferred by a certain amount of time. In contrast, flexibility has been defined as the maximum possible upward and downward deviation from a predefined power trajectory over a given time interval [38, 60–62]. A generalization of this concept, a so-called *asymmetric block bid* [63], captures the flexibility of a unit by means of possible deviations from a power reference, and also specifies the rebound effects that are characteristic for thermal loads [63, 64].

Other approaches to quantify flexibility focus on the temporal properties of energy units. The availability of units is taken into account explicitly in [41, 65] by specifying a time interval during which a unit is connected to the power grid and its flexibility can be accessed. Modeling the availability of flexibility is particularly important in the case of PEVs whose flexibility is available only when the vehicles are connected to the power grid. Similarly, the flexibility of various household appliances has been described by the possible deviations from a given load profile together with the range of admissible start times [66].

In this work, we seek a description of flexibility with the following properties. First, it must capture all the physical and operational characteristics and constraints of the underlying unit in order to guarantee the unit's safe and reliable operation. For instance, the proper consideration of power ramp-rate constraints has been identified as being crucial in power systems with rapidly fluctuating renewable generation [21, 25, 67]. Second, the description of flexibility should be generic so it can capture the flexibility of different types of energy units [41, 68, 69]. This is a necessary prerequisite for AGs to be able to aggregate the flexibility of multiple heterogeneous units and to exploit potential synergy effects [70, 71]. Finally, the flexibility of a unit should be described in a way that is independent of the purpose for which the flexibility will be used, i.e., independent of the type of flexibility services and products that are to be offered.

The approach followed here is to describe the flexibility of an individual energy unit by the set of all power trajectories that are feasible with regard to the dynamics and constraints of the unit over a finite time horizon. Such a set-theoretic description is widely adopted. It has been used to characterize the flexibility of a thermostatically controlled load (TCL) by describing it as a *generalized battery* with power and energy

constraints [33]. The concept has been extended to avoid short-cycling behavior [72], and has provided the basis for estimating the flexibility of energy units, such as heating and cooling systems [73], and for aggregating the flexibility of multiple units [74]. Related to the concept of generalized batteries is the idea of *load plasticity* introduced in [75]. It proposes a description of flexibility that is based on a simplified battery model that, in addition to the concept of a generalized battery, takes into account the temporal availability of a unit [75]. This makes the description particularly appropriate for characterizing the flexibility of PEVs.

Other set-based descriptions of flexibility focus on capturing the response dynamics of energy units with regard to external signals. This is crucial in cases in which flexibility is used to compensate for the rapidly fluctuating generation of VRES. *Flexibility envelopes* extend the generalized battery model by explicitly taking into account power ramp-rate limits and response delays of energy units [76, 77]. Similarly, the *power nodes* framework represents the dynamics of various types of DEUs by a battery that is subject to power, energy, and ramp-rate constraints [55]. The flexibility of a unit during a discrete time step is then described by the set of triples (ramp-rate, power, energy) that satisfy aforementioned constraints [56, 78]. This makes it possible to compare the flexibility available in a power system with the flexibility requirements imposed by particular portfolios of VRES [25].

The set-theoretic description of flexibility adopted in this work combines and extends several properties of the approaches outlined above.

2.2 MODEL OF AN ENERGY UNIT

Consider an individual energy unit over a finite time horizon comprising N time intervals each of duration t_s and indexed by k. The unit is subject to the following constraints.

1. *Power constraints:* In general, the power $p(t)$ the unit can draw from ($p > 0$) or feed into ($p < 0$) the power grid at time t is bounded. For all $t \in [kt_s, (k+1)t_s)$, $k = 0, \ldots, N-1$, we require that

$$\underline{p}_k \leq p(t) \leq \bar{p}_k, \tag{2.1}$$

where $\underline{p}, \bar{p} \in \mathbb{R}^N$ denote the piece-wise constant bounds on power.

2. *Power ramp-rate constraints:* Often, there are limits on the rate at which power can vary. Ramp-rate constraints play an important role, in particular when providing ancillary services for the purpose of balancing rapidly varying renewable generation. For all $t \in [kt_s, (k+1)t_s)$, $k = 0, \ldots, N-1$, we require that

$$\underline{r}_k \leq dp(t)/dt \leq \bar{r}_k, \tag{2.2}$$

where $\underline{r}, \bar{r} \in \mathbb{R}^N$ denote the piece-wise constant power ramp-rate limits.

3. *State constraints:* Various types of energy units have internal states that are driven by the electric power generation or consumption of the unit. Examples are the state of charge (SOC) of a battery, the temperature of a heating system, or the water level of the reservoir of a pumped hydro-power plant. We restrict ourselves to energy units whose single state $x(t) \in \mathbb{R}$ is governed by the linear time-invariant dynamics

$$dx(t)/dt = ax(t) + bu(t) + cp(t), \tag{2.3}$$

where $u(t)$ summarizes all exogenous (uncontrollable) inputs, such as the weather conditions in the case of heating systems, or the trips of a PEVs. The dynamics (2.3) are characterized by the parameters $a, b, c \in \mathbb{R}$. Many types of energy units, such as pumped hydro-power plants, batteries, heating and cooling systems, capacitors, and flywheels, owe their flexibility to an energy buffer. In this case, the state $x(t)$ represents the level of this buffer, $a \leq 0$ can be interpreted as the self-dissipation rate, and b, c determine the conversion efficiency of exogenous inputs and electric energy into buffered energy, respectively. Our sign convention implies $c \geq 0$. The generalization to multi-state systems is straightforward.

We consider piece-wise constant state constraints and require for all $t \in [kt_s, (k+1)t_s)$, $k = 0, \ldots, N-1$, that

$$\underline{x}_k \leq x(t) \leq \bar{x}_k, \tag{2.4}$$

where $\underline{x}, \bar{x} \in \mathbb{R}^N$ denote the physical limits of the system state. The initial state $x(0)$ is denoted by x_0.

In the following, we consider power trajectories $p(t)$ and exogenous inputs $u(t)$ that are piece-wise constant, i.e., $p(t) = p_k$ and $u(t) = u_k$ for $t \in [kt_s, (k+1)t_s)$, $k = 0, \dots, N-1$. In this case, the power, power ramp-rate, and state constraints (2.1), (2.2), and (2.4), respectively, are affine in p:

$$\underline{p}_k \leq p_k \leq \bar{p}_k, \; k = 0, \dots, N-1, \tag{2.5}$$

$$\underline{r}_k \leq (p_k - p_{k-1})/t_s \leq \bar{r}_k, \; k = 1, \dots, N-1, \tag{2.6}$$

$$\underline{x}_{k+1} \leq a_d x_k + b_d u_k + c_d p_k \leq \bar{x}_{k+1}, \; k = 0, \dots, N-1, \tag{2.7}$$

where a_d, b_d, c_d are the discrete-time counterparts of a, b, c. A broad family of energy units can be modeled according to (2.5)–(2.7). Among them are storage-like loads [79, 80], deferrable loads [43], and TCLs [81]. Because the majority of energy units cannot implement piece-wise constant power trajectories in practice, we study piece-wise affine and continuous power trajectories in Chapter 7. However, the description of flexibility introduced in the following is applicable to both types of power trajectories.

2.3 DESCRIPTION OF FLEXIBILITY

In this work, the electric flexibility of an energy unit is described by the set of discrete-time power trajectories $p \in \mathbb{R}^N$ that satisfy all the physical and operational constraints of the unit. We refer to this set as the *feasible set*, denote it by $\mathcal{P} \subset \mathbb{R}^N$, and make the following definition.

Definition 1 (Flexible Energy Unit (FEU))**.** *An energy unit with feasible set $\mathcal{P}$ is electrically flexible if and only if there exist at least two distinct and feasible power trajectories, i.e., $\exists p, q \in \mathcal{P} : p \neq q$.*

The description of flexibility by means of feasible sets offers a generic, quantitative, and application-independent description of the electric flexibility of different types of FEUs. The ability to properly model inter-temporal constraints, such as energy and power ramp-rate constraints, is one of its key features. This set-based description of flexibility is particularly convenient for the family of energy units described above whose power, power ramp-rate, and state constraints can be written according to (2.5)–(2.7), respectively. Because these constraints are affine

in the power values, the feasible set is a *compact convex polytope* and can be written as

$$\mathcal{P} = \{p \in \mathbb{R}^N : Ap \leq b\}, \tag{2.8}$$

where (A, b) summarize the constraint matrices and the limits of (2.5)–(2.7), respectively. Compact convex feasible sets have proven to be well suited for control [62, 82, 83] and optimization purposes [34, 79, 84]. In addition, they have served as the starting point for aggregating and disaggregating the flexibility of multiple energy units [40, 43, 80, 85, 86]. However, there also exist energy units whose feasible sets are not convex. Examples are units subjected to run-time or state-dependent constraints, units with discrete power levels, such as on/off-loads, or units whose state dynamics cannot be linearized.

An important family of energy units are those that can be modeled as an ideal battery. The battery's single state is its SOC and follows the dynamics (2.3) with parameters $a = b = 0$ and $c = 1$. That is, the state constraints (2.7) can be interpreted as energy constraints. In addition, the battery must respect the power constraints (2.5), but can neglect the ramp-rate constraints (2.6). Consequently, the feasible set of an ideal battery can be represented by a convex compact polytope (2.8). We refer to energy units behaving like ideal batteries as *PE-units*, and call their feasible sets *PE-polytopes*. The family of PE-units includes ideal batteries, capacitors, and flywheels, among others.

Figure 2.1 shows the feasible set of a typical PE-unit over two time steps. The power constraints define an axis-aligned box (dotted) whose lower left and upper right vertices are cut off by the energy constraints (dashed). The remaining feasible set $\mathcal{P}$ is a compact convex polytope. It properly incorporates inter-temporal constraints: the flexibility available at time step $k > 0$ depends on the power trajectory implemented by the unit up to time k, i.e., $\{p_0, \ldots, p_{k-1}\}$. For example, assume that during the first time step the unit consumed p_0^* units of power. Then, the power that can be consumed during the second time step, p_1, is restricted to the range $[\underline{p}_1, \bar{p}_1^*]$, where the new upper bound $\bar{p}_1^*$ is dictated by the upper energy bound.

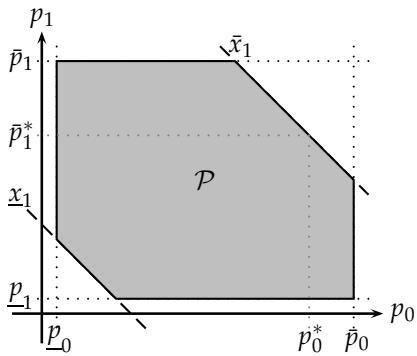

Figure 2.1: A typical feasible set $\mathcal{P}$ of a PE-unit in two dimensions. The power constraints define an axis-aligned box (dotted) whose lower left and upper right vertices are cut off by the state (energy) constraints (dashed). The resulting feasible set $\mathcal{P}$ is a compact convex polytope.

3 ESTIMATION OF FLEXIBILITY

In this chapter, we consider the problem of estimating the parameters of the flexibility description introduced in Chapter 2. We consider residential buildings with HP installations and identify the parameters that determine the dynamics and constraints of the energy unit based on smart meter (SM) and outdoor air-temperature measurement data. Given the parameters identified, the flexibility of an FEU can be estimated and predicted.

3.1 INTRODUCTION

The fundamental prerequisite for exploiting the flexibility of FEUs in direct-control schemes is that the flexibility of each participating FEU is known. The flexibility of a unit is determined by its dynamics and constraints, which are not fully known in general. This is true in particular for large numbers of DEUs, such as residential heating and cooling systems or PEVs, for which the collection of that data can be prohibitively expensive and time-consuming. Thus, flexibility estimation procedures are needed that can identify in an autonomous manner the relevant parameters of individual FEUs based on the measurement data available.

An entire field called *load monitoring* focuses on the detection and extraction of the power consumption of individual devices from aggregate load measurements, see [87, 88], for instance. However, mapping the extracted information onto explicit descriptions of flexibility is only rarely discussed. The flexibility estimation approach presented in [89] is based on decomposing historic consumption profiles of FEUs into flexible and non-flexible components. Similarly, a comparison of flexible and inflexible loads based on clustering techniques is used in [90] to infer the flexible fraction of a load profile. The model-based flexibility assessment strategies presented in [64, 91, 92] first identify appropriate models based on which the flexibility of FEUs is estimated subsequently by means of extensive simulations.

The work presented in this chapter attempts to bridge the gap between raw energy consumption data and applications relying on explicit characterizations of flexibility, such as direct-control DR schemes. The main contribution is an algorithm that estimates the flexibility of an FEU autonomously based only on energy-meter data and outdoor air-temperature measurements. The fact that the approach requires only limited, readily available measurement data and can easily be automated makes it an inexpensive and scalable tool for AGs and SOs to estimate and predict the flexibility of large populations of FEUs. The flexibility estimation algorithm is applied to real measurement data collected from residential buildings during the *Ecogrid2* project [93].

This chapter is organized as follows. A model of the thermal energy dynamics of a building and its heating system is introduced in Section 3.2. Section 3.3 identifies the model parameters based on which the flexibility of a buildings and its HP can be estimated in Section 3.4. The conclusions and an outlook on future work are given in Section 3.5.

3.2 SYSTEM MODELING

An individual building and its heating system are modeled as a single energy buffer that represents the storage of thermal energy in the building and in the heating circuit of the HP. The evolution of the building's energy level $e(t) \in \mathbb{R}_0^+$ is governed by the in- and out-flux of thermal energy:

$$de(t)/dt = q_c(t)u(t) - q_l(t), \tag{3.1}$$

where $q_c(t)$, $q_l(t) \in \mathbb{R}$ are referred to as the energy *charge* and *loss* rates, respectively. The charge rate $q_c(t)$ equals the rate at which the HP feeds thermal energy into the system, whereas the loss rate $q_l(t)$ corresponds to the rate at which the system loses thermal energy to the outside. The charge and loss rates depend on various factors and are discussed in more detail in Section 3.3.

The HP is modeled as an on-/off-load whose binary state is represented by $u(t) \in \{1, 0\}$. Other operational modes, such as electric boosting or defrosting, are neglected. The HP is controlled by a TC that switches the HP on (off) as soon as the energy level $e(t)$ reaches the user-defined lower (upper) bound $\underline{e}$ ($\bar{e}$). Thus, the amount of flexible

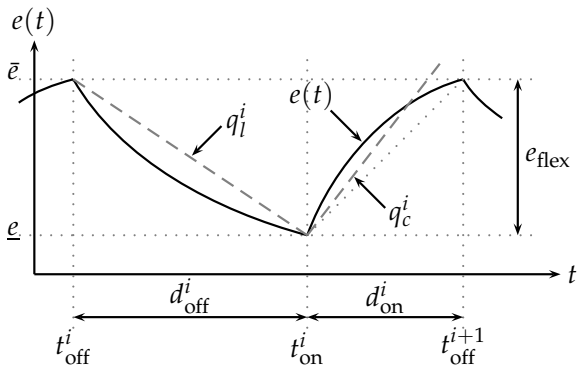

Figure 3.1: Evolution of the thermal energy content $e(t)$ (solid) of a thermostatically controlled heating system. The TC observes the energy limits $\underline{e}$, $\bar{e}$. The flexible energy capacity available during normal operation is e_{flex}. On- and off-switching times of the HP are denoted by t_{on}^i and t_{off}^i, with on- and off-durations d_{on}^i and d_{off}^i for heating cycle i. Also shown are the average charge and loss rates q_c^i and q_l^i (dashed).

energy available during normal operation is $e_{\text{flex}} := \bar{e} - \underline{e}$. Figure 3.1 shows a typical energy trajectory resulting from this type of control. The energy dynamics (3.1) depend on the time-varying charge and loss rates, and can be complex. However, in this work we are not interested in the exact temporal evolution of the energy level $e(t)$. Rather, our thermal model should be able to capture the on- and off-durations of the HP, denoted by d_{on}^i and d_{off}^i, respectively, for a heating cycle i. Thus, we approximate the time-varying charge and loss rates $q_c(t)$ and $q_l(t)$ for every heating cycle i by their average values q_c^i and q_l^i, defined as

$$q_c^i := \left(\frac{1}{d_{\text{on}}^i} + \frac{1}{d_{\text{off}}^i} \right) e_{\text{flex}}, \qquad q_l^i := \frac{e_{\text{flex}}}{d_{\text{off}}^i}. \tag{3.2}$$

The resulting energy dynamics are

$$de(t)/dt = q_c^i u(t) - q_l^i, \; t \in [t_{\text{off}}^i, t_{\text{off}}^{i+1}), \tag{3.3}$$

where t^i_{off} denotes the end time of a heating cycle. Because we are unable to measure the absolute amount of thermal energy $e(t)$ stored in the system, we consider the normalized state $x(t) := (e(t) - \underline{e})/(\bar{e} - \underline{e})$. From $\underline{e} \leq e(t) \leq \bar{e}$ guaranteed by the TC follows that $0 \leq x(t) \leq 1$. Thus, $x(t)$ can be interpreted as the SOC of the system and evolves according to

$$dx(t)/dt = r^i_c u(t) - r^i_l, \;\; t \in [t^i_{\text{off}}, t^{i+1}_{\text{off}}), \tag{3.4}$$

with the normalized average charge and loss rates

$$r^i_c := \frac{1}{d^i_{\text{on}}} + \frac{1}{d^i_{\text{off}}} \qquad \text{and} \qquad r^i_l := \frac{1}{d^i_{\text{off}}}. \tag{3.5}$$

The main advantage of the normalized energy dynamics (3.4)–(3.5) is that for each heating cycle i they are fully defined by the corresponding on- and off-durations d^i_{on} and d^i_{off}, respectively, and the initial state $x(t^i_{\text{off}}) = 1$.

3.3 SYSTEM IDENTIFICATION

3.3.1 *Identification of heat pump state and rated power*

The SM installed in each building measures the cumulative energy consumption of a building with a sampling time of $t_s = 5$ min. These energy measurements are translated into power values by computing their difference quotients. We denote by p_k the average power consumption of the building during the time interval $[kt_s, (k+1)t_s)$, where $k = 0, \ldots, N-1$ indexes discrete time and N is the number of time steps considered. A typical load profile of a building with predominant HP consumption is shown in the top plot in Figure 3.2.

Changes of the operational state of the HP usually involve significant changes in the power consumption of the HP. However, because the typical turn-on and turn-off procedures of an HP can cover more than one but not more than two sampling intervals, we consider power changes over two subsequent intervals computed as

$$\Delta_2 p_k := p_k - p_{k-2}, \; k = 2, \ldots, N-1, \tag{3.6}$$

and illustrated in the bottom plot in Figure 3.2. The local extrema of the time series $\Delta_2 p := \{\Delta_2 p_k\}_{k=2}^{N-1}$ are marked by red circles, and we

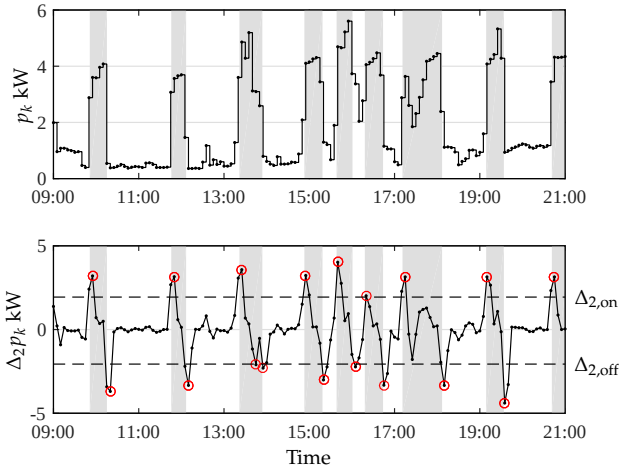

Figure 3.2: Aggregate power consumption p_k of a building and the corresponding changes over two subsequent time steps $\Delta_2 p_k$. The thresholds $\Delta_{2,\text{on}}$ and $\Delta_{2,\text{off}}$ are used to detect potential HP switching events. The shaded gray areas denote time periods during which the HP is believed to be ON.

summarize their indices by $\mathcal{I}_{\text{extr}} \subset \{2, \ldots, N-1\}$. These extrema can indicate switch-ons and -offs of electric devices such as HPs, plug-in electric vehicle chargers, or vacuum cleaners. Here we are interested in detecting switching actions of the HP. During the heating season, an HP exhibits a periodic heating behavior that results in several switching actions per day. Thus, the $\Delta_2 p$-data is expected to show at least three clusters: a cluster centered around zero comprises local extrema arising from small fluctuations of the aggregate load, and a cluster with positive and one with negative mean that correspond to changes in the load caused by switch-ons and -offs of the HP, respectively. Figure 3.3 shows the normalized histogram of the power changes $\Delta_2 p$ measured between September 1 2016 and March 15 2017. The two non-zero-mean clusters correspond to HP switching actions. Other electric devices that switch on and off frequently would show up as additional pairs of clusters.

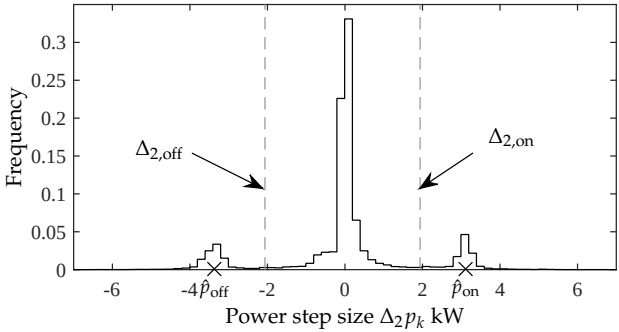

Figure 3.3: Distribution of power changes $\Delta_2 p_k$. The thresholds $\Delta_{2,\text{on}}$ and $\Delta_{2,\text{off}}$ are used for detecting HP switches. The estimated power consumption of the HP at the start and the end of a heating cycle is denoted by $\hat{p}_{\text{on}}$ and $\hat{p}_{\text{off}}$, respectively.

The HP detection procedure described here requires that the HP distinguishes itself in terms of its power rating from any other electric device installed in the building. In this case, the clusters of power changes arising from HP switching actions can be separated from potential clusters of other devices. A clustering technique, such as the k-means algorithm, can be used to group the data into these clusters. If the approximate value ranges of the power ratings of the individual devices are known, each device can be assigned to its corresponding pair of clusters.

We let $\mathcal{I}_{\text{on}}, \mathcal{I}_{\text{off}} \subset \mathcal{I}_{\text{extr}}$ denote the sets of indices of $\Delta_2 p$-values that belong to the switch-on and -off clusters of the HP, respectively. The thresholds

$$\begin{aligned}
\underline{\Delta}_{2,\text{on}} &:= \text{quantile}_{0.05}(\Delta_2 p_k,\ k \in \mathcal{I}_{\text{on}}),\\
\bar{\Delta}_{2,\text{on}} &:= \text{quantile}_{0.95}(\Delta_2 p_k,\ k \in \mathcal{I}_{\text{on}}),\\
\underline{\Delta}_{2,\text{off}} &:= \text{quantile}_{0.05}(\Delta_2 p_k,\ k \in \mathcal{I}_{\text{off}}),\\
\bar{\Delta}_{2,\text{off}} &:= \text{quantile}_{0.95}(\Delta_2 p_k,\ k \in \mathcal{I}_{\text{off}}).
\end{aligned} \tag{3.7}$$

are used to identify an HP switch-off if $\underline{\Delta}_{2,\text{off}} \leq \Delta_2 p_k \leq \bar{\Delta}_{2,\text{off}}$ and a switch-on if $\underline{\Delta}_{2,\text{on}} \leq \Delta_2 p_k \leq \bar{\Delta}_{2,\text{on}}$. We keep track of the operational state of the HP to enforce a series of strictly alternating switch-ons and -offs. For example, if we detect a switch-off while the HP is off already,

we check if we missed a switch-on or discard the current switch-off candidate. In addition, the minimum on- and off-duration of an HP are taken into account. If a switching event is detected at time k, the index k is stored in the set $\mathcal{I}_{\text{switch-on}}$ or $\mathcal{I}_{\text{switch-off}}$.

The $\Delta_2 p$-values that have been identified as a switch-on or -off of the HP are used to compute an estimate of the HP's rated power as

$$\begin{aligned}\hat{p}_r = (&\text{median}(\Delta_2 p_k,\, k \in \mathcal{I}_{\text{switch-on}}) \\ &- \text{median}(\Delta_2 p_k,\, k \in \mathcal{I}_{\text{switch-off}}))/2.\end{aligned} \tag{3.8}$$

3.3.2 *Estimation of exact switching times*

In general, the exact switching times of an HP state do not coincide with the 5-min sampling grid of the energy meter. For the purpose of illustration, Figure 3.4 shows an ideal switch-on of the heating system at time $t_{\text{on}} = 1.75 t_s$ reflected by a step in the power consumption from 0 kW to 3 kW. Also shown are the cumulative energy consumption of the HP as measured by the SM, and the power consumption time series derived from these measurements. Because the switching time $t_{\text{on}} = 1.75 t_s$ does not coincide with the sampling grid, the actual switch-on power step is not reproduced by the derived discrete-time power series as a single power step but by the two consecutive steps $p_1 - p_0$ and $p_2 - p_1$. Thus, the change in power over two sampling intervals $\Delta_2 p_k$ must be considered to properly capture an entire switching action. Here, the step $p_2 - p_0$ covers the entire switch-on, and the ratio $(p_2 - p_1)/(p_2 - p_0)$ is determined by the exact switching time of the HP.

Thus, if an HP switching action has been identified at time index k, i.e., $k \in \mathcal{I}_{\text{switch-on}}$ or $k \in \mathcal{I}_{\text{switch-off}}$, an estimate of the *exact* switching time $\hat{t}_{\text{switch}}$ is computed as

$$\hat{t}_{\text{switch}} = t_s(k-1) + t_s \left(\frac{p_k - p_{k-1}}{p_k - p_{k-2}} \right) \tag{3.9}$$

and stored as an estimated switch-on or switch-off time $\hat{t}^i_{\text{on}}$ or $\hat{t}^i_{\text{off}}$, respectively. Finally, estimates of the on- and off-durations of every HP heating cycle i are computed as

$$\hat{d}^i_{\text{on}} = \hat{t}^{i+1}_{\text{off}} - \hat{t}^i_{\text{on}} \qquad \text{and} \qquad \hat{d}^i_{\text{off}} = \hat{t}^i_{\text{on}} - \hat{t}^i_{\text{off}}. \tag{3.10}$$

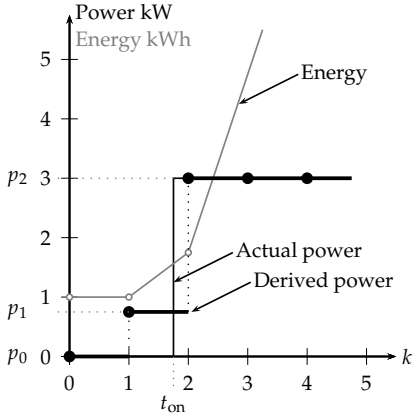

Figure 3.4: The heating system switches on at time $t_{on} = 1.75t_s$ and consumes 3 kW of power (thin solid black line). The cumulative energy consumption is shown in gray, and the corresponding power consumption is plotted in solid black. The power values start to rise before the switch occurs owing to discrete sampling of the cumulative energy. For the sake of simplicity, unit sampling time is assumed here.

The results of the HP state estimation are shown in Figure 3.2, where the shaded gray areas denote time periods during which the HP is believed to be running. The state estimation procedure is applicable to any type of on/off-load that switches its state frequently enough to appear as a peak in the $\Delta_2 p$-histogram in Figure 3.3. It can also be used to estimate the state of multiple loads simultaneously given that the loads can be distinguished by their power ratings. If multiple loads with similar power ratings are present, the state estimation fails because it cannot identify the device-specific clusters in the $\Delta_2 p$-data.

3.3.3 *Identification of the loss rate*

The loss rate r_l^i describes the average rate at which the normalized energy level $x(t)$ decreases over time. The loss rate depends on different factors, such as the the insulation of the building, the properties of the

heating system, and the ambient weather conditions. It is common in the literature to assume that the loss rate is proportional to the difference between the in- and the outdoor air temperature θ_{in} and θ_{out}, respectively, see [94–96]. Because the TC keeps θ_{in} within a temperature interval that is narrow compared with $\theta_{\text{in}} - \theta_{\text{out}}$, we only consider the loss rate's dependency on the average outdoor air temperature θ_{out} during a heating cycle and assume the piece-wise affine relationship

$$r_l(\theta_{\text{out}}) = \max\{0,\, a_l\theta_{\text{out}} + b_l\}, \tag{3.11}$$

with parameters $a_l, b_l \in \mathbb{R}$. The temperature θ^*_{out}, for which $r_l(\theta^*_{\text{out}}) = 0$, is referred to as the *zero-loss* temperature and can serve as an estimate of the building's indoor air temperature set point. The top plot in Figure 3.5 shows the normalized loss rate versus the mean outdoor air-temperature for all off-periods detected between September 1 2016 and March 15 2017. The red line depicts the robust least-squares fit of (3.11), with parameters $a_l = -0.084\ (^\circ\text{Ch})^{-1}$, $b_l = 1.722\ \text{h}^{-1}$, and a zero-loss temperature $\theta^*_{\text{out}} = 20.5\,^\circ\text{C}$. The data points above the fit correspond to heating cycles that are shorter than expected. The main reasons are *i)* irregular behavior of the inhabitants, such as the opening of windows or doors, and *ii)* errors in the detection of the HP switching actions. Samples below the fit indicate off-periods that are longer than expected. The main reason is that the heating activity can be shadowed by the consumption of other devices, and thus actual switch-ons of the HP are missed.

3.3.4 *Identification of the charge rate*

The charge rate r_c^i describes the rate at which the HP could increase the energy level $x(t)$ if there was no energy loss to the outside. Its inverse equals the time required for the HP to heat up the building from its lower energy bound $\underline{e}$ to its upper bound $\bar{e}$ if there are no thermal losses. The charge rate depends on the HP's rated power p_r and on its coefficient of performance (COP) η, i.e., $r_c = \eta(\cdot)p_r$. The COP is influenced by the in- and outdoor air temperatures, the HP outflow water temperature θ_{w}, and other factors, such as the run-time of the HP. The COP is determined by many factors. Here we only consider

the COP's dependency on the HP outflow water temperature θ_w and θ_{out} and use the formula for an HP's maximum theoretical efficiency

$$\eta = c_c \left(\frac{\theta_w + 273.15}{\theta_w - \theta_{out}} \right), \tag{3.12}$$

where c_c is an HP-specific scaling factor to incorporate all nonidealities [97]. The constant 273.15 is required to convert values from °C to K. The average outflow water temperature θ_w is determined by θ_{out} via the user-set *heat curve*. We assume the affine relationship $\theta_w(\theta_{out}) = a_c\theta_{out} + b_c$, which, together with (3.12), suggests that the dependency of r_c on θ_{out} is of the form

$$r_c(\theta_{out}) = c_c \left(\frac{a_c\theta_{out} + b_c + 273.15}{(a_c - 1)\theta_{out} + b_c} \right), \tag{3.13}$$

with parameters a_c, b_c, $c_c \in \mathbb{R}$. Thus, our COP model depends only on θ_{out}. The red line in the middle plot in Figure 3.5 depicts the robust nonlinear least-squares fit of (3.13). For the measurement data collected between September 1 2016 and March 15 2017, the parameter values $a_c = -17.85$, $b_c = 473.26\,°\mathrm{C}$, and $c_c = 1.6262\ \mathrm{h}^{-1}$ have been found.

3.3.5 *Identification of the duty cycle*

The duty cycle of the HP during the heating cycle i is defined as

$$d_c^i := d_{on}^i / (d_{on}^i + d_{off}^i). \tag{3.14}$$

According to (3.5), our estimate of the right-hand side of (3.14) is $\hat{r}_l^i / \hat{r}_c^i$. Thus, an estimate of the duty cycle is given by the ratio $\hat{r}_l(\theta_{out}) / \hat{r}_c(\theta_{out})$ and its inherent limitation to the interval $[0, 1]$:

$$\hat{d}_c(\theta_{out}) = \min\{1, \max\{0, \hat{r}_l(\theta_{out}) / \hat{r}_c(\theta_{out})\}\}. \tag{3.15}$$

The red line in the bottom plot in Figure 3.5 depicts the duty cycle of the HP estimated according to (3.14) based on measurement data collected between September 1 2016 and March 15 2017. The duty cycle takes the value 0 for $\theta_{out} \geq 20.5\,°\mathrm{C}$ and 1 for $\theta_{out} \leq -7.4\,°\mathrm{C}$.

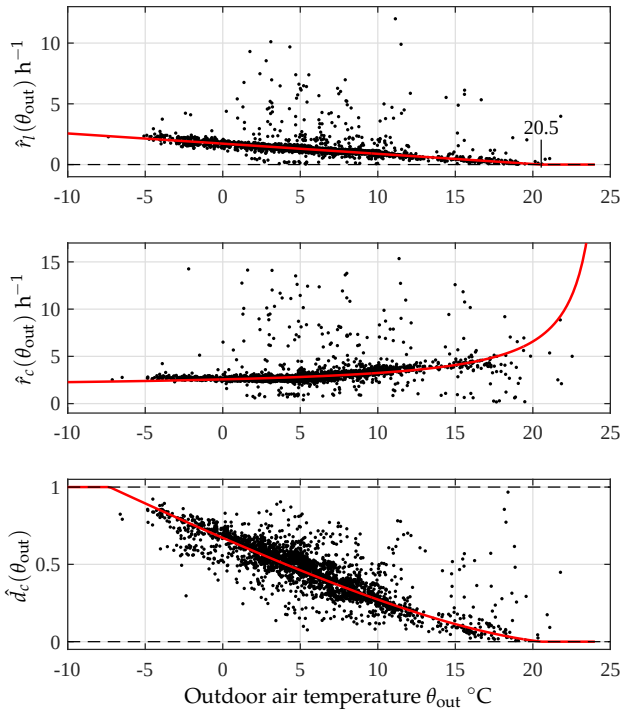

Figure 3.5: Estimates of the normalized average loss rate $\hat{r}_l$ (top), charge rate $\hat{r}_c$ (middle), and duty cycle $\hat{d}_c$ (bottom) versus outdoor air temperature. The solid lines represent least-squares fits according to (3.11) and (3.13), respectively, based on real SM and outdoor air-temperature measurements collected between September 1 2016 and March 15 2017.

3.3.6 *Flexibility level*

Under undisturbed operating conditions, the TC of the HP keeps the indoor air temperature θ_{in} within a range $[\underline{\theta}_{\text{in}}, \bar{\theta}_{\text{in}}]$ that is comfortable for the building inhabitants. The width of this range directly affects the potential of the system to provide DR services because it limits the maximum on- and off-durations of the HP. In most cases, however, the original temperature range is chosen conservatively, and the inhabitants accept wider ranges if they are awarded for it. We introduce a *flexibility level* as a simple means for inhabitants to express their willingness to accept wider indoor temperature ranges. Here we use a numeric flexibility level $f \geq 1$ that determines the maximum admissible off-duration of the HP as

$$f d_{\text{off}}(\theta_{\text{out}}) = f / r_l(\theta_{\text{out}}). \tag{3.16}$$

The AG in charge of controlling the HP must respect this limit. Because in our setup the AG does not know how long an HP has been off at the time it is throttled, the maximum allowed HP throttling duration is given as

$$T_{\max}(\theta_{\text{out}}) = (f - 1) / r_l(\theta_{\text{out}}). \tag{3.17}$$

Thus, the AG is not allowed to throttle the HP if $f = 1$. In this case, the HP is controlled only by the TC and the original indoor air temperature range is maintained. In contrast, flexibility levels $f > 1$ allow the AG to throttle the HP for longer time intervals compared to unperturbed operation. This can lead to indoor air temperatures below the original lower bound $\underline{\theta}_{\text{in}}$. Assuming the linear energy dynamics (3.1), the flexibility level f resets this lower bound to

$$\underline{\theta}_{\text{in}}(f) = f \underline{\theta}_{\text{in}} + (1 - f) \bar{\theta}_{\text{in}}. \tag{3.18}$$

The concept of the flexibility level can be generalized easily to also define the maximum on-duration of an HP. Moreover, categorical values such as {low, medium, high} can be offered to end-customers and then mapped to numeric values by the AG.

3.4 ESTIMATION OF FLEXIBILITY

Given the charge and loss rates identified above, the goal of this section is to estimate the parameters that define the flexibility of a residential buildings and its HP as represented by the feasible set introduced in Section 2.3.

3.4.1 *Power and state constraints*

An estimate of the rated heating power has been derived in Subsection 3.3.2. This value is interpreted as the average power the HP will consume during a heating period. Thus, the power consumption p_k of the HP is restricted to

$$p_k \in \{0,\, \hat{p}_r\},\, k = 0, \ldots, N-1. \tag{3.19}$$

As discussed in Section 3.2, a domestic building and its heating system are modeled as a single energy buffer. The discrete-time version of the normalized state dynamics (3.4) is

$$x_{k+1} = x_k + t_s(r_c^k u_k - r_l^k),\; k = 0, \ldots, N-1. \tag{3.20}$$

The temperature dependency of r_l and r_c is captured by the functions (3.11) and (3.13) which are used to predict the normalized rates based on forecasts of the outdoor air temperature. Let $\hat{\theta}_{\text{out}}^k$ denote the average outdoor temperature forecast for the time interval $[t_s k,\, t_s(k+1))$. Then, the forecast of r_c^k is $\hat{r}_c(\hat{\theta}_{\text{out}}^k)$, and similarly for $\hat{r}_l^k$. Thus, forecasts of the SOC are computed as

$$\hat{x}_{k+1} = \hat{x}_k + t_s[\hat{r}_c(\hat{\theta}_{\text{out}}^k)u_k - \hat{r}_l(\hat{\theta}_{\text{out}}^k)]. \tag{3.21}$$

An estimate of the initial SOC can be derived as follows:

$$\hat{x}_0 = \begin{cases} 1 - \hat{r}_l(\theta_{\text{out}})[t_0 - \hat{t}_{\text{off}}(\text{end})] & \text{if } \hat{t}_{\text{off}}(\text{end}) > \hat{t}_{\text{on}}(\text{end}) \\ \hat{r}_c(\theta_{\text{out}})[t_0 - \hat{t}_{\text{on}}(\text{end})] & \text{otherwise,} \end{cases} \tag{3.22}$$

with θ_{out} being the average outdoor air temperature measured since the last switching action of the HP.

3.4.2 *Estimation of flexibility*

At every time step k, the input $u_k \in \{0, 1\}$ can be chosen freely, subject to the constraint that the subsequent state value x_{k+1} satisfies the corresponding state constraints. Under normal operating conditions, i.e., with flexibility level $f = 1$, the state constraints are $0 \leq x_k \leq 1$ for all $k = 1, \ldots, N$. However, for general $f \geq 1$, which allow for a prolonged throttling of the HP, the lower bound $\underline{x}_k = 0$, $k = 1, \ldots, N$, does not hold anymore. Instead, it depends on the flexibility level according to

$$\underline{x}_k(f) = 1 - f, \; k = 1, \ldots, N. \tag{3.23}$$

Consequently, the state constraints are

$$1 - f \leq x_k \leq 1, \; k = 1, \ldots, N. \tag{3.24}$$

Summarize $\hat{x} := [\hat{x}_1, \ldots, \hat{x}_N]^\top \in \mathbb{R}^N$ and $u := [u_0, \ldots, u_{N-1}]^\top \in \mathbb{R}^N$. The predicted SOC evolution $\hat{x}$ for a particular input trajectory u can be written compactly as

$$\hat{x} = \hat{x}_0 \mathbf{1} + Bu - d, \tag{3.25}$$

with matrix $B \in \mathbb{R}^{N \times N}$, whose i^{th} row is given as

$$B_i = t_s[\hat{r}_c(\hat{\theta}^0_{\text{out}}), \hat{r}_c(\hat{\theta}^1_{\text{out}}), \ldots, \hat{r}_c(\hat{\theta}^{i-1}_{\text{out}}), 0, \ldots, 0],$$

and vector $d \in \mathbb{R}^N$, whose i^{th} element is given as

$$d_i = t_s \left(\hat{r}_l(\hat{\theta}^1_{\text{out}}) + \hat{r}_l(\hat{\theta}^2_{\text{out}}) + \cdots + \hat{r}_l(\hat{\theta}^{i-1}_{\text{out}}) \right).$$

A vector of ones of appropriate size is denoted by $\mathbf{1}$.

Using (3.25) and $p_k = \hat{p}_r u_k$, the state constraints (3.24) can be reformulated as constraints on the power consumption p of the heating system over a given planning horizon as

$$\begin{bmatrix} B \\ -B \end{bmatrix} p \leq \hat{p}_r \begin{bmatrix} (1 - \hat{x}_0)\mathbf{1} + d \\ (f - 1 + \hat{x}_0)\mathbf{1} - d \end{bmatrix}, \tag{3.26}$$

where $p := [p_0, \ldots, p_{N-1}]^\top \in \mathbb{R}^N$. Inequalities (3.26) define a convex compact polytope in $\mathbb{R}^N$ that contains all the power trajectories that are feasible with regard to the state limits.

The constraints (3.19) and (3.26) collectively describe the set of power trajectories that the system can follow over a given planning horizon while meeting all comfort constraints. That is, they represent the *feasible set*. Figure 3.6 visualizes the electric flexibility available from an individual building and its heating system under normal operating conditions ($f = 1$) over a time period of 12 h with sampling time $t_s = 5$ min. The initial SOC is $x_0 = 0.75$. For every time step, the maximum and the minimum possible cumulative energy have been computed. The shaded gray area between these lines is a measure of the flexibility because all feasible trajectories must lie inside of it. The flexibility of the system allows the heating to start immediately or to be deferred for at most 4 time steps, see the inset. Also shown is the outdoor air temperature, which influences the flexibility available.

