Preface

This thesis represents the last step of my academic journey at the University of Twente. It has always been a pleasure to study in Enschede, and I look back fondly at the last six years as a time full of amazing experiences. At the same time, I am looking forward to the opportunities and challenges that lie in my future, quite possibly in the field of energy management.

The work behind this thesis was carried out at TNO, the Hague in the period between September 2019 and May 2020. The past months have been both fun and challenging in many ways, and I could not have finished the thesis before you without the help of many people. First of all, I would like to thank my supervisors. I would of course like to thank Johann Hurink at the UT, for introducing me to the energy domain, and giving me the chance to combine Applied Mathematics with my passion for sustainability. Johann has been a great supervisor, and despite the physical distance was able to guide me through the many hurdles associated with the final project. Also at the UT, Jens Hönen was always available to help despite his busy PhD schedule. At TNO, Coen van Leeuwen got me through the day-to-day challenges and helped me feel at home at the organisation. Furthermore, Koen Kok was instrumental in the first part of my final project and provided interesting insights from his career in the energy domain. Lastly, I would like to thank my family and friends who, especially through the last weeks, have been of tremendous support. Special mention goes to Maaike who would always cheer me when it was needed most.

Abstract

Due to the ongoing energy transition and its associated trends, operation of the electricity grid is becoming more and more challenging. This is especially true for distribution grids where intermittent generation and increased demand coincide. In this thesis, energy management approaches for effective operation of the smart distribution grid are developed and compared. The energy management problem is first introduced in a general way, considering its main characteristics and developing a mathematical basis. The two main classes of energy management approaches are then represented by introducing a novel centralised model and a novel decentralised model. The fundamental mathematical features of the centralised and decentralised models are compared, focusing on the flow of information and availability of control. Furthermore, a case study is developed and tested to address the real-life performance of the two energy management approaches. In this case study, a dispatch problem in a day-ahead market is simulated using real-life data and realistic device specifications. The goal of the energy management approaches is to find a dispatch that follows a target demand profile in a grid feasible way while simultaneously being economic and sustainable.

The interaction of household loads, solar panels, household batteries, and heat pumps is investigated, comparing the centralised and decentralised models based on qualitative and quantitative performance criteria. The qualitative part of this comparison includes privacy and cybersecurity aspects, where we conclude that the decentralised model outperforms its centralised counterpart by the more local and anonymous nature of data processing. The quantitative part of this analysis considers robustness of the different modelling approaches by looking at the percentage of feasible runs for a number of different modelling instances. Furthermore, the degree too which the models are considered economic and sustainable is analysed by comparing the total energy loss results. For these quantitative measures, the centralised model was shown to significantly outperform the decentralised model.

Contents

1 Introduction **5**
1.1 Problem overview . . . 5
1.2 Organisation of this thesis . . . 5

2 Background **6**
2.1 The electricity system . . . 6
2.1.1 Historical centralised grid . . . 6
2.1.2 Market deregulation . . . 7
2.1.3 Energy transition . . . 8
2.2 Smart grids . . . 10
2.2.1 Definition and context . . . 10
2.2.2 Smart devices . . . 11
2.3 Energy management . . . 12
2.3.1 Definition and context . . . 12
2.3.2 Challenges . . . 13
2.3.3 Related work . . . 14

3 Centralised versus decentralised control **18**
3.1 Problem statement . . . 18
3.2 Model descriptions . . . 18
3.2.1 General model description . . . 19
3.2.2 Centralised optimisation model . . . 23
3.2.3 Decentralised optimisation model . . . 24
3.3 Theoretical comparison . . . 27

4 Modelling uncertainty **30**
4.1 Sources of uncertainty . . . 30
4.1.1 Intermittent generation . . . 30
4.1.2 Human behaviour . . . 30
4.1.3 General remarks . . . 31
4.2 Realistic data input . . . 32
4.2.1 Characterisation of input data . . . 32
4.2.2 Model data generation . . . 34
4.3 Receding horizon models . . . 35
4.3.1 Receding horizon definition . . . 35
4.3.2 General model description . . . 36
4.3.3 Receding horizon centralised model . . . 37
4.3.4 Receding horizon decentralised model . . . 38

5 Numerical experiments **40**
5.1 Experiment setup . . . 40
5.1.1 Models . . . 40
5.1.2 Performance criteria . . . 40
5.1.3 Case description . . . 41
5.2 Feasibility . . . 45
5.2.1 Device constraints . . . 46
5.2.2 Congestion constraints . . . 46
5.2.3 Market constraints . . . 47
5.2.4 Discussion . . . 47
5.3 Computation time . . . 48
5.3.1 Optimisation paths . . . 49
5.3.2 Stopping criteria . . . 50
5.3.3 Discussion . . . 51
5.4 Total loss . . . 52

5.4.1 Quality of solution 52
5.4.2 Device activation 53
5.4.3 Allocation 54
5.4.4 Discussion 56

6 Conclusion and recommendations **58**
6.1 Conclusion 58
6.1.1 Research question 1 58
6.1.2 Research question 2 58
6.2 Recommendations 60

References **64**

1 Introduction

Over the last decades, increasing concerns about man-induced global warming have called for a fundamental change in the energy system. The move from fossil based to renewable energy sources represents one of the biggest challenges of our time and is referred to as the *energy transition*. The energy transition requires an enormous amount of economical, political, and engineering developments. One such development is found in the field of *smart grids*, i.e. electricity grids where a smart combination of sensoring and automation aims to solve problems that arise from the inherently intermittent nature of renewable energy sources. Effective operation of a smart grid should not only be able to solve technical problems, it should also appeal to the involved stakeholders and function under a wide range of circumstances.

At TNO, and specifically in the Monitoring & Control Services (MCS) department, a great volume of research and many practical innovations have been added to the field of smart grids over the years. Some recent examples include the *PowerMatcher* and *ReFlex* algorithms, and the pilot smart grid in Heerhugowaard. Similarly, the University of Twente has made many contributions to the smart grid literature, notably in the Discrete Mathematics and Mathematical Programming (DMMP) department where energy and Applied Mathematics coincide. Under the banner of the *Agile project*, these developments now coincide through a collaboration between the University of Twente, the TU Eindhoven, and TNO. This thesis represent just a small element of this collaborative effort, aiming to combine expertise from both TNO and the University of Twente.

1.1 Problem overview

The smart grid problem that is considered in this thesis resides in the distribution grid, i.e. the low voltage portion of the electricity grid. The distribution network connects neighbourhoods and small businesses to the grid and is quickly transforming in conjunction with the energy transition. With increasing penetration of distributed energy resources such as solar panels, electric vehicles, and heat pumps on a household level, it becomes increasingly challenging to operate the distribution network. To address this challenge, three models are introduced that represent varying ways of preventing operational problems at the distribution level. One of these models is agent-based, or decentralised, and represent a novel approach for smart grid operation. The other two models are centralised and take a more conventional approach to smart grid operation. The three models are first compared theoretically, and later numerically through a case-study using real-life data. These comparisons play a central role in the analysis of the operations problem outlined in this thesis. The overarching problem can be summarised into two main research questions:

RQ1 How does the mathematical structure differ between the centralised and decentralised models and can we compare them?

RQ2 How effective are the centralised and decentralised models in operating the smart distribution grid?

1.2 Organisation of this thesis

This thesis consists of a total of 6 Chapters that together address the problem outlined above, and aim to answer the given research questions. The present Chapter gives an introduction to this process. Next, some background to the research is provided in Chapter 2 including context about the energy system, operation of smart grids, and the related literature. With the background firmly established, the base (mathematical) description of the problem and the centralised and decentralised models is introduced in Chapter 3 addressing the first part of **RQ1**. The chapter ends with a discussion on the compatibility of the models addressing the remainder of **RQ1**. In Chapter 4, the base models from Chapter 3 are further developed to more realistically reflect the presence of uncertainty in the operational problem. In Chapter 5, the numerical analysis is considered. In this chapter, a detailed case description is first given after which a number of experiments are carried out to help answer **RQ2**. The results and a brief discussion are both presented in the chapter. The thesis ends by outlining the main conclusions and recommendations in Chapter 6.

2 Background

To give context to the following analysis and provide the reader with the necessary tools, some background information is given in this Chapter. In Section 2.1, the most important features of the electricity system and its ongoing transformation are described briefly. Section 2.2 gives an introduction to the smart grid paradigm. And lastly, in Section 2.3 the concept of energy management is described together with an overview of the literature on this subject.

2.1 The electricity system

In this section, some background is provided to the energy system. In Subsection 2.1.1, the history of the current energy system is considered together with some aspects of the system that are still in place today. Next, Subsection 2.1.2 goes into the current deregulated energy market and general market dynamics and involved parties are introduced. Lastly, in Subsection 2.1.3, special attention is given to the energy transition and its most important trends.

2.1.1 Historical centralised grid

The first recorded instance of systematic energy distribution through an electrical grid happened in 1882. The grid was established by Thomas Edison in Manhattan, and connected 59 users using a direct current (DC) system. In the mid 1880's, Wettinghouse Electric introduced a new alternating current (AC) system that beat Edison's DC current in transport efficiency and ease of voltage transformation. Based on this technology, many localised networks were created in the following years, and their integration called for standardisation of the voltage profile. Eventually, AC emerged as the preferred technology and regional, national, and even international grids became connected [21].

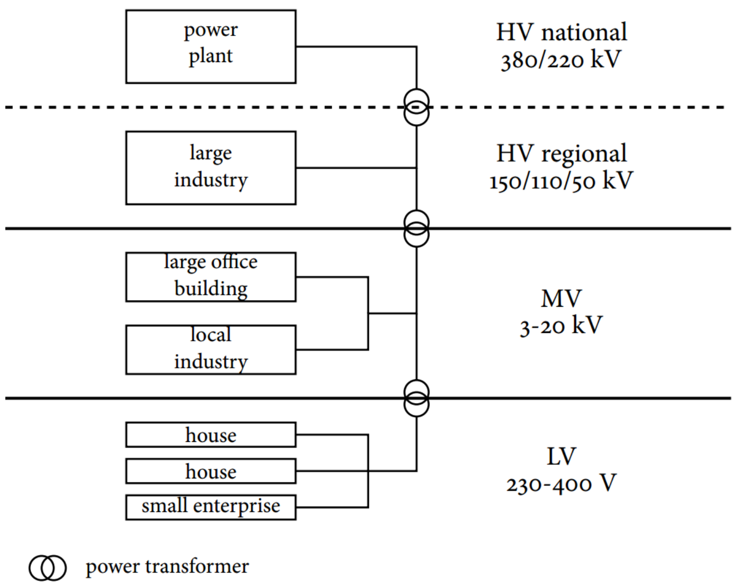

FIGURE 1: *Schematic representation of the electricity grid in the Netherlands. Solid lines are drawn to separate different voltage levels and the dotted line signifies the boundary between the generation and transmission networks [21].*

The resulting large-scale electricity grids have originally been operated centrally where large generators take care of the supply of electricity. The transport of electrical energy can roughly be subdivided into three main levels characterised by the used voltages: high voltage (HV), medium voltage (MV), and low voltage (LV). These levels are demonstrated for the Dutch power grid in Figure 1. At the highest level (HV), power plants are connected, which inject electrical energy into the grid. By this HV transmission network, electricity is carried from power plants to demand centres. In these centres, the distribution network (LV/MV) is used to supply power to end users including households, small offices, and local industrial complexes. Here, transformers are used to transfer power between different voltage levels. This hierarchical structure also corresponded with the flow of power: its direction was always top-down from power plant through transmission grid to distribution and finally ending at the user level. Storage of electricity played a negligible role which means that supply and demand needed to match at all times. This is referred to as *balancing* of the grid. To ensure grid balance, control strategies were implemented at the supply side because demand was assumed to be inflexible, i.e. "production follows demand" [45]. This balancing process follows a merit order: plants with the lowest operating cost should run first. A stable base load is provided by large, but typically inflexible, *base load power plants* such as coal and nuclear power plants. This base load is supplemented by *load-following plants* (e.g. combined-cycle gas plants) that mostly operate during the working day, and final balancing is done by flexible but less efficient *peak production plants* (e.g. gas turbine plants). In this context, flexibility refers to the ease at which generators may ramp their production level up or down. In the case of a coal power plant, going from zero to maximum production may take in the order of a full day. In contrast, a gas turbine plant can do the same in mere minutes. Power plants and system operators both benefit from a flat demand profile; variability results in power plants that run below or above their optimum working point and puts additional strain on the grid. Furthermore, sudden peak loads may lead to grid failure when the system does not have the time to adequately respond [44].

2.1.2 Market deregulation

Around the mid 1990's, legislation passed throughout Europe that lead to the privatisation/deregulation of the electricity supply chain. Operation of electricity generation, transmission, distribution, and retail was split between different parties. This led to a market mechanism which takes care of matching supply and demand of electricity. Ideally, this market resembles an optimisation process where (near) optimal allocation of resources is achieved. This optimum is called market equilibrium and is achieved through the economic principle of supply and demand. In the deregulated market, three main parties operate together:

- **TSOs:** *Transmission system operators* (TSOs) own the transmission sections (HV) of the grid. In balancing the grid, TSOs are tasked with assuring system stability by managing the AC frequency, which is a system-global property. TSOs also play a role as auctioneer in the electricity market. They collect bids from producers and retailers and determine a clearing price in a transparent manner. Participation in these regulated markets is mandatory for larger producers and consumers. In addition, some producers may be forced to dedicate part of their production as *reserve capacity* in exchange for a compensation to help balance the grid. These compensations may also be traded through *capacity markets* where value is assigned to increased generator flexibility [45].
- **DSOs:** *Distribution system operators* (DSOs) own the distribution sections (MV/LV) of the grid. Together with the TSOs, these operators are responsible for managing the proper functioning and connectivity of the electricity system. DSOs are responsible for maintaining voltage norms. In the centralised grid described in Section 2.1.1, this operational task was almost non-existent, consisting mostly of maintenance and incident response. However, as will be seen in Section 2.1.3 and Section 2.2, the role of DSOs has become increasingly complex in the last years. Because their shared role constitutes a natural monopoly, TSOs and DSOs are subject to strict state regulation (in Europe).
- **BRPs:** In addition to system operators, specific market participants are recognised under the name *balance responsible parties* (BRPs). These include power producers, retailers, and large consumers. Producers are generally private parties that take care of the generation of energy. Retailers (or aggregators) are parties that buy energy from producers and sell directly to users. Retailers collaborate with the system operators to ensure that it is possible to deliver the requested energy to the users. Furthermore, retailers may provide guarantees with respect to properties of their supplied power (e.g. at least above a certain

percentage renewable sources). All BRPs have some contractual obligations for energy consumption or production. Often in the Netherlands, a single BRP simultaneously represents a set of power plants (producers) together with a retail division. With permission of the system operators, these BRPs may regulate their dispatch internally (self-dispatch).

The operation of energy markets and resources can be divided into different timescales. Long term (2 - 20 years), *strategic*, decisions typically consist of large infrastructural projects, such as opening of new plants, and must consider projected changes in the relevant areas (e.g. due to new industrial areas). At the medium term (3 months - 2 years), decisions are made on when to schedule maintenance on e.g. the existing plants and grid resources. At the short term, different planning horizons are again considered. The day-ahead horizon refers to decisions that are made for the subsequent day. These decisions consist of three parts:

- **Unit commitment problems (UCPs)**: *Unit commitment problem* denotes the problem of scheduling what plants should be in operation at what time of the day [34].
- **Economic dispatch (ED)**: The *economic dispatch* problem consists of determining the production profiles of each plant during the commitment windows decided in the UCP [40].
- **Optimal power flow (OPF)**: In the *optimal power flow* problem, the schedule resulting from the UCP and ED may be altered to account for the constraints and losses of the physical grid [30].

Intraday planning (having a horizon of minutes - hours ahead) considers changes in behaviour and predictions during the day. At this time scale, operators can (manually) adapt the dispatch to new conditions. Finally, at the primary balancing time scale (also known as real-time control) automatic ICT systems take care of the finest balancing duties. This is done through *automatic generation control* (AGC) where strict balancing constraints and strong penalties for failure to respond are implemented. *Dynamic pricing* refers to dynamically updating electricity prices, and is often used for AGC at the transmission grid level. Note, that at the time of writing the Netherlands does not have a system in place that facilitates dynamic pricing for consumers. For consumers at the distribution level, there is only a day- and night-tariff.