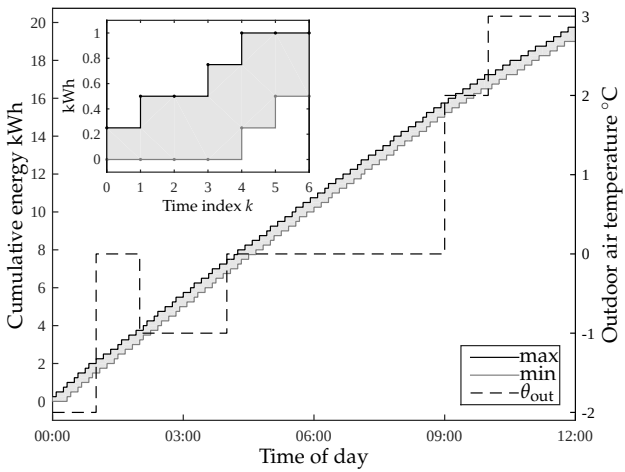

Figure 3.6: Minimum and maximum cumulative energy trajectories over 12 h with sampling time $t_s = 5$ min, $f = 1$, and $x_0 = 0.75$. The inset plot shows a detail view of the first 6 time slots. The flexibility of the system allows the HP to start immediately or to be deferred for at most 4 time steps.

3.5 CONCLUSION & FUTURE WORK

This chapter presented an algorithm to estimate the electric flexibility available from domestic buildings with HP installations. The approach requires only energy-meter data and measurements of the outdoor air temperature, which are readily available data sets. The fact that the estimation routine runs autonomously makes it well suited as an inexpensive and scalable tool to estimate the flexibility available from large populations of FEUs.

There are several future directions of interest. The robustness of the procedure to detect the HP consumption can be improved by considering additional HP states, such as the electric boost or defrost mode. In addition, learning the temporal operation pattern of the HP could help filter out the consumption of other electric devices. Further, the flexibility estimation presented here estimates the flexibility that is available during normal operation of the HP as determined by the settings of the TC, i.e., for a flexibility level $f = 1$. Future work will investigate strategies to iteratively increase the flexibility level with the aim of maximizing the amount of flexibility available without compromising the user comfort.

4 DEMONSTRATION OF DEMAND RESPONSE

In this chapter, we apply the flexibility estimation procedure presented in Chapter 3 to estimate the flexibility of a population of more than 200 residential buildings with HP installations. Based on the flexibility estimated, we show how the aggregate DR potential of the population can be quantified and predicted. We report on the results of various load-reduction and rebound-damping experiments to illustrate the effectiveness of the approach.

4.1 INTRODUCTION

Keeping the electricity demand and supply in a power system balanced at all times can be challenging for SOs, in particular with growing shares of uncertain and intermittent renewable generation. DR is one approach proposed to help balance the power grid by scheduling and adjusting the electricity consumption of various DEUs intelligently [98]. DR has been studied extensively, and there exist numerous projects implementing and evaluating DR schemes in practice, see the overviews provided in [99, 100].

DR schemes can be sorted into two main categories, namely, Direct Load Control (DLC) and Indirect Load Control (ILC) [35, 98]. ILC refers to a setup in which an incentive signal, such as a forecast of the electricity price, is sent to the consumers, which are expected to adjust their demand accordingly. The fact that only a single signal needs to be broadcast and that the control of the consumer devices is the responsibility of their operators makes ILC a simple, privacy-preserving and scalable approach for SOs. However, the demand sensitivity with regard to the incentive signal is unknown and must be estimated. Results from real-life implementations of ILC schemes can be found in [101–105], among others.

This work deals with DLC. In this setup, an AG controls the electricity consumption of each DEU in its control group individually via a control

signal. This individual control is the main advantage of DLC because it allows the AG to coordinate the DR of the units. For example, the DR of multiple units can be activated sequentially over time to achieve a prolonged response which would not be possible in an ILC setup in which all units react to a single control signal simultaneously. This makes DLC an accurate and versatile method to provide different types of DR services [35]. However, to efficiently and safely exploit the flexibility of a group of FEUs, their flexibility must be identified first.

We consider TCLs for which different flexibility identification approaches have been proposed. Top-down approaches attempt to capture the aggregate DR behavior of a group of DEUs directly based, for example, on the probabilistic properties of the underlying units [94–96]. Other approaches run extensive experimental tests on a population of units to identify their aggregate DR behavior [106]. In contrast, we consider a bottom-up approach in which the flexibility of each unit is identified first and aggregated subsequently. In [107], the deferrable energy consumption of an individual TCL is quantified based on high-resolution temperature and power measurements of the unit's previous heating cycles. Other approaches employ statistical models and artificial neural networks to identify the potential load shift of TCLs from temperature, set-point, and power consumption data [91, 108]. In cases in which all system parameters are known, the flexibility available from individual systems can be computed via simulations [73, 92] or optimization problems [59, 109].

The aforementioned approaches have several limitations that need to be addressed for real-life DR applications. First, many approaches, such as [59, 73, 94–96, 109], assume that the parameters governing the thermal and electric behavior of the units are known. However, this is rarely the case in practice. In particular for large populations of DEUs, it usually is prohibitively expensive and time-consuming to collect the nameplate and measurement data required by classical system identification techniques. Second, data-based flexibility identification procedures, such as [91, 107, 108, 110], typically require extensive, high-resolution, low-latency measurement data of system states, electricity consumption, and weather conditions, among others. These data sets are not easily available in general, and retrofitting existing systems with appropriate sensor equipment can be time-consuming and costly. Third,

some flexibility aggregation strategies, such as [74, 111, 112], do not properly take into account the heterogeneity of the units considered but resort to averaging or clustering techniques, for example. However, it has been shown that the heterogeneity among the units is important with regard to the accuracy [74] and the quantity of the flexibility available on the aggregate level [71]. Finally, most discussions of real-life DR demonstrations, such as [113, 114], focus on the size of the power or energy response achieved, or the costs saved. Results on the accuracy and reliability of the DR service prediction and delivery are reported rarely.

We address these challenges by making the following main contributions. First, we extend the flexibility estimation procedure introduced in Chapter 3 and use it to estimate the electric flexibility of each individual unit in a group of more than 200 residential buildings and HPs based on real measurement data. In contrast to the approaches mentioned above, our procedure requires only limited, readily available measurement data, and can easily be automated. These features make it an inexpensive and scalable tool for AGs. Second, we present a probabilistic characterization of the aggregate DR behavior of the population of buildings and HPs that accounts for heterogeneity by explicitly taking into account the thermal properties and constraints of individual units. Third, we describe a large-scale, real-life implementation of a DLC scheme and demonstrate by means of experiments that the aggregate load of a group of residential buildings can be reduced by 40–65% by throttling their HPs and that this load-reduction service can be delivered reliably with a median absolute percentage error of below 7%. In addition, we propose a rebound damping strategy that has lead to peak rebound power reductions of 50% in practice.

This chapter is organized as follows. Section 4.2 describes the experimental setup used. In Section 4.3, the aggregate DR behavior of a population of residential buildings with HPs is characterized based on the dynamics and constraints identified for each unit individually. Experimental results are presented in Section 4.4, and conclusions and an outlook are provided in Section 4.5.

4.2 EXPERIMENTAL SETUP

4.2.1 *Top-level demand response setup*

Figure 4.1 illustrates the direct-control DR setup implemented in the ECOGRID2 project [93] and used in this work. The setup comprises more than 300 inhabited residential buildings that are of different size, feature distinct thermal properties, and are equipped with HPs of different makes, types, and age. An AG manages the electricity consumption of a subset of J buildings and HPs indexed by $j \in \mathcal{J} := \{1, \ldots, J\}$. The AG concludes contracts with the unit owners, defining the set of admissible control signals and the details on how the units are reimbursed for their flexibility. The AG first estimates the electric flexibility of each unit based on measurement data, and aggregates the individual flexibilities subsequently. The aggregate flexibility can be used for different purposes, such as minimizing the energy procurement costs [62, 115] or offering ancillary services to SOs [32, 63]. Other approaches to exploit the flexibility are based on decentralized market designs incorporating the flexibility costs of end-customers [116], or integrating different types of energy carriers [117].

In this work, we consider a flexibility market setup in which buyers of flexibility, such as Distribution or Transmission SOs, can open various types of flexibility markets according to their particular needs. For example, after forecasting the imminent risk of a line congestion, a Distribution SO can open a market for load-reduction services that are to be delivered in the affected distribution grid zone during a predefined future time interval. The flexibility markets are hosted by a Flexibility Clearing House (FLECH) [118]. If a new market is opened, FLECH broadcasts the market and product specifications to all participating AGs, collects their flexibility bids, clears the market, and informs the parties involved about the clearing results, see Figure 4.1. FLECH is a market platform designed for future flexibility markets, i.e. , the processes of opening a market, submitting bids, and clearing a market occur prior to the time of delivery. During the time of delivery, FLECH acts as an intermediary between the sellers and buyers of flexibility. For example, an SO activates a DR service via FLECH which forwards the activation to

the corresponding AG that sends out individual control signals $\tau^{(j)}$ to all the units in its control group.

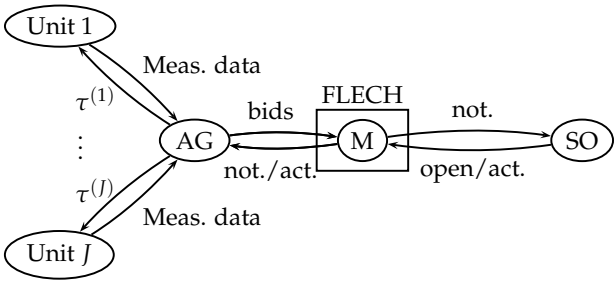

Figure 4.1: Schematic representation of the market-based direct-control DR setup used in this work. An AG collects measurement data from the individual units in its control group, estimates their flexibility, and may offer flexibility products on flexibility markets (M). Markets can be opened by potential buyers of flexibility, such as Distribution and Transmission SOs, and are hosted on FLECH. If an SO activates a flexibility product, FLECH notifies the AG which in turn sends out throttling commands $\tau^{(j)}$ to the individual units.

4.2.2 *Building-level communication and control setup*

An overview of the communication and control setup used in this work is provided in Figure 4.2. Each building is equipped with an off-the-shelf *Landis+Gyr E450* or *E650* SM that samples the total electric energy consumption and generation (photovoltaic or wind) of the building with a sampling time of $t_s = 5$ min rather than the 15 min or 60 min intervals commonly used. Once a day, the SMs upload their measurement data via a mobile internet connection to a central meter data management system by *Landis+Gyr*, from where the data is pushed to a data base hosted by the AG back-end. This data base also stores outdoor air-temperature values measured at a single location as well as corresponding forecasts provided by a weather service (WS). The AG back-end runs the system identification and the HP control algorithm. The identification procedure

estimates the flexibility of a building and is based exclusively on the data provided by the SM and the WS. It is executed only once for each building, and the results are stored in the data base. The HP control algorithm relies on the flexibility identified when controlling an HP j via the binary *throttling signal* $\tau^{(j)}$: If the value 1 is sent to an HP that is currently running, a rundown sequence is initiated, and the HP will cease operation within 2-10 min depending on its current state. As long as the signal remains 1, the HP cannot turn on. However, once the signal is reset to 0, the HP can operate freely according to its internal TC. Thus, the throttling signal can be interpreted as a "request to turn off" if set to 1, and as a "permit to operate according to TC" otherwise. It is important to note that with this control setup it is impossible to force an HP to turn on and consume power.

The control algorithm cannot access the HP directly but sends the throttling signal to a third-party home automation back-end, which communicates with an HP via several stations: It uses a DSL connection to the internet router (IR) of the building to access the home automation gateway (HAG) via Ethernet. The HAG controls a *Danfoss RXZ-1* relay (RE) via *Z-Wave* to apply the throttling signal to the HP. To do so, it uses the so-called *tariff input*, which is a feature common to all our HPs and can be used by SOs to prevent the HPs from consuming electricity during times of high demand.

The multi-state communication between the control algorithm and the HP proved to be inconvenient and unreliable for several reasons. First, finding a suitable position for the HAG can be tricky because it requires a cable connection to the IR but must at the same time be located such that the RE is reachable via *Z-Wave* radio signal. Power-line Ethernet bridges have been used to relax these restrictions. Second, the setup involves at least three on-premise devices, namely, the IR, HAG, and RE. If one of them is out of operation the control algorithm can no longer communicate with the HP. To prevent prolonged throttling of an HP in case of communication breakdown, the control algorithm always sends a throttling request together with a predefined throttling release time, which is stored locally on the HAG and is applied to the RE even if the communication link between the control algorithm and the HAG is lost.

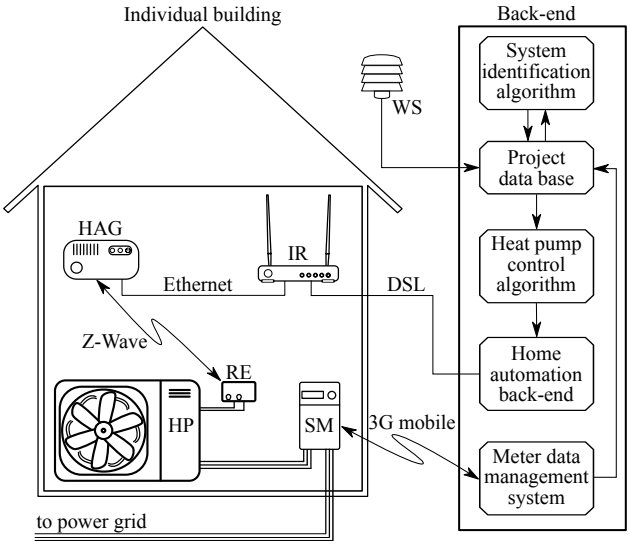

Figure 4.2: Communication and control setup used in this work: Each building is equipped with a smart meter (SM) that measures the energy consumption and generation of the building on a timescale of 5 min. The data is sent to the meter data management system and pushed to the aggregator data base, which also hosts outdoor air temperature data provided by a weather service (WS). The system identification uses this data to estimate the electric flexibility of a building. A control algorithm can send throttling commands to the HP via the home automation back-end and the home automation gateway (HAG).

4.3 POPULATION MODELING

In this section, we introduce a probabilistic characterization of the aggregate DR behavior of a population of FEUs. We characterize the aggregate flexibility of a population of J buildings and HP that are modeled as a TCL with a single energy buffer as described in Section 3.2. First, the estimation procedure discussed in Section 3.3 is applied to estimate the flexibility if individual buildings. Subsequently, the aggregate flexibility of the population of buildings is quantified by means of the aggregate load-reduction potential and the size of the corresponding rebound.

4.3.1 *Expected load reduction*

Because SM data can be delayed by several hours, the current state of an HP cannot be estimated by conventional methods such as [110]. Instead, we consider the binary operational state of HP j as a random variable whose probability of being ON equals its duty cycle $d_c^{(j)}(\theta_{\text{out}})$. Thus, if a throttling command is sent to this HP, the expected load reduction is

$$\hat{p}_{\text{red}}^{(j)}(\theta_{\text{out}}) = \hat{p}_r^{(j)} \hat{d}_c^{(j)}(\theta_{\text{out}}), \tag{4.1}$$

where $\hat{p}_r$ and $\hat{d}_c$ denote the estimates of the rated power and the duty cycle of the HP, respectively, derived in Section 3.3.

Experiments have shown that on average about 12% of all HPs fail to respond to our throttling commands. Thus, we include the empirical throttling success rate $\hat{\sigma}_{\text{red}} = 0.88$ when computing the *total* expected load reduction as

$$\hat{p}_{\text{red}}^{(\mathcal{J})}(\theta_{\text{out}}) = \hat{\sigma}_{\text{red}} \sum_{j \in \mathcal{J}} \hat{p}_{\text{red}}^{(j)}(\theta_{\text{out}}). \tag{4.2}$$

Figure 4.3 illustrates the theoretic and empirical total load reductions that can be achieved by our population of $J = 209$ buildings for outdoor air temperatures between $-30°$C and $30°$C. A maximum total load reduction of 454.4 kW can be reached at -30 °C.

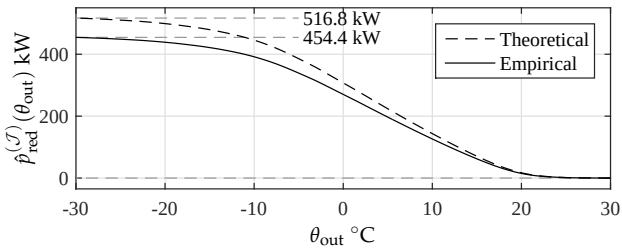

Figure 4.3: Expected maximum aggregate load reduction $\hat{p}_{\text{red}}^{(\mathcal{J})}(\theta_{\text{out}})$ that can be achieved by throttling all the 209 identified HPs simultaneously for different outdoor air temperatures θ_{out}.

Let $\mathcal{I}_T(\theta_{\text{out}}) \subseteq \mathcal{J}$ denote the set of HPs that can be throttled for a duration T or longer for a given temperature θ_{out}, i.e.,

$$\mathcal{I}_T(\theta_{\text{out}}) := \{j \in \mathcal{J} : T_{\text{max}}^{(j)}(\theta_{\text{out}}) \geq T\}, \tag{4.3}$$

where $T_{\text{max}}^{(j)}$ denotes the maximum throttling duration allowed for HP j, see (3.17). We characterize the aggregate load-reduction potential of population $\mathcal{J}$ by computing the maximum expected aggregate load reduction (4.2) for different outdoor air temperatures θ_{out} and throttling durations T as

$$\hat{p}_{\text{red}}^{(\mathcal{J})}(\theta_{\text{out}}, T) = \hat{\sigma}_{\text{red}} \sum_{j \in \mathcal{I}_T(\theta_{\text{out}})} \hat{p}_{\text{red}}^{(j)}(\theta_{\text{out}}). \tag{4.4}$$

Figure 4.4 shows the values of (4.4) obtained for flexibility factors $f^{(j)} = 4$, $j \in \mathcal{J}$. The achievable load reduction decreases for longer throttling times because the number of HPs that can be throttled for the full duration T declines. The gradual decline visible in Figure 4.4 is a result of the heterogeneity among the units with regard to the maximum throttle times (3.17). With decreasing θ_{out}, the load-reduction potential grows because the probability of an HP being ON increases according to its duty cycle, see Figure 3.5. The load-reduction curves for a constant throttling duration do not increase monotonically because the set of HPs admissible for throttling, $\mathcal{I}_T(\theta_{\text{out}})$, shrinks as θ_{out} decreases. The red line depicts the outdoor temperatures at which the largest load reduction is expected for a given throttling duration.

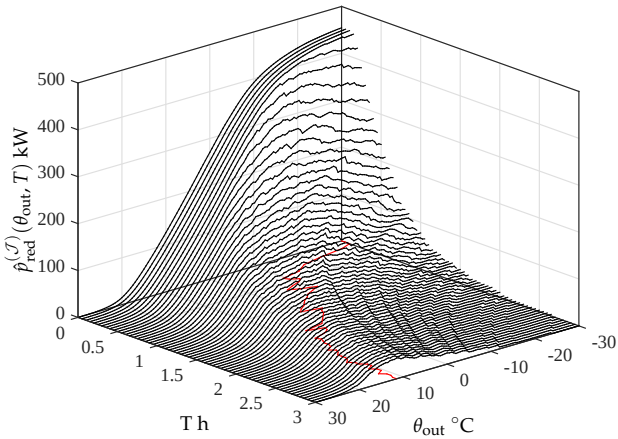

Figure 4.4: Maximum load reduction $\hat{p}_{\text{red}}^{(\mathcal{J})}(\theta_{\text{out}}, T)$ expected from our population of 209 HPs for different outdoor air temperatures θ_{out} and throttling durations T with flexibility factors $f^{(j)} = 4$, $j \in \mathcal{J}$. The red line depicts the outdoor temperatures at which the largest load reduction is expected for a given throttling duration.

4.3.2 *Expected rebound*

After releasing the throttling signal of an HP at the end time t_{end} of a throttling period of duration T, the TC turns the HP on if the energy level of the building has fallen below its lower bound, i.e., $x(t_{\text{end}}) \leq 0$. The probability of this happening is

$$P(x(t_{\text{end}}) \leq 0) = P(x(t_{\text{end}} - T) \leq T\hat{r}_l^{(j)}(\theta_{\text{out}})) = T\hat{r}_l^{(j)}(\theta_{\text{out}}), \quad (4.5)$$

given the linear state dynamics (3.4). This means that at time t_{end} the HP turns on with probability $T\hat{r}_l^{(j)}(\theta_{\text{out}})$ and consumes $\hat{p}_r^{(j)}$ units of power. If the HP had not been throttled, its expected average power consumption would be $\hat{d}_c^{(j)}(\theta_{\text{out}})\hat{p}_r^{(j)}$. We define the expected rebound power $\hat{p}_{\text{reb}}^{(j)}$ as the HP's *excess* consumption following a throttling period:

$$\hat{p}_{\text{reb}}^{(j)}(\theta_{\text{out}}) := T\hat{r}_l^{(j)}(\theta_{\text{out}})(1 - \hat{d}_c^{(j)}(\theta_{\text{out}}))\hat{p}_r^{(j)}. \quad (4.6)$$

Similar to (4.4), we take an empirical success rate $\hat{\sigma}_{\text{reb}}$ into account when computing the aggregate expected rebound power as

$$\hat{p}_{\text{reb}}^{(\mathcal{J})}(\theta_{\text{out}}, T) = \hat{\sigma}_{\text{reb}} \sum_{j \in \mathcal{I}_T(\theta_{\text{out}})} \hat{p}_{\text{reb}}^{(j)}(\theta_{\text{out}}) \tag{4.7}$$

which is shown in Figure 4.5 for flexibility factors $f^{(j)} = 4$, $j \in \mathcal{J}$. For short throttling durations, the rebound grows with T because longer throttling durations increase the probability of a rebound according to (4.5). However, as T grows further, fewer HPs can be throttled, see (4.3), and the rebound decreases again. For low θ_{out}, the rebound approaches zero because the HPs are ON most of the time, i.e., their duty cycles approach 1. Thus, the rebound diminishes in accordance with (4.6). For high θ_{out}, in contrast, the expected rebound is small because of the diminishing probability (4.5) that the system reaches its lower state limit and triggers a rebound. Peak rebound values (red line) lie in the range of 230–247 kW for $T > 1$ h and θ_{out} between 11–16 °C.

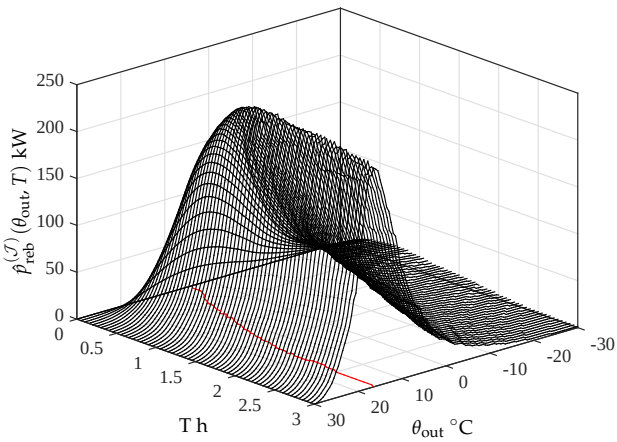

Figure 4.5: Rebound power $\hat{p}_{\text{reb}}^{(\mathcal{J})}(\theta_{\text{out}}, T)$ expected from our population of 209 HPs for different outdoor air temperatures θ_{out} and throttling durations T with flexibility factors $f^{(j)} = 4$, $j \in \mathcal{J}$. The red line depicts the outdoor temperatures at which the largest rebound power is expected for a given T.

4.4 EXPERIMENTAL RESULTS

We conducted more than 70 DR experiments involving 322 buildings with HPs. In each experiment, a subset $\mathcal{J}_{\text{ctrl}}$ of at most 209 HPs was used for DR. The remaining units served as the reference group $\mathcal{J}_{\text{ref}}$.

4.4.1 *Load reduction*

We focus on a particular DR service, namely *load reduction*, defined as the reduction of the total load of the units controlled relative to their undisturbed load profile referred to as the *baseline* [119]. We consider the *average* load reduction achieved during the response period, i.e.,

$$\bar{p}_{\text{red}}^{(\mathcal{J}_{\text{ctrl}})} := \bar{p}^{(\mathcal{J}_{\text{ctrl}})} - \bar{p}_{\text{base}}^{(\mathcal{J}_{\text{ctrl}})}, \tag{4.8}$$

where $\bar{p}^{(\mathcal{J}_{\text{ctrl}})}$ denotes the average power consumption of the group of controlled houses during the response period, and $\bar{p}_{\text{base}}^{(\mathcal{J}_{\text{ctrl}})}$ is the corresponding average baseline consumption. Because the baseline consumption cannot be measured, it must be estimated. Various techniques have been proposed to do this, see [119, 120]. Here, we derive an estimate of the baseline based on the consumption of the reference group. To make the power consumption of the controlled group and the reference group comparable, the reference consumption is scaled by a factor so as to optimally match (in the sense of least squares) that of the controlled group during the 4 h *preceding* the response period. A smoothing spline is then fitted to the scaled reference consumption data and serves as the baseline $\bar{p}_{\text{base}}^{(\mathcal{J}_{\text{ctrl}})}$ relative to which the load reduction is measured. The accuracy of our baseline estimation can be assessed by comparing the baseline with the consumption of the controlled group of houses measured under undisturbed conditions. In practice, our baseline was able to explain the 1-h average consumption of the controlled group with a mean absolute percentage error (APE) of 2.6%.

Figure 4.6 shows a typical load-reduction experiment: At 22.00 h, a throttling signal is sent to 141 out of 322 HPs and released one hour later. The remaining units were used as the reference group. The expected average aggregate load reduction $\hat{p}_{\text{red}}^{(\mathcal{J}_{\text{ctrl}})}(\hat{\theta}_{\text{out}}, T) = 177.5$ kW was computed in advance via (4.4) with $T = 1$ h and forecast $\hat{\theta}_{\text{out}} = 0.1\,^{\circ}$C. The

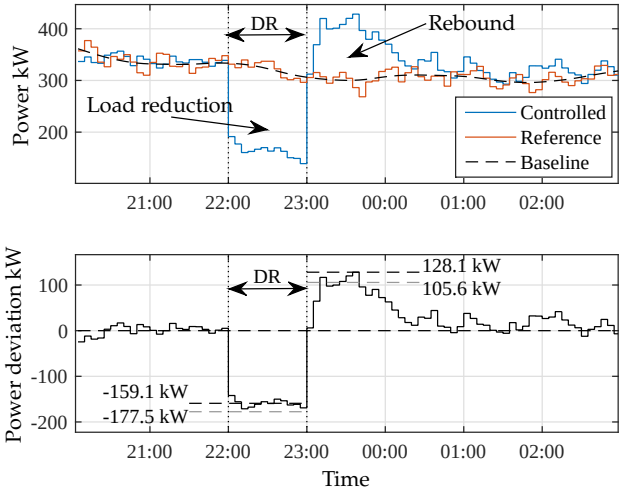

Figure 4.6: Results from a DR experiment involving 141 buildings with HPs. Top: Power consumption of the group of controlled units (blue), the scaled consumption of the reference group (red), and the fitted baseline (dashed black). A load-reduction service is delivered between 22.00 h and 23.00 h followed by a rebound. Bottom: The load reduction and rebound are defined relative to the baseline according to (4.8) and (4.10), respectively.

solid blue line in the top plot depicts the consumption of the controlled units measured by their SMs on the timescale of 5 min. The scaled reference consumption (solid red) is used to compute the baseline (dashed black) relative to which the load reduction is measured, see bottom plot in Figure 4.6. The data show that the majority of the HPs turns off within the first 5 min of the throttling period.

We evaluate the accuracy of the load-reduction service by means of its APE defined as

$$\mathrm{APE}_{\mathrm{red}} := \left| (\hat{\bar{p}}_{\mathrm{red}}^{(\mathcal{J}_{\mathrm{ctrl}})} - \bar{p}_{\mathrm{red}}^{(\mathcal{J}_{\mathrm{ctrl}})}) / \bar{p}_{\mathrm{red}}^{(\mathcal{J}_{\mathrm{ctrl}})} \right| . \tag{4.9}$$

In the experiment shown in Figure 4.6, $\bar{p}_{\text{red}}^{(\mathcal{J}_{\text{ctrl}})} = 159.1$ kW was achieved, which, compared with the predicted load reduction of $\hat{p}_{\text{red}}^{(\mathcal{J}_{\text{ctrl}})} = 177.5$ kW, yields an $\text{APE}_{\text{red}} = 11.6\%$.

4.4.2 *Load rebound*

In the case of TCLs, load reductions are typically followed by a rebound, during which the aggregate load is higher than usual because the devices are consuming additional power to restore nominal conditions as quickly as possible. The rebound power during time step k is defined relative to the baseline as

$$p_{\text{reb},k}^{(\mathcal{J}_{\text{ctrl}})} := p_k^{(\mathcal{J}_{\text{ctrl}})} - p_{\text{base},k}^{(\mathcal{J}_{\text{ctrl}})}, \tag{4.10}$$

for all time slots following a response period, i.e., $k \in \mathcal{N} : kt_s > t_{\text{end}}$. A typical load rebound is shown in Figure 4.6. After the HP throttling signal is released at 23.00 h, most HPs turn back on to restore their nominal energy levels. This synchronized power consumption leads to a load rebound. A key measure of the rebound is its peak value

$$p_{\text{reb,max}}^{(\mathcal{J}_{\text{ctrl}})} := \max_k \{p_{\text{reb},k}^{(\mathcal{J}_{\text{ctrl}})}\},\ k \in \mathcal{N} : kt_s > t_{\text{end}}. \tag{4.11}$$

We evaluate the accuracy of the peak rebound prediction (4.6) by means of the APE defined as

$$\text{APE}_{\text{reb,max}} := \left| (\hat{p}_{\text{reb,max}}^{(\mathcal{J}_{\text{ctrl}})} - p_{\text{reb,max}}^{(\mathcal{J}_{\text{ctrl}})}) / p_{\text{reb,max}}^{(\mathcal{J}_{\text{ctrl}})} \right|. \tag{4.12}$$

In the experiment in Figure 4.6, $p_{\text{reb,max}}^{(\mathcal{J}_{\text{ctrl}})} = 128.1$ kW was observed, which, compared with the predicted peak rebound of $\hat{p}_{\text{reb,max}}^{(\mathcal{J}_{\text{ctrl}})} = 105.6$ kW, yields an $\text{APE}_{\text{reb,max}} = 17.5\%$.

Different strategies have been proposed [113, 121] to avoid the undesirable features of rebounds, such as sharp load ramps and high demand peaks. We introduce a rebound damping strategy similar to the *sequential equipment recovery* described in [121], in which the throttling signals of individual HPs are released sequentially rather than simultaneously to avoid synchronized switch-ons of HPs. Consider a group of HPs $\mathcal{J}_T \subseteq \mathcal{J}$ that have been throttled during the time interval

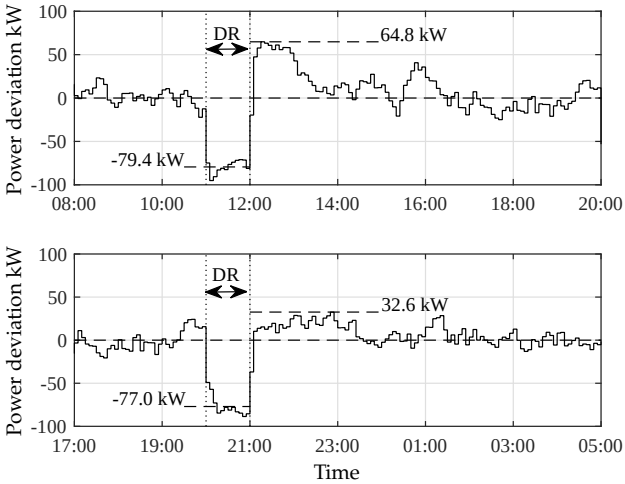

Figure 4.7: Results from two 1-h load-reduction experiments involving the same group of 74 buildings. Releasing the throttling signal simultaneously results in a pronounced rebound (top). The rebound damping strategy (4.13) can be used to significantly reduce the rebound by spreading out the throttling release times of individual HPs (bottom).

$[t_{\text{start}}, t_{\text{end}} = t_{\text{start}} + T]$. Instead of releasing the throttling signals for all HPs simultaneously at time t_{end}, the throttling signal applied to HP $j \in \mathcal{J}_T$ is released only at time

$$t_{\text{rel}}^{(j)} = \min\{t_{\text{end}} + \Delta T,\ t_{\text{start}} + T_{\text{max}}^{(j)}(\theta_{\text{out}})\}, \tag{4.13}$$

where the parameter $\Delta T \geq 0$ bounds the time interval over which the release times can be dispersed.

Figure 4.7 compares two load-reduction experiments involving the same group of 74 buildings and HPs. In the experiment shown in the top plot, the throttling signals are released simultaneously at 12.00 h, i.e., $\Delta T = 0$, which results in a sharp load ramp and a peak rebound of 64.8 kW. In the second experiment, shown in the bottom plot, $\Delta T = 45$ min was used to spread the individual release times (4.13) over a time interval of at most 45 min duration following the throttling period. This

reduced the peak rebound by 50% to values below 32.6 kW but also prolonged the duration of the rebound.

A prerequisite for our rebound damping strategy is the ability to control the HPs individually. This is possible in a DLC, but not in an ILC scheme. In addition, the rebound damping strategy benefits from the fact that we consider the flexibility of individual HPs and take their heterogeneity into account: a wide range of $T_{\max}^{(j)}(\theta_{\text{out}})$ values in (4.13) yields release times that are well spread out and, consequently, can result in significant rebound damping.

4.4.3 *Size and accuracy of load reduction*

Different groups of buildings have been used to achieve load reductions in the range of 35–288 kW, which correspond to 29–90% of the aggregate load, see Figure 4.8. 57 of the 67 load-reduction experiments resulted in a load reduction in the range of 40–65%, achievable over a broad air-temperature range. Four experiments yielded load reductions below 40%, and three experiments resulted in exceptionally high reductions (76–90%). The throttling periods of the latter experiments coincided with periods of high photovoltaic generation, which covered the lion's share of the buildings' aggregate load. Thus, the remaining load was dominated by the HP consumption.

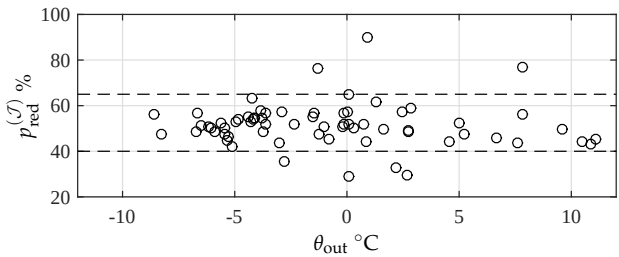

Figure 4.8: Aggregate load reductions $\bar{p}_{\text{red}}^{(\mathcal{J}_{\text{ctrl}})}$ in percent of the aggregate load achieved in 67 experiments for different outdoor air temperatures θ_{out}. 57 experiments resulted in load reductions in the range of 40–65% of the aggregate load.

The expected load reduction and rebound values computed via (4.4) and (4.7) were compared with the actual values observed in the experiments. We assess the load-reduction and rebound prediction accuracy by means of their APEs (4.9) and (4.12), respectively, whose statistics are summarized in Table 4.1. The average load reduction can be predicted accurately with a median APE of 6.7%. Predicting the peak rebound power is more challenging because it depends heavily on the degree of synchronization among the HPs when resuming operation.

Our experimental results show that a population of heterogeneous residential buildings and HPs can provide significant amounts of precisely predictable load-reduction services, and that rebound peaks can be reduced efficiently.

Comparing the results of different DR implementations is challenging for several reasons. First, many evaluations of DR report exclusively on the absolute size of the DR achieved. For example, in [113], the electricity consumption of 32 commercial buildings was controlled to reach a maximum load reduction of 35.9 kW. In [122], a population of 54 HPs was used to track a power references varying in the range of 15–35 kW. However, these numbers depend primarily on the type, size and number of the units used and do not allow us to draw any conclusions regarding the efficiency and the accuracy of a DR scheme. To provide a basis for future comparisons, we provide both absolute and relative load-reductions values. Second, the accuracy of the DR predictions and services is almost never assessed in the literature. A notable exception is [107], where the daily energy shifting potential of a population of 10 fridges is predicted with a median APE of 5.9%. We also use the APE to evaluate our load-reduction and rebound predictions, see Table 4.1.

Table 4.1: Load-reduction and rebound prediction error statistics.

	Avg. load reduction	Peak rebound
	$\mathrm{APE}_{\mathrm{red}}$	$\mathrm{APE}_{\mathrm{reb,max}}$
0.25-quantile	4.3%	12.5%
median	6.7%	22.9%
0.75-quantile	12.3%	40.5%

Finally, many DR measures, such as the load reduction and the rebound, are defined relative to a baseline which cannot be measured directly but must be estimated. Consequently, the DR evaluation depends on the accuracy of the baseline, and a DR estimate can only be as accurate as the baseline [119]. For this reason, we argue that any results on the prediction accuracy of DR measures defined relative to a baseline must be accompanied by corresponding measures of the baseline accuracy.

4.5 CONCLUSION & FUTURE WORK

This chapter presented the results from a large-scale demonstration of a direct-control DR scheme involving a population of more than 300 residential buildings with HPs. The electric flexibility of individual units was estimated autonomously from energy meter data and outdoor air temperature measurements only. Based on the flexibility of individual units, we developed a probabilistic characterization of the aggregate DR behavior of the group of units.

The results from various real-life experiments demonstrate that load reductions of 40–65% of the aggregate load can be achieved by throttling the HPs. The load-reduction services were predicted and delivered with a median APE of below 7%. In addition, we proposed a strategy to damp the rebound power and showed that it can reduce the rebound peak by 50% in practice. These results illustrate that with a minimalist and low-cost sensor and control equipment it is possible to efficiently exploit the flexibility of residential buildings and HPs to deliver significant amounts of precisely predictable load-reduction services.

Future work will improve the reliability of the communication and control setup by using the digital output of the SM to control the HP directly, for example. In addition, it will investigate how the flexibility level of individual units can be increased to maximize the flexibility available without compromising user comfort. Finally, the quantification of the rebound will be studied in more detail.

5 AGGREGATION OF FLEXIBILITY

To fully leverage the flexibility available from multiple FEUs, their aggregate flexibility must be quantified. This chapter introduces an approximation of the flexibility of individual energy units that is based on zonotopic sets. This approximation allows us to efficiently aggregate the flexibility of large numbers of units and to reduce the complexity of making control and market decisions on the aggregate level.

5.1 INTRODUCTION

The flexibility of small-scale DEUs, on both the demand and the supply side, can serve as an additional resource to the large-scale supply flexibility that has been traditionally used to balance the power grid. Sources of small-scale flexibility include heating, ventilation, and air-conditioning systems, PEVs and stationary batteries, as well as micro-generation units. However, quantifying and making use of the flexibility of large numbers of DEUs is challenging, in particular because of the computational complexity. To take full advantage of the flexibility available and to reduce the complexity of planning, trading and control, the aggregate flexibility of an entire group of units must be computed and represented in a concise and compact form. The importance of aggregation as a key enabler for incorporating large numbers of flexible units has been highlighted in numerous works, see [38, 41, 43, 123], among others. The tasks of collecting, aggregating and controlling the flexibility of a group of FEUs are performed by AGs [39, 42, 123]. The AG concludes contracts with each individual FEU, defining the way of communication and control between the unit and the AG, the specifics of the flexibility offered by the unit, as well as the details on how the unit is reimbursed for the flexibility provided [38, 39, 42, 45]. Thus, the AG acts as an intermediary between a group of FEUs and wholesale markets, molding individual flexibilities into tradable products.

Various approaches have been proposed to characterize the electric flexibility of individual FEUs, see Section 2.1. Regarding the description of the flexibility of an entire population of FEUs, two main types of approaches can be identified. Top-down approaches attempt to capture the aggregate flexibility directly based, for example, on the probabilistic properties of the underlying units [95, 124, 125]. In contrast, here we consider a bottom-up approach in which the flexibility of each individual unit is described first and aggregated subsequently. The particular way individual feasible sets are described determines the computational complexity of aggregating these sets. In general, computing the exact aggregate feasible set is computationally intractable [40, 43, 80]. Consequently, attempts to compute the exact aggregate feasible set are restricted to special cases, such as symmetric constraints [81] or polytopes with binary facet normals [40, 85]. Other approaches use set descriptions that provide inner- or outer-approximations of the exact aggregate feasible set. Outer-approximations have been derived based on polytopes [43, 79] or semi-definite constraints [126]. While some outer-approximations can be computed efficiently, the fact that they can contain infeasible power trajectories is a major drawback. In contrast, inner-approximations provide sufficient conditions for feasibility. Both inner- and outer-approximations for battery models are provided in [75, 81]. Moreover, approximation approaches based on clustering techniques have been proposed to group similar feasible sets and perform aggregation within clusters [75, 111].

This work builds on the above; in particular, it uses the polytopic feasible sets introduced in Section 2.3 as the basic description of flexibility for which zonotopic inner-approximations are computed. Zonotopes can be aggregated efficiently with regard to memory and computational complexity. Our zonotope approach allows for time-variant, asymmetric power and energy constraints, and is applicable for long planning horizons and large heterogeneous populations. We present different optimization problems that compute optimal zonotopic approximations.

This chapter is organized as follows. Section 5.2 introduces the aggregation concept considered. In Section 5.3, we discuss the challenges related to the aggregation of feasible sets and introduce zonotopic sets. They are used in Section 5.4 to approximate polytopic feasible sets in an optimal manner. We conclude our work in Section 5.5.

5.2 PROBLEM FORMULATION

We consider an AG that manages the electric energy generation and consumption of a population of FEUs indexed by $j \in \mathcal{J} := \{1, \ldots, J\}$ over a planning horizon comprising N time steps. Every unit provides the AG with a description of its flexibility in terms of its feasible set $\mathcal{P}^{(j)} \subset \mathbb{R}^N$. In addition, each unit specifies a cost function $T^{(j)} : \mathcal{P}^{(j)} \to \mathbb{R}$ that assigns to every feasible power trajectory the cost of implementing it. The AG collects this information from each individual unit and computes the amount of flexibility available on the aggregate level, $\mathcal{P}^{(\mathcal{J})}$, and the corresponding costs, $T^{(\mathcal{J})}$. We refer to this process as *aggregation*, and show how it can be executed efficiently in the case where the individual feasible sets $\mathcal{P}^{(j)}$ are approximated by zonotopic sets. Based on the aggregate flexibility and costs, the AG decides how to optimally make use of the flexibility available when bidding into wholesale energy and reserve markets. The aggregation setup is illustrated in Figure 5.1.

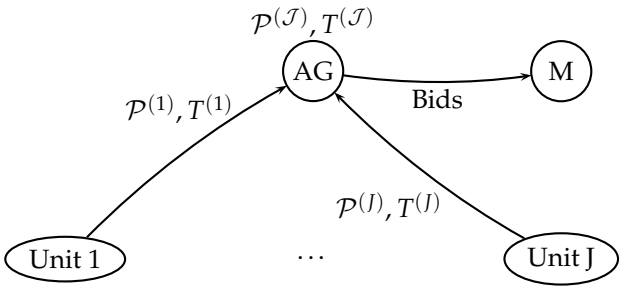

Figure 5.1: The aggregation setup considered in this work. The AG collects from each unit the description of its feasible set $\mathcal{P}^{(j)}$ and the corresponding tracking cost function $T^{(j)}$, and computes the aggregate descriptions $\mathcal{P}^{(\mathcal{J})}$ and $T^{(\mathcal{J})}$. The AG can interact with energy and reserve markets (M).

5.3 FEASIBLE SETS AND THEIR AGGREGATION

To reduce the complexity of making decisions on the aggregate level, it is desirable to aggregate the feasible sets $\mathcal{P}^{(j)}$ provided by the individual FEUs. A power trajectory $p^{(\mathcal{J})} \in \mathbb{R}^N$ that can be followed collectively by the population of J units is called *aggregate feasible*. The set of aggregate feasible power trajectories is denoted by $\mathcal{P}^{(\mathcal{J})}$ and is referred to as the *aggregate feasible set*. It is given by the *Minkowski* sum, denoted by $\oplus$, of all individual feasible sets [81, 126]:

$$\begin{aligned}\mathcal{P}^{(\mathcal{J})} &:= \mathcal{P}^{(1)} \oplus \cdots \oplus \mathcal{P}^{(J)} \\ &= \left\{ p \in \mathbb{R}^N : p = \sum_{j \in \mathcal{J}} p^{(j)},\ p^{(j)} \in \mathcal{P}^{(j)} \right\}. \end{aligned} \tag{5.1}$$

The complexity of computing $\mathcal{P}^{(\mathcal{J})}$ is determined by the particular type of description used. Here we consider the cases in which the $\mathcal{P}^{(j)}$ are given as polytopes or zonotopes.

5.3.1 *Feasible sets described by polytopes*

When the feasible sets $\mathcal{P}^{(j)}$ are given by polytopes, the aggregate feasible set (5.1) is also a polytope [127]. Different methods to compute the *Minkowski* sum of polytopes exist, for example by computing the convex hull of the pairwise sum of all vertices of the polytopes or by projecting an augmented polytope to a lower-dimensional subspace [80, 127, 128]. However, no efficient algorithm exists to compute the *Minkowski* sum of general polytopes in hyper-plane representation [127, 128]. Consequently, attempts to compute the *Minkowski* sum of polytopes are restricted to special types of polytopes (symmetric, binary hyperplane normal vectors) [81, 85], are limited to low dimensions N, or compute approximations of the true sum [75, 80].

5.3.2 *Feasible sets described by zonotopes*

Zonotopes are a subclass of polytopes. Their particular properties make them well suited for various applications where low memory

requirements of set descriptions and low time complexity of operations such as the *Minkowski* addition are crucial. Examples of zonotope applications include [129–131], where zonotopes are used for the representation and computation of reachable sets of dynamic systems, and [132], where zonotope describe bounding volumes for the purpose of collision detection.

We propose to use zonotopes to approximate polytopic feasible sets (2.8) with the purpose of making the description of individual feasible sets as well as their aggregation tractable for high dimensions and large populations of units. Zonotopes have several interesting characteristics. The ones most relevant to this work are the following:

1. Zonotopes are centrally symmetric polytopes in the sense that they are invariant under point reflection through their center. This allows a zonotope $\mathcal{Z}$ to be expressed in terms of its center $c \in \mathbb{R}^N$ and generator vectors $g^{(i)} \in \mathbb{R}^N$ with unit length, i.e., $\|g^{(i)}\|_2 = 1$, $i = 1, \dots, n_g$, as

$$\mathcal{Z} = \{x \in \mathbb{R}^N : x = c + G\beta,\ -\bar{\beta} \leq \beta \leq \bar{\beta}\}, \tag{5.2}$$

where the generators are summarized in the generator matrix

$$G := [g^{(1)}, \dots, g^{(n_g)}] \in \mathbb{R}^{N \times n_g},$$

where the number of generators is denoted by n_g. The vector of scaling factors $\beta \in \mathbb{R}^{n_g}$ is symmetrically bounded by $\bar{\beta}$. The shorthand $\mathcal{Z}(G, c, \bar{\beta})$ is used to denote such a zonotope. The generator vectors $g^{(i)}$ define the directions in which the zonotope can extend, whereas $\bar{\beta}$ determines its extent along the generator directions. Figure 5.2 provides a two-dimensional example of a zonotope that is constructed from the set of generators $\{g^{(1)}, g^{(2)}, g^{(3)}\}$ centered at c. Other examples of zonotope construction can be found in [129].

2. We refer to (5.2) as the *G-representation* of a zonotope. It is a particularly memory-efficient description compared to the vertex- and halfspace-representations of general (convex) polytopes.

3. Boxes are a subset of zonotopes: e.g., the two-dimensional axis-aligned box $\mathcal{B} := \{p \in \mathbb{R}^2 : l \leq p \leq u\}$ is a zonotope with center $c = (l + u)/2$, generators $g^{(1)} = [1, 0]^\top$ and $g^{(2)} = [0, 1]^\top$, and bounds $\bar{\beta} = (u - l)/2$.

4. The *Minkowski* sum of zonotopes can be computed efficiently with complexity $\mathcal{O}(N)$ with N being the system dimension [130]. The aggregate zonotope $\mathcal{Z}^{(\mathcal{J})}$ of the individual sets $\mathcal{Z}(G, c^{(j)}, \bar{\beta}^{(j)})$, $j \in \mathcal{J}$, is given explicitly by $\mathcal{Z}^{(\mathcal{J})} = \mathcal{Z}(G, c^{(\mathcal{J})}, \bar{\beta}^{(\mathcal{J})})$ with

$$c^{(\mathcal{J})} := \sum_{j \in \mathcal{J}} c^{(j)}, \qquad \bar{\beta}^{(\mathcal{J})} := \sum_{j \in \mathcal{J}} \bar{\beta}^{(j)}. \tag{5.3}$$

The structure of (5.3) is particularly convenient for AGs. First, $\mathcal{Z}^{(\mathcal{J})}$ is a zonotope because zonotopes are closed under *Minkowski* addition [131]. Second, if individual units are added to or removed from the population or if individual feasible sets change, the aggregate feasible set can be adjusted efficiently by updating the arithmetic sums (5.3) accordingly.

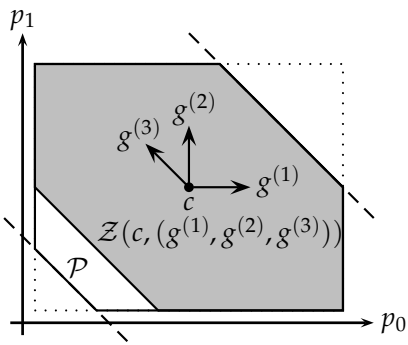

Figure 5.2: A typical polytopic feasible set $\mathcal{P}$ of a PE-unit in two dimensions. An inner-approximating zonotope $\mathcal{Z}(c, (g^{(1)}, g^{(2)}, g^{(3)}))$ is shown (shaded) with generators chosen according to (5.4).

In the following, we motivate the choice of a particular generator matrix such that the family of zonotopes it generates is able to approximate a given polytopic feasible set as accurately as possible. We restrict ourselves to the family of PE-units introduced in Section 2.3. Figure 5.2 shows a two-dimensional example. The power constraints define a box whose lower left and upper right vertices are

cut off by the energy constraints. By choosing three generators as $g^{(1)} = [1, 0]^\top$, $g^{(2)} = [0, 1]^\top$, $g^{(3)} = [-1/\sqrt{2}, 1/\sqrt{2}]^\top$, the zonotope $\mathcal{Z}$ can reproduce all the facets of the polytope $\mathcal{P}$. However, the zonotope cannot approximate the polytope perfectly due to symmetry restrictions. For arbitrary dimension $N \in \mathbb{N}$, a total of $2N - 1$ generators are constructed as

$$\begin{aligned} g^{(i)} &:= [0, \dots, 0, \overbrace{1}^{i}, 0, \dots, 0]^\top \in \mathbb{R}^N, \text{ and} \\ g^{(N+i')} &:= [0, \dots, 0, \underbrace{-1/\sqrt{2}}_{i'}, \underbrace{1/\sqrt{2}}_{i'+1}, 0, \dots, 0]^\top \in \mathbb{R}^N, \end{aligned} \tag{5.4}$$

with $i = 1, \dots, N$ and $i' = 1, \dots, N - 1$. The corresponding generator matrix is denoted by $G^{PE} \in \mathbb{R}^{N \times (2N-1)}$. It is shown in Section 5.6 that zonotopes with generator matrix G^{PE} can reconstruct every possible facet of a PE-polytope. However, the approximation quality is influenced by the fact that zonotopes, in general, will feature additional facets that do not exist in a PE-polytope.

5.4 COMPUTING ZONOTOPIC APPROXIMATIONS

In this section, it is shown how to compute zonotopes that inner-approximate a given polytope and that are optimal with regard to a specific objective. Given a polytopic feasible set $\mathcal{P}(A, b)$ and a generator matrix G, different objectives can be used to compute optimal zonotopes. The choice of objective depends on the purpose for which the aggregate feasible set will be used.

5.4.1 *Zonotope inclusion constraints*

A key ingredient to the optimization problems formulated below is a set of necessary and sufficient conditions under which a zonotope $\mathcal{Z}$ with generator matrix $G \in \mathbb{R}^{N \times n_g}$ is included in a polytope $\mathcal{P}$, i.e. $\mathcal{Z} \subseteq \mathcal{P}$.

The derivation that follows is based on the box-inclusion constraints used in [133]. We consider

$$\mathcal{P} = \{p \in \mathbb{R}^N : Ap \leq b\},\ A \in \mathbb{R}^{m \times N},\ b \in \mathbb{R}^m, \tag{5.5}$$

$$\mathcal{Z} = \{p \in \mathbb{R}^N : p = c + G\beta,\ -\bar{\beta} \leq \beta \leq \bar{\beta}\} \tag{5.6}$$

and denote the ith row of matrix A by A_i, and the ith element of vector b by b_i. The absolute value $|\cdot|$ is taken element-wise.

Claim 1. $\mathcal{Z} \subseteq \mathcal{P} \Leftrightarrow Ac + |AG|\bar{\beta} \leq b$.

Proof.

$$\begin{aligned}
&\mathcal{Z} \subseteq \mathcal{P} \\
\Leftrightarrow\ &\mathcal{Z} \subseteq \{p \in \mathbb{R}^N : A_i p \leq b_i\}, && \forall i = 1, \ldots, m \\
\Leftrightarrow\ &A_i(c + G\beta) \leq b_i, && \forall \beta : -\bar{\beta} \leq \beta \leq \bar{\beta} \\
& && \forall i = 1, \ldots, m \\
\Leftrightarrow\ &A_i c + |A_i G|\bar{\beta} \leq b_i, && \forall i = 1, \ldots, m \\
\Leftrightarrow\ &Ac + |AG|\bar{\beta} \leq b && (5.7)
\end{aligned}$$

□

Note that the zonotope inclusion constraints (5.7) are linear in the zonotope parameters $c, \bar{\beta}$.

5.4.2 *Objective 1: Constraints approximation*

Approximating the power and energy constraints of a PE-unit as accurately as possible is a natural objective for finding approximating zonotopes. The shape of a zonotope $\mathcal{Z}$ with given generator matrix G is determined by the offsets $\delta_{\mathcal{Z}}$ of the hyperplanes that define its facets, see Figure 5.3. Based on G, all the normal vectors $f^{(i)}$, $i = 1, \ldots, n_f$, of possible zonotope facets can be computed [129]. Because G^{PE} is fixed, this computation has to be executed only once. For every normal vector $f^{(i)}$, $i = 1 \ldots, n_f$, two Linear Programs (LPs) are solved to find the offsets of the corresponding hyperplanes that touch the original

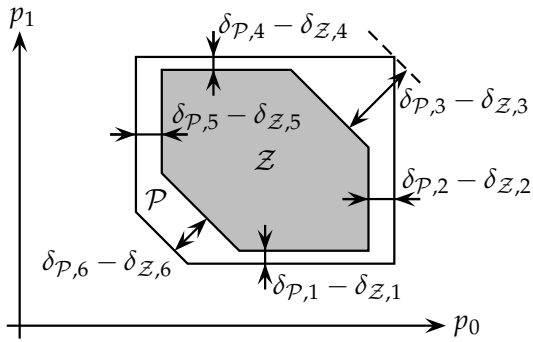

Figure 5.3: A polytopic feasible set $\mathcal{P}$ and inscribed zonotope $\mathcal{Z}$. The approximation quality measure introduced by (5.9) is to minimize the norm of the difference vector $(\delta_{\mathcal{P}} - \delta_{\mathcal{Z}})$. Its entries are the differences $\delta_{\mathcal{P},i} - \delta_{\mathcal{Z},i}$, $i = 1, \ldots, 2n_f$, between $\mathcal{Z}$ and corresponding parallel hyperplanes touching $\mathcal{P}$.

feasible set $\mathcal{P}$, as illustrated by the dashed line in Figure 5.3. For a zonotope $\mathcal{Z}$, the offsets $\delta_{\mathcal{Z}} \in \mathbb{R}^{2n_f}$ are given as

$$\delta_{\mathcal{Z}} = \begin{bmatrix} F^\top \\ -F^\top \end{bmatrix} c + \begin{bmatrix} W \\ W \end{bmatrix} \bar{\beta}, \tag{5.8}$$

where $F \in \mathbb{R}^{N \times n_f}$ is a matrix that contains the normal vectors of all possible zonotope facets as columns, i.e., $F := [f^{(1)}, \ldots, f^{(n_f)}]^\top \in \mathbb{R}^{N \times n_f}$, with $||f^{(i)}||_2 = 1$, $i = 1, \ldots, n_f$, and n_f denoting the number of pairs of parallel zonotope facets. For zonotopes with generators (5.4) it holds that $n_f = (N^2 + N)/2$, see Section 5.6. We define $W := |F^\top G^{\mathrm{PE}}| \in \mathbb{R}^{n_f \times n_g}$ [129].

Our objective is to find the zonotope that minimizes a norm of the offset differences $(\delta_{\mathcal{P}} - \delta_{\mathcal{Z}})$ between zonotope and polytope, see Figure 5.3. The corresponding optimization problem is:

$$\min_{\{c, \bar{\beta}\}} \quad \|\delta_P - \delta_Z\|_\ell \tag{5.9a}$$

$$\text{s.t.} \quad Ac + |AG|\bar{\beta} \leq b, \tag{5.9b}$$

$$\bar{\beta} \geq 0, \tag{5.9c}$$

where $\ell \in \{1, 2, \infty\}$. The objective (5.9a) can be extended by adding a vector of scaling factors to prioritize the approximation accuracy of certain types of constraints, such as for example the power constraints, or to focus on accurately approximating some constraints during select time steps. Problem (5.9) is convex and, depending on the choice of the norm ℓ, can be rewritten either as an LP or a Quadratic Program (QP), both of which can be solved efficiently by standard solvers.

How well a zonotope can approximate a given polytope depends on a number of factors, including the choice of generator matrix, the shape of the polytope (in particular its symmetry properties), and, in case the approximation is based on optimization, the objective function chosen. Moreover, the number of possible facets of the zonotope and the polytope grows with the number of dimensions N and also affects the approximation quality. Here, we investigate the approximation quality of zonotopes with generator matrix G^{PE} for different polytopes with regard to objective (5.9a). We consider a PE-unit with time-invariant power and energy constraints over a planning horizon of $N = 24$ time steps. The minimum and maximum rated power, $\underline{p}$, $\bar{p}$, as well as the minimum energy level, $\underline{e}$, of the unit are fixed, while the upper energy bound $\bar{e}$ and the initial energy level e_0 vary between experiments. In the simulations, different choices of the parameters

$$\epsilon := \frac{e_0 - \underline{e}}{(\bar{e} - \underline{e})} \qquad \text{and} \qquad \mu := \frac{(\bar{e} - \underline{e})}{T^H(\bar{p} - \underline{p})} \tag{5.10}$$

are used to generate various polytopes, see Table 5.2. The normalized initial energy level, i.e., the initial SOC, is given by $\epsilon \in [0, 1]$. Due to

Variable	Description	Value range
$\underline{p}$	Minimum rated power	-5 kW
$\bar{p}$	Maximum rated power	5 kW
T^H	Duration of planning horizon	24 h
t_s	Sampling time	1 h
ϵ	Normalized initial energy level (SOC)	$\{0, 0.5\}$
μ	Ratio of capacity and power rating	$\{0, 0.1, \ldots, 1\}$

Table 5.1: Parameters of PE-units.

symmetry, it is sufficient to consider $\epsilon \in [0, 0.5]$ only, where $\epsilon = 0.5$ results in a symmetric polytope. The ratio μ relates the capacity of the unit to its power rating. A measure of how much the unit is dominated by the energy constraints is given by

$$\lambda := \frac{\min\{e_0 + T^H \bar{p}, \bar{e}\} - \max\{e_0 + T^H \underline{p}, \underline{e}\}}{T^H(\bar{p} - \underline{p})}, \tag{5.11}$$

where coinciding energy bounds, i.e. $\underline{e} = \bar{e}$, yield $\lambda = 0$, and $\lambda = 1$ holds for units with redundant energy constraints.

For every combination of (ϵ, μ), problem (5.9) with $\ell = 1$ was solved twice: once with $G = G^{\mathrm{PE}}$ to compute the optimal zonotope, and once using $G = I_N$ to obtain the optimal axis-aligned box.

Let the optimal objective value of (5.9) scaled by the number of possible zonotope facets n_f be denoted by $\Delta_{\mathcal{BP}}$ for the box, and by $\Delta_{\mathcal{ZP}}$ for the zonotope:

$$\Delta_{\mathcal{BP}} := \frac{\|\delta_{\mathcal{P}} - \delta_{\mathcal{B}}\|_1}{n_f}, \qquad \Delta_{\mathcal{ZP}} := \frac{\|\delta_{\mathcal{P}} - \delta_{\mathcal{Z}}\|_1}{n_f}. \tag{5.12}$$

Above offset differences are a measure of how much, on average, the box and the zonotope underestimate the polytope constraints. Figure 5.4 shows the average offset differences obtained from boxes (dashed lines) and zonotopes (solid lines) for different initial energy levels ϵ and degrees of energy boundedness λ. For $(\epsilon = 0.5, \lambda = 1)$, the polytope is a hypercube and, thus, can be perfectly approximated by both the box and the zonotope. In the extreme case $\lambda = 0$ the energy bounds coincide and the polytope is approximated perfectly by the zonotope but not by the box. For $\mu = 0.5$, the maximum achievable λ-value is $\lambda = 0.5$.

Because the family of zonotopes considered includes axis-aligned boxes, the approximation quality of the zonotopes is always at least as good as for the boxes. Averaged over λ, the zonotopes yield offset errors that are approximately 26% lower compared to boxes for both choices of ϵ.

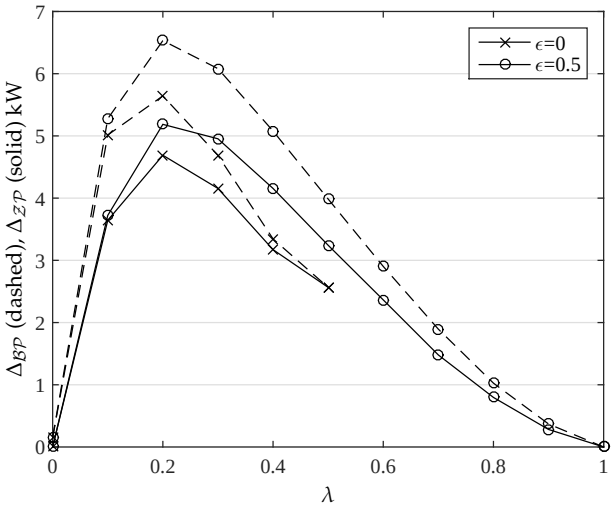

Figure 5.4: Average offset differences between box and polytope, $\Delta_{\mathcal{BP}}$ (dashed), and between zonotope and polytope, $\Delta_{\mathcal{ZP}}$, (solid) for different initial energy levels ϵ and degrees of energy boundedness λ.

5.4.3 *Objective 2: Shape approximation*

Another useful approximation criterion is to compare the dimensions of the zonotope along its generator vectors with the corresponding dimensions of the polytope, see Figure 5.5. The distances between parallel zonotope facets are given as $\Delta_{\mathcal{Z}} = 2|F^\top G|\bar{\beta}$. Similarly, the extension of the polytope in the direction of the ith zonotope facet-normal is denoted by $\Delta_{\mathcal{P},i}$, $i = 1, \ldots, n_f$, and computed in a preprocessing step by solving $2n_f$ LPs. Let $\Delta_{\mathcal{P}} := [\Delta_{\mathcal{P},1}, \ldots, \Delta_{\mathcal{P},n_f}]^\top$. The zonotope approximates the polytope perfectly if and only if $\Delta_{\mathcal{Z}} = \Delta_{\mathcal{P}}$. Thus, a convenient and informative approximation quality measure is

$$\Lambda_{\mathcal{Z}} := \frac{1}{n_f} \sum_{i=1}^{n_f} \frac{\Delta_{\mathcal{Z},i}}{\Delta_{\mathcal{P},i}} \in [0, 1], \tag{5.13}$$

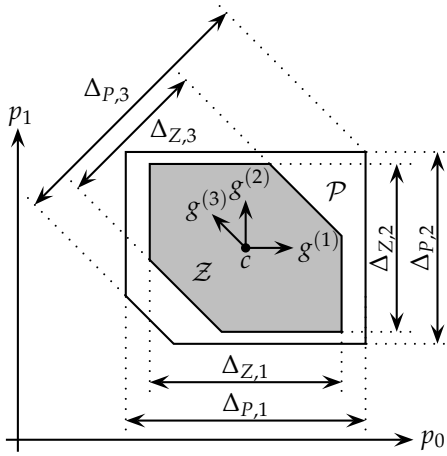

Figure 5.5: A PE-polytope $\mathcal{P}$ and inscribed zonotope $\mathcal{Z}$ (shaded gray) whose generator vectors $g^{(1)}$, $g^{(2)}$, and $g^{(3)}$ are chosen according to (5.4). The distances between parallel facets of the zonotope are referred to as $\Delta_{\mathcal{Z},i}$, $i = 1, \ldots, n_f$. The extension of the polytope in the direction of the ith zonotope facet-normal is denoted by $\Delta_{\mathcal{P},i}$.

where $\Lambda_{\mathcal{Z}} = 0$ implies that the zonotope is a point, and $\Lambda_{\mathcal{Z}} = 1$ is reached if and only if the polytope is approximated perfectly by the zonotope. The inner-approximating zonotope maximizing (5.13) can be found by solving the following LP:

$$\max_{\{c,\bar{\beta}\}} \quad \frac{2}{n_f}(1./\Delta_{\mathcal{P}})^\top |F^\top G|^\top \bar{\beta} \tag{5.14a}$$

$$\text{s.t.} \quad Ac + |AG|\bar{\beta} \leq b, \tag{5.14b}$$

$$\bar{\beta} \geq 0, \tag{5.14c}$$

where ./ denotes element-wise division. The objective (5.14a) expresses (5.13) as a linear function of $\bar{\beta}$. The constraints (5.14b)-(5.14c) guarantee that the zonotope is included in the polytope.

We assess the approximation quality for polytopic feasible sets of PEVs over a planning horizon of 24 h with 2 h sampling time. It is assumed

that the PEVs are subject to power and energy constraints only. Trip constraints can easily be translated into power constraints by setting both the upper and lower power bounds to zero for all time slots that fall into trip periods. Table 5.2 provides the parameter ranges from which the parameters of 100 PEVs are sampled uniformly. For every unit, an optimal inner-approximating axis-aligned box $\mathcal{B}^*$ and zonotope $\mathcal{Z}^*$ are computed via (5.14) using the generator matrices I_N and G^{PE}, respectively.

Averaged over the population, the quality of approximation (5.13) for the box and the zonotope are $\Lambda_{\mathcal{B}} = 0.28$ and $\Lambda_{\mathcal{Z}} = 0.62$, respectively. This shows that zonotopic approximations of the feasible set can capture a significantly larger portion of the flexibility of a PEV as compared to axis-aligned boxes. Zonotopes yield a good trade-off between an accurate but computationally complex description (in terms of polytopes) and an inaccurate but computationally simple description (in terms of axis-aligned boxes).

Variable	Description	Value range
N	Number of time steps	12
t_s	Duration of time steps	2 h
$\bar{p}$	Maximum rated power	3 kW
$\underline{p}$	Minimum rated power	-3 kW
$\bar{e} - \underline{e}$	Battery capacity	[20, 40] kWh
$(e_0 - \underline{e})/(\bar{e} - \underline{e})$	Initial SOC	[0.2, 0.8]

Table 5.2: PEV parameters.

5.4.4 *Extension: Including boxes in zonotopes*

Problems (5.9) and (5.14) may result in degenerate zonotopes with $\Delta_{\mathcal{Z},i} = 0$ for some $i \in \{1, \ldots, n_f\}$ even though $\Delta_{\mathcal{P},i} > 0$. For example, it might be preferable with regard to the objective (5.14a) to let the zonotope collapse in some directions, i.e., to set some $\bar{\beta}_i$ to zero, in favor of increasing the extent of the zonotope in other directions. This is undesirable because such zonotopes lack a significant portion of

the polytopic flexibility. To prevent such degeneracy, (5.9) and (5.14) can be augmented by constraints enforcing the zonotope to contain a full-dimensional object, e.g., a box. Here we consider axis-aligned boxes $\mathcal{B}$ which can be expressed as $\mathcal{B} = \{p \in \mathbb{R}^N : l \leq p \leq u\}$. The box edge lengths are $d = u - l \geq 0$. By augmenting (5.9) or (5.14) by the decision variable $\beta_{\mathcal{B}} \in \mathbb{R}^{n_g}$ and the affine constraints

$$c + G\beta_{\mathcal{B}} \geq d/2 \;\text{ and }\; -\bar{\beta} \leq \beta_{\mathcal{B}} \leq \bar{\beta} \tag{5.15}$$

it is guaranteed that $\mathcal{Z}(G, c, \bar{\beta})$ includes a box $\mathcal{B}$ with center c and edge lengths d. Consequently, the resulting zonotope $\mathcal{Z}(G, c, \bar{\beta})$ is full-dimensional for any full-dimensional $\mathcal{B} \subseteq \mathcal{P}(A, b)$.

Axis-aligned boxes inscribed in aggregate feasible sets play an important role for offering frequency-regulation services or for compensating disturbances affecting a balance group. Thus, it can be desirable to find zonotopic inner-approximations of feasible sets that contain a particular type of axis-aligned box, such as the *maximum-volume* axis-aligned box $\mathcal{B}^*$. We first compute the box $\mathcal{B}^* \subseteq \mathcal{P}$ and then find the zonotope $\mathcal{Z}^*$ that includes $\mathcal{B}^*$ and that is optimal with regard to a given objective.

By taking the logarithm of the product of the box edge lengths, the problem of finding the maximum-volume box contained in $\mathcal{P}(A, b)$ can be formulated as the convex optimization problem

$$\max_{\{l,u\}} \quad \sum_{i=1}^{N} \log(u_i - l_i) \tag{5.16a}$$

$$\text{s.t.} \quad A^+ u - A^- l \leq b, \tag{5.16b}$$

with $A^+_{ij} := \max\{0, A_{ij}\}$ and $A^-_{ij} := \max\{0, -A_{ij}\}$. We use *CVX*, a package for specifying and solving convex programs [134, 135], to solve problem (5.16). The constraints (5.16b) correspond to (5.7), with $G = I_N$ for an axis-aligned box.

5.5 CONCLUSION & FUTURE WORK

In this chapter, we used zonotopes to approximate the feasible sets of energy- and power-constrained units. We showed how optimal inner-approximations can be computed efficiently, and discussed the

tractability of the aggregation of zonotopic sets for long planning horizons and large populations of units.

There are several future directions of interest. We study the applicability of zonotopic approximations to units with additional constraints such as power ramp-rates or more general energy dynamics. The incorporation of uncertainty into the zonotope description is another topic of interest.

5.6 APPENDIX: FACETS OF PE-ZONOTOPES

In this section we characterize the facets of zonotopes with generators chosen according to (5.4). It follows that zonotopes can reconstruct every possible facet of a PE-polytope.

Theorem 1. *Every facet of a full dimensional zonotope $\mathcal{Z} \subseteq \mathbb{R}^N$ with generators (5.4) has normal vector $f_i = \sum\{e_i : j \leq i \leq k\}$ for some $1 \leq j \leq k \leq N$, with e_i being the unit vector of coordinate i.*

Proof. Let $\mathcal{F}$ be a facet of $\mathcal{Z}$, $\mathcal{H}$ the affine hyperplane generated by $\mathcal{F}$, and f a normal vector of $\mathcal{F}$ so that $\mathcal{F} = \{z \in \mathcal{Z} : z^\top f = M\}$ with $M = \max\{z^\top f : z \in \mathcal{Z}\}$. Let $1 \leq l \leq N$ be a coordinate which maximizes $|f_l|$. We know that $f_l \neq 0$. Let j be the minimum, and k the maximum value in $\{1, \ldots, N\}$ for which $f_i = f_l$ for all $j \leq i \leq k$. $\mathcal{Z}$ is centrally symmetric, so without loss of generality we can assume that $f_l > 0$. □

The following claim follows from (5.4).

Claim 2. *Let $z = c + \sum_{i=1}^{n_g} \beta_i g^{(i)} \in \mathcal{Z}$. If $z \in \mathcal{F}$ then*

1. $\beta_{N+j-1} = \overline{\beta}_{N+j}$ *if* $j > 1$,
2. $\beta_{N+k} = -\overline{\beta}_{N+k}$ *if* $k < N$.

For a set $\mathcal{X} \subseteq \mathbb{R}^N$ and $\mathcal{I} \subseteq \{1, \ldots, N\}$ we denote by $\mathcal{X}|_\mathcal{I}$ the projection of $\mathcal{X}$ onto $\mathbb{R}^\mathcal{I}$. Let $\mathcal{I}_1 = \{j, \ldots, k\}$ and $\mathcal{I}_2 = \{1, \ldots, N\} - \mathcal{I}_1$. $\mathcal{I}_2 = \emptyset$ is possible. For $x, y \in \mathcal{F}$ we define $\gamma(x, y) = x|_{\mathcal{I}_1} \times y|_{\mathcal{I}_2}$, the combination of x and y along the partition $\mathcal{I} = \mathcal{I}_1 \dot\cup \mathcal{I}_2$. Let

$$x = c + \sum_{i=1}^{n_g} \beta_i^{(x)} g^{(i)} \qquad \text{and} \qquad y = c + \sum_{i=1}^{n_g} \beta_i^{(y)} g^{(i)}.$$

By Claim 2 one can write $\gamma(x,y)$ as $c + \sum_{i=1}^{n_g} \beta_i g^{(i)} \in \mathcal{F}$ where

1. $\beta_i = \beta_i^{(x)}$ if $i \in \mathcal{I}_1$ and $\beta_i = \beta_i^{(y)}$ if $i \in \mathcal{I}_2$,
2. $\beta_{N+i} = \beta_{N+i}^{(x)}$ if $i \geq j$ and $i < k$,
3. $\beta_{N+i} = \beta_{N+i}^{(y)}$ if $i < j-1$ or $i > k$,
4. $\beta_{N+j-1} = \overline{\beta}_{N+j}$ if $j > 1$,
5. $\beta_{N+k} = -\overline{\beta}_{N+k}$ if $k < N$.

It follows that $\gamma(x,y) \in \mathcal{Z}$. By symmetry, $\gamma(y,x) \in \mathcal{Z}$. Clearly, $\gamma(x,y) + \gamma(y,x) = x + y$ and thus $\gamma(x,y)$, $\gamma(y,x) \in \mathcal{F}$.

Let $v_m \in \mathcal{F}|_{I_m} \subseteq \mathbb{R}^{\mathcal{I}_m}$ for $m = 1,2$. By definition, there exists a vector $z_m = c + \sum_{i=1}^{n_g} \beta_i^{(m)} g^{(i)} \in \mathcal{F}$ so that $z_m|_{\mathcal{I}_m} = v_m$ for $m = 1,2$. It is easy to see that $v_1 \times v_2 = \gamma(z_1, z_2) \in \mathcal{F}$. It follows that $\mathcal{F} = \mathcal{F}|_{\mathcal{I}_1} \times \mathcal{F}|_{\mathcal{I}_2}$, and that $\mathcal{F}|_{\mathcal{I}_m}$ is the set of points maximizing $c|_{\mathcal{I}_m}$ in $Z|_{\mathcal{I}_m}$.

$\dim(\mathcal{F}|_{\mathcal{I}_1}) \leq |\mathcal{I}_1| - 1$ because the vectors g_{N+i} for $j \leq i \leq k$ are generators, and clearly $\dim(\mathcal{F}|_{\mathcal{I}_2}) \leq |\mathcal{I}_2|$. Now $N - 1 = \dim(\mathcal{F}) = \dim(\mathcal{F}|_{I_1}) + \dim(\mathcal{F}|_{\mathcal{I}_2}) \leq |\mathcal{I}_1| + |\mathcal{I}_2| - 1 = N - 1$, and thus equality must hold throughout. $\mathcal{F}|_{\mathcal{I}_2}$, being the set of points maximizing $f|_{\mathcal{I}_2}$ in $Z_{\mathcal{I}_2}$, can only be full-dimensional if $f|_{\mathcal{I}_2} = 0$. After normalizing, we get that f has the form $\sum\{e_i : j \leq i \leq k\}$. □

Theorem 2. *Let $\mathcal{Z} \subseteq \mathbb{R}^N$ be a zonotope with generators (5.4). If $\overline{\beta} > 0$ then for all $j \leq i \leq k$ the objective vector $f = \sum\{e_i : j \leq i \leq k\}$ defines a facet.*

Proof. Let $\mathcal{F} = \{z \in \mathcal{Z} : z^\top f = M\}$ with $M = \max\{z^\top f : z \in \mathcal{Z}\}$, and $\mathcal{H}$ the affine hyperplane generated by $\mathcal{F}$. The following vectors are generators of $\mathcal{H}$:

1. $g^{(i)}$ for $i < j$ and $i > k$,
2. $g^{(N+i)}$ for $i \geq j$ and $i < k$.

The rank of this set of vectors is $N - 1$, so $\mathcal{H}$ is a hyperplane and $\mathcal{F}$ is a facet. □

It follows from Theorem 1 and Theorem 2 that the maximum number of facets of a zonotope with generators (5.4) is $N^2 + N$ and this bound is reached if $\overline{\beta} > 0$.

6 DISAGGREGATION OF FLEXIBILITY

This chapter considers the task of distributing control decisions made by an AG on the aggregate level among the individual units of a population. We present an algorithm that distributes control decisions in an economically optimal and computationally efficient way. The algorithm is applied to the problem of disaggregating reference schedules resulting from day-ahead energy markets.

6.1 INTRODUCTION

Although disaggregation is an essential problem faced by any AG of FEUs in a direct-control setting [39, 45], it has attracted limited attention. Hierarchical control schemes have been proposed in [40, 85]. Other approaches use a priority-stack [81] or randomized [95] controllers.

This work builds on the zonotopic approximation of feasible sets introduced in Chapter 5. First, it is shown how individual units can assign costs to their flexibility by means of flexibility cost functions. Second, we present an efficient disaggregation algorithm that distributes an aggregate-level signal among the units taking into account their individual flexibility cost functions. In contrast to existing approaches, our disaggregation method yields economically optimal solutions. The disaggregation algorithm proposed can be applied by AGs to check the feasibility of aggregate-level decisions, such as day-ahead energy schedules, and to compute the cost of offering regulation services.

This chapter is organized as follows. In Section 6.2, we outline the disaggregation setup considered and formulate the general disaggregation problem. The cost of flexibility is discussed in Section 6.3. In Section 6.4, a disaggregation algorithm is derived which is used in Section 6.5 to perform feasibility checks, distribute aggregate day-ahead energy schedules, and compute lower bounds on the prices of regulation power bids. A conclusion and outlook on future work are given in Section 6.6.

6.2 PROBLEM FORMULATION

6.2.1 *Disaggregation setup*

We consider the aggregation and disaggregation setup illustrated in Figure 6.1 in which an AG manages the electric energy consumption and generation of a population of FEUs indexed by $j \in \mathcal{J} := \{1, \ldots, J\}$. Every unit provides the AG with a description of its flexibility by means of its feasible set $\mathcal{P}^{(j)}$ and specifies a cost function $T^{(j)} : \mathcal{P}^{(j)} \to \mathbb{R}$ that assigns to every feasible power trajectory the cost of implementing it. Based on the aggregate flexibility and costs, the AG decides how to optimally make use of the flexibility available, see Chapter 5. For example, the AG can bid into wholesale energy and reserve markets. After the market clearing, the AG receives an aggregate reference schedule $p_{\text{ref}}^{(\mathcal{J})}$ that the population of units is committed to follow. The AG needs to check if this reference can indeed be implemented, i.e., check its feasibility, and is required to split it into the power references $p_{\text{ref}}^{(j)} \in \mathcal{P}^{(j)}$, $j \in \mathcal{J}$, that can be tracked by the individual units.

The process of distributing an aggregate signal among the individual units is referred to as *disaggregation*, and can be executed in different ways depending on the objective chosen by the AG. In this chapter, we present an efficient disaggregation algorithm that minimizes the total costs with regard to the flexibility cost functions specified. Given the electricity consumption and generation of an individual unit during the entire planning horizon, money flows either from the AG to the unit or vice versa according to the terms agreed upon by the unit and the AG.

6.2.2 *Economic disaggregation problem*

The AG concludes a contract with each individual FEU j defining the price $v_k^{(j)}$ at which the unit buys/sells electric energy from/to the AG for every time step k of a planning horizon comprising N time steps each of duration t_s. After receiving the aggregate reference schedule $p_{\text{ref}}^{(\mathcal{J})}$ from the market, the AG assigns a reference trajectory $p_{\text{ref}}^{(j)}$ to every unit, for which the unit pays the amount $t_s v^{(j)\top} p_{\text{ref}}^{(j)}$, $v^{(j)} := [v_1^{(j)}, \ldots, v_N^{(j)}]^\top$,

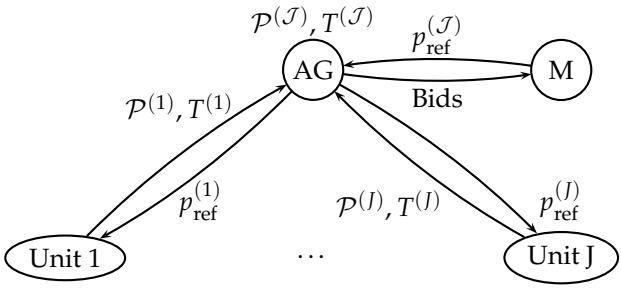

Figure 6.1: The aggregation and disaggregation setup considered in this work. Based on the aggregate description of flexibility, $\mathcal{P}^{(\mathcal{J})}$, and the aggregate cost function, $T^{(\mathcal{J})}$, the AG interacts with a market (M) from which it receives an aggregate reference trajectory $p_{\text{ref}}^{(\mathcal{J})}$ after clearing. The AG must then assign reference schedules $p_{\text{ref}}^{(j)} \in \mathcal{P}^{(j)}$ to the individual units.

to the AG. A contract design problem is studied in [45]. While the units guarantee that every $p_{\text{ref}}^{(j)} \in \mathcal{P}^{(j)}$ can be implemented, they might prefer certain trajectories over others. For example, the owner of a building might prefer those power trajectories that keep the indoor air-temperature within certain limits. Similarly, the operator of a gas power plant might favor constant power trajectories over ones that vary quickly because of ramping costs. To express particular preferences and additional operating costs, every unit defines a convex *flexibility cost function* $h^{(j)} : \mathcal{P}^{(j)} \to \mathbb{R}$. If we assume that the AG reimburses the individual units for their flexibility costs, the total costs $T^{(j)}(p^{(j)})$ of having unit j implementing the power trajectory $p^{(j)}$ are given by the difference of the flexibility costs and the energy costs:

$$T^{(j)}(p^{(j)}) := h^{(j)}(p^{(j)}) - t_s v^{(j)\top} p^{(j)}. \tag{6.1}$$

If $T^{(j)}$ is negative, money flows from the unit owner to the AG. If, in contrast, the costs are positive, the unit owner receives money from the AG.

The economic disaggregation problem the AG needs to solve is how to best distribute a given aggregate reference trajectory $p_{\text{ref}}^{(\mathcal{J})} \in \mathcal{P}^{(\mathcal{J})}$

into exactly J reference trajectories $p_{\text{ref}}^{(1)}, \dots, p_{\text{ref}}^{(J)}$ such that they are feasible and sum up to $p_{\text{ref}}^{(\mathcal{J})}$. This disaggregation is not unique in general, and the AG may want to find the one that is optimal with regard to a particular objective. A profit-oriented AG will try to find a disaggregation that minimizes the total costs by solving

$$\begin{aligned} \min_{\{p_{\text{ref}}^{(j)}\}_{j\in\mathcal{J}}} \quad & \sum_{j\in\mathcal{J}} T^{(j)}(p_{\text{ref}}^{(j)}) \\ \text{s.t.} \quad & \sum_{j\in\mathcal{J}} p_{\text{ref}}^{(j)} = p_{\text{ref}}^{(\mathcal{J})}, \\ & p_{\text{ref}}^{(j)} \in \mathcal{P}^{(j)}, \ \forall j \in \mathcal{J}. \end{aligned} \tag{6.2}$$

Problem (6.2) is convex and can, in theory, be solved efficiently. However, the problem grows with the population size J and the number of time steps N, and its solving time and memory requirements can be significant for many practical applications with polytopic feasible sets, see Subsection 6.4.2. In contrast, we consider feasible sets given by the zonotopes

$$\begin{aligned} \mathcal{Z}^{(j)} &= \mathcal{Z}(G, c^{(j)}, \bar{\beta}^{(j)}) \\ &= \{p^{(j)} \in \mathbb{R}^N : p^{(j)} = c^{(j)} + G\beta, \ -\bar{\beta}^{(j)} \leq \beta \leq \bar{\beta}^{(j)}\}, \ j \in \mathcal{J}. \end{aligned} \tag{6.3}$$

The particular structure of the *Minkowski* sum of zonotopes (5.3) allows an efficient disaggregation algorithm whose complexity scales particularly well with the number of units in the population.