2.1.3 Energy transition

The situation described in Section 2.1.1 has prevailed around the globe until fairly recently. Over the last decades, increased concerns about man-induced global warming have resulted in a need for rigorous re-engineering of the energy system. In accordance with the *Paris Agreement*, drafted by the UNFCCC in 2015 and signed by 195 member states, global average temperature rise should be limited to 2℃ compared to pre-industrial levels [1]. This goal is encapsulated in the so-called 20/20/20 targets, including the reduction of carbon dioxide (CO_2) emissions by 20%, the increase of the fraction of renewable energy sources (RES) by 20%, and a 20% increase in overall energy efficiency up to 2020. The process of altogether replacing fossil energy sources by renewables and the accompanying system alterations is colloquially referred to as the *energy transition*. Specifically, the energy transition represents the ambition to make the entire energy life cycle clean and renewable in order to halt the depletion of resources and to not restrict or harm future generations. This includes production and transportation of energy, but also sustainable storage technologies and end-of-life management of green energy generators. Following [44], this process is explained through four distinct trends: renewable generation, the electrification of demands, decentralisation of generation, and energy storage:

- **Renewable generation:** Renewable generation refers to the generation of energy through *renewable energy sources* (RES). A RES is an energy source that is naturally replenished in a human timescale [14]. Examples include wind, water, solar, and geothermal heat. Sometimes the terms *renewable* energy source and *sustainable* energy source are used interchangeably. However, this is not always correct as "sustainable" implies a stricter and often more subjective condition on the source of energy. In [15], sustainable energy is described as energy that "meets the needs of the present without compromising the ability of future generations to meet their own needs." For instance, some authors argue that hydro power does not constitute a sustainable energy source due to the damage to local ecosystems. An important aspect of RES is the typically stochastic and intermittent nature of production. Most notably, the solar and wind energy output depends directly on weather conditions. For these energy sources, supply side management is limited to curtailment.

In addition to directives such as those derived from the Paris Agreement, cost factors are (slowly) paving the way for economically attractive RES. Depletion of fossil fuel reserves increases the costs of "grey" energy while technological advances in renewable energy generation rapidly decrease costs of "green" alternatives. Most notably, solar power has experienced an exponential decrease in cost per unit power that has lead to a factor > 2000 reduction since 1956 [17]. On a global scale, RES represented 19.3% of the total energy production in 2017. The Netherlands is notably lagging behind globally and even more so compared to other EU states, with a renewable energy share of just 6% in 2015 [11]. An effort to catch up on the climate targets set by the EU is presented in the *Energieakkoord* response plan [44].

- **Electrification of demand:** Electricity as a carrier for sustainable energy has proven to be an indispensable part of the energy transition [26]. In recognition of this fact, a rapid development is seen towards more electrically powered devices, referred to as *electrification*. Coupled with the fact that most RES supply energy through electricity, these developments are drastically altering the way the power grid is used.

  A prominent example of such devices are electric vehicles (EVs). Instead of relying on carbon-based fuels, EVs run on electricity that is stored in a large battery. Depending on frequency and intensity of usage, EVs may more than double the electricity demands for an average household. Another example are heat pumps, which serve as an alternative to conventional natural gas fueled heating. Heat pumps are electrically powered and utilise the (vapor compression) refrigeration cycle to heat or cool households.

  As a consequence of the electrification, at all levels of the grid, higher total power transmission is seen with increased local peak production and demand. On the other hand, local RES such as privately owned photovoltaics (PV) also influences the power flow in the distribution grid and even results in temporary net back-flows of power. The effect of electrification on the load curve of an energy neutral house can be seen in Figure 2.

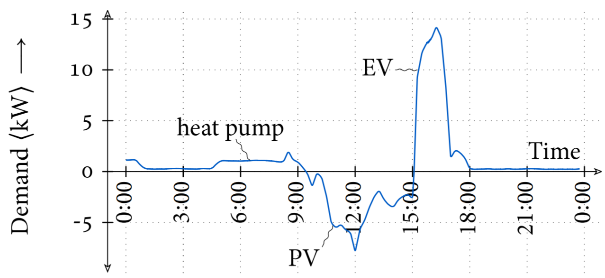

FIGURE 2: *Example of a power profile for a Dutch energy neutral house. Responsible devices are identified for each peak [33].*

- **Decentralisation of generation:** Decentralisation of generation, or *distributed generation* (DG), in the distribution grid was briefly touched upon in the previous paragraph when talking about household PV production. A similar effect is seen at the transmission level. In Section 2.1.1, the historical situation was described where centralised energy generation by means of fossil fuel power plants represented the cheapest and preferred means of production. Economies of scale enabled the large fossil fuel power plants to generate electrical power at a low market price. Because of their immense power output, these plants needed to be connected to the HV transmission network. With the current evolution of renewable generation, a trend towards smaller and more localised production is seen. Solar- and wind "parks" have less power output per square meter and are naturally more spread due to geographical and meteorological reasons. Because of these features, DGs are often connected to the MV grid as seen in the case of wind turbines or even the LV grid as seen in the case of household PV.

An important advantage of decentralisation is the reduction in average transport distance, minimising transport losses. A notable disadvantage was touched upon in regards to electrification. Net feed-in of power at the distribution grid level is not uncommon in the context of DGs [33]. As electrical grids and transformers are historically optimised for strict top-down transmission of electricity (see also Section 2.1.1), these feed-in peaks may result in local grid failures.

- **Storage:** The last trend is widespread storage of energy and electricity in particular. In the previous paragraphs, some challenges of the energy transition were described in the context of RES, electrification, and decentralisation. Furthermore, a fundamental change in overall grid usage was identified. This change manifests itself mostly in the form of uncertainty in generation, increase of total electricity demand, and higher variance in power demand characteristics. Storage has the potential to address these challenges by shifting demand in time; e.g. by charging the battery during feed-in peaks and discharging the battery at load peaks. Ideally, DGs should be coupled with decentralised storage to further minimise transport losses and locally manage congestion.

  Currently, storage technologies have substantial drawbacks and limitations. Storage devices are limited by their capacity, and minimum and maximum charge rates. Furthermore, batteries are still relatively expensive and deteriorate quickly depending on usage. Also, losses are significant, ranging somewhere between 10-70% for charging, with energy leakage over time further decreasing the overall efficiency. However, battery quality is improving every year while prices are declining steadily as seen in [29]. With these developments in mind, the rise of storage technologies represents an important trend in the energy transition.

2.2 Smart grids

This Section provides an introduction to the field of smart grids. A formal definition and some context are first provided in Subsection 2.2.1. Next, in Subsection 2.2.2, the different classes of smart devices are introduced.

2.2.1 Definition and context

In order to deal with the changes resulting from the energy transition, considerable changes in the grid need to be implemented (see also Section 2.1.3). At first glance, an obvious solution might be to upgrade transformers and transmission lines to be able to handle increased loads. These physical efforts are referred to by the term *copper plating*. Unfortunately, this conventional type of grid reinforcement typically does not result in balancing and is extremely expensive. Especially in the Netherlands, where most power lines are installed underground, replacing cables is very costly [21].

An emerging effort to mitigate the problems that the grid is facing is presented under the banner of *smart grid* technologies. The term smart grid is used to describe a variety of “smart” solutions that aim to simultaneously solve grid capacity- and balancing problems. A central concept is the integration of a communication layer in the process. This communication layer may be used to collect data at the distribution and transmission level and automatically control grid assets and even user appliances. The smart grid paradigm unites many technologies that communicate or are somehow connected through this communication layer. A more formal definition is expressed in [22] by the international energy agency (IEA):

Definition 1. *A smart grid is an electricity network system that uses digital technology to monitor and manage the transport of electricity from all generation sources to meet the varying electricity demands of end users. Such grids are able to coordinate the needs and capabilities of all generators, grid operators, end users and electricity market stakeholders in such a way that they can* *optimise asset utilisation and operation and, in the process, minimise both costs and environmental impacts while maintaining system reliability, resilience and stability.*

Smart grid technologies represent an immense saving potential over copper plating solutions but initial investments are significant [18]. Ideally, a high degree of electrification is coupled with a high penetration of smart devices. Take for example a household with solar panels, a heat pump, and an electric vehicle where all devices are connected through a smart meter. An automatic, smart, control algorithm may plan the charge period of the electric vehicle and pre-heating of the heat pump around the expected peak solar output to help flatten the power profile. A user can provide flexibility to the algorithm, e.g. by specifying that the vehicle does not need a full charge as soon as possible or by specifying a desired room temperature interval instead of a single point for the heat pump. In the context of smart grids, such *flexible* devices are referred to as *distributed energy resources* (DERs). DERs provide increased freedom of control, i.e. flexibility, to a coordination algorithm. Because of the demand side operation of DERs, smart grids constitute a shift from "generation-follows-demand" to "demand-follows-generation". To enable such a shift, sizable investments are required in smart devices, monitoring hardware, and control software.