6.3 COST OF FLEXIBILITY

Given a zonotope $\mathcal{Z}(G, c, \bar{\beta})$, a particular choice of the zonotope parameter $\beta \in [-\bar{\beta}, \bar{\beta}]$ uniquely defines the point $p = c + G\beta \in \mathcal{Z}$. To exploit the characteristic properties of zonotopic feasible sets, we restrict attention to flexibility costs that can be expressed as functions of β. In particular, we consider cost functions $h^{(j)}(\beta^{(j)}) : \mathbb{R}^{n_g} \to \mathbb{R}$ that can be written as

$$h^{(j)}(\beta^{(j)}) = h_1^{(j)}(\beta_1^{(j)}) + \cdots + h_{n_g}^{(j)}(\beta_{n_g}^{(j)}), \tag{6.4}$$

where each $h_i^{(j)} : \mathbb{R} \to \mathbb{R}$ is piece-wise affine convex, and $\beta_i^{(j)}$ refers to the ith-element of $\beta^{(j)}$. This family of functions can capture a wide range of flexibility costs. Figure 6.2 provides a two-dimensional example of a cost function that penalizes deviations from the zonotope center $c = [0, 0]^\top$. Every component $h_i^{(j)}(\beta_i^{(j)})$ of the cost function $h^{(j)}(\beta^{(j)})$ is piece-wise affine convex and is chosen such as to approximate a convex quadratic function. This is achieved by partitioning each range $[-\bar{\beta}_i, \bar{\beta}_i]$, $i = 1, \ldots, n_g$, into multiple segments and defining a linear cost for each segment, see Figure 6.3. This choice of cost function could be used by unit owners to express their preference on power trajectories that lie close to a desired trajectory. Many other types of flexibility costs can be modeled in this way. For example, the deviation of the energy content of a trajectory relative to a desired energy content can be penalized. This type of flexibility cost could be appropriate for heating and cooling applications where the temperature (i.e., energy) is required to stay close to its target value.

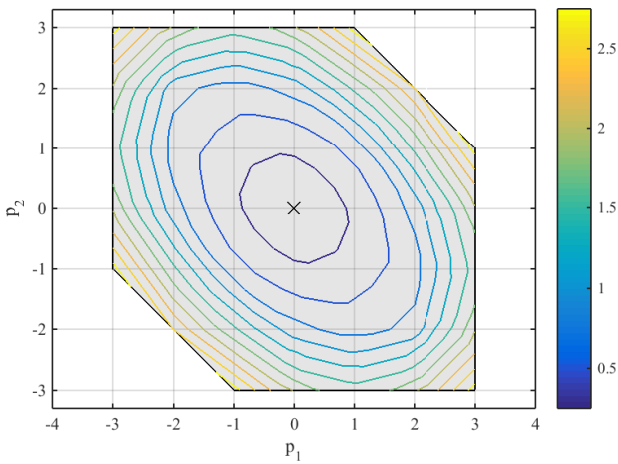

Figure 6.2: Example of a flexibility cost function $h^{(j)}(\beta^{(j)})$ that penalizes deviations from the zonotope center $c = [0, 0]^\top$ by approximating a convex quadratic function.

Substituting $p^{(j)}$ with $c^{(j)} + G\beta^{(j)}$ allows us to rewrite the implementation costs (6.1) in terms of β as

$$\begin{aligned} T^{(j)}(\beta^{(j)}) &= \sum_{i=1}^{n_g} \left(h_i^{(j)}(\beta_i^{(j)}) - t_s v^{(j)\top} g^{(i)} \beta_i^{(j)} \right) + T_{\text{fix}}^{(j)} \\ &= \sum_{i=1}^{n_g} T_i^{(j)}(\beta_i^{(j)}) + T_{\text{fix}}^{(j)}, \end{aligned} \tag{6.5}$$

where $g^{(i)}$ is the ith generator, and $T_{\text{fix}}^{(j)} := -t_s v^{(j)\top} c^{(j)}$ is a fixed cost independent of $\beta^{(j)}$. The components $T_i^{(j)}(\beta_i^{(j)})$ are piece-wise affine and convex. For the purpose of illustration, Figure 6.3 shows two example cost functions $T_i^{(1)}$ and $T_i^{(2)}$. The breakpoint interval lengths and the slopes of the line segments are denoted by $l_{i,i'}^{(j)}$ and $q_{i,i'}^{(j)}$, respectively, for $i' = 1, \ldots, m_i^{(j)}$ with $m_i^{(j)}$ being the number of line segments.

We define the *aggregate* cost $T^{(\mathcal{J})}(\beta^{(\mathcal{J})})$ as the total cost of the cheapest disaggregation of $\beta^{(\mathcal{J})}$ into $\beta^{(1)}, \ldots, \beta^{(J)}$:

$$\begin{aligned} T^{(\mathcal{J})}(\beta^{(\mathcal{J})}) := \min_{\{\beta^{(j)}\}_{j\in\mathcal{J}}} \quad & \sum_{j\in\mathcal{J}} T^{(j)}(\beta^{(j)}) \\ \text{s.t.} \quad & \sum_{j\in\mathcal{J}} \beta^{(j)} = \beta^{(\mathcal{J})}, \\ & -\bar{\beta}^{(j)} \leq \beta^{(j)} \leq \bar{\beta}^{(j)}, \ \forall j \in \mathcal{J}. \end{aligned} \tag{6.6}$$

The main reason for considering the cost functions (6.5) is the fact that this family of functions preserves its structure under aggregation, i.e., the aggregate implementation costs $T^{(\mathcal{J})}$ can be written as

$$T^{(\mathcal{J})}(\beta^{(\mathcal{J})}) = \sum_{i=1}^{n_g} T_i^{(\mathcal{J})}(\beta_i^{(\mathcal{J})}) + T_{\text{fix}}^{(\mathcal{J})}, \tag{6.7}$$

with $T_{\text{fix}}^{(\mathcal{J})} := T_{\text{fix}}^{(1)} + \cdots + T_{\text{fix}}^{(J)}$, which has the same structure as (6.5). Each component $T_i^{(\mathcal{J})}$ is piece-wise affine and convex and can be constructed efficiently by concatenating the line segments of $T_i^{(1)}, \ldots, T_i^{(J)}$

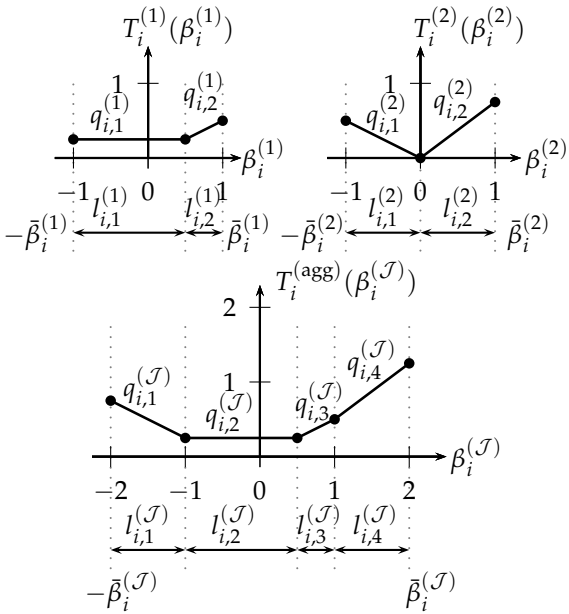

Figure 6.3: Examples of cost function components, $T_i^{(1)}$ and $T_i^{(2)}$, of two units (top). Aggregate cost function component $T_i^{(\mathcal{J})}$ derived from $T_i^{(1)}$ and $T_i^{(2)}$ by concatenating line segments ordered by ascending slopes (bottom).

sorted by ascending slopes $q_{i,i'}^{(j)}$, see Figure 6.3, where the offset of $T_i^{(\mathcal{J})}$ is

$$T_i^{(\mathcal{J})}(\bar{\beta}_i^{(\mathcal{J})}) = \sum_{j \in \mathcal{J}} T_i^{(j)}(\bar{\beta}_i^{(j)}) + T_{\text{fix}}^{(\mathcal{J})}. \tag{6.8}$$

A similar cost aggregation approach has been proposed in [59]. The fact that the aggregate costs $T^{(\mathcal{J})}$ can be computed efficiently and are parameterized by the aggregate variable $\beta^{(\mathcal{J})}$ allows to derive an efficient disaggregation algorithm to solve problem (6.2).

We introduce the auxiliary book-keeping variables $s_{i,i'}$ that store the unit index which belongs to line segment i' of the piece-wise affine function $T_i^{(\mathcal{J})}$. In the example shown in Figure 6.3, the first

line segment of $T_i^{(\mathcal{J})}$ originates from the cost function $T_i^{(2)}$ of unit 2. We encode this by setting $s_{i,1} = 2$. The second line segment of $T_i^{(\mathcal{J})}$ comes from the cost function $T_i^{(1)}$ of unit 1. Thus, we set $s_{i,2} = 1$. In this way, we know for every line segment of $T_i^{(\mathcal{J})}$ to which unit it belongs. This information is required in the disaggregation procedure discussed below, in particular when splitting the aggregate variable $\beta^{(\mathcal{J})}$ into individual $\beta^{(1)}, \dots, \beta^{(J)}$. Up to its offset, each cost component $T_i^{(\mathcal{J})}$, $i = 1, \dots, n_g$, is fully defined by the ordered list of triples $L_i := \{(l_{i,1}^{(\mathcal{J})}, q_{i,1}^{(\mathcal{J})}, s_{i,1}), \dots, (l_{i,m_i^{(\mathcal{J})}}^{(\mathcal{J})}, q_{i,m_i^{(\mathcal{J})}}^{(\mathcal{J})}, s_{i,m_i^{(\mathcal{J})}})\}$, with $m_i^{(\mathcal{J})} := m_i^{(1)} + \dots + m_i^{(J)}$. These lists are key prerequisites for the disaggregation algorithm discussed below.

6.4 DISAGGREGATION ALGORITHM

6.4.1 *Subgradient-based disaggregation algorithm*

Given an aggregate trajectory $p_{\text{ref}}^{(\mathcal{J})} \in \mathcal{Z}^{(\mathcal{J})}$, the individual feasible sets $\mathcal{Z}(G, c^{(j)}, \bar{\beta}^{(j)})$, $j \in \mathcal{J}$, and the aggregate cost function $T^{(\mathcal{J})}(\beta^{(\mathcal{J})})$, the economic disaggregation problem (6.2) can be expressed in terms of $\beta^{(J)}$ as

$$\begin{aligned} \min_{\beta^{(\mathcal{J})}} \quad & T^{(\mathcal{J})}(\beta^{(\mathcal{J})}) \\ \text{s.t.} \quad & p_{\text{ref}}^{(\mathcal{J})} = c^{(\mathcal{J})} + G\beta^{(\mathcal{J})}, \\ & -\bar{\beta}^{(\mathcal{J})} \leq \beta^{(\mathcal{J})} \leq \bar{\beta}^{(\mathcal{J})}. \end{aligned} \tag{6.9}$$

The key difference between the above problem and its original formulation (6.2) is that in (6.9) the number of decision variables (n_g), the number of equality constraints (N), and the number of inequality constraints ($2n_g$) are independent of the population size J. Problem (6.9) is convex with piece-wise affine objective and can be written as an LP using an epigraph reformulation, with the objective substituted by a set of linear inequalities. This is inconvenient for the application discussed here because the number of inequalities required to encode the objective function can be large. Instead, we solve problem (6.9) by

using a *projected subgradient method* [136]. The disaggregation algorithm proposed here comprises four steps:

1. *Initialization:* Find $\beta^{(\mathcal{J})}(0)$ in the feasible set of (6.9).
2. *Subgradient steps:* In iteration k, compute a subgradient of $T^{(\mathcal{J})}$ at $\beta^{(\mathcal{J})}(k)$ and execute a subgradient step to compute $\beta^{(\mathcal{J})}(k+1)$.
3. *Projection:* Project the new $\beta^{(\mathcal{J})}(k+1)$ back onto the feasible set of (6.9). Repeat 2) and 3) until converged, and denote the best $\beta^{(\mathcal{J})}$ found so far by $\beta^{(\mathcal{J})*}$.
4. *Assignment:* Distribute $\beta^{(\mathcal{J})*}$ into $\beta^{(1)*}, \dots, \beta^{(J)*}$.

In the following, we elaborate on the above steps. A summary of the subgradient-based disaggregation algorithm is provided in Algorithm 6.1.

1. *Initialization:* The subgradient method requires a feasible starting point $\beta^{(\mathcal{J})}(0) \in \mathbb{R}^{n_g}$. The choice $\beta^{(J)}$ is feasible if it lies in the feasible set $\mathcal{S}$ of (6.9), which is given as

$$\mathcal{S} := \{\beta \in \mathbb{R}^{n_g} : \beta \in [-\bar{\beta}^{(\mathcal{J})}, \bar{\beta}^{(\mathcal{J})}],\ p_{\text{ref}}^{(\mathcal{J})} = c^{(\mathcal{J})} + G\beta^{(\mathcal{J})}\}. \tag{6.10}$$

The starting point $\beta^{(\mathcal{J})}(0)$ can be obtained by projecting an arbitrary $\beta^{(\mathcal{J})} \in \mathbb{R}^{n_g}$ onto the feasible set $\mathcal{S}$ by using, for example, the *Euclidean* projection $Q_{\mathcal{S}}(\cdot)$.

2. *Subgradient step:* The subgradient method iteratively executes a subgradient step that updates the decision variable as

$$\beta^{(\mathcal{J})}(k+1) = \beta^{(\mathcal{J})}(k) - \alpha(k)\nabla(k), \tag{6.11}$$

where the subgradient and the step size of iteration k are denoted by $\nabla(k)$ and $\alpha(k)$, respectively. A subgradient of $T^{(\mathcal{J})}$ at $\beta^{(\mathcal{J})}(k)$ is computed element-wise according to line 6 in Algorithm 6.1.

Several convergence results for the projected sugradient method exist in the literature [136], including different step size rules that influence the convergence properties. Here, $\alpha(k) = a/k$, $a > 0$, is used, for which the subgradient method converges to the optimal value, i.e., $\lim_{k\to\infty} T^{(\mathcal{J})}(\beta^{(\mathcal{J})}(k)) = T^{(\mathcal{J})*}$ [136]. Setting $a = (\bar{\beta}_1^{(\mathcal{J})} + \cdots + \bar{\beta}_{n_g}^{(\mathcal{J})})/n_g$ proved to be effective.

3. *Projection:* The $\beta^{(\mathcal{J})}$ updated by the subgradient step in general does not satisfy all constraints of (6.9). To make it feasible, it is projected onto to the feasible set $\mathcal{S}$ using the *Euclidean* projection $Q_{\mathcal{S}}(\cdot)$, see line 9 in Algorithm 6.1.

The subgradient method is not a descent method. Thus, at every iteration, the best objective value found so far is stored. The subgradient algorithm is stopped once the termination criterion in line 15 in Algorithm 6.1 is met. The optimal $\beta^{(\mathcal{J})*}$ and the corresponding objective value $V(k^*)$ are returned.

4. *Distribution to individual values:* Once the subgradient method has terminated, the best choice of the decision variable $\beta^{(\mathcal{J})*}$ has to be distributed to individual $\beta^{(1)*}, \ldots, \beta^{(J)*}$. This is done by the algorithm presented in Algorithm 6.2. To illustrate the procedure, consider the bottom plot in Figure 6.3 and assume that $\beta_i^{(\mathcal{J})*} = 1.5$ results from Algorithm 6.1. To reach $\beta_i^{(\mathcal{J})*}$ starting from $-\bar{\beta}_i^{(\mathcal{J})}$, the full line segments $l_{i,1}^{(\mathcal{J})}, l_{i,2}^{(\mathcal{J})}, l_{i,3}^{(\mathcal{J})}$ and a fraction of the segment $l_{i,4}^{(\mathcal{J})}$ are required. The book-keeping variables $s_{i,1}, \ldots, s_{i,4}$ introduced above tell us that the line segments $l_{i,2}^{(\mathcal{J})}, l_{i,3}^{(\mathcal{J})}$ belong to unit 1. Thus, we set

$$\beta_i^{(1)*} = -\bar{\beta}_i^{(1)} + l_{i,2}^{(\mathcal{J})} + l_{i,3}^{(\mathcal{J})} = 1.$$

Similarly, from unit 2 we use the entire first line segment and half of the second segment, i.e.,

$$\beta_i^{(2)*} = -\bar{\beta}_i^{(2)} + l_{i,1}^{(\mathcal{J})} + 0.5 l_{i,4}^{(\mathcal{J})} = 0.5.$$

The distribution of $\beta_i^{(\mathcal{J})*} = 1.5$ into $\beta_i^{(1)*} = 1$ and $\beta_i^{(2)*} = 0.5$ is the minimum-cost distribution with regard to $T_i^{(\mathcal{J})}$. Finally, the corresponding individual minimum-cost power trajectories are given by $p^{(j)*} = c^{(j)} + G\beta^{(j)*}$, $j \in \mathcal{J}$.

Algorithm 6.1 Projected subgradient algorithm.

Require: $p^{(\mathcal{J})}$, $c^{(\mathcal{J})}$, $\beta^{(\mathcal{J})}(0)$, $\bar{\beta}^{(\mathcal{J})}$, $L_1, \ldots, L_{n_g}$, $T_{\text{fix}}^{(\mathcal{J})}$, a, δ, ϵ

$V(0) \leftarrow \text{INF}$, $k \leftarrow 0$
repeat
  $W \leftarrow T_{\text{fix}}^{(\mathcal{J})}$
  for $i = 1$ to n_g **do**
    Find largest $m : \sum_{i'=1}^{m} l_{i,i'}^{(\mathcal{J})} \leq \bar{\beta}_i^{(\mathcal{J})} + \beta_i^{(\mathcal{J})}(k)$
    $\nabla_i(k) \leftarrow q_{i,\min\{m+1,m_i^{(\mathcal{J})}\}}^{(\mathcal{J})}$
    $W \leftarrow W + \sum_{i'=1}^{m} l_{i,i'}^{(\mathcal{J})} q_{i,i'}^{(\mathcal{J})} + \ldots$
    $\left(\bar{\beta}_i^{(\mathcal{J})} + \beta_i^{(\mathcal{J})}(k) - \sum_{i'=1}^{m} l_{i,i'}^{(\mathcal{J})}\right) q_{i,\min\{m+1,m_i^{(\mathcal{J})}\}}^{(\mathcal{J})}$
  end for
  $\beta^{(\mathcal{J})}(k+1) \leftarrow Q_{\mathcal{S}}(\beta^{(\mathcal{J})}(k) - a/(k+1)\nabla(k))$
  $V(k+1) \leftarrow \min\{V(k), W\}$
  if $V(k+1) < V(k)$ **then**
    $k^* \leftarrow k+1$
    $\beta^{(\mathcal{J})*} \leftarrow \beta^{(\mathcal{J})}(k^*)$
  end if
  $k \leftarrow k+1$
until $[V(\max\{0, k-\delta\}) - V(k)]/V(k) \leq \epsilon$
return $\beta^{(\mathcal{J})*}$, $V(k^*)$

Algorithm 6.2 Distribution of $\beta^{(\mathcal{J})*}$ into $\beta^{(1)*}, \ldots, \beta^{(J)*}$.

Require: $\beta^{(\mathcal{J})*}$, $\bar{\beta}^{(1)}, \ldots, \bar{\beta}^{(J)}$, $L_1, \ldots, L_{n_g}$

$\beta^{(j)*} \leftarrow -\bar{\beta}^{(j)}$, $\forall j \in \mathcal{J}$
$\bar{\beta}^{(\mathcal{J})} \leftarrow \bar{\beta}^{(1)} + \cdots + \bar{\beta}^{(J)}$
for $i = 1$ to n_g **do**
  Find largest m : $\sum_{i'=1}^{m} l_{i,i'}^{(\mathcal{J})} \leq \bar{\beta}_i^{(\mathcal{J})} + \beta_i^{(\mathcal{J})*}$
  if $m \geq 1$ **then**
    for $i' = 1, \ldots, m$ **do**
      $\beta_i^{(s_{i,i'})*} \leftarrow \beta_i^{(s_{i,i'})*} + l_{i,i'}^{(\mathcal{J})}$
    end for
    $r \leftarrow \bar{\beta}_i^{(\mathcal{J})} + \beta_i^{(\mathcal{J})*} - \sum_{i'=1}^{m} l_{i,i'}^{(\mathcal{J})}$
    $i' \leftarrow \min\{m_i^{(\mathcal{J})}, m+1\}$
    $\beta_i^{(s_{i,i'})*} \leftarrow \beta_i^{(s_{i,i'})*} + r$
  else
    $\beta_i^{(s_{i,1})*} \leftarrow \beta_i^{(\mathcal{J})*}$
  end if
end for
return $\beta^{(1)*}, \ldots, \beta^{(J)*}$

6.4.2 *Performance of the subgradient-based disaggregation*

To assess the performance of the subgradient-based disaggregation, its solving time is compared with the solving time of the original LP formulation (6.2) for different populations of PEVs sampled uniformly from the parameter ranges provided in Table 5.2. The flexibility costs $h^{(j)}$ are assumed linear with random energy costs. The optimization models and the subgradient algorithm are implemented in *Matlab*®. All optimization problems are solved by *IBM ILOG CPLEX*®[1] v.12.6 run on a desktop PC featuring an *Intel*® Core i5-2400 CPU @ 3.1 GHz and 16 GB of RAM.

The original disaggregation problem (6.2) is convex and can, in theory, be solved efficiently. However, the problem grows with the population size J and the number of time steps N. In the case of PE-units, problem (6.2) comprises NJ decision variables, N equalities, and $J(4N-2)$ inequalities. Therefore, even in the simplest case of linear costs, the solving time and memory requirements of (6.2) can be significant. Table 6.1 summarizes the solving time for different problem sizes. For some combinations of N and J, the problem size exceeds the maximum memory that *Matlab* can allocate (abbreviated by M). We note that memory is exhausted for problem sizes of practical interest; for example, participation in the weekly SFR market with 15 min intervals would require the solution of a problem with $N = 672$. In addition, the problem size grows significantly with the number of units J in the population. Even for modest horizon lengths and population sizes, e.g., $N = 168$, $J = 5000 - 10000$, solving the disaggregation problem via the LP formulation can take over one hour.

The solving time of the subgradient-based disaggregation depends on the termination parameters δ and ϵ, among other things. The values $\delta = 5$ and $\epsilon = 10^{-3}$ have proved to yield good results in practice: the projected subgradient method terminates with a relative optimality gap of less than $2.5 \cdot 10^{-3}$ for all experiments executed. The maximum solving time for different choices of J and N are provided in Table 6.2. The subgradient-based disaggregation clearly outperforms the LP-based

1 *IBM*, *ILOG*, and *CPLEX* are trademarks of *International Business Machines Corp.*, registered in many jurisdictions worldwide. Other product and service names might be trademarks of *IBM* or other companies.

J \ N	24	48	96	168	288	672
10^3	2.9 s	15.9 s	34.0 s	88.9 s	871.7 s	M
$2 \cdot 10^3$	5.9 s	45.7 s	145.4 s	670.5 s	1603.7 s	M
$5 \cdot 10^3$	26.3 s	155.9 s	931.5 s	2647.4 s	M	M
10^4	99.8 s	449.9 s	3811.1 s	7319.2 s	M	M

Table 6.1: Solving time of LP-based disaggregation (6.2).

J \ N	24	48	96	168	288	672
10^3	0.13 s	0.13 s	0.17 s	0.24 s	0.32 s	0.70 s
$2 \cdot 10^3$	0.12 s	0.15 s	0.22 s	0.29 s	0.42 s	0.88 s
$5 \cdot 10^3$	0.15 s	0.20 s	0.27 s	0.47 s	0.76 s	1.71 s
10^4	0.19 s	0.29 s	0.48 s	0.82 s	1.30 s	2.97 s

Table 6.2: Solving time of SG-based disaggregation via Algorithm 6.1.

disaggregation (6.2). For example, all problem instances with $N = 672$ could not be solved by the LP-based disaggregation because of memory limitations. In contrast, the subgradient-based disaggregation solves these problems within a few seconds. It is also worth noting that the solving time of the subgradient-based method scales particularly well with regard to the population size J. This is true because J influences only the complexity of the objective function but not the number of decision variables and constraints of the subgradient-based optimization problem (6.9).

6.5 USE-CASES FOR ZONOTOPIC FLEXIBILITY

The particular properties of zonotopic feasible sets can be exploited in several problems that AGs face on a daily basis. In this section, we show how the characteristics of zonotopic flexibility can be used to check the feasibility of a given trajectory efficiently, and to compute regulation power bid prices that are required when bidding into reserve markets.

6.5.1 *Distributing day-ahead reference schedules*

An important use case for aggregation and disaggregation of flexibility arises for AGs participating in day-ahead energy markets. After the market has cleared, the AG receives an aggregate power reference schedule $p_{\text{ref}}^{(\mathcal{J})}$ that the units in the control group must follow collectively during the subsequent day. However, it is not guaranteed that $p_{\text{ref}}^{(\mathcal{J})}$ is aggregate feasible because *i)* some day-ahead energy bids submitted by the AG may not have been selected by the market, or *ii)* the aggregate feasible set $\mathcal{P}^{(\mathcal{J})}$ may have changed in the meantime due to varying operating conditions and states of individual FEUs. Thus, it is necessary to check if the reference schedule is aggregate feasible. However, checking the feasibility for polytopic feasible sets requires solving the disaggregation problem (6.2), which can be computationally expensive. In contrast, checking the feasibility in the case of zonotopes is easy because the aggregate feasible set can be computed explicitly via (5.3). It holds that $p_{\text{ref}}^{(\mathcal{J})} \in \mathcal{Z}^{(\mathcal{J})}$ if and only if $|F^\top(p_{\text{ref}}^{(\mathcal{J})} - c^{(\mathcal{J})})| \leq |F^\top G|\bar{\beta}^{(\mathcal{J})}$, where $F \in \mathbb{R}^{N \times n_f}$ is the matrix with the normal vectors of all possible zonotope facets as columns, see Subsection 5.4.2. That is, checking whether a given trajectory is aggregate feasible amounts to validating a set of n_f inequalities. For zonotopes with generator matrix G^{PE} according to (5.4), the number of possible zonotope facets is $n_f = (N^2 + N)$.

The particular structure of zonotope aggregation (5.3) makes it easy to update the aggregate feasible set to reflect changes in the individual feasible sets. The same applies for checking the feasibility of a given trajectory. If, for instance, the flexibility of an individual unit changes or a unit becomes unavailable, above inequalities can be adjusted easily by making appropriate changes to $\bar{\beta}^{(\mathcal{J})}$.

If the reference $p^{(\mathcal{J})}$ is infeasible, the AG can buy a corrective schedule $p^{(\text{corr})}$ from other AGs or from intra-day markets to make the sum $p^{(\mathcal{J})} + p^{(\text{corr})}$ feasible again. The cheapest correction $p^{(\text{corr})*}$ can be found by solving

$$\begin{aligned} p^{(\text{corr})*} := \underset{p^{(\text{corr})}}{\arg\min} \quad & t_s \hat{v}^\top p^{(\text{corr})} \\ \text{s.t.} \quad & p^{(\mathcal{J})} + p^{(\text{corr})} \in \mathcal{Z}(G, c^{(\mathcal{J})}, \bar{\beta}^{(\mathcal{J})}), \end{aligned} \tag{6.12}$$

where $\hat{v} \in \mathbb{R}^N$ denotes the expected energy prices. Note that the problem size is determined by the time horizon N only, and is independent of the population size J.

Once an aggregate feasible power reference is available, it has to be distributed among the individual units of the population. Algorithm 6.1 can be used to efficiently find the disaggregation that minimizes the total cost of the AG and that is economically optimal with regard to the flexibility cost functions defined by the individual units.

6.5.2 *The cost of offering regulation power*

In addition to trading energy, AGs can use flexibility to offer ancillary services to the power grid. Here, we consider a symmetric and constant regulation power service, such as SFR [137], for example. Offering $\bar{r}$ units of regulation power requires that in every time step of the planning horizon, a deviation of $\pm\bar{r}$ units of power from the aggregate baseline $p_{\text{ref}}^{(\mathcal{J})}$ is possible. The requirement to keep the reserves $\bar{r}$ available, restricts the possible choices for the baseline $p_{\text{ref}}^{(\mathcal{J})}$, and, consequently, results in higher energy procurement costs compared to the case where no reserves are offered. This trade-off between energy procurement costs and offering reserves is illustrated in Figure 6.4. The aggregate feasible set is given by $\mathcal{Z}^{(\mathcal{J})}$ and the expected energy costs are denoted by $\hat{v}$. If no reserves are offered, the baseline $p_{\text{ref}}^{(\mathcal{J})}$ can be chosen everywhere in $\mathcal{Z}^{(\mathcal{J})}$. The choice minimizing the energy costs is labeled $p_{\text{ref}}^{(\mathcal{J})*}(0)$. In contrast, if $\bar{r}$ units of regulation power are offered, the possible choices for $p_{\text{ref}}^{(\mathcal{J})}$ are restricted to the set

$$\mathcal{Z}^{(\text{rem})} := \{p \in \mathbb{R}^N : p \pm w \in \mathcal{Z}^{(\mathcal{J})} \ \forall w \in [0, \bar{r}]^N\}, \tag{6.13}$$

and the corresponding minimum-cost baseline is $p_{\text{ref}}^{(\mathcal{J})*}(\bar{r})$.

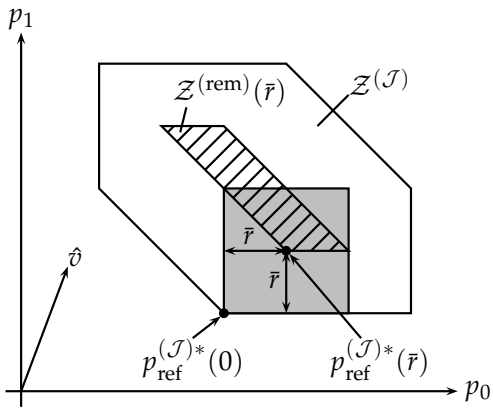

Figure 6.4: Minimum-cost baselines $p_{\mathrm{ref}}^{(\mathcal{J})*}(0)$ and $p_{\mathrm{ref}}^{(\mathcal{J})*}(\bar{r})$ in the aggregate feasible set $\mathcal{Z}^{(\mathcal{J})}$ for 0 and $\bar{r}$ units of symmetric constant regulation power, respectively. The set of baselines that allow for $\bar{r}$ units of regulation power is given by $\mathcal{Z}^{(\mathrm{rem})}(\bar{r})$. The expected price of energy is $\hat{v}$.

A lower bound $\bar{r}_{\max}$ on the maximum regulation power available from the population of flexible units is

$$\max_{\bar{r},\{p^{(j)},\, r^{(j)}\}_{j\in\mathcal{J}}} \quad \bar{r} \tag{6.14a}$$

$$\text{s.t.} \quad \mathcal{B}(p^{(j)}, r^{(j)}) \subseteq \mathcal{P}^{(j)}, j \in \mathcal{J}, \tag{6.14b}$$

$$r^{(j)} \geq 0,\ j \in \mathcal{J}, \tag{6.14c}$$

$$\sum_{j\in\mathcal{J}} r^{(j)} \geq \bar{r}\mathbf{1} \tag{6.14d}$$

with box $\mathcal{B}(p^{(j)}, r^{(j)}) := \{p \in \mathbb{R}^N : p^{(j)} - r^{(j)} \leq p \leq p^{(j)} + r^{(j)}\}$.

For a fixed amount of regulation power $\bar{r} \in [0, \bar{r}_{\max}]$ and expected day-ahead energy prices $\hat{v} \in \mathbb{R}^N$, the AG can find the cheapest baseline by solving

$$\hat{R}(\bar{r}) := \min_{\{p^{(j)},r^{(j)}\}_{j\in\mathcal{J}}} \sum_{j\in\mathcal{J}} \left\{ T^{(j)}(p^{(j)}) + \hat{v}^\top p^{(j)} \right\}, \tag{6.15}$$

subject to the constraints (6.14b)-(6.14d). The expected cost of offering $\bar{r} > 0$ units of regulation power is $\hat{R}(\bar{r}) - \hat{R}(0)$. This quantity can serve as a lower bound on the regulation power bid price.

However, problem (6.15) with polytopic feasible sets has the same complexity as the disaggregation problem (6.2) and thus is inconvenient for large populations and long planning horizons. In contrast, for zonotopic feasible sets, an approximative solution to (6.15) can be found efficiently by using the subgradient-based disaggregation algorithm. In a first step, $\bar{r}_{\max}$ is computed by solving (6.14) reformulated for zonotopes with generator matrices G^{PE} as

$$\begin{aligned} \bar{r}_{\max} := \quad & \underset{\{\beta^{(j)}, \bar{r}\}_{j \in \mathcal{J}}}{\arg\min} \quad \bar{r} \\ & \text{s.t.} \qquad G^{\mathrm{PE}} \beta^{(j)} \geq \mathbf{0}, \ \forall j \in \mathcal{J}, \\ & \qquad G^{\mathrm{PE}} \sum_{j \in \mathcal{J}} \beta^{(j)} \geq \bar{r}\mathbf{1}, \\ & \qquad -\bar{\beta}^{(j)} \leq \beta^{(j)} \leq \bar{\beta}^{(j)}, \ \forall j \in \mathcal{J}. \end{aligned} \tag{6.16}$$

Denote by $\beta^{(j)}_{\max}$, $j \in \mathcal{J}$, the optimizers of (6.16). The constraints guarantee that every feasible set $\mathcal{Z}(G^{\mathrm{PE}}, c^{(j)}, \bar{\beta}^{(j)})$ includes a box with semi-edge lengths given by $G^{\mathrm{PE}} \beta^{(j)}_{\max}$. To keep a certain amount $\bar{r} = \eta \bar{r}_{\max}$, $\eta \in [0, 1]$, of regulation power available, every unit reserves a portion $\eta |\beta^{(j)}_{\max}|$. The remaining $\beta^{(j)}_{\mathrm{rem}} := \bar{\beta}^{(j)} - \eta |\beta^{(j)}_{\max}|$ defines a zonotope within which the baseline $p^{(j)}_{\mathrm{ref}}$ can be chosen freely. The cheapest aggregate baseline is $p^{*}_{\mathrm{ref}} = c^{(\mathcal{J})} + G^{\mathrm{PE}} \beta^{(\mathcal{J})*}$ with $\beta^{(\mathcal{J})*}$ minimizing

$$\begin{aligned} \min_{\beta^{(\mathcal{J})}} \quad & T^{(\mathcal{J})}(\beta^{(\mathcal{J})}) \\ \text{s.t.} \quad & -\sum_{j \in \mathcal{J}} \beta^{(j)}_{\mathrm{rem}} \leq \beta^{(\mathcal{J})} \leq \sum_{j \in \mathcal{J}} \beta^{(j)}_{\mathrm{rem}}. \end{aligned} \tag{6.17}$$

This problem is solved efficiently by Algorithm 6.1.

Table 6.3 provides the expected baseline costs $\hat{R}$ for different amounts of regulation power $\bar{r}$ reserved. The costs are computed for a population of 100 PEVs with the parameters sampled from Table 5.2. Historic *EPEX SPOT* day-ahead prices are used for $\hat{v}$. Considered are the

cases where *i)* problem (6.15) is solved using the original polytopic feasible sets, *ii)* problem (6.17) is solved for zonotopic feasible sets, and *iii)* problem (6.17) is solved for feasible sets approximated by axis-aligned (A-A) boxes. While the zonotopic approximation is suboptimal compared with the solution found on the original polytopic sets, it still yields significantly lower baseline costs than feasible sets approximated by boxes do. These results indicate that zonotopes are a reasonable compromise between the complex, but accurate polytopic feasible sets and the simple, but inaccurate approximations by boxes.

	$\bar{r} = 0$ kW	20 kW	40 kW	60 kW
Polytopes	−77.9 EUR	−59.5 EUR	−41.0 EUR	−22.2 EUR
Zonotopes	−77.7 EUR	−58.6 EUR	−39.1 EUR	−18.5 EUR
A-A Boxes	−43.1 EUR	−25.8 EUR	−8.6 EUR	8.7 EUR

Table 6.3: Baseline costs for different types of feasible sets.

6.6 CONCLUSION & FUTURE WORK

In this chapter we introduced a subgradient-based disaggregation algorithm based on zonotopic feasible sets that can be used to efficiently distribute aggregate-level signals in an economically optimal way among the FEUs of a balance group. The disaggregation algorithm was applied to the problems of distributing day-ahead reference schedules and computing the expected cost of offering regulation services.

Future work will investigate the choice of the termination parameters of the subgradient-based disaggregation algorithm, and the applicability of distributed optimization methods for solving the disaggregation problem. Further, other flexibility reimbursement schemes will be studied to better reflect to contribution of individual energy units to the aggregate flexibility product.

7 ENERGY AND RESERVE BIDDING

This chapter investigates how an AG can exploit the flexibility available from the energy units in its balance group to maximize the economic profit from trading on energy and SFR markets. To guarantee the accurate provision of SFR services, we take power ramp-rate constraints into account explicitly when computing optimal adjustable energy-trading policies that are robust against uncertain activation of reserves. In addition, we model all the different timescales of the day-ahead, intra-day, and the SFR market to make the approach applicable to real market settings. Finally, it is shown that the option to trade energy sequentially over time can free up a significant portion of an energy unit's flexibility.

7.1 INTRODUCTION

To guarantee the safe and reliable operation of the electricity grid, an SO depends on several types of ancillary services. In particular, the SFR service is invoked by the SO to reestablish the nominal grid frequency and to compensate for mismatches of energy schedules on the timescale of seconds to minutes [137, 138]. SFR has traditionally been provided by controllable large-scale generating units, such as fossil-fired steam turbines, combustion turbines, nuclear power plants, and hydro power units. With a growing share of VRES in the power grid, the need for rapid and accurate SFR is increasing. To support the traditional providers of SFR, various types of electric energy units have been identified whose flexibility in producing and consuming electricity can be used for SFR. In particular, the balancing potential of heating and cooling systems [70, 125, 139–141], plug-in electric vehicles [142, 143], and units behaving like batteries [81] has been studied extensively.

AGs can act as intermediaries between electric energy units, electricity markets, and SOs, see Chapter 5. AGs represent a single or a group of electric energy units, referred to as their *balance group*. Throughout

this work, we assume that AGs are *balance-responsible*, i.e., they are responsible for trading electric energy on energy markets to satisfy the needs of their balance group and to adjust energy schedules to match actual electricity generation and consumption as accurately as possible. The AG is held accountable for any mismatches between the actual and planned energy schedule [144]. In this sense, we assume that an AG acts as a Balance Responsible Party (BRP). If the balance group of an AG comprises FEUs, the AG can offer their flexibility to the SO in the form of ancillary services that are traded on dedicated flexibility markets [138]. The provision of SFR has been studied for heating, ventilation, and air-conditioning systems [139, 145, 146], PEVs [142, 143], and units behaving like generalized batteries [74, 147], among others.

Here, we consider an AG of a single FEU that wishes to maximize the economic profit made from trading energy on day-ahead and intra-day energy markets and offering SFR services to the SO via reserve markets. These trading decisions are mutually dependent and constitute a trade-off between the revenues from offering reserves and the costs of procuring energy. We refer to the problem of making optimal decisions on energy and reserve markets as the *energy and reserve bidding problem*. The problem has been approached by means of Model Predictive Control [148], Stochastic Optimization [149, 150], and Robust Optimization [34, 70, 151, 152], among others. The robust formulations rely on the concept of affinely adjustable robust control (AARC) policies, which have first been applied to the regulation power context in [151, 153–155], building on the initial work in [156] and the references therein.

The provision of SFR services requires the tracking of an activation signal broadcast by the SO on the timescale of a few seconds. The ability of an FEU to accurately follow such signals is limited by its physical properties, in particular by power, power ramp-rate, and energy constraints. While power and energy constraints are considered by most AARC approaches, all works mentioned above neglect power ramp-rate constraints. However, the crucial importance of ramping constraints has been recognized in the context of grid balancing [147, 157] and unit commitment [158, 159], and has lead to the adoption of new performance-based schemes for the remuneration of SFR by several SOs [160]. But the methods developed in [157–159] cannot be directly applied to the energy and reserve bidding problem considered here.

In our formulation of the bidding problem, we adopt the AARC formulation used in [145, 151–153], and make two main contributions. First, in contrast to the aforementioned works, we explicitly incorporate constraints on power ramp-rates in our modeling framework to guarantee the feasibility and accurate provision of SFR. Instead of the commonly used piece-wise constant power schedules, we base our formulation on piece-wise affine and continuous power trajectories, which makes it possible to accurately model power ramp-rates on arbitrary timescales. We then derive a reformulation of the AARC optimization problem which is necessary to guarantee the feasibility of the piece-wise affine solutions. Second, in our problem formulation we include all the different timescales and market lead times involved in the bidding problem. This allows us to study how different energy and reserve market parameters affect the amount of SFR capacity an FEU can provide. These two contributions make our approach versatile and widely applicable, in particular in real-world settings, where accurate modeling of system constraints and market specifications is indispensable.

This chapter is organized as follows. The energy and reserve bidding problem of an AG is introduced in Section 7.2. In Section 7.3, we illustrate how piece-wise affine power schedules can account for finite power ramp-rates. Section 7.4 provides a robust reformulation of the flexibility of an energy unit. It is used in Section 7.5 to formulate and solve the energy and reserve bidding problem. A conclusion and outlook on future work are provided in Section 7.6.

7.2 PROBLEM FORMULATION

We consider an AG that participates in energy and reserve markets and whose balance group comprises a single FEU. The AG strives to schedule the electric energy generation and consumption of the unit so as to maximize the profit made from trading energy and from offering SFR services over a planning horizon of duration T^{H}. The bidding problem involves different timescales and requires the AG to take decisions on multiple stages both offline, i.e., prior to the time of delivery, and online, i.e., at delivery.

7.2.1 *Energy markets*

To satisfy the energy requirements of the units in its balance group, the AG can trade electric energy prior to delivery.

On the day-ahead market, energy is traded once a day for every time interval of duration T^{DA} of the next day. The Swiss day-ahead market runs on an hourly timescale, i.e., $T^{\mathrm{DA}} = 1$ h, and closes at 11 a.m. on the day before delivery [161]. The outcome of the day-ahead markets is the energy schedule $e^{\mathrm{DA}} \in \mathbb{R}^{N^{\mathrm{DA}}}$ of energy quantities to be produced or consumed during each of the $N^{\mathrm{DA}} = 7 \cdot 24 = 168$ time intervals in the planning horizon T^{H}. The energy procurement costs are $C^{\mathrm{DA}} := {c^{\mathrm{DA}}}^{\top} e^{\mathrm{DA}}$, where $c^{\mathrm{DA}} \in \mathbb{R}^{N^{\mathrm{DA}}}$ denotes the day-ahead market clearing prices, which are unknown at the time of bidding.

To make adjustments to the day-ahead energy schedule during a particular day, the intra-day markets allow trading energy continuously for time intervals of durations T^{ID} (usually $T^{\mathrm{ID}} = 5 - 15$ min). We focus on quarter-hourly energy block products, i.e., $T^{\mathrm{ID}} = 15$ min, as traded on the Swiss intra-day market with a lead time of 1 hour [161] and denote by N^{ID} the number of intra-day time slots in the planning horizon T^{H}, i.e., $N^{\mathrm{ID}} = 7 \cdot 24 \cdot 4 = 672$. All intra-day energy trades are summarized by $e^{\mathrm{ID}} \in \mathbb{R}^{N^{\mathrm{ID}}}$, $N^{\mathrm{ID}} := T^{\mathrm{H}} / T^{\mathrm{ID}}$. The total energy procurement costs are $C^{E} := {c^{\mathrm{DA}}}^{\top} e^{\mathrm{DA}} + {c^{\mathrm{ID}}}^{\top} e^{\mathrm{ID}}$, where $c^{\mathrm{DA}} \in \mathbb{R}^{N^{\mathrm{DA}}}$ and $c^{\mathrm{ID}} \in \mathbb{R}^{N^{\mathrm{ID}}}$ denote the energy prices, which are unknown at the time of bidding.

7.2.2 *Reserve markets*

The SO constantly observes the state of the power grid and, in case of imbalances, takes measures to reestablish the nominal operating state. To be able to compensate fluctuating electricity generation and consumption and other unexpected disturbances, the SO procures ancillary services on reserve markets ahead of time [138]. We focus on SFR and adopt the product specifications that apply to the Swiss SFR market, see [137]: every week on Tuesday not later than at 1 p.m., an AG can submit bids to the SFR market covering the subsequent week, i.e., $T^{\mathrm{RES}} = 1$ week [137]. Each bid consists of a constant and symmetric SFR capacity

$\gamma \in \mathbb{R}_0^+$ and price c^{RES}. The approach can be generalized to asymmetric reserves [34]. For simplicity, we will use $T^{\mathrm{RES}} = T^{\mathrm{H}}$.

If the capacity bid γ of an AG is accepted, the SO has the right, but not the obligation, to ask the AG to deviate from the planned power reference by $\gamma w(t)$ units of power at any time during the planning horizon. The SFR activation signal $w(t)$ is limited to the interval $\mathcal{W} := [-1, 1]$ and is unknown ahead of time. In return for keeping the reserve capacity γ available over the horizon T^{H}, the AG receives the capacity reservation payment $R^{\mathrm{RES}} := c^{\mathrm{RES}}\gamma$, where c^{RES} is the reserve capacity price, which is unknown at the time of bidding [137].

Power deviations due to SFR activation can result in energy deviations from the planned energy schedule of an AG. This so-called up- and down-regulation energy is usually measured over each time interval T^{ID} separately for positive and negative activation, and is denoted by e^{up} and e^{dn}, respectively [137]. The AG is remunerated for providing SFR based on the amount of up- and down-regulation energy delivered, i.e., the remuneration is $R^{\mathrm{REG}} := c^{\mathrm{up}\top} e^{\mathrm{up}} - c^{\mathrm{dn}\top} e^{\mathrm{dn}}$, where $c^{\mathrm{up}}, c^{\mathrm{dn}} \in \mathbb{R}^{N^{\mathrm{ID}}}$ are the corresponding regulation energy prices [137]. However, recent regulatory changes in the USA require all SOs to base their SFR compensation payment not only on the amount of regulation power reserved and regulation energy delivered, but also on the achieved accuracy of up- and down-regulation [160, 162]. This motivates us to impose strict requirements on SFR performance, see Subsection 7.3.3.

7.2.3 *The bidding problem*

The energy and reserve bidding problem of the AG is to decide how to trade on the energy markets and how much reserve capacity to offer on the SFR market so as to maximize its profit. The profit made by the AG is given by

$$V(\Omega, \Psi) := R^{\mathrm{RES}}(\Omega) + R^{\mathrm{REG}}(\Omega, \Psi) - C^{E}(\Omega, \Psi) \tag{7.1}$$

and is a function of the AG's decision variables $\Omega := \{\gamma, e^{\mathrm{DA}}, e^{\mathrm{ID}}\}$, and of all the uncertain parameters $\Psi := \{c^{\mathrm{DA}}, c^{\mathrm{ID}}, c^{\mathrm{RES}}, c^{\mathrm{up}}, c^{\mathrm{dn}}, e^{\mathrm{up}}, e^{\mathrm{dn}}\}$, which are unknown at the time of bidding. However, an empirical estimate of the probability density of Ψ is assumed to be known. We

consider the AG to be a price taker in both the energy and the reserve markets and consider the expected profit $\mathbf{E}_\Psi[V(\Omega, \Psi)]$. The problem of finding optimal bids in terms of price-quantity pairs is out of the scope of this work. It has been studied in [163, 164], among others.

A particular choice of Ω together with a realization of the SFR activation signal w results in a unique power trajectory $p(\Omega, w)$ that the balance group of the AG is committed to implement. The AG must choose Ω subject to the constraints that the energy needs of the units in the balance group are satisfied and the offered SFR capacity is available over the entire planning horizon *for any realization* of the uncertain SFR activation $w \in \mathcal{W}$. The energy and reserve bidding problem can be written as

$$\max_{\Omega} \quad \mathbf{E}_\Psi[V(\Omega, \Psi)] \tag{7.2a}$$

$$\text{s.t.} \quad p(\Omega, w) \text{ feasible } \forall w \in \mathcal{W}. \tag{7.2b}$$

The more reserve capacity γ the AG offers, the more restrictive become the constraints (7.2b) with regard to possible energy trades e^{DA}, e^{ID}. Consequently, offering SFR reserves and trading energy involves a trade-off between the reserve reward R^{RES}, the potential remuneration of regulation energy R^{REG}, and the energy procurement costs C^E.

Problem (7.2) is a robust optimization problem because the constraints (7.2b) must hold for *any* admissible realization of the SFR activation signal w, see [165]. In Section 7.3, we provide the details of the mapping $p(\Omega, w)$ based on which we derive a robust reformulation of the constraint (7.2b). This allows us to rewrite the robust bidding problem (7.2) as a standard LP in Section 7.5.

7.2.4 *Different timescales*

The bidding problem takes place on different timescales. Trading decisions are made on the timescales T^{RES}, T^{DA}, and T^{ID} corresponding to the reserve market, and the day-ahead and the intra-day energy market, respectively. Online activation of SFR occurs at the faster timescale T^{C}. The dynamics of individual systems are time-discretized on an intermediate timescale T^{S} commonly chosen in the range of 5–15 min in the case of SFR. In general, it holds that $T^{\mathrm{H}} \geq T^{\mathrm{RES}} \geq T^{\mathrm{DA}} \geq T^{\mathrm{ID}} \geq T^{\mathrm{S}} \gg T^{\mathrm{C}}$.

For simplicity, we assume that longer time horizons are integer multiples of shorter ones, and for $* \in \{\mathrm{RES}, \mathrm{DA}, \mathrm{ID}, \mathrm{S}, \mathrm{C}\}$ define $N^* := T^{\mathrm{H}}/T^*$, and $\mathcal{T}^*(s)$ the continuous time interval $[(s-1)T^*, sT^*)$. Finally, let $\mathcal{N}^s := \{1, 2, \ldots, s\}$ for $s \in \mathcal{N}^+$. Different timescales are used in different market regions. All timescales T^* are used as parameters in our approach making it versatile and applicable to the various market settings. Here we consider the Swiss market setting and use $T^{\mathrm{H}} = T^{\mathrm{RES}} = 1$ week, $T^{\mathrm{DA}} = 1$ h, $T^{\mathrm{ID}} = 15$ min, $T^{\mathrm{S}} = 5$ min, and $T^{\mathrm{C}} = 1$ s.

7.3 ENERGY BASELINE AND POWER REFERENCE

7.3.1 *Energy baseline*

The day-ahead and intra-day energy trades of an AG result in the energy market baseline $e^{\mathrm{base}} \in \mathbb{R}^{N^{\mathrm{ID}}}$, which comprises the total energy trades in each interval of duration T^{ID} in the planning horizon. The energy baseline is defined as

$$e_k^{\mathrm{base}} := e_j^{\mathrm{DA}} T^{\mathrm{ID}}/T^{\mathrm{DA}} + e_k^{\mathrm{ID}},\ j := \lceil kT^{\mathrm{ID}}/T^{\mathrm{DA}} \rceil,\ k \in \mathcal{N}^{\mathrm{ID}}, \tag{7.3}$$

and is written in vector form as $e^{\mathrm{base}} = Me^{\mathrm{DA}} + e^{\mathrm{ID}}$ with $e^{\mathrm{base}} := [e_1^{\mathrm{base}}, \ldots, e_{N^{\mathrm{ID}}}^{\mathrm{base}}]^\top$ and $M \in \mathbb{R}^{N^{\mathrm{ID}} \times N^{\mathrm{DA}}}$.

7.3.2 *Power reference*

It has been common in recent works, e.g. [139, 145], to associate with the energy baseline e^{base} the piece-wise *constant* continuous-time power reference

$$p(t) := e_k^{\mathrm{base}}/T^{\mathrm{ID}},\ t \in \mathcal{T}^{\mathrm{ID}}(k),\ k \in \mathcal{N}^{\mathrm{ID}}, \tag{7.4}$$

that the AG is committed to follow. Whereas the energy content of $p(t)$ over each time interval $\mathcal{T}^{\mathrm{ID}}(k)$ matches e_k^{base} as required by market contracts, implementing the instantaneous changes of $p(t)$ at the interface of subsequent market time intervals would require infinite ramping capabilities. Inaccurate tracking of the power reference (7.4) can cause the grid frequency to deviate from its nominal value when switching between market time intervals, making it necessary to activate frequency

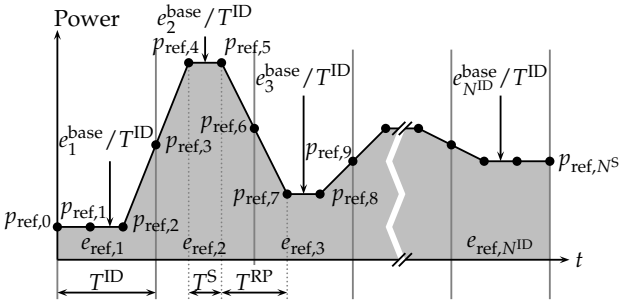

Figure 7.1: The piece-wise affine continuous-time power reference $p_{\text{ref}}(t)$ (solid black line) defined by its breakpoints $p_{\text{ref}} := [p_{\text{ref},0}, \ldots, p_{\text{ref},N^{\text{S}}}]^\top$. The energy content (shaded gray area) of $p_{\text{ref}}(t)$ over the interval $\mathcal{T}^{\text{ID}}(k)$ is $e_{\text{ref},k}$. The power ramps of duration T^{RP} are located symmetrically at the intersections of market time intervals.

reserves [157, 158]. To alleviate this difficulty, we propose to translate the energy baseline e^{base} into a continuous-time *piece-wise affine* and continuous power reference $p_{\text{ref}}(t)$ whose breakpoint intervals are each of duration T^{S} [158]. The power reference $p_{\text{ref}}(t)$ is fully defined by its breakpoint values $p_{\text{ref}} := [p_{\text{ref},0}, \ldots, p_{\text{ref},N^{\text{S}}}]^\top$ according to

$$p_{\text{ref}}(t) := p_{\text{ref},s-1} + (p_{\text{ref},s} - p_{\text{ref},s-1})(t - (s-1)T^{\text{S}})/T^{\text{S}}, \tag{7.5}$$

for $t \in \mathcal{T}^{\text{S}}(s)$, $s \in \mathcal{N}^{\text{S}}$. An example of such a piece-wise affine and continuous power reference is given in Figure 7.1.

Our choice of a piece-wise affine power trajectory is in full accordance with ENTSO-E standards [138]. The power reference comprises N^{ID} periods during which the power is constant and equal to $e_k^{\text{base}}/T^{\text{ID}}$, $k \in \mathcal{N}^{\text{ID}}$. These periods of constant power are connected by linear ramps, each of duration T^{RP} as illustrated in Figure 7.1. We assume that T^{RP} is an integer multiple of T^{S} and that $0 \leq T^{\text{RP}} \leq T^{\text{ID}}$. For the

Swiss case of $T^{\mathrm{ID}} = 15$ min, $T^{\mathrm{RP}} = 10$ min, and $T^{\mathrm{S}} = 5$ min, the relation between e^{base} and p_{ref} is the following:

$$p_{\mathrm{ref},s} = \begin{cases} e_1^{\mathrm{base}}/T^{\mathrm{ID}}, & s = 0, \\ e_k^{\mathrm{base}}/T^{\mathrm{ID}}, & s = 3k-2, 3k-1, \\ & k \in \mathcal{N}^{\mathrm{ID}}, \\ (e_k^{\mathrm{base}} + e_{k+1}^{\mathrm{base}})/(2T^{\mathrm{ID}}), & s = 3k,\ k \in \mathcal{N}^{\mathrm{ID}-1}, \\ e_{N^{\mathrm{ID}}}^{\mathrm{base}}/T^{\mathrm{ID}}, & s = N^{\mathrm{S}}. \end{cases} \tag{7.6}$$

Thus, the power reference can be expressed compactly as

$$p_{\mathrm{ref}} = Re^{\mathrm{base}} = RMe^{\mathrm{DA}} + Re^{\mathrm{ID}}, \tag{7.7}$$

where we encode (7.6) by the matrix $R \in \mathbb{R}^{(N^{\mathrm{S}}+1)\times N^{\mathrm{ID}}}$. Note that (7.6) reverts to (7.4) in the limit case $T^{\mathrm{RP}} = 0$. For all $T^{\mathrm{RP}} > 0$, our choice (7.5) is continuous and, thus, allows us to model the power ramp-rates. We make the assumption that an AG with energy baseline e^{base} will use the piece-wise affine continuous-time power trajectory $p_{\mathrm{ref}}(t)$ (7.5) defined by its breakpoint values p_{ref} (7.7) as its power reference.

7.3.3 *Activation of SFR*

The SFR service is activated by the SO by sending out an activation signal to all AGs that offer SFR for the corresponding time interval. The activation signal $w := [w_1, \ldots, w_{N^{\mathrm{C}}}]^\top$ is a discrete-time signal with $w_l \in \mathcal{W} := [-1, 1]$, $l = 1, \ldots, N^{\mathrm{C}}$. The signal is computed by the SO and broadcast *sequentially* on the timescale T^{C} ($T^{\mathrm{C}} = 1 - 5$ s in Continental Europe [138], $T^{\mathrm{C}} = 1$ s in Switzerland [150], and $T^{\mathrm{C}} = 4 - 6$ s in the USA [160]). That is, the activation w_l becomes available to the AG only at time $(l-1)T^{\mathrm{C}}$. We interpret the discrete signal w as the continuous-time piece-wise affine and continuous activation signal

$$w(t) := w_{l-1} + (w_l - w_{l-1})(t - (l-1)T^{\mathrm{C}})/T^{\mathrm{C}}, \tag{7.8}$$

where $t \in \mathcal{T}^{\mathrm{C}}(l)$, $l \in \mathcal{N}^{\mathrm{C}}$.

The target power level $p_{\mathrm{tgt}}(t)$ that the AG is responsible to track is a function of the time t and the activation $w(t)$ and is given by

$$p_{\mathrm{tgt}}(t, w) := p_{\mathrm{ref}}(t) + \gamma w(t). \tag{7.9}$$

Compared with the power reference $p_{\text{ref}}(t)$, which is piece-wise affine on the timescale T^{S}, the activation $w(t)$ can vary at the higher rate T^{C}. Thus, accurate tracking of the target power level (7.9) is challenging. The minimum required tracking accuracy of $p_{\text{tgt}}(t)$ is specified by the SOs individually, see [166] for details on the Swiss requirements. In the USA, SOs are required to base their regulation compensation payment R^{REG} not only on the quantity of regulation provided, but must also take into account how accurately the activation signal has been tracked [160]. This motivates the use of a continuous power reference according to (7.7) and the explicit consideration of ramp-rate constraints.

7.3.4 *Adjustable energy schedules*

The decision on the amount of SFR to be offered is made once for each reserve tendering period and cannot be changed afterwards. In contrast, energy schedules can be adjusted from day to day or intra-day by trading energy on respective markets. These adjustments can depend on past SFR activation, for example. Let $\tilde{w} \in \mathcal{W}^{N^{\text{S}}}$ be the activation signal $w(t)$ averaged over time intervals of duration T^{S}, i.e.,

$$\tilde{w}_s := \frac{1}{T^{\text{S}}} \int_{(s-1)T^{\text{S}}}^{T^{\text{S}}} w(t)dt, \; s \in \mathcal{N}^{\text{S}}. \tag{7.10}$$

Following the work in [32, 118, 151–153], we express the adjustable energy schedules as affine functions of the averaged activation signal:

$$\begin{aligned} e^{\text{DA}}(\tilde{w}) &= Q^{\text{DA}} A^{\text{DA}} \tilde{w} + q^{\text{DA}}, \\ e^{\text{ID}}(\tilde{w}) &= Q^{\text{ID}} A^{\text{ID}} \tilde{w} + q^{\text{ID}}, \end{aligned} \tag{7.11}$$

with parameters $Q^{\text{DA}} \in \mathbb{R}^{N^{\text{DA}} \times N^{\text{DA}}}$, $q^{\text{DA}} \in \mathbb{R}^{N^{\text{DA}}}$, $Q^{\text{ID}} \in \mathbb{R}^{N^{\text{ID}} \times N^{\text{ID}}}$, and $q^{\text{ID}} \in \mathbb{R}^{N^{\text{ID}}}$. The matrices $A^{\text{DA}} \in \mathbb{R}^{N^{\text{DA}} \times N^{\text{S}}}$ and $A^{\text{ID}} \in \mathbb{R}^{N^{\text{ID}} \times N^{\text{S}}}$ are used to average $\tilde{w}$ over time intervals of duration T^{DA} and T^{ID}, respectively. For instance, A^{ID} is block-diagonal with each block given by a row-vector of length $T^{\text{ID}}/T^{\text{S}}$ whose elements all equal $T^{\text{S}}/T^{\text{ID}}$. By plugging (7.11) into (7.7) we can write the power reference as an affine function of the averaged activation:

$$p_{\text{ref}}(\tilde{w}) = R(MQ^{\text{DA}} A^{\text{DA}} + Q^{\text{ID}} A^{\text{ID}})\tilde{w} + R(Mq^{\text{DA}} + q^{\text{ID}}). \tag{7.12}$$

The policies (7.11) describe how much energy should be traded on the day-ahead and intra-day energy markets depending on the average amount of SFR activated during past time intervals. The policies must be causal, i.e., they can only depend on past activation, and must respect the different lead times of the markets. These requirements can be incorporated by imposing a particular structure onto Q^{DA} and Q^{ID}. Figure 7.2 illustrates which elements of Q^{DA} and Q^{ID} can be nonzero to guarantee causality of (7.11). Because the day-ahead market has a lead time of 13 h, the trading decisions for the second day can only depend on the first 11 h of the first day. Similarly, the $4 \cdot 15$ min lead time of the intra-day market forces the first 5 rows of Q^{ID} to be zero. The size of the bidding problem can be reduced by not taking into account all the available historic activation data in (7.11) but only considering the $N_{\mathrm{lb}}^{\mathrm{DA}}$ and $N_{\mathrm{lb}}^{\mathrm{ID}}$ most recent measurements. The policy parameters $N_{\mathrm{lb}}^{\mathrm{DA}}$ and $N_{\mathrm{lb}}^{\mathrm{ID}}$ determine the width of the blocks in Q^{DA} and the number of off-diagonals in the lower-triangular part of Q^{ID}, respectively.

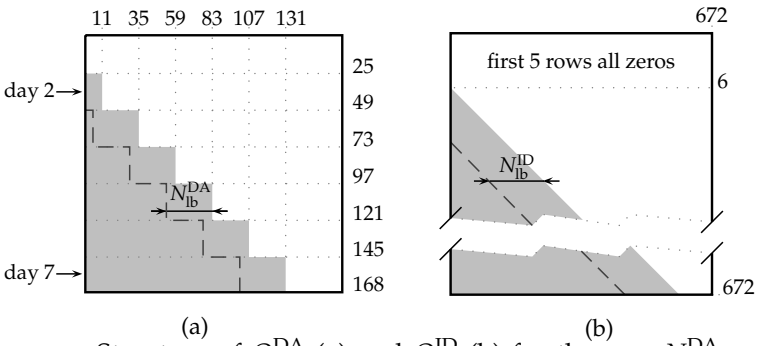

Figure 7.2: Structure of Q^{DA} (a) and Q^{ID} (b) for the case $N^{\mathrm{DA}} = 168$ and $N^{\mathrm{ID}} = 672$. The shaded grey areas denote the location of potentially nonzero elements. The amount of historic activation data considered by the policies (7.11) is governed by the parameters $N_{\mathrm{lb}}^{\mathrm{DA}}$ and $N_{\mathrm{lb}}^{\mathrm{ID}}$.

7.4 DESCRIPTION OF FLEXIBILITY

The decisions an AG makes on the different energy and reserve markets are restricted by the requirement that the target power trajectory (7.9) must be feasible for all realizations of the unknown activation signal and must satisfy all physical constraints of the unit. Thus, the AG must know the exact conditions under which a particular combination of energy trading policies and SFR capacity is feasible. We consider a unit with power, power ramp-rate, and state constraints.

7.4.1 *Power constraints*

The power the unit can draw from or feed into the power grid is limited. For all $t \in \mathcal{T}^{\mathrm{S}}(s)$, $s \in \mathcal{N}^{\mathrm{S}}$, we require that

$$\underline{p}_s \leq p_{\mathrm{tgt}}(t, w) \leq \bar{p}_s, \ \forall w \in \mathcal{W}^{N^{\mathrm{C}}}, \tag{7.13}$$

where $\underline{p}, \bar{p} \in \mathbb{R}^{N^{\mathrm{S}}}$ denote the piece-wise constant bounds on power. The constraints above are satisfied if and only if for all $\tilde{w} \in \mathcal{W}^{N^{\mathrm{S}}}$ it holds that

$$\begin{aligned} p_{\mathrm{ref},s}(\tilde{w}) + \gamma &\leq \begin{cases} \bar{p}_1, & s = 0, \\ \min\{\bar{p}_s, \bar{p}_{s+1}\}, & s = 1, \ldots, N^{\mathrm{S}} - 1, \\ \bar{p}_{N^{\mathrm{S}}}, & s = N^{\mathrm{S}}, \end{cases} \\ p_{\mathrm{ref},s}(\tilde{w}) - \gamma &\geq \begin{cases} \underline{p}_1, & s = 0, \\ \max\{\underline{p}_s, \underline{p}_{s+1}\}, & s = 1, \ldots, N^{\mathrm{S}} - 1, \\ \underline{p}_{N^{\mathrm{S}}}, & s = N^{\mathrm{S}}. \end{cases} \end{aligned} \tag{7.14}$$

The robust counterparts of the above constraints are derived according to [165] by computing the maximum and minimum of $p_{\mathrm{ref}}(\tilde{w})$ with regard to $\tilde{w}$ according to (7.12). An analytic solution exists because each element of $\tilde{w}$ is restricted to $\mathcal{W}$: the power constraints (7.13) are satisfied if and only if for all $\tilde{w} \in \mathcal{W}^{N^{\mathrm{S}}}$

$$\begin{aligned} p_{\mathrm{ref}}(\tilde{w}) &\leq \|R(MQ^{\mathrm{DA}}A^{\mathrm{DA}} + Q^{\mathrm{ID}}A^{\mathrm{ID}})\|_1 + R(Mq^{\mathrm{DA}} + q^{\mathrm{ID}}), \\ p_{\mathrm{ref}}(\tilde{w}) &\geq -\|R(MQ^{\mathrm{DA}}A^{\mathrm{DA}} + Q^{\mathrm{ID}}A^{\mathrm{ID}})\|_1 + R(Mq^{\mathrm{DA}} + q^{\mathrm{ID}}), \end{aligned} \tag{7.15}$$

where $\mathbf{1}$ denotes a column vector of ones of appropriate size.

7.4.2 *Power ramp-rate constraints*

Limits on the rate at which power can vary over time play an important role, in particular for providing ancillary services with high accuracy. For all $t \in \mathcal{T}^{\mathrm{S}}(s)$, $s \in \mathcal{N}^{\mathrm{S}}$, we require that the rate of change of the target power be bounded, i.e.,

$$\underline{r}_s \leq \partial p_{\mathrm{tgt}}(t, w)/\partial t \leq \bar{r}_s, \ \forall w \in \mathcal{W}^{N^{\mathrm{C}}}, \tag{7.16}$$

where $\underline{r}, \bar{r} \in \mathbb{R}^{N^{\mathrm{S}}}$ denote the piece-wise constant power ramp-rate limits. The above inequalities are satisfied if and only if for all $\tilde{w} \in \mathcal{W}^{N^{\mathrm{S}}}$ it holds that

$$\begin{aligned} &(p_{\mathrm{ref},s}(\tilde{w}) - p_{\mathrm{ref},s-1}(\tilde{w}))/T^{\mathrm{S}} + 2\gamma/T^{\mathrm{C}} \leq \bar{r}_s, \\ &(p_{\mathrm{ref},s}(\tilde{w}) - p_{\mathrm{ref},s-1}(\tilde{w}))/T^{\mathrm{S}} - 2\gamma/T^{\mathrm{C}} \geq \underline{r}_s, \ s \in \mathcal{N}^{\mathrm{S}}. \end{aligned} \tag{7.17}$$

The above inequalities illustrate that ramp-rate constraints limit the sum of the rate of change of the power reference and the rate of change of activated regulation power. They are equivalent to

$$\begin{aligned} &\|(R_{s+1} - R_s)(MQ^{\mathrm{DA}}A^{\mathrm{DA}} + Q^{\mathrm{ID}}A^{\mathrm{ID}})\|_1 \\ &+ (R_{s+1} - R_s)(Mq^{\mathrm{DA}} + q^{\mathrm{ID}}) + 2\gamma T^{\mathrm{S}}/T^{\mathrm{C}} \leq \bar{r}_s T^{\mathrm{S}}, \\ &- \|(R_{s+1} - R_s)(MQ^{\mathrm{DA}}A^{\mathrm{DA}} + Q^{\mathrm{ID}}A^{\mathrm{ID}})\|_1 \\ &+ (R_{s+1} - R_s)(Mq^{\mathrm{DA}} + q^{\mathrm{ID}}) - 2\gamma T^{\mathrm{S}}/T^{\mathrm{C}} \geq \underline{r}_s T^{\mathrm{S}}, \end{aligned} \tag{7.18}$$

with $s \in \mathcal{N}^{\mathrm{S}}$. The case of maximum negative ramp-rates is illustrated in Figure 7.3, where SFR activation changes from maximum positive activation ($w_{l-1} = 1$) to maximum negative activation ($w_l = -1$) within one control time step T^{C}. In this case, maximum negative ramp-rates are required during the time interval $[(l-1)T^{\mathrm{C}}, lT^{\mathrm{C}}]$.

7.4.3 *State constraints*

Many types of flexible energy units owe their flexibility to an energy buffer with limited capacity. Examples of such units are pumped hydropower plants, batteries, heating and cooling systems, capacitors, and

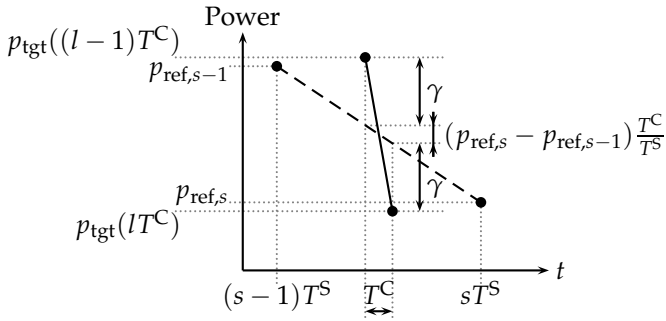

Figure 7.3: The target power trajectory $p_{\text{tgt}}(t)$ (solid line) requires maximum negative ramp-rates if $w_{l-1} = -w_l = 1$.

flywheels. Here we consider units whose energy buffer level $x(t) \in \mathbb{R}$ is governed by the linear time-invariant dynamics

$$\partial x(t,w)/\partial t = ax(t,w) + bu_s + cp_{\text{tgt}}(t,w), \tag{7.19}$$

with $t \in \mathcal{T}^{\text{S}}(s)$, $s \in \mathcal{N}^{\text{S}}$. Exogenous (uncontrollable) inputs, such as the weather conditions in the case of heating systems, or trips in the case of electric-vehicle batteries, are summarized by $u := [u_1, \ldots, u_{N^{\text{S}}}]^\top$ and are assumed to be known. Because the bidding problem has to be solved ahead of time, the initial state $x(0)$ is not known precisely, but is known to lie in the interval $[x_0^{\min}, x_0^{\max}]$. The dynamics (7.19) are characterized by the parameters $\{a, b, c\}$, where $a \leq 0$ can be interpreted as the self-dissipation rate of the energy buffer, and b, c determine the conversion efficiencies of exogenous inputs and electric energy into buffered energy, respectively. Our sign convention implies $c \geq 0$. The behavior of an ideal battery, for instance, can be modeled by setting $a = b = 0$ and $c = 1$.

We consider the piece-wise constant state constraints

$$\underline{x}_s \leq x(t,w) \leq \bar{x}_s, \ \forall w \in \mathcal{W}^{N^{\text{C}}}, \tag{7.20}$$

for all $t \in \mathcal{T}^{\text{S}}(s)$, $s = \mathcal{N}^{\text{S}}$, where $\underline{x}, \bar{x} \in \mathbb{R}^{N^{\text{S}}}$ denote the physical limits of the system state. The initial state is assumed to be feasible, i.e., $\underline{x}_1 \leq x_0^{\min} \leq x_0^{\max} \leq \bar{x}_1$. The exact evolution of the state (7.19) is unknown because it depends on the uncertain activation $w \in \mathcal{W}^{N^{\text{C}}}$. For

our case of piece-wise affine power references, the state $x(t,w)$ can evolve non-monotonically during a discretization interval $\mathcal{T}^{\mathrm{S}}(s)$. To ensure that (7.20) holds would require bounding the extreme values of $x(t,w)$ for $t \in \mathcal{T}^{\mathrm{S}}(s)$. This, however, would result in non-convex constraints. Instead, we approximate the right-hand side of (7.19) by affine functions of time and define for every interval $\mathcal{T}^{\mathrm{S}}(s)$, $s \in \mathcal{N}^{\mathrm{S}}$, the auxiliary dynamics

$$\begin{aligned}\dot{y}^{(s)}(\tau,\tilde{w}) := {} & a\bar{x}_s + bu_s + c(p_{\mathrm{ref},s-1}(\tilde{w}) - \gamma) \\ & + [\min\{0, a\bar{x}_s + bu_s + c(p_{\mathrm{ref},s}(\tilde{w}) - \gamma)\} \\ & - (a\bar{x}_s + bu_s + c(p_{\mathrm{ref},s-1}(\tilde{w}) - \gamma))]\tau/T^{\mathrm{S}},\end{aligned} \tag{7.21}$$

$$\begin{aligned}\dot{z}^{(s)}(\tau,\tilde{w}) := {} & a\underline{x}_s + bu_s + c(p_{\mathrm{ref},s-1}(\tilde{w}) + \gamma) \\ & + [\max\{0, a\underline{x}_s + bu_s + c(p_{\mathrm{ref},s}(\tilde{w}) + \gamma)\} \\ & - (a\underline{x}_s + bu_s + c(p_{\mathrm{ref},s-1}(\tilde{w}) + \gamma))]\tau/T^{\mathrm{S}},\end{aligned} \tag{7.22}$$

with $\tau \in [0, T^{\mathrm{S}}]$ and initial conditions

$$\begin{aligned} y^{(1)}(0) &= x_0^{\min} \\ z^{(1)}(0) &= x_0^{\max}, \\ y^{(s)}(0,\tilde{w}) &= z^{(s)}(0,\tilde{w}) = x((s-1)T^{\mathrm{S}},\tilde{w}). \end{aligned} \tag{7.23}$$

For brevity we use $\dot{y}^{(s)}(\tau,\tilde{w})$ to denote $\partial y^{(s)}(\tau,\tilde{w})/\partial\tau$, and similarly for x and z.

Claim 3. *If for all $s \in \mathcal{N}^S$ and $\tilde{w} \in \mathcal{W}^{N^S}$*

$$y^{(s)}(\tau,\tilde{w}) \geq \underline{x}_s \text{ and } z^{(s)}(\tau,\tilde{w}) \leq \bar{x}_s,\ \tau \in [0, T^S], \tag{7.24}$$

then the state constraints (7.20) are satisfied for all $t \in \mathcal{T}^S(s)$.

Proof. Consider the time interval $\mathcal{T}^{\mathrm{S}}(s)$ and divide it into $\kappa \in \mathcal{N}^+$ subintervals, each of duration $\delta := T^{\mathrm{S}}/\kappa$. For $\tau = 0$, the constraints (7.24) are equivalent to $\underline{x}_s \leq x((s-1)T^{\mathrm{S}}, w) \leq \bar{x}_s$, $\forall \tilde{w} \in \mathcal{W}^{N^{\mathrm{S}}}$. By definitions (7.21) and (7.22), it follows that

$$\dot{y}^{(s)}(0,\tilde{w}) \leq \dot{x}((s-1)T^{\mathrm{S}}, w) \leq \dot{z}^{(s)}(0,\tilde{w}),\ \forall \tilde{w} \in \mathcal{W}^{N^{\mathrm{S}}}. \tag{7.25}$$

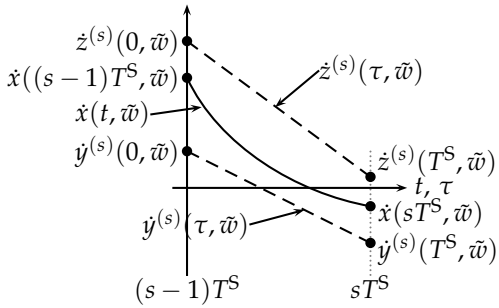

Figure 7.4: Illustration of $\dot{y}^{(s)}(\tau,\tilde{w}) \leq \dot{x}(t) \leq \dot{z}^{(s)}(\tau,\tilde{w})$ for $t \in \mathcal{T}^{\mathrm{S}}(s)$, $\tau \in [0, T^{\mathrm{S}}]$.

Approximate

$$x((s-1)T^{\mathrm{S}}+\delta, w) \approx x((s-1)T^{\mathrm{S}}, w) + \delta\dot{x}((s-1)T^{\mathrm{S}}, w), \tag{7.26}$$

$$y^{(s)}(\delta,\tilde{w}) \approx x((s-1)T^{\mathrm{S}}, w) + \delta\dot{y}^{(s)}(0,\tilde{w}), \tag{7.27}$$

$$z^{(s)}(\delta,\tilde{w}) \approx x((s-1)T^{\mathrm{S}}, w) + \delta\dot{z}^{(s)}(0,\tilde{w}), \tag{7.28}$$

and note that $y^{(s)}(\delta,\tilde{w}) \leq x((s-1)T^{\mathrm{S}}+\delta, w) \leq z^{(s)}(\delta,\tilde{w})$, which, via (7.24), implies that $\underline{x}_s \leq x((s-1)T^{\mathrm{S}}+\delta, w) \leq \bar{x}_s$. The same line of arguments holds for all subsequent subintervals. Consequently, in the limit case $\delta \to 0$, we have that $\underline{x}_s \leq x(t,w) \leq \bar{x}_s$ $\forall t \in \mathcal{T}^{\mathrm{S}}(s)$, $s \in \mathcal{N}^{\mathrm{S}}$, $w \in \mathcal{W}^{N^{\mathrm{C}}}$. □

Note that $\dot{y}^{(s)}(\tau,\tilde{w})$, $\dot{z}^{(s)}(\tau,\tilde{w})$ are affine functions of τ that can be integrated explicitly. In addition, the $\min\{\cdot\}$ and the $\max\{\cdot\}$ in the definitions (7.21) and (7.22) ensure that the minimum value of $y^{(s)}(\tau,\tilde{w})$ and the maximum value of $z^{(s)}(\tau,\tilde{w})$ over $\tau \in [0, T^{\mathrm{S}}]$ are reached at the edges of the interval, i.e., at $\tau \in \{0, T^{\mathrm{S}}\}$. Thus, (7.24) holds if for all $s \in \mathcal{N}^{\mathrm{S}}$ and $\tilde{w} \in \mathcal{W}^{N^{\mathrm{S}}}$

$$y^{(s)}(\tau,\tilde{w}) \geq \underline{x}_s,\ z^{(s)}(\tau,\tilde{w}) \leq \bar{x}_s,\ \tau \in \{0, T^{\mathrm{S}}\}, \tag{7.29}$$

which is equivalent to requiring that for all $s \in \mathcal{N}^{\mathrm{S}}$, $w \in \mathcal{W}^{N^{\mathrm{C}}}$, and $\tilde{w} \in \mathcal{W}^{N^{\mathrm{S}}}$

$$\underline{x}_s \leq x((s-1)T^{\mathrm{S}}, w) \leq \bar{x}_s, \tag{7.30}$$

$$y^{(s)}(T^{\mathrm{S}}, \tilde{w}) \geq \underline{x}_s, \tag{7.31}$$

$$z^{(s)}(T^{\mathrm{S}}, \tilde{w}) \leq \bar{x}_s, \tag{7.32}$$

Thus, (7.30) and (7.32) are sufficient conditions for the state constraints (7.20). The robust counterparts of (7.30) are derived in Section 7.7. The robust counterparts of (7.31) and (7.32) can be found similarly.

7.4.4 *Description of flexibility*

The power, ramp-rate, and state constraints (7.13), (7.16), and (7.20), respectively, are satisfied if the offered SFR capacity γ together with the energy trading policies $(Q^{\mathrm{DA}}, q^{\mathrm{DA}})$, $(Q^{\mathrm{ID}}, q^{\mathrm{ID}})$ meet the constraints (7.15), (7.18), and the robust counterparts of (7.30)-(7.32). All these constraints are linear in the decision variables $\Omega := \{Q^{\mathrm{DA}}, q^{\mathrm{DA}}, Q^{\mathrm{ID}}, q^{\mathrm{ID}}, \gamma\}$ and thus define a convex polytope $\mathcal{P}(\Phi)$ parameterized by the system constraints parameters $\Phi := \{\underline{p}, \bar{p}, \underline{r}, \bar{r}, \underline{x}, \bar{x}, x_0^{\min}, x_0^{\max}\}$. This polytope can serve as a concise and convenient description of the flexibility the unit offers with regard to trading energy in the day-ahead and intra-day markets and offering regulation power on the SFR market.

7.5 SOLVING THE BIDDING PROBLEM

7.5.1 *Formulation of the bidding problem*

Based on our model of the adjustable power references and the reformulation of the robust power, power ramp-rate and state constraints provided above, we can rewrite the robust energy and reserve bidding problem (7.2) as a standard LP. The energy and reserve bidding problem consists of finding the energy trading policies (7.11) and the reserve capacity γ that maximize the expected profit while satisfying all system constraints, i.e. ,

$$\begin{aligned} \max_{\Omega} \quad & \mathbf{E}_{\Psi}[V(\Omega, \Psi)] \\ \text{s.t.} \quad & \Omega \in \mathcal{P}(\Phi), \\ & \text{Causality of (7.11).} \end{aligned} \tag{7.33}$$

The results of the economic bidding problem (7.33) do not only depend on the market characteristics and technical constraints of the unit, but are also influenced by market prices. An economic analysis of (7.33) can be found in [139]. To investigate the effects that ramp-rate constraints, different market timescales, lead times, and trading policies have on the maximum available SFR capacity γ^*, we solve

$$\begin{aligned} \max_{\Omega} \quad & \gamma \\ \text{s.t.} \quad & \Omega \in \mathcal{P}(\Phi), \\ & \text{Causality of (7.11)} \end{aligned} \tag{7.34}$$

instead of (7.33).

It has been shown in [72, 146], among others, that the flexibility of various types of units can be characterized by the dynamics and constraints of a generalized battery. We consider a generic system modelled as an ideal battery according to (7.19) with parameters $a = 0$, $b = c = 1$, and no exogenous input, i.e., $u = 0$. The physical constraints of the battery are similar to those of a *Tesla Powerwall 2* battery [167]: $-\underline{p}_s = \bar{p}_s = 5$ kW, $\underline{x}_s = 0$ kWh, $\bar{x}_s = 15$ kWh, $s = 1, \ldots, N^{\mathrm{S}}$, and $x_0 = 7.5$ kWh. However, in contrast to [72, 146], we explicitly include the power ramp-rate constraints (7.16) in our model and investigate how different ramp-rate limits affect the result of (7.33). We emphasize that above battery model is used as an illustrative example and to facilitate the interpretation of the results. Our approach is applicable to arbitrary units (7.19) subject to power, power ramp-rate, and energy constraints (7.13), (7.16), and (7.20), respectively.