Efforts to control the demand profile are often denoted by *demand response* (DR) or *demand side management* (DSM). In this work, any approach that aims to provide safe and effective management of DERs is referred to as an *energy management* (EM) approach (see Section 2.3). This term reflects the fact that both demand and supply have an important role to play. In the context of smart grids, EM describes the overarching (ICT) technique that handles the coordination. By proper leveraging of EM techniques, balancing of the grid may still be achieved at high penetrations of RES. It is mainly this potential that motivates the growing volume of smart grid related research and applications [21].

2.2.2 Smart devices

In the previous subsection, a definition of smart grids was given and the concept of DERs, or smart devices, was introduced. As explained, DERs play a central role in the smart grid paradigm by providing *flexibility* to an EM approach. Currently, many such devices exist that can be categorised into three main classes [45, 46]. These classes are described below together with potential scheduling decisions that an EM approach could make for the specific classes:

- **Uncontrollable** This class of devices encompasses all "conventional" devices that cannot be controlled and therefore do not offer any measure of flexibility. Typical examples of uncontrollable devices are televisions, water cookers, and lighting. In current EM approaches, these devices together make up the static loads around which flexible devices need to be scheduled. This implies that no scheduling decisions are made for this class of device. A subclass of uncontrollable devices are the *curtailable devices*. Curtailable devices are typically power generating, having an uncontrollable input power process, and offer flexibility in the option to reduce the power output. For some of these devices, it may be possible to linearly reduce the power, which means that the fraction of curtailment can be chosen freely. Other smart devices offer curtailment in several discrete steps of which binary curtailment is a special case. It is important to note that maximum power generation is still an uncontrollable process. The most prevalent curtailable DER is the solar (PV) cell, it converts (sun)light into DC electrical power and through an inverter converts this DC into AC for residential use and potential grid feed-in. The available curtailment options differ between PV units. Scheduling decisions consist therefore of specifying (partial) curtailment over time.
- **Time-shiftable** Time-shiftable devices are DERs that provide flexibility by having a fixed power profile that may be shifted in time. Examples include smart washing machines, automatic vacuum cleaners, and pool pumps. The flexibility of such devices is typically characterised by a time window in which the device needs to schedule its "programme". In the case of the smart washing machine, a user may provide flexibility by choosing a latest time at which the programme needs be finished. Any additional time in this interval on top of the fixed programme time provides flexibility for shifting in time. Some time-shiftable devices increase this flexibility by having the option to pause the programme during operation.

Scheduling decisions then consist of setting the starting time of the programme and possible pausing intervals when this flexibility is available.

- **Buffer** The buffer class consists of two subclasses: pure storage devices such as electric and heat batteries, and devices that require their internal storage for primary functions such as EVs and heat pumps. The name buffer refers to the buffering capacity that such devices may have in a grid. For instance, a battery may absorb part of the feed-in power of a PV cell during peak solar irradiation and redistribute it to a consuming device when production is low. This allows buffer class devices to decouple energy consumption/production of other devices and their output to the grid. Buffer devices have constraints on the limits of their internal storage level and on their maximum ramping rates (charge and discharge in the case of a battery). Devices that require their internal storage for core functions typically have constraints that require them to be charged above a certain point at a user specified time (e.g. departure to work for EVs). Scheduling decisions for this class of devices consists of what amount of power to charge/discharge at different time periods while respecting the given constraints.

2.3 Energy management

Energy management is used to describe various approaches that take care of the operation of smart grids. In this section, energy management is first introduced formally in Subsection 2.3.1. In subsection 2.3.2, some current energy management challenges are outlined. Finally, in Subsection 2.3.3, related work is described in the field of energy management.

2.3.1 Definition and context

In the previous sections, some elements are described that relate to the planning and scheduling within electricity grids. Unit commitment, economic dispatch, and optimal flow problems were mentioned briefly in Subsection 2.1.2 to describe the different problem structures in the electricity market. Energy management (EM) was introduced in Subsection 2.2.1 as an overarching term to describe any approach that aims to provide safe and effective management of DERs. The following, slightly more general definition is used in [45] and [21] and adapted to the current work:

Definition 2. *An EM approach takes care of the complete planning and operation of devices in the grid. It may simultaneously encompass the challenge of solving unit commitment problems, economic dispatch problems, and optimal flow problems while using flexibility from DERs to obtain a solution/schedule.*

In order to create schedules using EM, predictions and a control structure need to be in place. The accuracy of predictions and the nature of these control structures may differ between approaches. In this subsection, it is assumed that accurate predictions and a valid control structure are given.

A simple example of an EM problem is the classical UCP wherein the production planning of a set of power plants is optimised over a certain time horizon (see e.g. [28] for a mathematical description). In this approach, user demands are predicted beforehand and control is exerted directly at the generators. A more general situation can be described by modelling a set of DERs instead of power plants. To give an idea of the structure of a general EM problem, a brief and general mathematical model is introduced here. Consider a time horizon $\mathcal{T}$ that is discretised into a finite number of N time slots with a fixed duration τ. Furthermore, let $\mathcal{I}$ describe a set of M DERs. For each device $i \in \mathcal{I}$, schedules are described in the form of vectors $\boldsymbol{x}_i \in \mathbb{R}^N$ of power values. This implies that power values are assumed to be constant for a given time slot. The complete schedule is then given by the set $\boldsymbol{x} = \{\boldsymbol{x}_i : \forall i \in \mathcal{I}\}$. The system goal is described in the form of a function $f(\boldsymbol{x})$ that indicates to what extent the chosen schedules achieve this goal. A typical example is the 2-norm of the sum of all schedules, i.e. $f(\boldsymbol{x}) = \|\sum_{i\in\mathcal{I}} \boldsymbol{x}_i\|$. Minimising this function reflects the system goal of flattening the aggregated power schedule. Constraints on the schedule x can be divided into *local constraints* and *system constraints* (also referred to as grid constraints). The local constraints are specified through an allowed set X^i of schedules for each device $i \in \mathcal{I}$. These constraints limit the schedules of DERs to reflect realistic device behaviour. For instance, a household battery will have constraints that specify the minimum and maximum state of charge (SoC), and mimimum and maximum charging rate. System constraints describe limits on the behaviour on an aggregated scale. These may, for instance, include line-ratings and balancing requirements at a transformer. The system constraints are described in the form of an allowed set of states X for the complete dispatch $\boldsymbol{x}$.

Using the above structure and notation, the most general form of the *EM scheduling problem* (EMS) is then described in the following way [45]:

$$\min_{\boldsymbol{x}} \quad f(\boldsymbol{x}) \tag{1a}$$
$$\text{s.t.} \quad \boldsymbol{x}_i \in X^i \quad \forall i \in \mathcal{I} \tag{1b}$$
$$\boldsymbol{x} \in X \tag{1c}$$

The optimisation of the system goal is given in (1a). The local constraints are given in (1b) and the system constraint is given in (1c). The objective function $f(\boldsymbol{x})$ is interpreted as a cost function and minimised. The constraints may be integer, which means that the EMS problem is, in general, NP-hard.

Note that the general description of the EMS in (1) to some extent describes the EM problem structure but not necessarily the control or communication structures. As is described in Section 2.3.2, some aspects of the physical problem may restrict these control and communication structures. As an example, an EM approach that has a centralised decision maker with full control leads typically to computation times with exponential complexity and quickly becomes intractable for real-life situations [44]. The class of decentralised EM approaches has the potential to alleviate some of these problems using different communication and control structures. Some of these approaches are outlined in Section 2.3.3. For an in-depth discussion of centralised versus decentralised EM, the reader is referred to Section 3.3.