7.5.2 *Effect of ramp-rate constraints*

Consider the battery system described above with timescales $T^{\mathrm{S}} = 5$ min, $T^{\mathrm{C}} = 1$ s, $T^{\mathrm{H}} = T^{\mathrm{RES}} = 1$ d, and the Swiss intra-day energy market

where $T^{\mathrm{ID}} = 15$ min and $T^{\mathrm{ID}}_{ld} = 1$ h. If power ramp-rate constraints are neglected, the maximum SFR capacity γ^* that the unit can provide independently, i.e., without interacting with energy markets, can be computed analytically by dividing the available up- and down energy buffer capacity by the duration of the planning horizon: $\gamma^* = 7.5$ kWh / 24 h ≈ 0.31 kW, corresponding to 6.25 % of the rated power $\bar{p}$. In the case where the power reference can be adjusted via the intra-day market, the available SFR capacity computed via (7.34) increases to $\gamma^* \approx 2.59$ kW, corresponding to 51.87 % of $\bar{p}$. However, if the unit is subject to power ramp-rate limits, the feasible SFR capacity can differ from these limit-case values. Figure 7.5 illustrates the values of γ^* in percent of $\bar{p}$ for different power ramp-rate limits given in (% of $\bar{p}$)/s. The dashed lines show the SFR capacity computed above for the ramp-rate-free cases with and without access to energy markets. In the range $\bar{r} \in [0, 103.7]$ (% of $\bar{p}$)/s, the solution of (7.34) is dominated by the ramp-rate constraints (7.18). The maximum available SFR capacity γ^* increases almost linearly with $\bar{r}$ because the left-hand sides of (7.18) are dominated by the term $2\gamma T^{\mathrm{S}}/T^{\mathrm{C}}$. For ramp-rate limits $\bar{r} > 103.7$ (% of $\bar{p}$)/s, γ^* remains constant because it is bounded by the battery's energy constraints. The results in Figure 7.5 show how crucial it is to include ramp-rate limits in the formulation of the problems (7.33) and (7.34) to guarantee that the SFR service offered by the unit can be delivered accurately.

7.5.3 *Effects of market parameters*

In addition to power ramp-rate constraints, our formulation of the bidding problem takes into account all the different timescales and lead times of the energy and reserve markets. This further distinguishes our approach from existing work, such as [145, 151–153]. In the following we show that the market parameters can significantly influence the amount of SFR capacity a unit can offer. For all simulations discussed below, we consider the battery system introduced above and compute the maximum available SFR capacity γ^* via (7.34) with timescales $T^{\mathrm{S}} = 5$ min, $T^{\mathrm{C}} = 1$ s, $N^{\mathrm{DA}}_{\mathrm{lb}} = 13$ h, and $T^{\mathrm{H}} = T^{\mathrm{RES}}$. Table 7.1 provides the different market and trading-policy parameter settings considered and

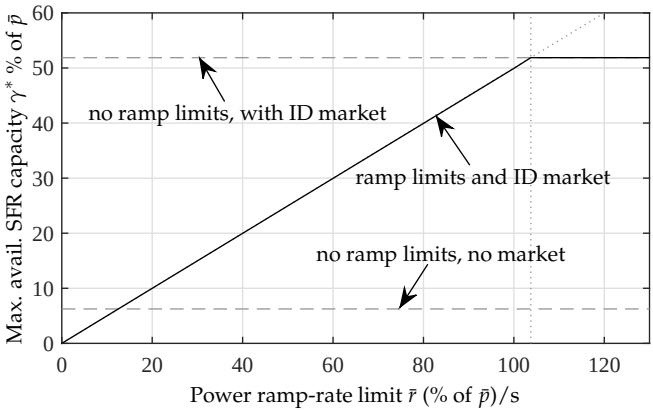

Figure 7.5: Maximum available SFR capacity γ^* (solid black) computed via (7.34) for different power ramp-rate limits $\bar{r}$. For $\bar{r} \leq 103.7\%$ of $\bar{p}/s$, the ramp-rate constraints (7.18) dominate the solution. Thus, taking them into account in (7.33) and (7.34) is crucial. The dashed lines depict γ^* for the cases in which the ramp-rate limits are neglected and *1)* the unit can trade on the intra-day market ($\gamma^* = 51.87\%$), and *2)* the unit does not trade on energy markets ($\gamma^* = 6.25\%$).

summarizes the results obtained. The minimum power ramp-rates required to offer γ^* units of SFR capacity are denoted by $\bar{r}^*$.

As a reference case, the first setting in Table 7.1 considers the case where the power reference cannot be adjusted via day-ahead or intra-day energy trades; it remains fixed regardless of potential SFR activation. This setting results in the smallest γ^*. More SFR capacity can be made available by adjusting the power reference via day-ahead energy trades depending on past activation. The capacity γ^* increases as trading decisions are based on more data on past activation, see settings 2–6, where the value of $N_{\text{lb}}^{\text{DA}}$ is gradually increased. Considering more than 24 h of past activation data, i.e., $N_{\text{lb}}^{\text{DA}} > 24$, does not yield more SFR capacity because the activation that occurred more than 24 h in the past has already been compensated for in the previous day-ahead market.

Whereas day-ahead adjustments of the power reference yield only minor increases of SFR capacity over the reference setting, intra-day adjustments make more than 50% of the rated power available for SFR, see settings 7–9. Reducing the intra-day market lead time T_{ld}^{ID} from 1 h to 1/4 h has only minor effects on γ^*, see settings 11–13. Reducing the planning horizon from 7 d to 1 d results in an increase of γ^* from 50.26% to 51.87%, see settings 7 and 11. Considering more than 15 min of past activation data, i.e., $N_{\mathrm{lb}}^{\mathrm{ID}} > 1$, does not yield more SFR capacity because the activation that occurred more than 15 min in the past has already been compensated for in the previous intra-day market step.

The results highlight the importance of the intra-day market which operates at faster timescales than the day-ahead market. It allows energy-constrained units to compensate for the regulation energy provided. The introduction of energy markets operating at even shorter time intervals, e.g. 5 min, could further increase the amount of available reserves. Additional reserves can be made available by reducing the length of the tendering period T^{RES} and by relaxing the requirement that a constant amount of SFR must be offered.

#	T^{RES}	T_{ld}^{ID}	N_{lb}^{DA}	N_{lb}^{ID}	γ^*	$\bar{r}^*$	Sol. time
	d	h			% of $\bar{p}$	(% of $\bar{p}$)/s	
1	7	1	0	0	0.89	1.78	8 min
2	7	1	1	0	0.93	1.87	9 min
3	7	1	2	0	0.96	1.94	12 min
4	7	1	6	0	1.14	2.29	40 min
5	7	1	12	0	1.55	3.13	100 min
6	7	1	≥ 24	0	4.05	8.25	270 min
7	7	1	0	≥ 1	50.26	100.69	30 min
8	7	1/2	0	≥ 1	50.34	100.84	30 min
9	7	1/4	0	≥ 1	50.37	100.91	30 min
10	1	1	0	0	6.25	12.50	5 s
11	1	1	0	≥ 1	51.87	103.90	25 s
12	1	1/2	0	≥ 1	52.38	104.92	25 s
13	1	1/4	0	≥ 1	52.63	105.42	25 s

Table 7.1: Maximum SFR capacity and minimum required ramp-rates for different market settings and trading policies.

7.5.4 *Computational complexity*

The bidding problem (7.34) has been solved with *IBM ILOG® CPLEX®*, v12.6 on a system featuring an *Intel® Xeon®* 12-core CPU @2.4 GHz and 20 GB of RAM. The solving times are provided in the rightmost column of Table 7.1. The number of decision variables and constraints of (7.34) is determined by various factors, such as the planning horizon, the different market timescales, the system discretization time, and also the trading policies used. An efficient means to reduce the problem size is to limit the number of free variables in the trading policies (7.11) via the parameters N_{lb}^{DA} and N_{lb}^{ID}.

7.6 CONCLUSION & FUTURE WORK

The provision of accurate and reliable SFR services is important to compensate for imbalances in the power grid. We introduced piece-wise affine and continuous power trajectories compliant with ENTSO-E requirements, and derived a reformulation of the energy and reserve bidding problem as a robust linear program that takes into account power, power ramp-rate and state constraints of an FEU and all the different timescales and lead times involved in energy and reserve markets. Our approach is versatile and applicable to various market settings. Simulation results highlight the importance of power ramp-rate limits and illustrate the effect of different market settings on the amount of available SFR capacity a unit can deliver.

There are several directions for future work. First, the outcomes of the energy and SFR markets are unknown at the time of bidding because they are both auction-based. Thus, strategies to compute competitive energy and reserve bids will be studied. In addition, other objectives than the expected profit can be used in the bidding problem, such as risk-aware measures, for example. Further, the effect of using piece-wise affine as opposed to piece-wise constant power trajectories should be studied with regard to real, performance-based SFR reimbursement schemes and the residual power imbalance.

7.7 APPENDIX: BOUNDS ON DISCRETE-TIME STATE DYNAMICS

Consider the linear time-invariant state dynamics (7.19):

$$\partial x(t,w)/\partial t = ax(t,w) + bu_s + cp_{\text{tgt}}(t,w),$$

$t \in \mathcal{T}^{\text{S}}(s)$, $s \in \mathcal{N}^{\text{S}}$, with $x_0 := x(0) \in [x_0^{\min}, x_0^{\max}]$ and $p_{\text{tgt}}(t,w)$ according to (7.9). Standard integration techniques can be used to express the state $x_s(w) := x(sT^{\text{S}}, w)$ as

$$x_s(w) = fx_{s-1}(w) + gu_s + h_1 p_{\text{ref},s-1}(\tilde{w}) + h_2 p_{\text{ref},s}(\tilde{w}) + \gamma v_s, \quad (7.35)$$

where we have used

$$f := e^{aT^{\text{S}}}, \qquad g := b\int_0^{T^{\text{S}}} e^{a\tau} d\tau, \quad (7.36)$$

$$h_1 := \frac{c}{T^{\text{S}}}\int_0^{T^{\text{S}}} e^{a\tau}\tau d\tau, \quad h_2 := \frac{c}{T^{\text{S}}}\int_0^{T^{\text{S}}} e^{a\tau}(T^{\text{S}} - \tau) d\tau, \quad (7.37)$$

$$v_s := c\int_0^{T^{\text{S}}} e^{a\tau} w(sT^{\text{S}} - \tau) d\tau. \quad (7.38)$$

Note that the power reference is determined by the averaged activation $\tilde{w}$, whereas the term v_s depends on the unfiltered activation signal $w(t)$. Because $\tilde{w}$ and v are computed from w, we write $x(w)$ only. The discrete-time state evolution can be written in vector form as

$$x(w) = Fx_0 + Gu + Hp_{\text{ref}}(\tilde{w}) + \gamma Kv, \quad (7.39)$$

with $F \in \mathbb{R}^{N^{\text{S}}}$, $G, K \in \mathbb{R}^{N^{\text{S}} \times N^{\text{S}}}$, and $H \in \mathbb{R}^{N^{\text{S}} \times (N^{\text{S}}+1)}$.

By plugging (7.7) into (7.39), we have

$$x_s(w) = \sum_{i=1}^{s} \{Q_{s,i}\tilde{w}_i + \gamma K_{s,i} v_i\} + \kappa_s, \quad (7.40)$$

where we have used

$$Q := HR(MQ^{\text{DA}}A^{\text{DA}} + Q^{\text{ID}}A^{\text{ID}}) \quad (7.41)$$

$$\kappa := Fx_0 + Gu + HR(Mq^{\text{DA}} + q^{\text{ID}}) \quad (7.42)$$

for brevity. Using the definitions of $\tilde{w}$ and v, we can write

$$x_s(w) = \sum_{i=1}^{s} \left\{ \int_0^{T^S} \left(\frac{Q_{s,i}}{T^S} + c\gamma e^{a\tau} K_{s,i} \right) w(\tau) d\tau \right\} + \kappa_s. \tag{7.43}$$

Note that in case of self-dissipation, i.e. , $a < 0$, the integrand is time-dependent, making it difficult to evaluate the integral analytically. To circumvent this, we substitute $e^{a\tau}$ by the constant $e^{a\hat{\tau}}$ which is chosen such that for all $\tau \in [0, T^S]$

$$|e^{a\tau} - e^{a\hat{\tau}}| \leq \epsilon := (1 - e^{aT^S})/2. \tag{7.44}$$

Because

$$\begin{aligned} \max_{w \in \mathcal{W}} \int_0^{T^S} & c\gamma K_{s,i}(e^{a\tau} - e^{a\hat{\tau}}) w(\tau) d\tau \\ & = \int_0^{T^S} c\gamma K_{s,i} |e^{a\tau} - e^{a\hat{\tau}}| d\tau \leq c\gamma\epsilon T^S K_{s,i}, \end{aligned} \tag{7.45}$$

the integral in (7.43) is bounded from above:

$$\begin{aligned} & \int_0^{T^S} \left(\frac{Q_{s,i}}{T^S} + c\gamma e^{a\tau} K_{s,i} \right) w(\tau) d\tau \\ \leq & \int_0^{T^S} \left(\frac{Q_{s,i}}{T^S} + c\gamma e^{a\hat{\tau}} K_{s,i} \right) w(\tau) d\tau + c\gamma\epsilon T^S K_{s,i} \\ \leq & |Q_{s,i} + c\gamma e^{a\hat{\tau}} T^S K_{s,i}| + c\gamma\epsilon T^S K_{s,i} \quad \forall w \in \mathcal{W}^{N^C}. \end{aligned} \tag{7.46}$$

Bounds on $x_s(w)$ can be found by substituting the summands in (7.43) with (7.46): For all $w \in \mathcal{W}^{N^C}$ and $s \in \mathcal{N}^S$

$$\begin{aligned} x_s(w) &\leq \|Q_s + c\gamma e^{a\hat{\tau}} T^S K_s\|_1 + c\gamma\epsilon T^S K_s \mathbf{1} + \kappa_s, \\ x_s(w) &\geq -\|Q_s + c\gamma e^{a\hat{\tau}} T^S K_s\|_1 - c\gamma\epsilon T^S K_s \mathbf{1} + \kappa_s. \end{aligned} \tag{7.47}$$

Inequalities (7.47) are the robust counterparts of (7.30).

8 SYNERGY EFFECT OF AGGREGATION

In the present chapter we show how the flexibility of groups of energy units can be exploited more efficiently by considering units with complementary physical properties and controlling them in a coordinated way. Based on the robust optimization framework introduced in Chapter 7, we demonstrate that aggregations of select energy units can offer significantly more regulation capacity than the units could provide individually. We call this the *synergy* effect of aggregation.

8.1 INTRODUCTION

The potential of individual units to provide SFR can be severely limited by their physical constraints. For instance, batteries, residential heating and cooling systems, and flywheels are mainly limited by the size of their energy buffers [141, 168], while steam turbines and industrial cooling systems must adhere to restrictive power ramp-rate limits [169]. Moreover, operational constraints, such as the trips of PEVs, run-time constraints, and control delays, can have significant impact on the ability of a unit to provide SFR. In addition to its physical and operational constraints, a unit offering SFR services must satisfy the terms and conditions dictated by the energy and reserve markets, such as market timescales and lead times, as well as particular requirements on bids, such as time-invariant and symmetric reserve bids [137].

Approaches to remedy these issues fall into two main groups. The approaches of the first group propose new SFR products that are better suited for the kind of units listed above, for example by limiting the energy content [34, 150, 170] or the frequency bandwith of SFR products [32, 140, 157], or by reducing the time period over which SFR capacity must be available [70]. The second group of approaches relies on aggregating the flexibility of multiple units and controlling them in a coordinated way. This has proven useful to alleviate the difficulties related to the duration and symmetry requirements of certain SFR mar-

kets [70, 71]. The same aggregation techniques can be used to couple energy storage systems with energy markets, allowing the former to adjust their energy schedules via the markets and, thereby, to exploit their flexibility more efficiently [139, 168].

While considering power and energy constraints of individual units when solving the energy and reserve bidding problem, power ramp-rate constraints are commonly neglected in the literature. This is a major shortcoming because the accurate provision of SFR services can be compromised by insufficient ramping capabilities. The crucial importance of incorporating ramping constraints has already been recognized in the context of grid balancing [157] and unit commitment [158, 159], and has even lead to the adoption of new performance-based schemes for the remuneration of SFR provision by several SOs [160].

We consider an AG that controls the electricity generation or consumption of a group of energy units and study the energy and reserve bidding problem introduced in Chapter 7. In our formulation of the bidding problem, we adopt the robust formulations of [70, 118, 152, 153], and make two main contributions. First, building on the foundations developed in Chapter 7, we formulate the energy and reserve bidding problem for an aggregation of multiple units. In contrast to existing approaches, such as [70, 118, 152, 153], we properly model and aggregate the flexibility of energy units that are subject to power ramp-rate constraints and control delays, and consider all the different timescales of the energy and reserve markets. Second, we demonstrate that aggregating multiple FEUs with complementary physical properties can unlock significant amounts of SFR capacity. Simulations show that aggregating fast-responding but energy-constrained units, such as batteries or capacitors, and slowly-responding units with large energy buffers, such as industrial heating and cooling systems or steam turbines, yield an increase in the available SFR capacity in the range of 18-1190% depending on the system parameters.

This chapter is organized as follows. Section 8.2 introduces the control setup that underlies the energy and reserve bidding problem formulated for multiple FEUs in Section 8.3. In Section 8.4 and Section 8.5, we solve the bidding problem for aggregations of FEUs with complementary physical properties and discuss the synergy effects that can be observed. A conclusion and outlook on future work are given in Section 8.6.

8.2 ADJUSTABLE POWER REFERENCE

8.2.1 *Power reference and power target*

Every unit $j \in \mathcal{J} := \{1, \ldots, J\}$ in the balance group of the AG implements a continuous-time piece-wise affine and continuous power reference which is fully defined by its breakpoints $p_{\text{ref}}^{(j)} := [p_{\text{ref},0}^{(j)}, \ldots, p_{\text{ref},N^{\text{S}}}^{(j)}]^\top$, according to

$$p_{\text{ref}}^{(j)}(t) := p_{\text{ref},s-1}^{(j)} + (p_{\text{ref},s}^{(j)} - p_{\text{ref},s-1}^{(j)})(t - (s-1)T^{\text{S}})/T^{\text{S}}, \tag{8.1}$$

for $t \in \mathcal{T}^{\text{S}}(s)$, $s \in \mathcal{N}^{\text{S}}$. Thus, the units in the balance group collectively implement the aggregate reference $p_{\text{ref}}^{(\mathcal{J})}(t) := p_{\text{ref}}^{(1)}(t) + \cdots + p_{\text{ref}}^{(J)}(t)$.

We assume that an AG offering the SFR capacity $\gamma^{(\mathcal{J})} \in \mathbb{R}^{N^{\text{S}}}$ is responsible for its balance group to continuously track the aggregate target power level

$$p_{\text{tgt}}^{(\mathcal{J})}(t, w) := p_{\text{ref}}^{(\mathcal{J})}(t) + \gamma_s^{(\mathcal{J})} w(t), \tag{8.2}$$

with $w(t)$ denoting the continuous SFR activation signal (7.8), and $t \in [0, T^{\text{H}}]$, $s \in \mathcal{N}^{\text{S}} : t \in \mathcal{T}^{\text{S}}(s)$. The SFR delivered by an AG is the result of a collective effort of all the units in its balance group, i.e., $\gamma^{(\mathcal{J})} := \gamma^{(1)} + \cdots + \gamma^{(J)}$. Consequently, each unit is required to follow its own target power trajectory

$$p_{\text{tgt}}^{(j)}(t, w) := p_{\text{ref}}^{(j)}(t) + \gamma_s^{(j)} w(t). \tag{8.3}$$

Compared with $p_{\text{ref}}^{(j)}(t)$, which is piece-wise affine on the timescale T^{S}, the activation $w(t)$ can vary at the higher rate T^{C}. Thus, accurate tracking of (8.3) puts higher requirements on the ramping capabilities and control delays of a unit.

8.2.2 *Adjustable power reference*

To exploit the flexibility of energy units more efficiently, it has been proposed to adjust the reference power schedules (8.1) in response to past activation, see [34, 152–154], among others. Here we consider

power references (8.1) whose breakpoints $p_{\text{ref}}^{(j)}$ are affinely adjustable based on past activation, i.e.,

$$p_{\text{ref}}^{(j)}(\tilde{w}) = Q^{(j)}\tilde{w} + q^{(j)},\ j \in \mathcal{N}^{J}, \tag{8.4}$$

where $\tilde{w} \in \mathcal{W}^{N^{\text{S}}}$ denotes the activation signal $w(t)$ averaged over time intervals of duration T^{S}, and $Q^{(j)} \in \mathbb{R}^{(N^{\text{S}}+1)\times N^{\text{S}}}$ and $q^{(j)} \in \mathbb{R}^{N^{\text{S}}+1}$ are parameters. Equation (8.4) is a control policy describing the power reference a unit should follow depending on a certain SFR activation.

To be applicable in practice, the policy (8.4) must be *causal*, i.e., it must consider the fact that the value of $\tilde{w}_s$ becomes known only at time sT^{S} and may be used exclusively to adjust future breakpoint values $p_{\text{ref},s+1}^{(j)}, \ldots, p_{\text{ref},N^{\text{S}}}^{(j)}$. Thus, the first two breakpoints, $p_{\text{ref},0}$ and $p_{\text{ref},1}$, must be set ahead of time without any knowledge about $\tilde{w}$. That is, $p_{\text{ref},0} = q_1^{(j)}$ and $p_{\text{ref},1} = q_2^{(j)}$. Consequently, the first two rows of the matrix $Q^{(j)}$ must be all zeros. At time T^{S}, the first measurement $\tilde{w}_1$ is available and can be used to compute $p_{\text{ref},2}^{(j)}$. In general, the conditions for the policy (8.4) to be causal are

$$Q_{m,n}^{(j)} = 0 \ \ \forall m \in \mathcal{N}^{\text{S}+1},\, n \in \mathcal{N}^{\text{S}} : m \leq n+1, \tag{8.5}$$

where $Q_{m,n}^{(j)}$ denotes the mth element in the nth column of the matrix $Q^{(j)}$.

8.2.3 *Energy balance*

To comply with the energy contracts concluded on the day-ahead and intra-day markets, the balance-responsible AG must ensure that its balance group follows the net power references p^{DA} and p^{ID} resulting from the day-ahead and intra-day energy markets, respectively, i.e.

$$p_{\text{ref}}^{(\mathcal{J})}(\tilde{w}) = p^{\text{DA}} + p^{\text{ID}}\ \forall \tilde{w} \in \mathcal{W}^{N^{\text{S}}}. \tag{8.6}$$

Given our choice of affine policies (8.4), the balance equation (8.6) suggests that the day-ahead and intra-day trading decisions made on the aggregate level be affine policies of $\tilde{w}$ also, i.e.

$$\begin{aligned} p^{\mathrm{DA}}(\tilde{w}) &= Q^{\mathrm{DA}}\tilde{w} + q^{\mathrm{DA}}, \\ p^{\mathrm{ID}}(\tilde{w}) &= Q^{\mathrm{ID}}\tilde{w} + q^{\mathrm{ID}}. \end{aligned} \tag{8.7}$$

In this case, the balance equation (8.6) is satisfied if and only if

$$\begin{aligned} Q^{(\mathcal{J})} &:= \sum_{j \in \mathcal{N}^{\mathcal{J}}} Q^{(j)} = Q^{\mathrm{DA}} + Q^{\mathrm{ID}}, \text{ and} \\ q^{(\mathcal{J})} &:= \sum_{j \in \mathcal{N}^{\mathcal{J}}} q^{(j)} = q^{\mathrm{DA}} + q^{\mathrm{ID}}, \end{aligned} \tag{8.8}$$

with parameters $Q^{\mathrm{DA}}, Q^{\mathrm{ID}} \in \mathbb{R}^{(N^{\mathrm{S}}+1)\times N^{\mathrm{S}}}$ and $q^{\mathrm{DA}}, q^{\mathrm{ID}} \in \mathbb{R}^{N^{\mathrm{S}}+1}$. The power policies (8.7) correspond to the energy trading policies

$$\begin{aligned} e^{\mathrm{DA}}(\tilde{w}) &= E^{\mathrm{DA}}(Q^{\mathrm{DA}}\tilde{w} + q^{\mathrm{DA}}), \\ e^{\mathrm{ID}}(\tilde{w}) &= E^{\mathrm{ID}}(Q^{\mathrm{ID}}\tilde{w} + q^{\mathrm{ID}}), \end{aligned} \tag{8.9}$$

where $E^{\mathrm{DA}} \in \mathbb{R}^{N^{\mathrm{DA}}\times(N^{\mathrm{S}}+1)}$ and $E^{\mathrm{ID}} \in \mathbb{R}^{N^{\mathrm{ID}}\times(N^{\mathrm{S}}+1)}$ map power to energy. The AG can use the policies (8.9) to decide on the amount of energy to trade on the corresponding markets based on past SFR activation data. To be applicable, the policies must, however, be causal and take into account the lead times of the different energy markets. These requirements can be incorporated by imposing particular structures onto the matrices $(E^{\mathrm{DA}}Q^{\mathrm{DA}})$ and $(E^{\mathrm{ID}}Q^{\mathrm{ID}})$ as discussed in Subsection 7.3.4 and illustrated in Figure 7.2.

8.3 THE BIDDING PROBLEM

8.3.1 *Description of flexibility*

The decisions the AG makes on the different energy and reserve markets are restricted by the requirement that the individual target power trajectories (8.3) must satisfy all the physical constraints of the corresponding FEU for all realizations of the unknown activation signal. We

consider the same power, power ramp-rate, and state constraints as in Section 7.4. All their robust counterparts are affine in the AG's decision variables $\zeta^{(j)} := \{Q^{(j)}, q^{(j)}, \gamma^{(j)}\}$. Thus, the feasible set of each unit j can be represented by a convex compact polytope $\mathcal{P}(\Phi^{(j)})$ parameterized by the system constraints and dynamics parameters $\Phi^{(j)} := \{\underline{p}^{(j)}, \bar{p}^{(j)}, \underline{r}^{(j)}, \bar{r}^{(j)}, \underline{x}^{(j)}, \bar{x}^{(j)}, x^{(j)}_{0,\min}, x^{(j)}_{0,\max}, a^{(j)}, b^{(j)}, c^{(j)}\}$.

8.3.2 *Problem formulation*

Based on our model of the adjustable power references and the reformulation of the robust power, power ramp-rate and state constraints provided above, we can rewrite the robust energy and reserve bidding problem (7.2) as a standard LP. The problem consists of finding the power reference policies (8.4), the trading policies (8.7), and the SFR capacities $\gamma^{(j)}$ that maximize the expected profit while satisfying all the system and market constraints, i.e.,

$$\max_{\Omega} \quad \mathbf{E}_{\Psi}[V(\Omega, \Psi)] \tag{8.10a}$$

$$\text{s.t.} \quad \zeta^{(j)} \in \mathcal{P}(\Phi^{(j)}),\ j \in \mathcal{N}^J, \tag{8.10b}$$

$$\sum_{j \in \mathcal{N}^J} Q^{(j)} = Q^{\mathrm{DA}} + Q^{\mathrm{ID}}, \tag{8.10c}$$

$$\sum_{j \in \mathcal{N}^J} q^{(j)} = q^{\mathrm{DA}} + q^{\mathrm{ID}}, \tag{8.10d}$$

$$\text{Causality of (8.4) via (8.5)}, \tag{8.10e}$$

$$\text{Market constraints}, \tag{8.10f}$$

where $\Omega := \{\zeta^{(1)}, \ldots, \zeta^{(J)}, Q^{\mathrm{DA}}, q^{\mathrm{DA}}, Q^{\mathrm{ID}}, q^{\mathrm{ID}}\}$ summarizes all the decision variables. The constraints (8.10b) ensure the feasibility of the target power trajectory (8.3) with respect to the constraints of the individual unit $j \in \mathcal{N}^J$. The constraints (8.10c) and (8.10d) guarantee that the AG complies with all energy market trades. The power reference policies must be causal (8.10e), and respect all market constraints (8.10f), such as market lead times or the requirement of time-invariant reserve bids.

The constraints of the bidding problem (8.10) can easily be augmented. For example, special requirements of energy or reserve markets

can make it necessary to add market-specific constraints. For example, some reserve markets, such as the Swiss SFR market [137], require the reserve bids to be constant over the tendering period. This can be considered by adding the auxiliary decision variable $\gamma^{(agg)}$ and the constraints $\gamma_s^{(1)} + \cdots + \gamma_s^{(J)} = \gamma^{(agg)}$, $s \in \mathcal{N}^{\mathrm{S}}$. In addition, regulations in Continental Europe require that energy schedules are implemented as piece-wise affine and continuous power trajectories that alternate between linear ramps and periods of constant power [138]. These requirements can be formulated as linear constraints on $\{Q^{\mathrm{DA}}, q^{\mathrm{DA}}, Q^{\mathrm{ID}}, q^{\mathrm{ID}}\}$.

8.3.3 *Complexity of the bidding problem*

The number of decision variables and constraints of problem (8.10) is determined by various factors, such as the length of the planning horizon, the system discretization time, the different market timescales, and also the structure of the Q-matrices that determine the power reference and trading policies (8.4) and (8.7), respectively. For instance, the problem size can be reduced significantly, if the power reference policy (8.4) considers only the $r \geq 1$ most recent measurements of the SFR activation signal instead of its entire history. That is, in addition to (8.5) we require that

$$Q_{m,n}^{(j)} = 0 \; \forall m \in \mathcal{N}^{\mathrm{S}+1}, \, n \in \mathcal{N}^{\mathrm{S}} : m > n + r + 2, \tag{8.11}$$

which imposes a band structure onto $Q^{(j)}$. The effects of different policies on the solution of problem (8.10) and the required solving time have been investigated in Chapter 7.

Solving problem (8.10) for aggregations of a large number of units can be inefficient. However, the particular structure of the problem, that is its decomposable objective and the linear coupling constraints (8.10c) and (8.10d), allow us to apply decentralized optimization methods, such as the Alternating Direction Method of Multipliers (ADMM), see [71, 171] for more details.

8.3.4 *The cost of flexibility*

Exploiting the flexibility of electric energy units can have undesirable impacts on the units and their owners. For example, heating and cooling systems deviating from set-point profiles can lead to discomfort, and forcing batteries or gas turbines to track volatile SFR activation signals can cause battery degradation [172], put high thermal stress on turbine components, and entail higher plant operations and maintenance costs [173].

Thus, AGs conclude contracts with individual unit owners, defining the specifics of the flexibility offered by the units, as well well as the details on how the units are reimbursed for the flexibility provided [39, 45]. For example, the flexibility reimbursement of individual units could be based on flexibility cost functions $C^{(j)}$ [48] that assign a cost to the family of possible power references (8.4) and power target trajectories (8.3) that a unit j might be required to follow for a given choice of the decision variables $\{Q^{(j)}, q^{(j)}, \gamma^{(j)}\}$. The AG can consider the cost of flexibility when solving the energy and reserve bidding problem by augmenting its profit function (7.1) by the negative flexibility costs $-C^{(j)}$, $j \in \mathcal{N}^J$.

8.4 SYNERGY I: BATTERY AND FREEZER

Aggregating the flexibility of multiple energy units has been shown to alleviate some of the challenges faced by individual units, such as limits on the minimum bid size and the requirement of time-invariant SFR bids [70, 71]. Here, we show that aggregating the flexibility of units with complementary physical properties can significantly increase the amount of SFR capacity that can be offered by the aggregation.

For all the simulations discussed below we consider a planning horizon of duration $T^{\mathrm{H}} = 1$ d, an intra-day market operating on the timescale $T^{\mathrm{ID}} = 15$ min with a lead time of $T^{\mathrm{ID}}_{\mathrm{ld}} = 1$ h, and an SFR market that accepts only *symmetric and time-invariant* capacity bids [137] for a tendering period of $T^{\mathrm{RES}} = 1$ d. The dynamics of the units are time-discretized on the timescale $T^{\mathrm{S}} = 5$ min, and the SFR service is activated on the timescale $T^{\mathrm{C}} = 10$ s.

We consider two flexible energy units, an ideal battery (B), and an industrial freezer warehouse (F).

8.4.1 *Battery model*

The battery is modeled as an ideal energy buffer whose state dynamics are governed by (7.19) with parameters $a^{(B)} = b^{(B)} = 0$ and $c^{(B)} = 1$. The ramping capabilities of the battery are assumed to be fast enough so that the power ramp-rate constraints can be neglected. However, the battery is subjected to power constraints (7.13) with symmetric limits $-\underline{p}^{(B)} = \bar{p}^{(B)}$, and state constraints (7.20) with limits $\underline{x}^{(B)} = 0$, $\bar{x}^{(B)} > 0$ and $x^{(B)}_{0,\min} = x^{(B)}_{0,\max} = \bar{x}^{(B)}/2$.

The maximum SFR capacity the battery can offer is

$$\gamma^{(B)}_{\max} = \min\{|\bar{p}^{(B)} - \underline{p}^{(B)}|/2, (\bar{x}^{(B)} - x^{(B)}_{0,\max})/T^{\mathrm{SFR}}, \ldots \\ (x^{(B)}_{0,\min} - \underline{x}^{(B)})/T^{\mathrm{SFR}}\}. \quad (8.12)$$

The maximum constant and symmetric SFR capacity $\gamma^{(B)}_{\max}$ that an individual battery can provide is determined by the most restrictive of three limits given on the right-hand side of (8.12). First, the SFR capacity is limited by the battery's power constraints. Because the battery must be able to provide the same amount of up- and down regulation, $\gamma^{(B)}_{\max} \leq |\bar{p}^{(B)} - \underline{p}^{(B)}|/2$ is required. The second limiting factor is the battery's energy capacity. In the extreme case where the full up-regulation $-\gamma^{(B)}_{\max}$ is activated over the entire period T^{RES}, the battery must be able to deliver the corresponding energy to the power grid, i.e., $\gamma^{(B)}_{\max} T^{\mathrm{RES}} \leq x^{(B)}_{0,\min} - \underline{x}^{(B)}$. Similarly, if the full down-regulation $\gamma^{(B)}_{\max}$ is activated over the entire period T^{RES}, the battery must be able to absorb the corresponding amount of energy, i.e., $\gamma^{(B)}_{\max} T^{\mathrm{RES}} \leq \bar{x}^{(B)} - x^{(B)}_{0,\max}$.

In general, the SFR potential of individual batteries is limited by these capacity constraints because their energy capacity is usually small compared to their power rating. Consider for example the battery of a *Tesla Model-S* with power rating $\bar{p}^{(B)} = -\underline{p}^{(B)} = 17.2$ kW, energy capacity $\bar{x}^{(B)} = 100$ kWh [174], and initial state $x^{(B)}_0 = 50$ kWh providing constant and symmetric SFR over a time horizon of $T^{\mathrm{RES}} = 24$ h. According

to (8.12), the battery can provide $\gamma_{\max}^{(B)} = 2.08$ kW of SFR which corresponds to only 12.1% of the SFR capacity it could provide if the capacity constraints are neglected (namely, $\bar{p}^{(B)} = 17.2$ kW).

8.4.2 *Freezer model*

The freezer is modeled as a single energy buffer whose state evolves according to (7.19). It must respect the power, power ramp-rate, and state constraints (7.13), (7.16), and (7.20), respectively. The parameters of the freezer are similar to the ones of the freezer warehouse described in [52] and are given in Table 8.1. The freezer consumes electric energy to keep its indoor air temperature θ_{in} within the admissible range $[\underline{\theta}_{\text{in}}, \bar{\theta}_{\text{in}}]$. The time required to cool the freezer from $\bar{\theta}_{\text{in}}$ down to $\underline{\theta}_{\text{in}}$ at maximum cooling power $\bar{p}^{(F)}$ is assumed to be 6 h. That is, the freezer has an energy storage capacity of $\bar{x}^{(F)} = 6 \cdot \bar{p}^{(F)} = 1.8$ MW. Given the constant outdoor air temperature $\theta_{\text{out}} = 5°\text{C}$, we assume that without cooling the indoor temperature rises from $\underline{\theta}_{\text{in}}$ to $\bar{\theta}_{\text{in}}$ in $T^{\text{dis}} = 10$ h. The parameters of the freezer state dynamics (7.19) are

$$\begin{aligned} a^{(F)} &= \frac{1}{T^{\text{dis}}} \ln\left(\frac{\bar{\theta}_{\text{in}} - \underline{\theta}_{\text{in}}}{\underline{\theta}_{\text{in}} - \theta_{\text{out}}}\right), \\ b^{(F)} &= a^{(F)} \bar{x}^{(F)} / (\underline{\theta}_{\text{in}} - \bar{\theta}_{\text{in}}), \\ c^{(F)} &= 1. \end{aligned} \tag{8.13}$$

We assume that the power consumption of the freezer's cooling system can be set to any value in the continuous range $[\underline{p}^{(F)}, \bar{p}^{(F)}]$. However, power set point changes are not implemented immediately but only after a constant delay $\delta^{(F)}$ which can be determined by measuring the change in the unit's power consumption in response to a power set point change. Here we assume that $T^{\text{C}} < \delta^{(F)} < T^{\text{S}}$ which makes it impossible for the freezer to track an SFR target power trajectory (8.3) because it cannot react on the timescale T^{C}. Consequently, the freezer cannot offer any SFR, i.e. $\gamma_{\max}^{(F)} = 0$ kW. A delay can also impact the ability of a unit to adjust its power reference based on past SFR activation

Symbol	Value	Description
$\underline{p}^{(F)}$, $\bar{p}^{(F)}$	0, 300 kW	Min., max. power consumption
$\underline{r}^{(F)}$, $\bar{r}^{(F)}$	−100, 100 kW/min	Min., max. power ramp-rate
$\underline{x}^{(F)}$, $\bar{x}^{(F)}$	0, 1.8 MWh	Min., max. state
$\underline{x}^{(F)}_{\min} = \bar{x}^{(F)}_{\max}$	0.9 MWh	Limits on initial state
$\underline{\theta}_{\text{in}}$, $\bar{\theta}_{\text{in}}$	−29, −27 °C	Admissible indoor temp. range
θ^{out}	5 °C	Constant outdoor air temp.
$a^{(F)}$	$-0.0061\ \text{h}^{-1}$	Self-dissipation rate
$b^{(F)}$	5.4562 kW/°C	Heat flux to outside
$c^{(F)}$	41	Efficiency of cooling system

Table 8.1: Freezer warehouse parameters.

according to (8.4). This can be incorporated explicitly in the bidding problem by forcing additional elements of $Q^{(F)}$ to zero:

$$Q^{(F)}_{m,n} = 0 \ \forall m \in \mathcal{N}^{\text{S}+1},\ n \in \mathcal{N}^{\text{S}} : n \geq m - 1 - \left\lceil \frac{\delta^{(F)}}{T^{\text{S}}} \right\rceil . \tag{8.14}$$

8.4.3 *Maximizing the aggregate SFR potential*

Individual FEUs can adjust their power trajectories only by trading energy on day-ahead or intra-day markets. In contrast, the units in a balance group can make bilateral adjustments to their power references while keeping the aggregate power reference unchanged, i.e., without the need for trading energy via markets.

To illustrate the beneficial effect of aggregation, we compute the maximum constant and symmetric SFR capacity $\gamma^{(\mathcal{J})}_{\max} \in \mathbb{R}^{+}_{0}$ that can be

offered by the freezer and the battery collectively by solving the slightly modified bidding problem:

$$\gamma_{\max}^{(\mathcal{J})} := \max_{\Omega} \quad \min_{k \in \mathcal{N}^{\mathrm{S}}} \{\gamma_k^{(B)} + \gamma_k^{(F)}\} \tag{8.15a}$$

$$\text{s.t.} \quad \zeta^{(j)} \in \mathcal{P}(\Phi^{(j)}),\ j \in \{B, F\}, \tag{8.15b}$$

$$Q^{(B)} + Q^{(F)} = Q^{\mathrm{DA}} + Q^{\mathrm{ID}}, \tag{8.15c}$$

$$q^{(B)} + q^{(F)} = q^{\mathrm{DA}} + q^{\mathrm{ID}}, \tag{8.15d}$$

$$Q^{\mathrm{DA}} = Q^{\mathrm{ID}} = 0,\ q^{\mathrm{ID}} = 0, \tag{8.15e}$$

$$\text{Causality of (8.4) via (8.5)}, \tag{8.15f}$$

$$\text{Control delay constraint (8.14)}. \tag{8.15g}$$

By adding the constraints (8.15e), we model the case where the *aggregate* power reference (8.6) is fixed and cannot be adjusted via trades on energy markets. However, the individual power references of the battery and the freezer, $p_{\mathrm{ref}}^{(B)}$ and $p_{\mathrm{ref}}^{(F)}$, respectively, are still adjustable according to (8.4) as long as all adjustments cancel out on the aggregate level. Any market constraints (8.10f) can be omitted. The control delay of the freezer is considered by (8.15g).

Solving (8.15) using the *Tesla Model-S* battery parameters yields $\gamma_{\max}^{(\mathcal{J})} = \gamma^{(B)*} = 9.61$ kW, which is 55.9% of $\bar{p}^{(B)}$ and represents an increase by a factor of 4.6 compared to the amount of SFR the battery can provide on its own. The optimal power reference policy parameters are

$$\begin{aligned} q_k^{(B)*} &= 0,\ q_k^{(F)*} = q_k^{\mathrm{DA}*} = 180.05,\ k \in \mathcal{N}^{\mathrm{S}}, \\ Q_{m,n}^{(F)*} &= \begin{cases} 7.59 \text{ for } m \in \mathcal{N}^{\mathrm{S}+1}, n \in \mathcal{N}^{\mathrm{S}} : m = n+3, \\ 0 \text{ for } m \in \mathcal{N}^{\mathrm{S}+1}, n \in \mathcal{N}^{\mathrm{S}} : m \neq n+3, \end{cases} \\ Q^{(B)*} &= -Q^{(F)*}. \end{aligned} \tag{8.16}$$

Figure 8.1 illustrates how the battery and the freezer adjust their power references in response to SFR activation. The first plot shows a real activation signal w which was downsampled from 1 s to 10 s resolution over the course of 1 day. The signal is biased toward positive values, in particular during the first half of the planning horizon. The second plot depicts the power reference $p_{\mathrm{ref}}^{(B)}$ of the battery for the

case where no SFR is activated (horizontal dashed grey line at 0 kW), and for the particular activation shown in the first plot (solid grey line). The solid black line denotes the target power level (8.2). The power trajectories of the freezer are provided in the third plot. Because $\gamma^{(F)*} = 0$, the power reference and power target trajectories of the freezer coincide (solid black line). Note that $p_{\text{ref}}^{(F)} = -p_{\text{ref}}^{(B)}$ as required by the balance constraints (8.15c) and (8.15d). The evolution of the battery and freezer states $x^{(B)}$ and $x^{(F)}$ are shown in the two bottom plots. In the case of no activation, the states remain constant at their initial values (dashed grey lines). However, for the activation given, both the battery and the freezer must adjust their power references (solid grey line for $x^{(B)}$ and solid black line for $x^{(F)}$). In this way, the battery is able to keep its SOC (solid black line) within its feasible range despite a biased activation signal. The energy content of the SFR activation is absorbed by the freezer, whose reference and target state trajectories coincide (solid black line). By acting as a responsive energy buffer, the freezer contributes to providing SFR on the aggregate level even though it cannot offer any SFR on its own.

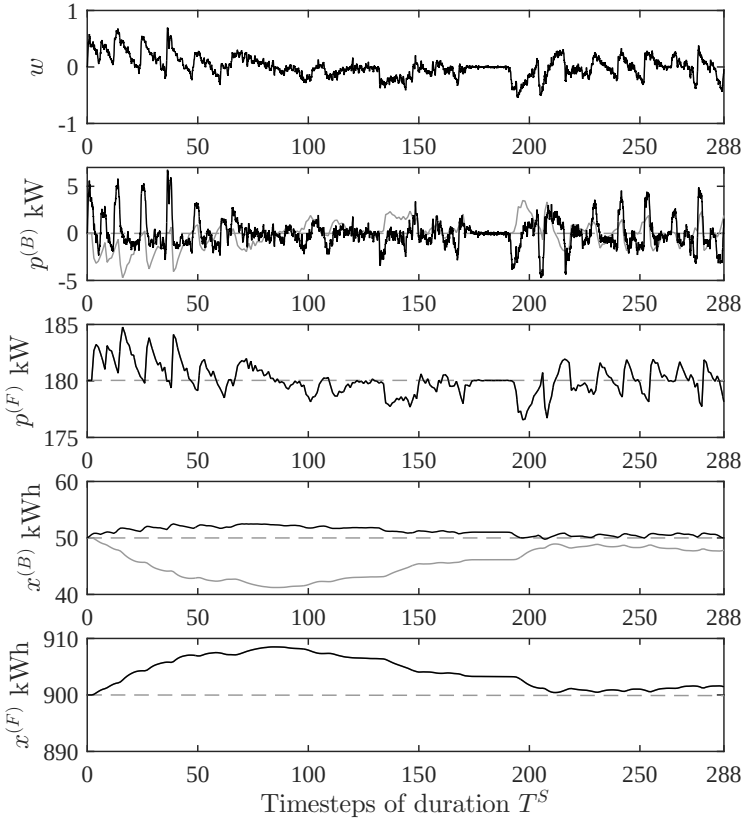

Figure 8.1: Power and state trajectories of the battery and the freezer in response to a given activation w. Trajectories in case of no activation are drawn as dashed grey lines. The reference and target trajectories for the given activation are shown as solid gray and solid black lines, respectively. The reference and target trajectories for the freezer coincide because $\gamma^{(F)*} = 0$.

8.4.4 *Synergy effect of aggregation*

To quantify the benefit of aggregating multiple units, we define the *SFR synergy factor*

$$\sigma^{\mathrm{SFR}} := \gamma_{\max}^{(\mathrm{agg})} \Big/ \sum_{j \in \mathcal{N}^I} \gamma_{\max}^{(j)} - 1 \in \mathbb{R}_0^+. \tag{8.17}$$

The synergy factor measures the amount of *additional* SFR capacity that is available from the aggregation of FEUs relative to the sum of SFR capacities the units could provide individually. A value $\sigma^{\mathrm{SFR}} = 0$ means that aggregating the units does not yield any benefit with regard to SFR capacity, i.e., the SFR capacity available from the aggregation of units equals the sum of the SFR capacities the units can provide individually. In contrast, a value $\sigma^{\mathrm{SFR}} > 0$ indicates that aggregation results in a synergy effect and is able to free up SFR capacity that the units cannot offer individually. The value $100\sigma^{\mathrm{SFR}}$ is the increase in available SFR capacity in %. Figure 8.2 provides the σ^{SFR}-values for the aggregation of a battery and a freezer for different choices of the battery parameters. The circles indicate the battery parameters corresponding to aggregations of 1–5 *Tesla Model-S* batteries [174], 1–2 *Tesla Powerpack* batteries [175], and 2–10 *Tesla Powerwall* batteries [167]. The battery parameters and the corresponding individual and aggregate maximum SFR capacities and synergy factors are summarized in Table 8.2.

Aggregating batteries with a freezer can unlock a significant amount of SFR capacity. Depending on the battery parameters, synergy factors in the range of 3.61-11.90 are achieved. This corresponds to a relative increase in the available SFR capacity by 361–1190% compared with the case where the units provide SFR individually.

The synergy factor obtained from coupling batteries with a freezer remains constant for the same power-to-capacity ratio $\bar{p}^{(B)}/\bar{x}^{(B)}$ of the battery. For example, aggregations of 1–5 *Model-S* batteries and a freezer all result in the same synergy factor $\sigma^{\mathrm{SFR}} = 3.61$, see Table 8.2. The same is true for aggregations of 2–10 *Powerwall* batteries, which all reach synergy factors $\sigma^{\mathrm{SFR}} = 11.90$. In these cases, the SFR capacity available on the aggregate level is limited by the battery's power constraints. The capacity constraints are not active because the battery can balance its energy level via the freezer. However, at some point the energy

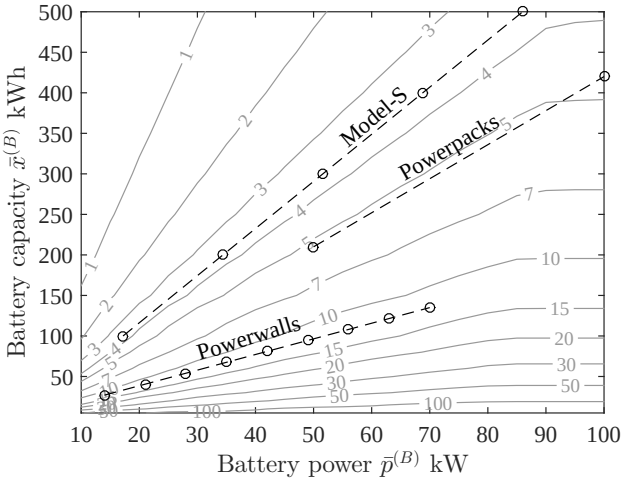

Figure 8.2: Values of the synergy factor σ^{SFR} for different values of $\bar{p}^{(B)}$, $\bar{x}^{(B)}$ for the aggregation of a freezer with batteries.

balancing requirement of the battery exceeds the energy capacity of the freezer. This is when the energy capacity constraints of both the battery and the freezer become active simultaneously and limit the maximum available SFR capacity and the corresponding synergy factor. This effect can be seen using the example of the *Powerpacks*. Coupling a single *Powerpack* battery with the freezer yields $\sigma^{\text{SFR}} = 5.19$, whereas the aggregation of 2 *Powerpacks* with a freezer yields only $\sigma^{\text{SFR}} = 4.65$. At the points where the energy capacity constraints of the freezer become active, the isolines in Figure 8.2 turn horizontal and increasing the number of batteries does not lead to larger synergy factors.

The data shown in Figure 8.2 can be used by AGs to infer what type of battery should be added to an existing aggregation to increase the synergy factor the most. Consider for example the aggregation of a single *Model-S* battery and a freezer. The synergy factor can be increased the most by adding more power capacity, i.e., moving horizontally to the right. In the case of two *Powerpacks*, however, adding more power capacity would not result in a higher synergy factor because the isolines

are horizontal. Here, a larger synergy factor can be reached only by augmenting the energy capacity of the freezer.

Battery type	#	$\bar{p}^{(B)}$ kW	$\bar{x}^{(B)}$ kWh	$\gamma_{\max}^{(\mathrm{B})}$ kW	$\gamma_{\max}^{(\mathcal{J})}$ kW	σ^{SFR}
Model-S	1	17.2	100	2.08	9.61	3.61
	5	86	500	10.42	48.04	3.61
Powerpack	1	50	210	4.38	27.09	5.19
	2	100	420	8.75	49.47	4.65
Powerwall	2	14	27	0.56	7.25	11.90
	10	70	135	2.81	36.26	11.90

Table 8.2: Results of battery-freezer aggregation.

The results summarized in Figure 8.2 and Table 8.2 show that aggregating FEUs with complementary physical properties – a fast-responding but energy-constrained battery and a delayed-responding freezer with a comparatively large energy buffer – can unlock significant amounts of SFR capacity. Depending on the battery parameters, a relative increase in the available SFR capacity by 361–1190% is achieved.

8.5 SYNERGY II: BATTERY AND TURBINE

A synergy effects of aggregation can also be observed when coupling a battery with a unit whose ability to offer SFR is restricted by power ramp-rate limits. Consider a steam turbine with the time-invariant power limits $\underline{p}^{(ST)} = 0$ MW and $\bar{p}^{(ST)} = 250$ MW, similar to the *Siemens SST-3000* turbine [176]. We assume that the turbine can be operated at full power over the entire planning horizon, i.e. the state constraints are omitted. In contrast to the freezer discussed in Section 8.4, the turbine is assumed to be running at an operating point where it can react to power set point changes without delays. However, the rate at which the turbine power output can be varied is limited by $-\underline{r}^{(ST)} = \bar{r}^{(ST)} = 4.5$ MW/min [169]. These ramp-rate limits determine the maximum SFR capacity the turbine can offer. The accurate tracking of the target power

trajectory (8.2) requires that, in the extreme case, the power output of the turbine can be changed by $\pm 2\gamma^{(ST)}$ within one SFR time step $T^{\mathrm{RES}} = 10$ s. Thus, the turbine's power ramp-rate limits must satisfy

$$2\gamma_{\max}^{(ST)} / T^{\mathrm{RES}} \leq \min\{\bar{r}^{(ST)}, -\underline{r}^{(ST)}\}, \tag{8.18}$$

from which follows that the maximum SFR capacity the turbine can provide is $\gamma_{\max}^{(ST)} = 375$ kW, which corresponds to 0.15% of $\bar{p}^{(ST)}$.

We solve the bidding problem (8.15) for an aggregation of a steam turbine and different types of batteries. The control delay constraints (8.15g) are neglected. The resulting aggregate SFR capacities and synergy factors are given in Table 8.3. The results show that the aggregation of a steam turbine and a battery can offer significantly more SFR than the two units could provide individually: the synergy effects amount to 18–125%. The synergy factor grows as the power and capacity of the battery increases. The synergy effects are achieved because the two units complement each other. If needed, the turbine can exchange energy with the battery via adjustments of the reference schedules (8.4) to keep the SOC of the battery feasible. The battery, on the other hand, takes over the high-frequency component of the activation signal and thereby reduces the thermal stress and operational costs of the turbine.

Battery type	#	$\bar{p}^{(B)}$ kW	$\bar{x}^{(B)}$ MWh	$\gamma_{\max}^{(\mathrm{B})}$ kW	$\gamma_{\max}^{(\mathcal{J})}$ kW	σ^{SFR}
Model-S	10	172	1	20.83	468.70	0.18
	50	860	5	104.17	843.50	0.76
	100	1720	10	208.33	1312.00	1.25
Powerpack	5	250	1.05	21.88	506.84	0.28
	10	500	2.1	43.75	638.68	0.53
	20	1000	4.2	87.50	902.35	0.95
Powerwall	50	350	0.675	14.06	551.00	0.42
	100	700	1.35	28.13	726.99	0.80

Table 8.3: Results of battery-steam turbine aggregation.

8.6 CONCLUSION & FUTURE WORK

This chapter considered the energy and reserve bidding problem of an AG trading energy and offering SFR services on behalf of a group of FEUs. The problem was formulated as a robust optimization problem whose solution comprises optimal affine control policies that the AG can apply directly. The problem formulation explicitly takes into account power ramp-rate limits and control delays of individual units, as well as all the different timescales involved in the bidding problem. This makes the approach versatile and applicable to various market settings.

Simulation results demonstrate that aggregating FEUs with complementary physical properties can free up additional SFR capacity in the range of 18-1190% compared with the amount of reserves the units could provide individually. This synergy effect can unlock the regulation potential of different types of FEUs with control delays or restrictive energy or power ramp-rate constraints.

There are several future directions of interest. First, the energy and reserve scheduling problem can be augmented to consider additional types of ancillary services, such as spinning reserves and ramping capacity services. In the bidding problem, the AG must consider all the different timescales of the services and the corresponding markets and ensure the feasibility of all individual target power trajectories even if multiple services are activated simultaneously. Second, the AG is currently assumed to be a price taker. Future work could investigate how to find optimal price-capacity bids and evaluate the synergy effect with regard to the AG's economic profit. The pricing of SFR capacity could be based on flexibility cost functions that reflect the costs of individual units deviating from their reference schedules. Third, the requirement to track SFR activation signals exactly can be relaxed to track the signals only as accurately as required by the SO, for example with regard to regulation mileage constraints or other performance-based measures [162, 177]. Further, the incorporation of distribution network constraints [178, 179] into the bidding problem should be studied.

9 DISCUSSION

In this last chapter, we summarize and consolidate the work presented in this thesis. The chapter is divided into the four main topics covered by this thesis. The first deals with the description and estimation of electric flexibility, the second considers the aggregation and disaggregation of flexibility, the third is concerned with demonstrating the use of flexibility for DR purposes, and the fourth focuses on how flexibility can be exploited to offer SFR services to the power grid. For each topic, we discuss the lessons learned and make suggestions for future work.

9.1 DESCRIPTION AND ESTIMATION OF FLEXIBILITY

In Chapter 2, a generic and quantitative description of the electric flexibility of an individual energy units was proposed based on the unit's feasible set. Relying on this concept of flexibility, we derived a flexibility estimation procedure in Chapter 3 that can be used to identify the parameters of a simple thermal model which defines the flexibility of a thermostatically controlled FEU. The procedure was applied in practice to estimate the flexibility of more than 200 residential buildings with HP installations from measurement data.

9.1.1 *Lessons learned*

1. *Feasible sets are useful.* The description of the flexibility based on feasible sets has proven useful because it allows us to characterize the flexibility of an FEU in a way that is generic, quantitative, and application-independent. Feasible sets can serve as a common approach to quantify, compare and trade the flexibility of various types of energy units. Flexibility cost functions can harness this set-based description by assigning a cost to each power trajectory comprised in the feasible set. In addition, feasible sets can serve as the starting point for the aggregation and disaggregation of flexibility.

2. *Polytopic feasible sets are convenient.* In particular, convex polytopic feasible sets have proven convenient. First, they arise naturally when describing the power, power ramp-rate, and energy constraints of FEUs that can continuously adjust their electricity generation and/or consumption in volume and/or time. Second, the polytopic feasible sets of most FEUs are determined by a few parameters only, such as the time-invariant rated power and energy capacity bounds of PE-system. These properties of polytopic feasible sets simplify their identification and allow for a concise description. Third, polytopic sets are well suited for control and optimization purposes because they can be represented by a set of linear inequalities that are handled by all standard control and optimization frameworks.

3. *Polytopic feasible sets are not perfect.* Although the flexibility of many types of energy units can be characterized or at least approximated by convex polytopic feasible sets, there exist classes of FEUs for which polytopic sets are inappropriate. Examples are energy units with run-time constraints, the modeling of which requires auxiliary binary variables, units whose dynamics cannot be appropriately captured by their linearized counterparts, or units with a fixed power profile that can be shifted only in time but not in volume.

4. *Autonomous estimation of flexibility is needed and feasible.* Because the flexibility available from DEUs in general is small compared with that of large-scale power plants, large numbers of DEUs are required to accumulate a sufficient amount of aggregate flexibility to have an impact on the grid level. However, exploiting the flexibility of numerous units requires that the flexibility of all these DEUs is known. Because this is rarely the case in practice, flexibility estimation procedures are needed that can identify the relevant flexibility parameters autonomously based on measurement data. We showed that it is possible to autonomously identify a simple thermal model of a residential building and its HP based on limited measurement data. The model identified can be used to estimate the flexibility of an individual building and its HP. Our system identification and flexibility estimation procedure can be applied to any other type of TCL as long as the energy dynamics of the system can be described as a single energy buffer with time-invariant capacity and power rating.

9.1.2 *Suggestions for future work*

1. *Improving the estimation of flexibility.* The flexibility estimation procedure offers room for improvements in several directions. First, the extraction of the HP behavior from aggregate energy consumption data can be made more accurate and robust. The HP state detection accuracy could be improved by taking into account additional operational states of the HP, such as the boost and defrost mode. The robustness of the approach could be increased further by learning the cyclic operation pattern of an HP and using this knowledge to distinguish between the power consumption of the HP and that of other electric devices with a similar power rating. Second, the process of fitting the charge- and loss-rate functions to the measurement data should be investigated in more detail. More advanced methods for outlier detection could make the fits more robust. Third, solar irradiation is known to have a significant effect on the heating need, and, thus, the flexibility of a building. The thermal model could be augmented to include solar irradiation as an additional input.

2. *Learning the flexibility factor.* We introduced the flexibility factor as a means for building owners to express their willingness to accept indoor air temperatures that are lower compared with undisturbed HP operation and make a larger amount of flexibility available. Future work could study the acceptance rate of the flexibility factor concept, and look into how unit owners can be incentivized to choose larger flexibility factors. Another approach could be to try to learn the maximum accepted flexibility factor of a building by increasing it iteratively until certain system constraint become active or the building inhabitants intervene.

3. *Considering uncertain flexibility.* The accuracy of the flexibility estimated for individual energy units can be corrupted by modeling and parameter identification errors. In addition, FEUs in general are affected by disturbances, such as uncertain weather conditions and user behavior. These sources of uncertainty should be considered in the description of flexibility. For example, the flexibility of an uncertain FEUs could be described by a collection of nested feasible sets each corresponding to a certain confidence level.

9.2 DEMONSTRATION OF DEMAND RESPONSE

Based on the flexibility estimation discussed in Chapter 3, we demonstrated in Chapter 4 that the flexibility of a population of residential buildings with HP installations can be *i)* described by means of the polytopic feasible sets introduced in Chapter 2, *ii)* identified autonomously by the flexibility estimation procedure presented in Chapter 3, and *iii)* aggregated to offer DR services reliably.

9.2.1 *Lessons learned*

1. *Feasible sets prove themselves in practice.* We showed that polytopic feasible sets can capture a significant portion of the flexibility of residential buildings with HP installations. We modeled an individual building and its HP as a single energy buffer with time-invariant capacity and power rating. The corresponding feasible set can be constructed from the power and state constraints only. The parameters defining the feasible sets were identified by our flexibility estimation procedure based on outdoor air-temperature measurements and SM data only.

2. *Residential buildings with HPs can offer significant DR.* We demonstrated that the flexibility of populations of residential buildings with HP installations can be exploited to offer load-reduction services to the grid in a reliably way. Based on a minimalist and inexpensive sensor and actuator setup, we were able to accurately predict the average load reduction of a select group of HPs. The load reductions achieved amount to 40–60% of the aggregate load of the group of controlled buildings. The load-reduction values were predicted with an APE of less than 7% for 50%, and less than 13% for 75% of all experiments conducted.

3. *"Dirty reality".* The estimation of flexibility and the demonstration of DR were complicated by various problems and difficulties that occur in reality but not in theory. First, the extraction of the HP power consumption from aggregate SM data can be challenging because *i)* the measurement data can be incomplete, *ii)* the characteristic HP consumption pattern can be corrupted by the consumption of other electric devices, such as fridges, vacuum cleaners, and baking ovens, and *iii)* the HP consumption can be shadowed completely by the electricity

generation of solar PV panels or wind turbines installed on-premise. In addition, our flexibility estimation procedure tries to identify the relationship between the duty cycle of an HP and the outdoor air temperature. However, this relationship can be blurred by unobservable user actions, such as opening doors and windows or using a woodstove for heating. Second, our communication and control infrastructure proved unreliable: several HPs did not react to our throttling requests, some HAGs went offline repeatedly, and some SMs would stop reporting measurement data. We learned that setting up and maintaining a reliable and robust DR scheme in a real-life setting is far from trivial, even if it only relies on a limited set of hardware devices, such as an HAG, an RE, and an SM. DR setups that involve more sophisticated measuring equipment, such as indoor air temperature sensors and clamp readers, make the situation even more challenging, although such additional measurement data would be highly valuable for system modeling and identification purposes.

4. *A stochastic model of aggregate flexibility.* We learned that the description of the flexibility of populations of FEUs should incorporate uncertainty. To take into account system modeling and identification errors as well as the unreliability of our communication and control setup mentioned above, we applied a stochastic approach to characterize the aggregate flexibility available from the population of buildings and HPs. For example, we interpreted the duty cycle of an HP as the probability of the HP being ON, and introduced an empirical success rate when predicting load-reduction and rebound values.

9.2.2 *Suggestions for future work*

1. *Improving the communication and control setup.* We strongly recommend that future DR schemes put emphasis on creating a robust and reliable communication and control setup. For example, the digital output of a modern SM could be used to control an external RE or the HP directly. In addition, it would be desirable to have the SM measurements available as close to real time as possible because this would allow us to apply state-estimation methods to infer the current operational state of an HP.

2. *Understanding the rebound.* Rebound effects are characteristic for TCLs. While we succeeded in accurately predicting the load reduction from a population of HP, forecasting the corresponding rebound proved challenging. Future work should investigate the rebound effect of TCLs in more detail, in particular with regard to its duration, peak-power value, and energy content. In addition, our rebound avoidance strategies can be improved further. Similar sequential HP throttling and release strategies could be developed to reduce the peak power ramp-rates of the aggregate load occurring at the beginning and at the end of a load-reduction interval.

3. *Validation of flexibility.* Our estimation of the flexibility of a building and its HP depends on the dynamic model used, the parameters identified, the flexibility factor chosen, and the forecasts of the outdoor air temperature. We assessed the accuracy of this description of flexibility with regard to offering load-reduction services by a group of buildings. Unfortunately, no indoor air-temperature measurements were available in the *Ecogrid2* project that would have allowed us to validate if the description of flexibility is valid also on the level of individual units, i.e., whether or not the bounds on the indoor air temperature of a building are respected. By installing multiple air temperature sensors and HP clamp readers in select buildings, future work should validate our description of flexibility and could yield further insights into how it could be improved further.

4. *No short cycling.* Our HP control algorithm should be extended with the aim of avoiding short cycling of individual units. For this purpose, we need to keep track of past activations of the flexibility of each individual unit and identify the minimum recovery period that is required until the same unit can provide flexibility again.

5. *Coordination of flexibility.* In our current estimation of the aggregate flexibility of buildings and HPs we assume that the throttling start times of all HPs coincide. However, more flexibility can be made available on the aggregate level if this assumption is dropped and the flexibilities of individual units can be concatenated over time. For example, when computing the maximum aggregate load reduction that can be maintained over a given throttling duration, we currently only consider those

buildings whose HPs can be throttled over the entire duration. This can be overly restrictive because the throttling intervals of individual HPs could be concatenated such that in total they cover the entire throttling interval. This motivates the development of more advanced flexibility aggregation and coordination methods. Future work could study the problem of how to optimally schedule the use of individual flexibilities such as to achieve a certain amount of load-reduction capacity that can be maintained over a given time duration.

9.3 AGGREGATION AND DISAGGREGATION OF FLEXIBILITY

The aggregation of polytopic feasible sets is computationally intractable in general. To alleviate this difficulty, we proposed in Chapter 5 to use zonotopes as approximations of the original polytopic feasible sets. We formulated optimization problems to find optimal zonotopic approximations and evaluated their approximation quality. We showed that the zonotopic flexibility can be aggregated efficiently and that it is well suited for making market and control decisions on the aggregate level. In Chapter 6 we presented a disaggregation algorithm that distributes aggregate-level decisions among the individual FEUs in an economically optimal way.

9.3.1 *Lessons learned*

1. *Aggregation of flexibility is not trivial.* To reduce the complexity of managing populations of FEUs, aggregating their flexibility is desirable. However, a major drawback of describing flexibility by means of their polytopic feasible sets is the fact that computing their aggregation is computationally intractable in general. To alleviate this problem, we proposed to inner-approximate polytopic feasible sets by zonotopes, which can be aggregated efficiently. Unfortunately, finding zonotopic approximations entails its own difficulties. First, the set of zonotope generators has to be chosen to make the approximation problem well defined. Because the generators determine the possible shapes of a zonotope, their choice significantly affects the approximation quality. For feasible sets given by PE-polytopes, we proposed a set of generators

that can reconstruct all possible polytope facets. Second, we formulated the zonotope approximation problem as an optimization problem. Consequently, the extension of the optimal zonotope along its generators as well as its approximation quality depend on the objective function chosen. Third, it is not obvious which objective functions and quality measures should be used to compute zonotopes and assess their approximation accuracy.

2. *Zonotopic approximations are a good trade-off between accuracy and complexity.* Our simulations show that zonotopic feasible sets can capture a significantly larger portion of the flexibility of an FEU compared with axis-aligned boxes. Zonotopes yield a good trade-off between an accurate but computationally complex description in terms of polytopes and an inaccurate but computationally simple description in terms of axis-aligned boxes. The approximation quality of zonotopes is determined by several factors, such as the choice of zonotope generators, the symmetry properties of the polytopic set that is to be approximated, as well as the objective function used to compute the optimal zonotope. This can make it difficult to know in advance if zonotopes will yield good approximations for a particular FEU.

3. *Explicit aggregations might not always be necessary.* Although an explicit description of the aggregate feasible set or an approximation thereof can be desirable and convenient, it might not be necessary for particular use-cases of flexibility. For example, several top-down approaches have been proposed to model the DR potential of populations of DEUs based on their statistical properties and without explicitly identifying and aggregating their flexibility. Similarly, the control framework we used in Chapter 7 and Chapter 8 to offer SFR services does not require an explicit characterization of the aggregate feasible set.

4. *Efficient economic disaggregation.* The particular properties of zonotopes can also be used for the purpose of disaggregating flexibility. We showed that flexibility cost functions defined on zonotopes can be aggregated efficiently. We developed a disaggregation algorithm that distributes aggregate-level control decisions in a computationally efficient and economically optimal way based on the aggregate flexibility cost function. The option to assign costs to the flexibility of individual

FEUs paves the way for new market-based flexibility aggregation and disaggregation schemes in which individual FEUs place flexibility bids and are reimbursed for the flexibility reserved and delivered.

9.3.2 *Suggestions for future work*

1. *Approximation quality of zonotopes.* The approximation quality of zonotopic feasible sets should be studied in more detail for different types of FEUs in real-life settings. For instance, initial results suggest that the approximation quality of zonotopes depends, among other things, on how strongly the original polytopic feasible set is dominated by state (energy) constraints. Understanding this dependency in more detail could yield insights about which classes of FEUs can be approximated appropriately by zonotopes.

2. *Generalization of zonotope generators.* We proposed a set of zonotope generators that are able to reproduce all possible facets of a PE-polytope. A more general set of generators could be found which is able to reproduce all possible facets of a polytopic feasible set that is created by the power, power ramp-rate, and state constraints of an arbitrary single-state linear time-invariant system. This would extend the applicability of zonotopic approximations to a wider family of FEUs.

3. *Cost and reimbursement of flexibility.* We presented a disaggregation algorithm that distributes an aggregate-level control signal among the energy units of a population in an economically optimal way based on flexibility cost functions defined on individual feasible sets. Future work could study more complex disaggregation objectives. For example, an AG might decide to reimburse an individual FEU based on how much its flexibility contributed to an aggregate-level flexibility product, such as SFR. Another option could be for an AG to open a market to which the individual FEUs can submit flexibility bids in terms of flexibility cost functions defined over their feasible sets. Subsequently, a market-based process reserves and activates the flexibility of individual units and the FEUs are reimbursed for the flexibility reserved and delivered. Such market-based aggregation and disaggregation schemes could foster competition between the individual FEUs, and make more flexibility available to AGs and SOs at a lower price.

9.4 USING FLEXIBILITY FOR REGULATION SERVICES

In the last part of this thesis we investigated how flexibility can be used to provide SFR services to the power grid. Chapter 7 adopts a control approach that is based on affinely adjustable control policies to formulate the energy and reserve bidding problem as a robust optimization problem. We focus on properly taking into account different timescales, lead times, and power ramp-rate constraints and show that trading energy on intra-day markets is crucial for energy-constrained FEUs to efficiently exploit their flexibility. In addition, it was demonstrated in Chapter 8 that aggregating the flexibility of a heterogeneous group of FEUs can lead to a synergy effect with regard to the aggregate SFR capacity available.

9.4.1 *Lessons learned*

1. *Power ramp-rates are important.* The incorporation of power ramp-rate constraints into the description of flexibility is of particular importance for two main reasons. First, higher shares of VRES lead to larger, more frequent, and more rapid variations of the electricity supply. Thus, there is a growing need for fast-responding and accurate regulation power. To reflect these changing needs, SOs have recently started to adjust their reserve requirements and reimbursement schemes to reward the timely and accurate provision of SFR. This requires that power ramp-rate constraints of individual FEUs are considered in the description and the control of flexibility. Second, power ramp-rate constraints and control delays are distinguishing characteristics of energy units. Their consideration is a prerequisite for AGs to build up portfolios of FEUs whose characteristics complement each other and can give rise to synergy effects.

2. *Flexibility benefits from diversity.* We showed that aggregating the flexibility of a group of heterogeneous FEUs and controlling them in a coordinated way can make a significantly larger quantity of flexibility available for SFR services compared with the case where the units are considered individually. The synergy effect of aggregation can be observed in situations where the diverse characteristics of individual

units complement each other. This is the case for example when energy-constrained FEUs are allowed to balance their energy levels by trading on intra-day energy markets, or when traditional large-scale and slowly-responding suppliers of flexibility are supported by groups of small-scale fast-responding FEUs.

3. *New flexibility products are desirable.* Current SFR products are tailored to the capabilities of traditional providers of regulation services, such as gas turbines and pumped hydro-power plants. They usually require that a constant amount of regulation power is kept available over several hours and they are procured up to several days before the time of delivery. Such requirements dramatically reduce the amount of SFR services that small-scale FEUs can offer. First, predicting the flexibility available over long planning horizons is difficult because the behavior of the units can be subject to uncertainty. Second, for most types of FEUs the energy constraints will limit the amount of SFRs capacity that can be offered. These limitations can be relaxed by coupling energy-constrained FEUs with energy markets to allow them to rebalance their energy content via energy trades if necessary.

9.4.2 *Suggestions for future work*

1. *Considering ramping costs.* SFR activation signals vary on the timescale of seconds. Thus, tracking them accurately can require significant ramping, which for units like thermal and pumped hydro-power plants, can be costly and entail high thermal and mechanical stress. For other types of FEUs, such as capacitors and flywheels, ramping might be less of a challenge. Future work could include the cost of ramping in the objective of the energy and reserve bidding problem. The resulting affine controllers would implicitly consider these ramping costs and split the SFR signal into components with different frequency contents.

2. *Studying synergy effects.* The synergy effects that can occur when the flexibility of multiple FEUs is aggregated should be studied in more detail, in particular the conditions under which synergy arises. We have shown that aggregating the flexibility of heterogeneous FEUs can lead to a synergy effect with regard to SFR capacity. However, other types of synergy effects exist. For example, in [47] we study two types of

synergies, one with regard to the volume, and the other with regard to the position of an aggregate feasible axis-aligned box.

3. *Developing new flexibility products.* As mentioned above, we believe that current flexibility products, in particular SFR products, are inappropriate for small-scale FEUs and the challenges that come with integrating VRES. We suggest to develop new flexibility products whose specifications are chosen in accordance with the stochastic properties of the uncertainty introduced by VRES and that are fitted to the needs and capabilities of FEUs to efficiently exploit their flexibility. Several steps have been made in this direction. For example, SFR products with energy or activation frequency limits have been proposed. In addition, new SFR products should be defined over shorter time intervals (in the range of minutes or a few hours) and should be traded closer to the time of delivery to facilitate the prediction of (uncertain) flexibility.

4. *Creating new flexibility markets.* New flexibility products require new flexibility markets that are open to transmission and distribution SOs, BRPs, and AGs. These flexibility markets bring together suppliers of flexibility, such as individual or groups of FEUs, and potential buyers of flexibility, such as SOs or BRPs, that wish to hedge against uncertain supply or demand of electricity or other contingencies. For example, new flexibility markets could run on sub-hourly timescales with short lead times, enable BRPs to trade flexibility bilaterally in order to avoid imbalance penalties, or allow distribution SOs to purchase dedicated flexibility products that are delivered in certain topological areas of the distribution grid.

BIBLIOGRAPHY

[1] *Kyoto Protocol to the United Nations framework convention on climate change*. Tech. rep. United Nations, 1998, pp. 1–21. URL: http://unfccc.int/resource/docs/convkp/kpeng.pdf.

[2] *Paris Agreement*. Tech. rep. United Nations, 2015, pp. 1–27. URL: http://unfccc.int/resource/docs/2015/cop21/eng/l09r01.pdf.

[3] *Energy roadmap 2050*. Tech. rep. European Commission, 2011, pp. 1–24. URL: https://ec.europa.eu/energy/sites/ener/files/documents/2012{_}energy{_}roadmap{_}2050{_}en{_}0.pdf.

[4] *A roadmap for moving to a competitive low carbon economy in 2050*. Tech. rep. European Commission, 2011, pp. 1–34. URL: http://ec.europa.eu/clima/documentation/roadmap/docs/com{_}2011{_}112{_}en.pdf.

[5] *Renewable energy policies in a time of transition*. IRENA, IEA, REN21, 2018, pp. 1–112. URL: http://www.irena.org/-/media/Files/IRENA/Agency/Publication/2018/Apr/IRENA{_}IEA{_}REN21{_}Policies{_}2018.pdf.

[6] *Energy efficiency directive*. Tech. rep. European Union, 2012, pp. 1–56. URL: http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32009L0028{\&}from=EN.

[7] *Renewable energy directive*. Tech. rep. European Union, 2009, pp. 1–47. URL: http://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32009L0028{\&}from=EN.

[8] T. B. Johansson, K. McCormick, L. Neij, and W. Turkenburg. "The potentials of renewable energy." In: *Int. Conf. Renew. Energies*. REN21, 2004, pp. 1–40.

[9] *Renewables global futures report*. Tech. rep. REN21, 2017, pp. 1–98. URL: http://www.ren21.net/wp-content/uploads/2017/10/GFR-Full-Report-2017{_}webversion{_}3.pdf.

[10] S. Koch. "Demand response methods for ancillary services and renewable energy integration in electric power systems." PhD Thesis. ETH Zurich, 2012, pp. 1–299.

[11] J. H. R. Enslin. "Network impacts of high penetration of photovoltaic solar power systems." In: *IEEE PES Gen. Meet.* IEEE, 2010, pp. 1–5.

[12] C. Ziems, S. Meinke, H. Weber, E. Hassel, and J. Nocke. "Influence of intermittent wind and photovoltaic power production onto thermal power plants in Germany until 2020." In: *17th Power Syst. Comput. Conf.* PSCC, 2011, pp. 1–9.

[13] *What the duck curve tells us about managing a green grid*. Tech. rep. CAISO, 2016, pp. 1–4. URL: https://www.caiso.com/Documents/FlexibleResourcesHelpRenewables{_}FastFacts.pdf.

[14] A. Gilbert. *Solar to wreck economics of existing power markets*. 2016. URL: https://www.sparklibrary.com/solar-wreck-economics-existing-power-markets/ (visited on 05/16/2018).

[15] *Levelized cost of energy analysis*. Tech. rep. Lazard, 2016, pp. 1–21. URL: https://www.lazard.com/media/438038/levelized-cost-of-energy-v100.pdf.

[16] U.S. Energy Information Administration. *Electric generator dispatch depends on system demand and the relative cost of operation*. 2012. URL: https://www.eia.gov/todayinenergy/detail.php?id=7590 (visited on 05/16/2018).

[17] A. Ulbig, T. S. Borsche, and G. Andersson. "Impact of low rotational inertia on power system stability and operation." In: *IFAC Proc. Vol.* Vol. 47. 3. IFAC, 2014, pp. 7290–7297.

[18] *The power of transformation - Wind, sun and the economics of flexible power systems*. Tech. rep. IEA, 2014, pp. 1–238. URL: https://webstore.iea.org/the-power-of-transformation.

[19] E. Lannoye, D. Flynn, and M. O'Malley. "Evaluation of power system flexibility." In: *IEEE Trans. Power Syst.* 27.2 (2012), pp. 922–931.

[20] P. D. Lund, J. Lindgren, J. Mikkola, and J. Salpakari. "Review of energy system flexibility measures to enable high levels of variable renewable electricity." In: *Renew. Sustain. Energy Rev.* 45 (2015), pp. 785–807.

[21] F. Bouffard and M. Ortega-Vazquez. "The value of operational flexibility in power systems with significant wind power generation." In: *IEEE PES Gen. Meet.* IEEE, 2011, pp. 1–5.

[22] H. Holttinen, P. Meibom, A. Orths, B. Lange, M. O'Malley, J. O. Tande, A. Estanqueiro, E. Gomez, L. Söder, G. Strbac, J. C. Smith, and F. van Hulle. "Impacts of large amounts of wind power on design and operation of power systems, results of IEA collaboration." In: *Wind Energy* 14.2 (2011), pp. 179–192.

[23] J. Ma, V. Silva, R. Belhomme, D. S. Kirschen, and L. F. Ochoa. "Evaluating and planning flexibility in sustainable power systems." In: *IEEE Trans. Sustain. Energy* 4.1 (2013), pp. 200–209.

[24] Y. Dvorkin, D. S. Kirschen, and M. A. Ortega-Vazquez. "Assessing flexibility requirements in power systems." In: *IET Gener. Transm. Distrib.* (2014), pp. 1–11.

[25] Y. V. Makarov, C. Loutan, J. Ma, and P. D. Mello. "Operational impacts of wind generation on California power systems." In: *IEEE Trans. Power Syst.* 24.2 (2009), pp. 1039–1050.

[26] M. Khederzadeh. "Defining and realizing flexibility in distribution grid." In: *CIRED Work.* IET, 2016, pp. 1–4.

[27] T. Noll, M. Greve, C. Rehtanz, S. Dierkes, A. Moser, P. Schäfer, and H. Vennegeerts. "Flexibility of the distribution grid for system stability." In: *19th Power Syst. Comput. Conf.* IEEE, 2016, pp. 1–7.

[28] J. Cochran, M. Miller, O. Zinaman, M. Milligan, D. Arent, B. Palmintier, M. O. Malley, S. Mueller, E. Lannoye, A. T. Epri, B. Kujala, N. Power, M. Sommer, H. Holttinen, J. Kiviluoma, and S. Soonee. *Flexibility in 21st century power systems*. Tech. rep. NREL, 2014, pp. 1–14. URL: https://www.nrel.gov/docs/fy14osti/61721.pdf.

[29] *Empowering variable renewables*. Tech. rep. IEA, 2008, pp. 1–35. URL: https://www.iea.org/publications/freepublications/publication/Empowering{_}Variable{_}Renewables.pdf.

[30] G. T. Costanzo, J. Kheir, and G. Zhu. "Peak-load shaving in smart homes via online scheduling." In: *IEEE Int. Symp. Ind. Electron.* IEEE, 2011, pp. 1347–1352.

[31] N. Chen, C. W. Tan, and T. Q. Quek. "Electric vehicle charging in smart grid: Optimality and valley-filling algorithms." In: *IEEE J. Sel. Top. Signal Process.* 8.6 (2014), pp. 1073–1083.

[32] L. Fabietti, T. T. Gorecki, F. A. Qureshi, A. Bitlislioglu, I. Lymperopoulos, and C. N. Jones. "Experimental implementation of frequency regulation services using commercial buildings." In: *IEEE Trans. Smart Grid* (2016), pp. 1–9.

[33] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent. "Frequency regulation from flexible loads: Potential, economics, and implementation." In: *Americ. Control Conf.* IEEE, 2014, pp. 65–72.

[34] E. Vrettos, F. Oldewurtel, and G. Andersson. "Robust energy-constrained frequency reserves from aggregations of commercial buildings." In: *IEEE Trans. Power Syst.* 31.6 (2016), pp. 4272–4285.

[35] D. S. Callaway and I. A. Hiskens. "Achieving controllability of electric loads." In: *Proc. IEEE* 99.1 (2011), pp. 184–199.

[36] F. L. Müller and B. Jansen. "Large-scale demonstration of precise demand response provided by residential heat pumps." In: *Submitt. to Energy* (2018), pp. 1–8. URL: http://arxiv.org/abs/1806.07670.

[37] *2010 Smart Grid system report*. Tech. rep. U.S. Department of Energy, 2012, pp. 1–120. URL: https://www.energy.gov/sites/prod/files/2010SmartGridSystemReport.pdf.

[38] C. Eid, P. Codani, Y. Chen, Y. Perez, and R. Hakvoort. "Aggregation of demand side flexibility in a SmartGrid: A review for European market design." In: *12th Int. Conf. Eur. Energy Mark.* IEEE, 2015, pp. 1–5.

[39] L. Gkatzikis, I. Koutsopoulos, and T. Salonidis. "The role of aggregators in smart grid demand response markets." In: *IEEE J. Sel. Areas Commun.* 31.7 (2013), pp. 1247–1257.

[40] K. Trangbaek and J. Bendtsen. "Exact constraint aggregation with applications to smart grids and resource distribution." In: *51st IEEE Conf. Decis. Control*. IEEE, 2012, pp. 4181–4186.

[41] A. Nayyar, J. Taylor, A. Subramanian, K. Poolla, and P. Varaiya. "Aggregate flexibility of a collection of loads." In: *52nd IEEE Conf. Decis. Control*. IEEE, 2013, pp. 5600–5607.

[42] P. Koponen, J. Ikaheimo, A. Vicino, A. Agnetis, G. De Pascale, N. Ruiz Carames, J. Jimeno, E. F. Sanchez-Ubeda, P. Garcia-Gonzalez, and R. Cossent. "Toolbox for aggregator of flexible demand." In: *IEEE Int. Energy Conf.* IEEE, 2012, pp. 623–628.

[43] S. Barot and J. A. Taylor. "A concise, approximate representation of a collection of loads described by polytopes." In: *Int. J. Elect. Power Energy Syst.* 84.1 (2016), pp. 55–63.

[44] D. Pudjianto, D. Pudjianto, C. Ramsay, C. Ramsay, G. Strbac, and G. Strbac. "Virtual power plant and system integration of distributed energy resources." In: *Renew. Power Gener.* 1.1 (2007), pp. 10–16.

[45] M. Balandat, F. Oldewurtel, M. Chen, and C. Tomlin. "Contract design for frequency regulation by aggregations of commercial buildings." In: *52nd Allert. Conf. Comm., Control. Comput.* IEEE, 2014, pp. 38–45.