2.3.2 Challenges

The system goal for EM approaches may differ between implementations. Recently, EM approaches have mostly focussed on finding schedules that are economically attractive while simultaneously facilitating a high penetration of renewable energy sources. Especially the second part of this objective has become increasingly challenging with the current trends of electrification and decentralisation of generation (see Subsection 2.1.3). The EM system goal is becoming more complex as an increasing number of stakeholders become involved. Currently, objectives for EM approaches include: balancing, congestion mitigation, minimising transport losses, maximising self-consumption etc. An EM approach needs to consider different time-scales, varying from seasonal to sub-second scales, where control strategies may differ immensely. In the pursuit of these goals lies an immense coordination and control challenge. Specifically at the distribution level where many heterogeneous clusters of DERs and ordinary devices are concentrated, complicated problems arise. Several main themes can be distinguished in the coordination problem:

- **Flexibility:** Flexibility is provided by users through DERs. Flexibility as a resource for EM leads to a number of issues. The most important challenge is social: device flexibility may be plentiful, but lack of user flexibility may still restrict control. In general, DERs may have a preconfigured flexibility that can always be (instantly) overridden by the user. For example, a user may have expressed a need for the smart washing machine to run its programme within a specified time interval. This data is sent upward where it is used by the EM algorithm. However, if the user decides to prematurely load and start the machine, the EM algorithm will not be able to initiate the cycle as planned. To maximise available flexibility, a balance needs to be found between device and user autonomy while guaranteeing comfort [44].

- **Uncertainty:** Part of the EM challenge stems from the uncertainty typically found in systems with high penetration of flexible devices. Smart devices have significant uncertainties in production profile for RES, and demand profile for DERs. These uncertainties are partially systematic (e.g. in the prediction of intermittent energy sources) and partially social (e.g. in the unscheduled behaviour of users). Uncertainty is especially relevant when looking at local behaviour where variances often are at their highest. When aggregating multiple households, variances cancel each other out in accordance with the law of large numbers; this is referred to as *load diversification* [45]. However, most power lines have been optimised for usage that was typical for the "pre energy transition era". Because of the trends introduced in Subection 2.1.3, their tolerances are not sufficient anymore. To mitigate this, coordination needs to consider the full scope of the distribution grid and exert control locally. This motivates the use of DSM and the merit of smart grids in general. For an in-depth discussion of uncertainty in the distribution grid, the reader is referred to Chapter 4.
- **Physical robustness:** System operators have the responsibility to ensure physical resilience of the grid. This is embodied by the $n-1$ *criterion*: if any grid resource fails, this should never lead to a system failure [4]. Satisfying this requirement is becoming more challenging as smart devices increase the variance and uncertainty of power demands significantly, as was explained in the previous paragraph. To ensure physical robustness, the electrical system should be as independent, i.e. self-reliant, as possible.
- **Privacy and cyber-security:** Increasing the amount of smart devices in distribution grids results in a large volume of user data that is analysed for control purposes. Users may have objections about the handling and storage of this data, dissuading them from installing smart devices. Furthermore, any ICT system may be vulnerable to cyber-attacks that could potentially leak this data. Sensitive information about users may be inferred such as the occupancy of the home over time. In the opposite direction, control signals are sent down to homes that may similarly be intercepted, potentially leading to unwanted device behaviour and congestion. To prevent such security breaches, EM systems need to be resilient not only in a physical sense, but also in the context of cyber-security. As a consequence, it is desirable to analyse data as close to source as possible to minimise these risks [45].
- **Scalability of optimisation:** Large amounts of all sort of smart devices, each with user specific flexibility constraints and uncertainties lead to a complex optimisation problem. Compared to more traditional EM approaches such as the unit commitment problems (UCPs) mentioned in Subsection 2.1.2, this optimisation problem has a high level of heterogeneity (in size, constraints, granularity, etc.). The resulting mathematical programmes have a high level of complexity and therefore lack scalability, i.e. increased problem size quickly leads to intractable models. Especially in medium time-scales where *operational planning* relies on many predictions of RES conditions and user behaviours, solving large-scale central optimisation problems quickly becomes unpractical. Decentralised EM approaches may be used to solve a set of smaller optimisation problems locally to improve scalibility. Heuristics are employed to further decrease solution times.

2.3.3 Related work

A large body of work exists on the topic of EM. One of the very first discussions on the topic was published in the 1980s that foresaw aspects of the energy transition developments that are only now unfolding [36]. In the study, it is argued that economic principles can be used to balance supply and demand in the electricity grid. A methodology is introduced that is very similar to the balancing strategies employed by current day smart devices. Since 1980, many extensions have been made to this concept. A rapid increase in the volume of publication has only been observed in the last 20 or so years with the emergence of enabling technologies and the accelerating progression of the energy transition. This trend is visualised in Figure 3. An overview of the many EM related publications is given in a number of recent surveys (see for instance [47, 3, 5, 13, 23]).

In the overview paper [26], the different EM approaches are characterised in one of four clusters

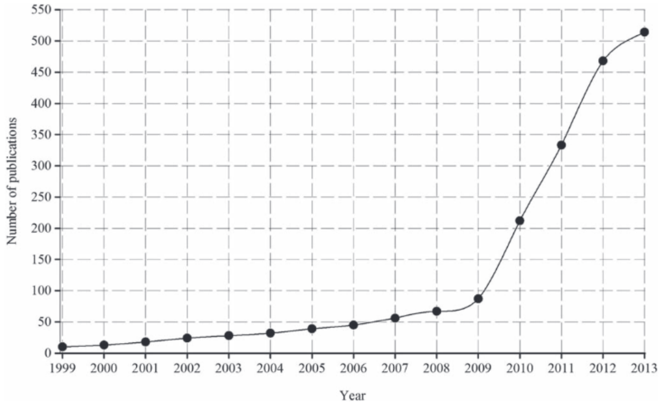

FIGURE 3: *Trend of the amount of published EM optimisation models between 1999 and 2013 [47].*

Decisions on local issues made locally	**Price Reaction** – No (certain) global optimum + Full use of response potential – Uncertain system reaction – Market Inefficiency + No privacy issues	**Transactive Control** – No (certain) global optimum + Full use of response potential + Certain system reaction + Efficiency market + No privacy issues
Decisions on local issues made centrally	**Top-Down Switching** – No (certain) global optimum – Partial use of response potential – Uncertain system reaction – Autonomy issues + No privacy issues	**Centralised optimisation** + Certain global optimum + Full use of response potential + Certain system reaction – Privacy and autonomy issues – Low scalability
	One-way communications	Two-way communications

TABLE 1: *The energy management matrix: comparison of four different energy management approach with degree of decentralisation on the y-axis and communication approach on the x-axis. Adapted from [26].*

based on the used decision-making approach (centralised versus decentralised) and the communication approach (one-way versus two-way) . An *energy management matrix* is introduced and adapted here (see Table 1) where the different EM clusters are judged on a number of metrics. The four methods are outlined below:

- **Top-down switching:** As seen in Section 2.1, earlier EM approaches tended to have all decisions made centrally in accordance with the structure of the grid. The oldest out of these approaches is *top-down switching*. In top-down switching, a signal may be broadcasted by a utility company when demand exceeds a critical level prompting a connection or even an entire grid area to switch off. Although simple and effective, top-down switching overwrites user autonomy completely ignoring user comfort and does not deal well with distributed generation.
- **Centralised optimisation:** The second approach is *centralised optimisation*. Compared to top-down switching, *centralised optimisation* introduces two-way communication within EM. Decisions are still made centrally, but local data is used as much as possible to aid in finding an optimal dispatch. At the top, a complex optimisation engine oversees all (smart) RES and DERs in the relevant grid cluster. All available local data is communicated to this central controller and, accounting for local and/or global control goals as described in Section 2.3.2, an optimal solution is sought. This optimal solution, or dispatch, is communicated back down to

the devices where control is exerted directly. With complete access to local data and complete control over the smart devices, centralised optimisation is able to make full use of the smart grid response potential. The main disadvantage lies in the scalability of the approach; because the optimisation is done centrally accounting for all device states, the size of the problem may quickly become intractable with increased number of devices. Furthermore, as with top-down switching, user autonomy is not upheld completely. Lastly, a privacy issue and additional security risk is incurred as detailed local information is communicated all the way to the top level.