[46] F. L. Müller, B. Jansen, and O. Sundström. "Autonomous estimation of the energetic flexibility of buildings." In: *Amer. Control Conf.* IEEE, 2017, pp. 2713–2718.

[47] F. L. Müller, O. Sundström, J. Szabo, and J. Lygeros. "Aggregation of energetic flexibility using zonotopes." In: *54th IEEE Conf. Decis. Control*. IEEE, 2015, pp. 6564–6569.

[48] F. L. Müller, O. Sundström, J. Szabo, and J. Lygeros. "Aggregation and disaggregation of energetic flexibility from distributed energy resources." In: *IEEE Trans. Smart Grid* (2017), pp. 1–10. arXiv: arXiv:1705.02815.

[49] F. L. Müller, S. Woerner, and J. Lygeros. "Ramp-rate-constrained bidding of energy and frequency reserves in real market settings." In: (2018), pp. 1–8. URL: https://arxiv.org/abs/1804.03892.

[50] F. L. Müller, S. Woerner, and J. Lygeros. "Unlocking the potential of flexible energy resources to help balance the power grid." In: *Submitt. to IEEE Trans. Smart Grids* (2018), pp. 1–8. URL: http://arxiv.org/abs/1806.08237.

[51] A. Caracas, F. L. Müller, O. Sundström, C. Binding, and B. Jansen. "VillaSmart : Wireless sensors for system identification in domestic buildings." In: *IEEE PES Innov. Smart Grid Technol. Eur.* IEEE, 2013, pp. 1–5.

[52] F. L. Müller, O. Sundström, C. Binding, and D. Dykeman. "Power reference tracking of a large-scale industrial freezer system for ancillary service delivery." In: *IEEE PES Innov. Smart Grid Technol. Eur.* IEEE, 2013, pp. 1–5.

[53] F. L. Müller, C. Binding, O. Sundström, and M. Bengsch. "Minimum-cost charging of electrical storage heaters." In: *IEEE Int. Conf. Smart Grid Comm.* IEEE, 2014, pp. 740–745.

[54] J. Zhao, T. Zheng, and E. Litvinov. "A unified framework for defining and measuring flexibility in power system." In: *IEEE Trans. Power Syst.* 31.1 (2016), pp. 339–347.

[55] A. Ulbig and G. Andersson. "Analyzing operational flexibility of electric power systems." In: *Power Syst. Comput. Conf.* IEEE, 2014, pp. 1–8.

[56] A. Ulbig and G. Andersson. "On Operational Flexibility in Power Systems." In: *IEEE Power Energy Soc. Gen. Meet.* IEEE, 2012, pp. 1–8.

[57] M. Huber, D. Dimkova, and T. Hamacher. "Integration of wind and solar power in Europe: Assessment of flexibility requirements." In: *Energy* 69 (2014), pp. 236–246.

[58] M. I. Alizadeh, M. Parsa Moghaddam, N. Amjady, P. Siano, and M. K. Sheikh-El-Eslami. "Flexibility in future power systems with high renewable penetration: A review." In: *Renew. Sust. Energy Rev.* 57 (2016), pp. 1186–1193.

[59] R. De Coninck and L. Helsen. "Bottom-up quantification of the flexibility potential of buildings." In: *13th Conf. Int. Build. Perform. Sim. Assoc.* IBPSA, 2013, pp. 3250–3258.

[60] R. D'hulst, W. Labeeuw, B. Beusen, S. Claessens, G. Deconinck, and K. Vanthournout. "Demand response flexibility and flexibility potential of residential smart appliances: Experiences from large pilot test in Belgium." In: *Appl. Energy* 155 (2015), pp. 79–90.

[61] S. Stinner, K. Huchtemann, and D. Müller. "Quantifying the operational flexibility of building energy systems with thermal energy storages." In: *Appl. Energy* 181 (2016), pp. 140–154.

[62] J. L. Mathieu, M. Kamgarpour, J. Lygeros, and D. S. Callaway. "Energy arbitrage with thermostatically controlled loads." In: *Eur. Control Conf.* IEEE, 2013, pp. 2519–2526.

[63] N. O'Connell, P. Pinson, H. Madsen, and M. O'Malley. "Economic dispatch of demand response balancing through asymmetric block offers." In: *IEEE Trans. Power Syst.* 31.4 (2015), pp. 2999–3007.

[64] D. Fischer, T. Wolf, J. Wapler, and R. Hollinger. "Model-based flexibility assessment of a residential heat pump pool." In: *Energy* 118 (2016), pp. 853–864.

[65] H. Hao and W. Chen. "Characterizing flexibility of an aggregation of deferrable loads." In: *53rd IEEE Conf. Decis. Control*. IEEE, 2014, pp. 4059–4064.

[66] E. Valsomatzis, T. B. Pedersen, and A. Abell. "Aggregating energy flexibilities under constraints." In: *IEEE Int. Conf. Smart Grid Comm.* IEEE, 2016, pp. 484–490.

[67] Q. Wang and B. M. Hodge. "Enhancing power system operational flexibility with flexible ramping products: A review." In: *IEEE Trans. Ind. Informatics* 13.4 (2017), pp. 1652–1664.

[68] K. Heussen, S. Koch, A. Ulbig, and G. Andersson. "Unified system-level modeling of intermittent renewable energy sources and energy storage for power system operation." In: *IEEE Syst. J.* 6.1 (2012), pp. 140–151.

[69] L. Šikšnys, M. E. Khalefa, and T. B. Pedersen. "Aggregating and disaggregating flexibility objects." In: *24th Int. Conf. Sci. Stat. Database Manag.* Springer, 2012, pp. 379–396.

[70] E. Vrettos, F. Oldewurtel, F. Zhu, and G. Andersson. "Robust provision of frequency reserves by office building aggregations." In: *19th World Congr. Int. Fed. Autom. Control*. Vol. 19. IFAC, 2014, pp. 12068–12073.

[71] F. Rey, X. Zhang, S. Merkli, V. Agliati, M. Kamgarpour, and J. Lygeros. "Strengthening the group: Aggregated frequency reserve bidding with ADMM." In: *IEEE Trans. Smart Grid* (2018), pp. 1–10.

[72] B. M. Sanandaji, H. Hao, K. Poolla, and T. L. Vincent. "Improved battery models of an aggregation of thermostatically controlled loads for frequency regulation." In: *Amer. Control Conf.* IEEE, 2014, pp. 38–45.

[73] J. T. Hughes, A. D. Dominguez-Garcia, and K. Poolla. "Identification of virtual battery models for flexible loads." In: *IEEE Trans. Power Syst.* 31.6 (2016), pp. 4660–4669.

[74] H. Hao, B. M. Sanandajib, K. Poolla, and T. L. Vincent. "Aggregate flexibility of thermostatically controlled loads." In: *IEEE Trans. Power Syst.* 30.1 (2014), pp. 1–10.

[75] M. Alizadeh, A. Scaglione, A. Goldsmith, and G. Kesidis. "Capturing aggregate flexibility in demand response." In: *53rd IEEE Conf. Decis. Control*. IEEE, 2014, pp. 6439–6445.

[76] H. Nosair and F. Bouffard. "Flexibility envelopes for power system operational planning." In: *IEEE Trans. Sust. Energy* 6.3 (2015), pp. 800–809.

[77] H. Nosair and F. Bouffard. "Reconstructing operating reserve: Flexibility for sustainable power systems." In: *IEEE Trans. Sust. Energy* 6.4 (2015), pp. 1624–1637.

[78] M. A. Bucher, S. Delikaraoglou, K. Heussen, P. Pinson, and G. Andersson. "On quantification of flexibility in power systems." In: *IEEE Eindhoven PowerTech*. IEEE, 2015, pp. 1–6.

[79] M. Gonzalez Vaya, L. Baringo, and G. Andersson. "Integration of PEVs into power markets: A bidding strategy for a fleet aggregator." In: *Plug Electr. Veh. Smart Grids*. Springer, 2015. Chap. 9, pp. 233–260.

[80] L. Zhao, H. Hao, and W. Zhang. "Extracting flexibility of heterogeneous deferrable loads via polytopic projection approximation." In: *55th IEEE Conf. Decis. Control*. IEEE, 2016, pp. 6651–6656.

[81] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent. "A generalized battery model of a collection of thermostatically controlled loads for providing ancillary service." In: *51st Allert. Conf. Commun., Control Comput.* IEEE, 2013, pp. 551–558.

[82] T. T. Gorecki, A. Bitlislioglu, G. Stathopoulos, and C. N. Jones. "Guaranteeing input tracking for constrained systems: Theory and application to demand response." In: *Proc. Amer. Control Conf.* IEEE, 2015, pp. 232–237.

[83] B. Jin, M. Maasoumy, P. Nuzzo, and A. Sangiovanni-Vincentelli. "Online computation of polytopic flexibility models for demand shifting applications." In: *IEEE Conf. Autom. Sci. Eng.* IEEE, 2017, pp. 1–6.

[84] O. Sundström, C. Binding, D. Gantenbein, D. Berner, and W. C. Rumsch. "Aggregating the flexibility provided by domestic hot-water boilers to offer tertiary regulation power in Switzerland." In: *IEEE PES Innov. Smart Grid Technol. Conf. Eur.* IEEE, 2012, pp. 1–7.

[85] K. Trangbaek, M. Petersen, J. Bendtsen, and J. Stoustrup. "Exact power constraints in smart grid control." In: *50th IEEE Conf. Decis. Control Eur. Control Conf.* IEEE, 2011, pp. 6907–6912.

[86] L. Zhao, W. Zhang, H. Hao, and K. Kalsi. "A geometric approach to aggregate flexibility modeling of thermostatically controlled loads." In: *IEEE Trans. Power Syst.* 32.6 (2017), pp. 4721 –4731.

[87] A. Zoha, A. Gluhak, M. Imran, and S. Rajasegarar. "Non-intrusive load monitoring approaches for disaggregated energy sensing: A survey." In: *Sensors* 12.12 (2012), pp. 16838–16866.

[88] D. He, W. Lin, N. Liu, R. G. Harley, and T. G. Habetler. "Incorporating non-intrusive load monitoring into building level demand response." In: *IEEE Trans. Smart Grid* 4.4 (2013), pp. 1870–1877.

[89] D. Kaulakienė, L. Šikšnys, and Y. Pitarch. "Towards the automated extraction of flexibilities from electricity time series." In: *Jt. EDBT/ICDT Work.* ACM, 2013, pp. 267–272.

[90] K. Kouzelis, Z. H. Tan, B. Bak-Jensen, J. R. Pillai, and E. Ritchie. "Estimation of residential heat pump consumption for flexibility market applications." In: *IEEE Trans. Smart Grid* 6.4 (2015), pp. 1852–1864.

[91] E. C. Kara, M. D. Tabone, J. S. MacDonald, D. S. Callaway, and S. Kiliccote. "Quantifying flexibility of residential thermostatically controlled loads for demand response." In: *Proc. 1st ACM Conf. Embed. Syst. Energy-Efficient Build.* ACM, 2014, pp. 140–147.

[92] R. Yin, E. C. Kara, Y. Li, N. DeForest, K. Wang, T. Yong, and M. Stadler. "Quantifying flexibility of commercial and residential loads for demand response using setpoint changes." In: *Appl. Energy* 177 (2016), pp. 149–164.

[93] *EcoGrid2.0*. 2018. URL: http://www.ecogrid.dk/en/home{_}uk (visited on 04/24/2018).

[94] D. S. Callaway. "Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy." In: *Energy Conv. Manag.* 50.5 (2009), pp. 1389–1400.

[95] M. Kamgarpour, C. Ellen, S. E. Z. Soudjani, S. Gerwinn, J. L. Mathieu, N. Muellner, A. Abate, D. S. Callaway, M. Fraenzle, and J. Lygeros. "Modeling options for demand side participation of thermostatically controlled loads." In: *IREP Symp. Bulk Power Syst. Dyn. Control*. IEEE, 2013, pp. 1–15.

[96] J. L. Mathieu, S. Koch, and D. S. Callaway. "State estimation and control of electric loads to manage real-time energy imbalance." In: *IEEE Trans. Autom. Control* 28.1 (2012), pp. 430–440.

[97] H. D. Baehr and S. Kabelac. *Thermodynamik*. 13th ed. Springer, 2006, pp. 1–661.

[98] N. O'Connell, P. Pinson, H. Madsen, and M. Omalley. "Benefits and challenges of electrical demand response: A critical review." In: *Renew. Sust. Energy Rev.* 39 (2014), pp. 686–699.

[99] N. G. Paterakis, O. Erdinç, and J. P. Catalão. "An overview of demand response: Key-elements and international experience." In: *Renew. Sust. Energy Rev.* 69.September 2015 (2017), pp. 871–891.

[100] U. Obinna, P. Joore, L. Wauben, and A. Reinders. "Insights from stakeholders of five residential smart grid pilot projects in the Netherlands." In: *Smart Grid Renew. Energy* 7.1 (2016), pp. 1–15.

[101] P. Lund, P. Nyeng, R. D. Grandal, S. H. Sorensen, M. F. Bendtsen, G. LeRay, E. M. Larsen, J. Mastop, F. Judex, F. Leimgruber, K. Kok, and P. MacDougall. *EcoGrid EU - Overall evaluation and conclusion*. Tech. rep. 2015, pp. 1–88. URL: http://www.eu-ecogrid.net/.

[102] W. Cardinaels and I. Borremans. *LINEAR - Demand response for families*. Tech. rep. LINEAR consortium, 2014, pp. 1–223. URL: http://www.linear-smartgrid.be.

[103] S. De Francisci, R. Attema, O. Franz, V. Sebastian, P. Frias, M. Valles, P. Lewis, T. Schmid, T. Mikkelsen, G. Bodea, S. Goudet, and N. Becuwe. *Consolidated report on the key ADVANCED conclusions*. Tech. rep. ADVANCED consortium, 2014, pp. 1–135. URL: https://www.tno.nl/media/6221/d6-5{_}key-advanced-results.pdf.

[104] S. Widergren, J. Fuller, C. Marinovici, and A. Somani. "Residential transactive control demonstration." In: *IEEE PES Innov. Smart Grid Technol. Conf.* IEEE, 2014, pp. 1–5.

[105] D. P. Chassin, D. J. Hammerstrom, and J. G. De Steese. "The Pacific Northwest demand response market demonstration." In: *IEEE PES Gen. Meet.* IEEE, 2008, pp. 1–6.

[106] J. L. Mathieu, M. Kamgarpour, J. Lygeros, G. Andersson, and D. S. Callaway. "Arbitraging intraday wholesale energy market prices with aggregations of thermostatic loads." In: *IEEE Trans. Power Syst.* 30.2 (2015), pp. 763–772.

[107] V. Lakshmanan, K. Gudmand-Hoyer, M. Marinelli, A. M. Kosek, and P. Norgard. "Energy shift estimation of demand response activation on refrigerators - A field test study." In: *49th Int. Univ. Power Eng. Conf.* IEEE, 2014, pp. 1–6.

[108] E. A. M. Klaassen, J Frunt, and J. G. Slootweg. "Experimental validation of the demand response potential of residential heating systems." In: *Proc. 19th Power Syst. Comput. Conf.* IEEE, 2016, pp. 1–7.

[109] E. Georges, B. Cornélusse, D. Ernst, V. Lemort, and S. Mathieu. "Residential heat pump as flexible load for direct control service with parametrized duration and rebound effect." In: *Appl. Energy* 187 (2017), pp. 140–153.

[110] E. Vrettos and J. L. Mathieu. "Demand response with moving horizon estimation of individual thermostatic load states from aggregate power measurements." In: *Amer. Control Conf.* IEEE, 2014, pp. 4846–4853.

[111] B. S. Palmintier and M. D. Webster. "Heterogeneous unit clustering for efficient operational flexibility modeling." In: *IEEE Trans. Power Syst.* 29.3 (2014), pp. 1089–1098.

[112] K. Hreinsson and A. Scaglione. "On aggregating thermostatically controlled loads based on energy losses." In: *IEEE PES Gen. Meet.* IEEE, 2017, pp. 1–5.

[113] M. Wrinch, G. Dennis, T. H. El-Fouly, and S. Wong. "Demand response implementation for improved system efficiency in remote communities." In: *IEEE Elect. Power Energy Conf.* IEEE, 2012, pp. 105–110.

[114] O. De Somer, A. Soares, T. Kuijpers, K. Vossen, K. Vanthournout, and F. Spiessens. "Using reinforcement learning for demand response of domestic hot water buffers: a real-life demonstration." In: *IEEE PES Innov. Smart Grid Technol. Conf. Eur.* IEEE, 2017, pp. 1–7.

[115] R. Henriquez Auba, G. Wenzel, D. Olivares, and M. Negrete-Pincetic. "Participation of demand response aggregators in electricity markets: Optimal portfolio management." In: *IEEE Trans. Smart Grid* (2017), pp. 1–11.

[116] S. Bahrami, M. H. Amini, M. Shafie-khah, and J. P. Catalao. "A decentralized electricity market scheme enabling demand response deployment." In: *IEEE Trans. Power Syst.* 33.4 (2018), pp. 4218–4227.

[117] S. Bahrami and A. Sheikhi. "From demand response in smart grid toward integrated demand response in smart energy hub." In: *IEEE Trans. Smart Grid* 7.2 (2016), pp. 650–658.

[118] L. Zhang, F. Jabbari, T. Brown, and S. Samuelsen. "Coordinating plug-in electric vehicle charging with electric grid: Valley filling and target load following." In: *J. Power Sources* 267 (2014), pp. 584–597.

[119] C. K. Woo and K Herter. *Residential demand response evaluation: A scoping study*. Tech. rep. DRRC, 2006, pp. 1–24. URL: https://drrc.lbl.gov/publications/residential-demand-response.

[120] J. L. Mathieu, P. N. Price, S. Kiliccote, and M. A. Piette. "Quantifying changes in building electricity use, with application to demand response." In: *IEEE Tran. Smart Grid* 2.3 (2011), pp. 507–518.

[121] N. Motegi, M. A. Piette, D. S. Watson, S. Kiliccote, and P. Xu. *Introduction to commercial building control strategies and techniques for demand response*. Tech. rep. Lawrence Berkeley Nat. Lab., 2007, pp. 1–107. URL: http://gaia.lbl.gov/btech/papers/59975.pdf.

[122] B. Biegel, P. Andersen, J. Stoustrup, L. H. Hansen, and A. Birke. "Sustainable reserve power from demand response and fluctuating production - Two Danish demonstrations." In: *Proc. IEEE* 104.4 (2016), pp. 780–788.

[123] A. M. Carreiro, H. M. Jorge, and C. H. Antunes. "Energy management systems aggregators: A literature survey." In: *Renew. Sust. Energy Rev.* 73 (2017), pp. 1160–1172.

[124] I. Sajjad, G. Chicco, and R. Napoli. "Definitions of demand flexibility for aggregate residential loads." In: *IEEE Trans. Smart Grid* 7.6 (2016), pp. 2633–2643.

[125] S. Koch, J. L. Mathieu, and D. S. Callaway. "Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services." In: *14th Power Syst. Comput. Conf.* 2011, pp. 1–7.

[126] S. Barot and J. A. Taylor. "An outer approximation of the Minkowski sum of convex conic sets with application to demand response." In: *55th IEEE Conf. Decis. Control*. IEEE, 2016, pp. 4233–4238.

[127] C. Weibel. "Minkowski sums of polytopes: Combinatorics and computation." PhD Thesis. Lausanne: Ecole Polytechnique Federale de Lausanne, 2007, pp. 1–114. URL: https://infoscience.epfl.ch/record/109334/files/EPFL{_}TH3883.pdf.

[128] M. Kvasnica. *Minkowski Addition of Convex Polytopes*. 2005. URL: http://www.researchgate.net/publication/249915078{_}Minkowski{_}addition{_}of{_}convex{_}polytopes.

[129] M. Althoff, O. Stursberg, and M. Buss. "Computing reachable sets of hybrid systems using a combination of zonotopes and polytopes." In: *Nonlinear Anal.: Hybrid Syst.* 4.2 (2010), pp. 233–249.

[130] M. Althoff and B. H. Krogh. "Zonotope bundles for the efficient computation of reachable sets." In: *50th IEEE Conf. Decis. Control Eur. Control Conf.* IEEE, 2011, pp. 6814–6821.

[131] A. Girard and C. Le Guernic. "Zonotope/hyperplane intersection for hybrid systems reachability analysis." In: *Hybrid Syst. Comput. Control*. Vol. 4981. Springer, 2008, pp. 215–228.

[132] L. J. Guibas, A Nguyen, and L Zhang. "Zonotopes as bounding volumes." In: *Proc. 14th Annu. ACM-SIAM Symp. Discret. Algorithms*. SIAM, 2003, pp. 803–812.

[133] L. Vandenberghe, S. Boyd, and S.-P. Wu. "Determinant maximization with linear matrix inequality constraints." In: *SIAM J. Matrix Anal. Appl.* 19.2 (1998), pp. 499–533.

[134] M. Grant and S. Boyd. *CVX: Matlab Software for Disciplined Convex Programming, version 2.1*. 2014.

[135] M. Grant and S. Boyd. "Graph implementations for nonsmooth convex programs." In: *Recent Adv. Learn. Control*. Springer, 2008, pp. 95–110.

[136] S. Boyd, L. Xiao, and A. Mutapcic. "Subgradient methods." In: *Lect. Notes EE3920*. Stanford University. 2003, pp. 1–21. URL: http://xxpt.ynjgy.com/resource/data/20100601/U/stanford201001010/02-subgrad{_}method{_}notes.pdf.

[137] Swissgrid. *Basic principles of ancillary service products*. Tech. rep. Swissgrid, 2017, pp. 1–10. URL: https://www.swissgrid.ch/dam/swissgrid/experts/ancillary{_}services/Dokumente/D170214{_}AS-Products{_}V9R2{_}en.pdf.

[138] ENTSO-E. *Load-frequency control and performance*. ENTSO-E, 2004, pp. 1–28. URL: https://www.entsoe.eu/resources/publications/system-operations/operation-handbook/.

[139] F. A. Qureshi, I. Lymperopoulos, A. Ahmadi Khatir, and C. N. Jones. "Economic advantages of office buildings providing ancillary services with intraday participation." In: *IEEE Trans. Smart Grid* 99 (2016), pp. 1–10.

[140] H. Hao, A. Kowli, Yashen Lin, P. Barooah, and S. Meyn. "Ancillary service for the grid via control of commercial building HVAC systems." In: *Americ. Control Conf.* IEEE, 2013, pp. 467–472.

[141] Y. Lin, P. Barooah, and S. Meyn. "Low-frequency power-grid ancillary services from commercial building HVAC systems." In: *IEEE Int. Conf. Smart Grid Comm.* IEEE, 2013, pp. 169–174.

[142] S. I. Vagropoulos and A. G. Bakirtzis. "Optimal bidding strategy for electric vehicle aggregators in electricity markets." In: *IEEE Trans. Power Syst.* 28.4 (2013), pp. 4031–4041.

[143] H. Liu, Z. Hu, Y. Song, J. Wang, and X. Xie. "Vehicle-to-grid control for supplementary frequency regulation considering charging demands." In: *IEEE Trans. Power Syst.* 30.6 (2014), pp. 3110–3119.

[144] Swissgrid. *Balance group model introduction*. Tech. rep. Swissgrid, 2008, pp. 1–13. URL: https://www.swissgrid.ch/dam/swissgrid/experts/bgm/bg{_}documents/en/balance{_}group{_}model{_}intro{_}en.pdf.

[145] E. Vrettos and G. Andersson. "Scheduling and provision of secondary frequency reserves by aggregations of commercial buildings." In: *IEEE Trans. Sust. Energy* 7.2 (2016), pp. 850–864.

[146] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent. "Potentials and economics of residential thermal loads providing regulation reserve." In: *Energy Policy* 79 (2015), pp. 115–126.

[147] B. M. Sanandaji, T. L. Vincent, and K. Poolla. "Ramping rate flexibility of residential HVAC loads." In: *IEEE Trans. Sust. Energy* 7.2 (2016), pp. 865–874.

[148] S. Koch and G. Andersson. "Assessment of revenue potentials of ancillary service provision by flexible unit portfolios." In: *IEEE PES Gen. Meet.* IEEE, 2012, pp. 1–8.

[149] X. Zhang, E. Vrettos, M. Kamgarpour, G. Andersson, and J. Lygeros. "Stochastic frequency reserve provision by chance-constrained control of commercial buildings." In: *Eur. Control Conf.* 2015, pp. 1134–1140.

[150] I. Lymperopoulos, F. A. Qureshi, T. Nghiem, A. A. Khatir, and C. N. Jones. "Providing ancillary service with commercial buildings: The Swiss perspective." In: *Int. Fed. Autom. Control*. Vol. 48. 8. Elsevier, 2015, pp. 6–13.

[151] J. Warrington, P. Goulart, S. Mariéthoz, and M. Morari. "Policy-based reserves for power systems." In: *IEEE Trans. Power Syst.* 28.4 (2013), pp. 4427–4437.

[152] T. T. Gorecki, L. Fabietti, F. A. Qureshi, and C. N. Jones. "Experimental demonstration of buildings providing frequency regulation services in the Swiss market." In: *Energy Build.* 144 (2017), pp. 229–240.

[153] J. Warrington, P. J. Goulart, S Mariethoz, and M Morari. "Robust reserve operation in power systems using affine policies." In: *51st IEEE Conf. Decis. Control*. IEEE, 2012, pp. 1111–1117.

[154] D. Bertsimas, E. Litvinov, X. A. Sun, J. Zhao, and T. Zheng. "Adaptive robust optimization for the security constrained unit commitment problem." In: *IEEE Trans. Power Syst.* 28.1 (2013), pp. 52–63.

[155] X. Zhang, M. Kamgarpour, P. Goulart, and J. Lygeros. "Selling robustness margins: a framework for optimizing reserve capacities for linear systems." In: *53rd IEEE Conf. Decis. Control*. IEEE, 2014, pp. 6419–6424.

[156] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. "Adjustable robust solutions of uncertain linear programs." In: *Math. Program.* 99.2 (2004), pp. 351–376.

[157] T. Borsche, A. Ulbig, and G. Andersson. "A new frequency control reserve framework based on energy-constrained units." In: *Power Syst. Comput. Conf.* IEEE, 2014, pp. 1–7.

[158] G. Morales-Espana, A. Ramos, and J. Garcia-Gonzalez. "An MIP formulation for joint market-clearing of energy and reserves based on ramp scheduling." In: *IEEE Trans. Power Syst.* 29.1 (2014), pp. 476–488.

[159] M. Parvania and A. Scaglione. "Unit commitment with continuous-time generation and ramping trajectory models." In: *IEEE Trans. Power Syst.* 31.4 (2016), pp. 3169–3178.

[160] FERC. *Frequency regulation compensation in the organized wholesale power markets*. Tech. rep. U.S. Federal Energy Regulatory Commission, 2011, pp. 1 –123. URL: https://www.ferc.gov/whats-new/comm-meet/2011/102011/e-28.pdf.

[161] J. Abrell. *The Swiss wholesale electricity market*. Tech. rep. 2016, pp. 1–26. URL: https://www.ethz.ch/content/dam/ethz/special-interest/mtec/cer-eth/econo-mics-energy-economics-dam/documents/people/jabrell/Abrell{_}Swiss{_}Wholesale{_}Electricity{_}Market.pdf.

[162] A. Sadeghi-Mobarakeh and H. Mohsenian-Rad. "Performance accuracy scores in CAISO and MISO regulation markets: A comparison based on real data and mathematical analysis." In: *IEEE Trans. Power Syst.* 33.3 (2018), pp. 3196–3198.

[163] T. Wei, B. Zheng, Q. Zhu, and S. Hu. "Proactive demand participation of smart buildings in smart grid." In: *IEEE Trans. Comput.* 65.5 (2016), pp. 465–472.

[164] M. Gonzalez Vaya and G. Andersson. "Optimal bidding strategy of a plug-in electric vehicle aggregator in day-ahead electricity markets." In: *IEEE Trans. Power Syst.* 30.5 (2015), pp. 2375–2385.

[165] D. Bertsimas, D. B. Brown, and C. Caramanis. "Theory and applications of robust optimization." In: *Soc. Ind. Appl. Math.* 53.3 (2011), pp. 464–501.

[166] D. Schlipf, M. Scherer, and M. Haller. *Test for secondary control capability*. Tech. rep. Swissgrid, 2012, pp. 1–5. URL: https://www.swissgrid.ch/dam/swissgrid/experts/ancillary{_}services/prequalification/en/D130422{_}Test-for-secon-dary-control-capability{_}V2R1{_}EN.pdf.

[167] Tesla. *Power Wall 2 AC*. 2016. URL: http://www.energymatters.com.au/wp-content/uploads/2016/11/tesla-powerwall-2-datasheet.pdf (visited on 06/26/2018).

[168] T. Borsche, A. Ulbig, M. Koller, and G. Andersson. "Power and energy capacity requirements of storages providing frequency control reserves." In: *IEEE PES Gen. Meet.* IEEE, 2013, pp. 1–5.

[169] Y. V. Makarov, S. Lu, J. Ma, and T. B. Nguyen. *Assessing the value of regulation resources based on their time response characteristics*. Tech. rep. Pacific Northwest National Laboratory, 2008, pp. 1–83. URL: https://www.pnnl.gov/main/publications/external/technical{_}reports/PNNL-17632.pdf.

[170] H. Nosair and F. Bouffard. "Energy-centric flexibility management in power systems." In: *IEEE Trans. Power Syst.* 31.6 (2016), pp. 5071–5081.

[171] A. Bitlislioglu, T. T. Gorecki, and C. N. Jones. "Robust tracking commitment." In: *IEEE Trans. Autom. Control* 62.9 (2017), pp. 4451–4466.

[172] B. Foggo and N. Yu. "Improved battery storage valuation through degradation reduction." In: *IEEE Trans. Smart Grid* 3053 (2017), pp. 1–13.

[173] N. Kumar, B. Peter, S. Lefton, D. Agan, and D Hilleman. *Power plant cycling costs*. Tech. rep. NREL, 2012, pp. 1–83. URL: https://www.nrel.gov/docs/fy12osti/55433.pdf.

[174] *Tesla Model S*. 2018. URL: https://www.tesla.com/models (visited on 04/26/2018).

[175] *Tesla Powerpack*. 2018. URL: https://www.tesla.com/powerpack (visited on 04/26/2018).

[176] Siemens. *Siemens steam turbine portfolio*. Tech. rep. Siemens, 2017, pp. 1–88. URL: https://www.energy.siemens.com/mx/pool/hq/power-generation/steam-turbines/downloads/new/siemens-steam-turbines-product-overview-EN-updated.pdf.

[177] C. Hinman. *Pay for performance regulation (FERC Order 755): Updated with year one design changes*. Tech. rep. CAISO, 2015, pp. 1–41. URL: https://www.caiso.com/Documents/Pay-PerformanceRegulationFERC{_}Order755Presentation.pdf.

[178] W. Wang and N. Yu. "Chordal conversion based convex iteration algorithm for three-phase optimal power flow problems." In: *IEEE Trans. Power Syst.* 33.2 (2018), pp. 1603–1613.

[179] E. Dall'Anese, Hao Zhu, and G. B. Giannakis. "Distributed optimal power flow for smart microgrids." In: *IEEE Trans. Smart Grid* 4.3 (2013), pp. 1464–1475. arXiv: 1211.5856.

NOTATION

DESCRIPTION OF FLEXIBILITY

SYMBOL	DESCRIPTION
a	System dynamics parameter, energy self-dissipation rate
b, c	System dynamics parameters, energy conversion factors
a_d, b_d, c_d	Discrete-time counterparts of a, b, c
(A, b)	Constraints matrix and limits of a polytopic feasible set
k	Discrete-time index
N	Number of time steps in the planning horizon
p_k	Average power consumption during time step k
$\underline{p}_k$, $\bar{p}_k$	Bounds on power during time step k
$\mathcal{P}$	Feasible set of an energy resource
r_k	Power ramp-rate during time step k
$\underline{r}_k$, $\bar{r}_k$	Bounds on power ramp-rate during time step k
t_s	Sampling time
u_k	Constant exogenous input during time step k
x_k	System state at the end of time step k
$\underline{x}_k$, $\bar{x}_k$	Bounds on state during time step k

ESTIMATION OF FLEXIBILITY

SYMBOL	DESCRIPTION
$\hat{d}_c$	Estimated duty cycle of an HP
$\hat{d}^i_{\text{on}}$, $\hat{d}^i_{\text{off}}$	Estimated on-, off-duration of an HP in heating cycle i
$d_{\text{off,max}}$	Max. allowed off-time for an HP
f	Flexibility level
i	Index of the heating cycle of an HP
$\mathcal{I}_{\text{extr}}$	Indices of local extrema of $\Delta_2 p_k$
$\mathcal{I}_{\text{on}}$, $\mathcal{I}_{\text{off}}$	Indices of switch-on, -off candidates
$\mathcal{I}_{\text{switch-on}}$, $\mathcal{I}_{\text{switch-off}}$	Indices of switch-ons, -offs
$\hat{p}_{\text{on}}$, $\hat{p}_{\text{off}}$	Estimated switch-on, -off power of an HP
$\hat{p}_r$	Estimate of the rated power of an HP
$\Delta_2 p_k$	Change in power over two time steps
$\hat{q}^i_c$, $\hat{q}^i_l$	Avg. charge, loss rate during heating cycle i
$\hat{r}^i_c$, $\hat{r}^i_l$	Avg. normalized charge, loss rate during heating cycle i
$\hat{t}^i_{\text{on}}$, $\hat{t}^i_{\text{off}}$	Estimated switch-on, -off times of the HP in heating cycle i
T^H	Duration of planning horizon
Δ_{on}, Δ_{off}	Thresholds for detection of switch-ons, -offs
θ^*_{out}	Zero-loss temperature

RESIDENTIAL BUILDING WITH HEAT PUMP

SYMBOL	DESCRIPTION
a_l, b_l	Parameters of loss-rate function
a_c, b_c, c_c	Parameters of charge-rate function
d_c	Duty cycle of an HP
e_k	Energy level at the end of time step k
$\underline{e}$, $\bar{e}$	Constant energy bounds
e_{flex}	Flexible energy
R	Lumped heat-transfer coefficient of the building
$T_{\max}$	Max. allowed HP throttling time
η	COP of the HP
θ_{in}, θ_{out}	In-, outdoor air temperature
$\underline{\theta}_{\text{in}}$, $\bar{\theta}_{\text{in}}$	Bounds on indoor air temperature
θ_w	Heating circuit water temperature
ζ_{HP}	Exergetic efficiency of an HP

POPULATION OF BUILDINGS WITH HEAT PUMPS

SYMBOL	DESCRIPTION
$\mathcal{I}_T$	Indices of HPs that can be throttled for duration T
j	System index
J	Total number of systems
$\mathcal{J}$	Set of system indices
$\hat{p}_{\text{reb}}^{(j)}$	Expected rebound power from system j
$\hat{p}_{\text{red}}^{(j)}$	Expected load reduction from system j
$P(\cdot)$	Probability mass function
$T_{\max}^{(j)}$	Maximum throttling time of system j
$\hat{\sigma}_{\text{red}}$	Estimated load-reduction success rate
$\hat{\sigma}_{\text{reb}}$	Estimated rebound occurrence rate

AGGREGATION OF FLEXIBILITY

SYMBOL	DESCRIPTION
$\mathcal{B}(l, u)$	Axis-aligned box with bounds l, u
e_i	Unit vector of coordinate i
$f^{(i)}$	Normal vector of the ith zonotope facet
F	Matrix of zonotope facet normal vectors
$\mathcal{F}$	Zonotope facet
$g^{(i)}$	Generator vector i
G	Zonotope generator matrix
G^{PE}	Special choice of G for PE-systems
$\mathcal{H}$	Affine hyperplane
I_N	Identity matrix of dimension N
n_f	Number of zonotope facets
n_g	Number of generator vectors
$\mathcal{P}^{(j)}$, $\mathcal{P}^{(\mathcal{J})}$	Feasible set of system j, aggregate feasible set
$T^{(j)}$, $T^{(\mathcal{J})}$	Tracking cost function of system j, aggregate tracking cost function
$\mathcal{Z}(G, c, \bar{\beta})$	Zonotope with generator matrix G, center c, and symmetric bounds $\bar{\beta}$
β	Zonotope parameter, bounded by $\bar{\beta}$
$\delta_{\mathcal{P}}$, $\delta_{\mathcal{Z}}$	Offsets of hyperplanes with touching $\mathcal{P}$, $\mathcal{Z}$
$\Delta_{\mathcal{BP}}$, $\Delta_{\mathcal{ZP}}$	Approximation quality measure for box, zonotope
$\Delta_{\mathcal{P}}$, $\Delta_{\mathcal{Z}}$	Extension of $\mathcal{P}$, $\mathcal{Z}$ along zonotope facet normals
ϵ	Initial SOC of a PE-system
λ	Measure of dominance of energy constraints
$\Lambda_{\mathcal{Z}}$	Measure of zonotope approximation quality
μ	Ratio of capacity and power rating of a PE-system

DISGGREGATION OF FLEXIBILITY

SYMBOL	DESCRIPTION
a	Subgradient step-size update parameter
$h^{(j)}$	Flexibility cost function of system j
k	Iteration index of subgradient algorithm
$l_{i,i'}^{(j)}$	Length of line segment i' of $T_i^{(j)}$
L_i	List of ordered triples $(l_{i,i'}^{(\mathcal{J})}, q_{i,i'}^{(\mathcal{J})}, s_{i,i'})$
$m_i^{(j)}$	Number of line segments of $T_i^{(j)}$
$p_{\text{ref}}^{(j)}$, $p_{\text{ref}}^{(\mathcal{J})}$	Individual, aggregate reference power trajectory
$q_{i,i'}^{(j)}$	Slope of line segment i' of $T_i^{(j)}$
$Q_{\mathcal{S}}(\cdot)$	*Euclidean* projection onto set $\mathcal{S}$
$\bar{r}$	Units of regulation power offered
$\hat{R}(\bar{r})$	Expected costs of offering $\bar{r}$ units of regulation power
$s_{i,i'}$	Book-keeping variable storing the system index
$\mathcal{S}$	Feasible set of the disaggregation problem
$T^{(j)}(p)$	Implementation costs of system j for trajectory p
$T_i^{(j)}$	Cost component along direction of generator $g^{(j)}$
v	Vector of energy prices agreed upon in AG contract
$V(k)$	Objective value in iteration k
$\alpha(k)$	Subgradient step size in iteration k
$\beta^{(j)*}$, $\beta^{(\mathcal{J})*}$	Optimal individual, aggregate β found
δ	Number of iterations considered by termination criterion
ϵ	Relative convergence tolerance
η	Auxiliary scaling factor
$\nabla(k)$	Subgradient in iteration k

ENERGY AND RESERVE BIDDING

SYMBOL	DESCRIPTION
c^{DA}, c^{ID}	Day-ahead, intra-day energy prices
c^{RES}	SFR reserve capacity price
c^{up}, c^{dn}	Up-, down-regulation energy prices
C^{DA}, C^{ID}	Day-ahead, intra-day energy procurement costs
e^{base}	Energy market baseline
e^{DA}, e^{ID}	Day-ahead, intra-day energy schedule
e^{up}, e^{dn}	Up-, down-regulation energy
$\mathbf{E}[\cdot]$	Expectation
k, l, s	Indices of intra-day, SFR activation, system time intervals
N^{DA}, N^{ID}	Number of DA-, ID-time intervals
$N^{\mathrm{DA}}_{\mathrm{lb}}$, $N^{\mathrm{ID}}_{\mathrm{lb}}$	Number of activation samples considered by day-ahead, intra-day trading policies
p_{ref}, p_{tgt}	Reference,target power trajectory
Q^{DA}, q^{DA}	Parameters of day-ahead control policy
Q^{ID}, q^{ID}	Parameters of intra-ahead control policy
R^{REG}	Remuneration for regulation energy
R^{RES}	SFR capacity reservation payment
T^{DA}, T^{ID}	Tendering period of day-ahead, intra-day energy market
$\mathcal{T}^{\mathrm{DA}}$, $\mathcal{T}^{\mathrm{ID}}$	Day-ahead, intra-day market time interval
T^{C}, T^{S}	Discretization time of SFR activation, system dynamics
T^{H}	Duration of planning horizon
T^{RES}	SFR tendering period
T^{RP}	Power schedule ramp duration
V	AG profit made from selling energy and SFR
w	SFR activation signal
$\mathcal{W}$	Set of possible SFR activations
y, z	Auxiliary state dynamics
γ	Constant and symmetric SFR capacity
Ω, ζ	Set of decision variables of the bidding problem
Φ	Set of system parameters
Ψ	Set of uncertain parameters of the bidding problem

SYNERGY EFFECTS OF AGGREGATION

SYMBOL	DESCRIPTION
$a^{(B)}$, $a^{(B)}$, $a^{(B)}$	System parameters of battery
$a^{(F)}$, $a^{(F)}$, $a^{(F)}$	System parameters of freezer
$\underline{\theta}$, $\bar{\theta}$	Lower, upper bound on freezer temperature
$\gamma_{\max}^{(B)}$, $\gamma_{\max}^{(F)}$, $\gamma_{\max}^{(ST)}$	Max. SFR capacity available from battery, freezer, steam turbine
$\gamma^{(\mathcal{J})}$	Max. SFR capacity available from population $\mathcal{J}$
δ	Control delay
θ_{out}	Outdoor air temperature
σ^{SFR}	SFR synergy factor

ACRONYMS

AARC	affinely adjustable robust control
ADMM	Alternating Direction Method of Multipliers
AG	Aggregator
APE	absolute percentage error
BRP	Balance Responsible Party
COP	coefficient of performance
DEU	Distributed Energy Unit
DLC	Direct Load Control
DR	Demand Response
ENTSO-E	European Network of Transmission System Operators
EU	European Union
FEU	Flexible Energy Unit
FLECH	Flexibility Clearing House
GHG	greenhouse gas
HAG	home automation gateway
HP	heat pump
ILC	Indirect Load Control
IR	internet router
LP	Linear Program
PEV	plug-in electric vehicle
PV	photovoltaic
QP	Quadratic Program
RE	relay
RES	renewable energy sources
SFR	Secondary Frequency Regulation
SG	subgradient
SM	smart meter
SO	System Operator
SOC	state of charge
TC	thermostat controller
TCL	thermostatically controlled load
VRES	variable renewable energy sources

WS weather service