- **Price reaction:** Recently, decentralised EM approaches have gained popularity. The trend towards decentralised generation combined with the ICT developments have paved the way for more intricate EM systems. In *price-reaction* approaches, a retailer communicates prices that differ over time in order to achieve some control goal, this is referred to as *dynamic pricing* (see for instance [50, 43]). Prices vary over discrete time intervals that typically have a length of 15 minutes or one hour. Dynamic pricing is used to steer users or devices towards more desirable behaviour, e.g. balancing and peak reduction. With recent improvements and the spread of smart-metering, *direct load control* has become available on a residential scale [16, 37]. In direct load control, users can be represented by automated *agents* that react to price signals in real-time. Because of this agent interaction, decentralised EM approaches are often referred to as *multi-agent systems* in the literature. Due to the local nature of the decision making, price reaction methods are decentralised EM approaches. Advantages of price reaction approaches are the simple communication structure and robustness with respect to privacy and cyber-security issues. In [3], EM benefits for users are outlined and an overview is given of the different approaches that have been studied on this topic. Disadvantages of price reaction methods are addressed in the survey paper [7] where the authors show that many such approaches only shift peaks instead of reducing them. The authors further show that intervals of highest consumption are simply moved to time periods with the lowest price. Moreover, power quality issues may arise from synchronous user behaviour resulting from identical incentives.

- **Transactive control:** By combining two-way communication and decentralised EM, it is shown in [26] that *transactive control* offers distinct advantages over the previous approaches. Similarly to price reaction approaches, direct load control is used where agents represent flexible devices *locally*. Local agents have the option to communicate available flexibility and demand preferences based on dynamic pricing communications from the retailer. Both parties can then interact until a stable solution is found. Compared to price reaction approaches, the pool reaction to price signals is known because of the additional communication layer. In [51, 2], it is shown that transactive control approaches are capable of using the full flexibility potential of DERs and result in schedules that are globally optimal under a, quite restrictive, set of conditions. Autonomy is preserved as agents may be configured to represent user preferences. Furthermore, transactive control offers distinct advantages in robustness, cyber-security, and scalability [31, 10]. In [25], transactive control techniques are used to interact in an instantaneous electricity market. Major drawbacks of this approach are the limited incorporation of future predictions. Another contribution is made in [49], where transactive control is applied to a day-ahead market using a *receding horizon scheme* (see Section 4.3 for more information) allowing the decision maker to dynamically update decisions with the most recent future predictions. A market operator aggregates power profiles and changes prices using a gradient descent approach in order to optimise the dispatch. Furthermore, local pricing is introduced as a way to further increase stability and mitigate congestion locally. However, the approach is applied to the transmission level whereas most congestion problems are expected to arise at the distribution level. A similar transactive control approach was introduced in [9].

Some EM approaches do not strictly fall into any of the above four categories. For example, in [45], a market operator sends preferred profiles to device nodes instead of price signals. In this approach, called “*profile steering*”, agents do not receive monetary incentives but are instead steered towards a cluster goal which corresponds to balancing the grid. Authors show that this approach does not have the instability and synchronisation issues that are typically associated with using price signals. A disadvantage of this approach is that user incentives are not as evident as for transactive control approaches. This may result in a lack of motivation for users to sacrifice private rewards for adapting their power demand. In [19], an approach is proposed based on *DCOPs* (Distributed Constraint Optimisation Problems) to solve the economic dispatch problem using a combination of decentralised optimisation and dynamic programming. The resulting approach finds optimal solutions but scales very poorly.

3 Centralised versus decentralised control

In this chapter, an EM problem is considered for a distribution network with a high penetration of smart devices. The problem is first described in detail in Section 3.1. Two optimisation models are then introduced in Section 3.2, one using a centralised and one using a decentralised approach. The comparison between these two models plays a central role in this chapter. Therefore, a general description of the common structure between the two models is first introduced in Section 3.2.1. The features that are specific to the centralised and decentralised models are introduced in Section 3.2.2, and 3.2.3 respectively. Finally, an in-depth discussion on the differences between centralised and decentralised modelling is presented in Section 3.3.

3.1 Problem statement

In this section, we give a detailed description of the concrete problem that is treated in this chapter. The problem at hand is a combination of the unit commitment (UC), economic dispatch (ED), and optimal power flow (OPF) problems and is a specific case of the EMS problem (see (1) in Section 2.3.1). The operation of flexible devices has to be scheduled in a way that minimises some cost function while balancing production and demand, and respecting local and system constraints. In the following, this problem is referred to as "the dispatch problem" for brevity. Historically, the dispatch problem included controllable generators and static loads and was almost exclusively solved on the TSO level in a regulated market (see Section 2.1.2). However, the current energy system features intermittent generation (RES), flexible devices (DERs), and the possibility to control these resources at the distribution level (DSM). This increased complexity combined with the current deregulated electricity market has resulted in a complicated situation with many stakeholders and involved parties. Arguably the most interesting case is encountered in the distribution grid (LV) where users with flexible resources interact with retailers and system operators. These LV grids are expected to face many challenges in the (near) future with the ongoing energy transition (see Section 2.1.3).

In an effort to confront some of these challenges, in this work a dispatch problem is considered for a neighbourhood LV grid consisting of a number of households, each equipped with a set of flexible devices. This problem is simulated in the context of a *day-ahead market*. Because of the multi-stakeholder aspect and the involved flexibility, the current problem does not have the single goal of minimising generator costs as with the historical problem. DSOs require a solution that respects all grid constraints and that does not lead to unnecessary wear or even damage to the grid. Retailers, (external) producers, and TSOs need to have a reliable prediction of the power demand (target demand profile) to be able to interact efficiently in the day-ahead market. Furthermore, the underlying sustainability concerns have to be considered, meaning that carbon emissions should be minimised. Finally and most importantly, users need an incentive for offering flexibility to an outside controller. This last objective can be achieved most easily and directly by ensuring some financial gain [26, 5]. Accounting for all of these aspects, the problem statement can be described as follows:

> "*To find a dispatch that follows a target demand profile in a grid feasible way while simultaneously being economic and sustainable*"

With this goal in mind, two distinct EM approaches are presented in the following subsections. First, a *centralised optimisation* approach is introduced in Subsection 3.2.2 that serves as an ideal *benchmark* for the dispatch problem. Then in Subsection 3.2.3, a novel *decentralised optimisation* approach is presented that uses transactive control strategies to solve the problem.

3.2 Model descriptions

In this section, the mathematical models for the considered problems are introduced. First, the general structure that is identical for the models is described. This forms the basis of the analysis and gives a framework for developing the specific features for the individual models. The centralised and decentralised approaches are then considered respectively.

3.2.1 General model description

The models that will be developed are both used to solve the same dispatch problem, so naturally they share some common features. Most importantly, the network structure of the grid, and all constraints and decision variables should be the same. These factors determine the feasible set $\mathcal{F}$ of the problem and should be compatible to make a meaningful comparison between the different models. Furthermore, time discretisation, and some general parameters and specifications are mentioned here.

NETWORK STRUCTURE

The most obvious shared feature between the two models is the structure of the considered grid. Mathematically, power grids can be modelled as a *rooted tree*, which is an acyclic, connected, undirected graph. Note that in a rooted tree, each node except for the root node has exactly one parent node. Let the LV grid be described by a graph G using the conventional graph notation $G = (V, E)$, where V is the set of *nodes*, and E is the set of connecting *edges*. In the tree model of a LV network, three different types of node are given:

- $\mathcal{I} \subset V$ The leaves of the tree. These represent the devices connected to the grid and are called *device nodes*, whereby each individual device has its own leaf in the tree. Devices are the only source/sink of power in the grid.
- $\mathcal{J} \subset V$ The inner nodes of the tree. These nodes correspond to locations in the grid where congestion is deemed most likely to occur and may represent e.g. field stations or specific cable segments.
- $v_{mo} \in V$ The root node of the tree. This node contains the transformer that connects the LV network to the transmission network above it and is called the *market operator*.

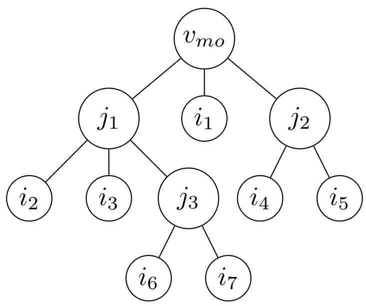

FIGURE 4: *Schematic example of possible graph G.*

In the modelled problem, *flexible* power demand of devices is considered as the main *decision variable*. Devices are the only nodes that directly consume or produce power, and they are represented by the set of leaves $\mathcal{I}$. These devices are connected to inner nodes or the market operator node via *edges*. The edge set E describes hierarchical relations, and does not directly specify a physical part of the system. Therefore, edges do not have further characteristics like capacity. Instead, grid capacity is modelled by a set of nodes $\mathcal{J}$ called the congestion nodes. They are called congestion nodes because they correspond to “capacity bottlenecks” in the network where congestion is most likely to occur. These *inner nodes* of the tree may have parent nodes in the set $v_{mo} \cup \mathcal{J}$ and child nodes in the set $\mathcal{J} \cup \mathcal{I}$. Inner nodes may be seen as concentrator nodes because they somehow represent the power demands of all their child nodes. At the root of the LV network lies the transformer v_{mo} that connects the neighbourhood distribution grid to the transmission grid. It is referred to as the *market operator* because it forms the connection with the outside market. The market operator can also be seen as an concentrator, taking the power demand of all its children as input. To model this process, the notation $D(\cdot)$ is introduced. For each parent node $v \in v_{mo} \cup \mathcal{J}$, $D(v)$ gives the set of all direct child nodes. Furthermore, $P(\cdot)$ gives the parent node $P(v)$ for each child node or inner node $v \in \mathcal{I} \cup \mathcal{J}$. An example of a simple LV graph is given in Figure 4 including the market operator (v_{mo}), three congestion nodes ($j_1, j_2, j_3 \in \mathcal{J}$), and seven device nodes ($i_1, \ldots, i_7 \in \mathcal{I}$).

As was mentioned in Section 3.1, it is our aim to model a day-ahead market. This means that power demands are not decided instantaneously but in a day-ahead fashion. In the day-ahead market a schedule has to be specified, the *power profile*, that spans a time horizon of exactly 24 hours and describes the power demand/supply values during this horizon. To model this, the time horizon is discretised into a finite number of time slots. Let N be the number of such time slots, and let $\mathcal{T} = \{t_1, t_2, \ldots, t_N\}$ describe the ordered set of time intervals for the complete horizon. In the context of electricity trading, these time slots are referred to as *programme time units* (PTUs), and typically span a fixed time interval τ of one hour or fifteen minutes [21]. With the chosen discretisation, the power profile of a node $v \in V$ may then be defined by a vector $\boldsymbol{p}_v \in \mathbb{R}^N$ with elements p_{vt} for each $t \in \mathcal{T}$ representing an *average* power demand value for each time slot t. To emphasise this distinction, vectors in the following will be described by boldface symbols. The complete dispatch may now be represented by the power output matrix $\boldsymbol{p} \in \mathbb{R}^{|V| \times N}$, wherein the power profile for each node is given at every time slot.

DEVICE NODES

Device nodes represent the devices and are modelled to reflect their behaviour. For each type of device node, the relevant modelling features are explained here. For each device $i \in \mathcal{I}$ a decision $\boldsymbol{x}_i$ has to be made, and this decision leads to an output power profile $\boldsymbol{p}_i$. The allowed values of $\boldsymbol{x}_i$ are limited by the accompanying *device constraints* and the constraints on the realised power profile. For each device category, the input/output relation is first described followed by the relevant device constraints:

- **Static loads**: Static loads, given by the set $\mathcal{I}_{load} \subset \mathcal{I}$, represent the summed power demand of all inflexible loads (e.g. television, lights, etc.) for a given household and are modelled as a single *uncontrollable* device node. Note that static loads may also model other inflexible loads such as street lamps and traffic lights. A load profile $\boldsymbol{\chi}_i$ is taken from a data set for every load device node $i \in \mathcal{I}_{load}$ (more about the data sets in Section 4.2). Because the static load device is uncontrollable, no decisions have to be taken and its output power profiles $\boldsymbol{p_i}$ for $i \in \mathcal{I}_{load}$ are equal to the input data as defined by

$$\boldsymbol{p}_i := \boldsymbol{\chi}_i, \qquad \forall i \in \mathcal{I}_{load}. \tag{2}$$

- **Solar cells**: Solar cells, or PVs for short, form the set $\mathcal{I}_{PV} \subset \mathcal{I}$. A power generation profile $\boldsymbol{\xi}_i$ is taken from a data set and assigned to a PV node $i \in \mathcal{I}_{PV}$ in the model. Solar cells are *curtailable* devices, i.e. the decision maker has the choice to completely turn off the device for each time slot $t \in \mathcal{T}$. This curtailment is characterised by a binary decision variable $x_{it} \in \{0, 1\}$. The output power profile is a combination of the input generation data profile $\boldsymbol{\xi}_i$ and this decision variable, defined by

$$p_{it} := \xi_{it} \cdot x_{it}, \qquad \forall i \in \mathcal{I}_{PV}, t \in \mathcal{T}. \tag{3}$$

  Note that (3) expresses that curtailment may be specified independently for each individual PTU.

- **Batteries**: Batteries, specified in the set $\mathcal{I}_{bat} \subset \mathcal{I}$, are able to charge and discharge up to their power limits; i.e., they are very *flexible* devices. The storage level or SoC s_{it} represents the *energy* contained in battery i at time slot t and results from the charge and discharge decisions since the first period and its starting level s_{it_0}. The decision variable x_{it} for $i \in \mathcal{I}_{bat}$ and $t \in \mathcal{T}$ describes the power that is "experienced" by battery i in time slot t. When multiplied with the time duration τ of the time slot, this is equivalent to the increment in SoC during a time period. An energy conversion efficiency of $\eta_i \in [0, 1]$ is assumed specific to each battery. This means that, when discharging, a decrease of the SoC of $x_{it} \cdot \tau$ is accompanied by a power output p_{it} of $\eta_i \cdot x_{it}$. When charging, an increase of the SoC of $x_{it} \cdot \tau$ requires a power input p_{it} of $\frac{1}{\eta_i} \cdot x_{it}$ as described by

$$p_{it} := \begin{cases} \eta_i \cdot x_{it}, & \text{for } x_{it} < 0, \\ \dfrac{1}{\eta_i} \cdot x_{it}, & \text{for } x_{it} \geq 0. \end{cases} \qquad \forall i \in \mathcal{I}, t \in \mathcal{T} \tag{4}$$

The decision variable x_{it} thus represents the charge/discharge rate of battery i over time period t. It is limited by the (negative) minimum discharge rate X_i^{min} and (positive) maximum charge rate X_i^{max} that are specific to the battery. These constraints can be described by

$$x_{it} \in [X_i^{min}, X_i^{max}], \qquad \forall i \in \mathcal{I}_{bat}, t \in \mathcal{T}. \tag{5}$$

In addition to these bounds on charge and discharge rate, there exist bounds on the maximum and minimum SoC of the batteries. Based on the above, the SoC s_{it} is given by

$$s_{it} := s_{it_0} + \sum_{t'=t_0}^{t'=t} x_{it'} \cdot \tau, \qquad \forall i \in \mathcal{I}_{bat}, t \in \mathcal{T}. \tag{6}$$

The SoC of the battery is further limited by the minimum SoC S_i^{min} and the maximum SoC S_i^{max}, i.e. we must have

$$s_{it} \in [S_i^{min}, S_i^{max}], \qquad \forall i \in \mathcal{I}_{bat}, t \in \mathcal{T}. \tag{7}$$

Without loss of generality, S_i^{min} is assumed to be 0 for all batteries.

- **Heat pumps**: Heat pumps, specified in the set $\mathcal{I}_{HP} \subset \mathcal{I}$ are in many ways similar to batteries. A heat pump creates a certain "heat storage level" s_{it} in a building/household that is comparable to the SoC of a battery. The user can use the heat pump to "charge", which corresponds to injecting heat into the building/household. This heat transfer is assumed to be lossless, but a constant power leakage rate of β_i is modelled specifically for each household. The definition for the power output is then given by

$$p_{it} := x_{it} - \beta_i, \qquad \forall i \in \mathcal{I}_{HP}, t \in \mathcal{T} \tag{8}$$

The heat storage level of the heat pump is given by

$$s_{it} := s_{it_0} + \sum_{t'=t_0}^{t'=t} \big(x_{it'} - \beta_i\big) \cdot \tau, \qquad \forall i \in \mathcal{I}_{HP}, t \in \mathcal{T} \tag{9}$$

The possible heat transfer x_{it} has a lower bound $X_i^{min} \geq 0$ and an upper bound of X_i^{max} given by

$$x_{it} \in [X_i^{min}, X_i^{max}], \qquad \forall i \in \mathcal{I}_{HP}, t \in \mathcal{T}. \tag{10}$$

Because the heat pump can only be charged and not discharged, we have that $X_i^{min} = 0$. The heat storage level of the heat pump is limited by lower bound S_i^{min} and upper bound S_i^{max} by

$$s_{it} \in [S_i^{min}, S_i^{max}], \qquad \forall i \in \mathcal{I}_{HP}, t \in \mathcal{T}, \tag{11}$$

where $S_i^{min} = 0$ can again be assumed without loss of generality.

CONGESTION NODES

As was mentioned in the beginning of this subsection, grid capacity is not modelled through line ratings but instead at the congestion nodes $j \in \mathcal{J}$ where upper and lower bounds on the power throughput are imposed. The congestion nodes are also seen as concentrators, as they sum the power profiles of the child nodes directly below them through

$$\boldsymbol{p}_j = \sum_{v \in D(j)} \boldsymbol{p}_v, \qquad j \in \mathcal{J}. \tag{12}$$

The *grid constraints* include all limitations that are imposed on the dispatch by capacity limits of grid assets and connections. These are modelled at the congestion nodes as constraints on the summed power profiles. More specifically, the summed power profiles $\boldsymbol{p}_j$ are subject to bounds P_{jt}^{min} and P_{jt}^{max} that are specific to each congestion node and may differ over time. Each element p_{jt} of the congestion node power profile must respect these bounds and the resulting grid constraints are thus defined by

$$p_{jt} \in [P_{jt}^{min}, P_{jt}^{max}], \qquad \forall j \in \mathcal{J}, t \in \mathcal{T}. \tag{13}$$

MARKET OPERATOR

Finally, the central *market operator* deals with the connection to the outside market at the market operator node v_{mo}. Analogously to the congestion nodes, the market operator sums its underlying power profiles. The summed power profile at the market operator is described by

$$\boldsymbol{p}_{v_{mo}} = \sum_{v \in D(v_{mo})} \boldsymbol{p}_v. \tag{14}$$

Note that $\boldsymbol{p}_{v_{mo}}$ also matches exactly the sum of the power profiles of all device nodes $\mathcal{I}$ since these represent the only sources/sinks of power and no loss is modelled. In other words, all power production/consumption happens at the devices and the congestion and market operator nodes only serve to sum this information, apply the given bound constraints, and send it, if feasible, upwards. The power profile at the market operator might therefore equally well be written as

$$\boldsymbol{p}_{v_{mo}} = \sum_{i \in \mathcal{I}} \boldsymbol{p}_i \tag{15}$$

At the market operator, the *market constraints* are evaluated. These constraints describe how well the achieved profile $\boldsymbol{p}_{v_{mo}}$ matches a given target profile $\boldsymbol{\Theta}$. In this study, we assume that this is modelled by an absolute error tolerance of α, i.e. we must have

$$p_{v_{mo}t} \in [\Theta_t - \alpha, \Theta_t + \alpha], \qquad \forall t \in \mathcal{T}. \tag{16}$$

SUMMARY

In this subsection, the general model was introduced, i.e. all features that are identical for the centralised and the decentralised approach. This included the physical structure of the grid: device nodes, congestion nodes, the market operator node, and the connections between them. Furthermore, the behaviour and constraints at each of the nodes have been described. These constraints together bound the feasible set $\mathcal{F}$ of the dispatch problem. This feasible set $\mathcal{F}$ is part of the general model structure and therefore completely identical for the different modelling approaches. Referring back to the problem description of Section 3.1, the constraints take care of "*a dispatch that follows a target demand profile in a grid feasible way*" part of the goal. What has not been addressed yet is the "*while simultaneously being economic and sustainable*" part of the goal. It turns out that the centralised and decentralised models approach this in a fundamentally different way as will become apparent in Sections 3.2.2 and 3.2.3. An in depth discussion about this distinction is included in Section 3.3.

3.2.2 Centralised optimisation model

The centralised optimisation model introduced here aims to solve the dispatch problem described in Section 3.1. The model is part of the *centralised optimisation* cluster of EM methods as classified in [26] and described in 2.3.3. All available information about each device and grid asset, i.e. about each node in the rooted tree, is assumed to be known perfectly in the centralised optimisation problem and used by a single controller to find a dispatch. This dispatch is subject to the full set of constraints that together bound the feasible set $\mathcal{F}$ as was explained in Section 3.2.1. In the centralised case, the optimisation consists of minimising a given global cost function that somehow represents the goal: to find a dispatch that is "*economic and sustainable*".

We denote the global cost function by $f_{cost}(\boldsymbol{x}) : \mathbb{R}^{|\mathcal{I}|\times N} \mapsto \mathbb{R}$. Using this, a more formal definition of the optimisation problem can be obtained. The goal of the centralised optimisation is to minimise f_{cost}, i.e. to find a decision vector $\boldsymbol{x}^* \in \mathcal{F}$ with $f_{cost}(\boldsymbol{x}^*) = \min_{\boldsymbol{x}\in\mathcal{F}} f_{cost}(\boldsymbol{x})$. Such a dispatch is said to be *globally optimal* in the specified time horizon.

The set of equations (2) - (16) are all reformulated as linear and integer constraints. Depending on the nature of the function f_{cost}, the resulting model may either be a *mixed integer linear programme* (MILP/MIP), a *mixed integer quadratic programme* (MIQP), or even a *mixed integer nonlinear programme* (MINLP). A summary of complete mathematical description of the centralised model is given by (17).

$$
\begin{aligned}
\min_{\boldsymbol{x},\boldsymbol{y}} \quad & f_{cost}(\boldsymbol{x},\boldsymbol{y}) && && \text{(17a)}\\
\text{s.t.} \quad & \textit{Device constraints} && && \\
& x^d_{it} \geq y^d_{it} \cdot X^{min}_i, && \forall i \in \mathcal{I}_{bat}, t \in \mathcal{T}, && \text{(17b)}\\
& x^c_{it} \leq y^c_{it} \cdot X^{max}_i, && \forall i \in \mathcal{I}_{bat}, t \in \mathcal{T}, && \text{(17c)}\\
& y^c_{it} + y^d_{it} \leq 1, && \forall i \in \mathcal{I}_{bat}, t \in \mathcal{T}, && \text{(17d)}\\
& x_{it} \geq X^{min}_i, && \forall i \in \mathcal{I}_{HP}, t \in \mathcal{T}, && \text{(17e)}\\
& x_{it} \leq X^{max}_i, && \forall i \in \mathcal{I}_{HP}, t \in \mathcal{T}, && \text{(17f)}\\
& s_{it} \geq S^{min}_i, && \forall i \in \mathcal{I}_{bat} \cup \mathcal{I}_{HP}, t \in \mathcal{T}, && \text{(17g)}\\
& s_{it} \leq S^{max}_i, && \forall i \in \mathcal{I}_{bat} \cup \mathcal{I}_{HP}, t \in \mathcal{T}, && \text{(17h)}\\
& \textit{Congestion constraints} && && \\
& \boldsymbol{p}_{jt} \geq P^{min}_{jt}, && \forall j \in \mathcal{J}, t \in \mathcal{T}, && \text{(17i)}\\
& \boldsymbol{p}_{jt} \leq P^{max}_{jt}, && \forall j \in \mathcal{J}, t \in \mathcal{T}, && \text{(17j)}\\
& \textit{Market constraints} && && \\
& \boldsymbol{p}_{v_{mo}t} \geq \Theta_t - \alpha, && \forall t \in \mathcal{T}, && \text{(17k)}\\
& \boldsymbol{p}_{v_{mo}t} \leq \Theta_t + \alpha, && \forall t \in \mathcal{T}, && \text{(17l)}
\end{aligned}
$$

Further constraints

$$\boldsymbol{p}_i = \boldsymbol{\chi}_i, \qquad \forall i \in \mathcal{I}_{load}, \tag{17m}$$

$$p_{it} = \xi_{it} \cdot x_{it}, \qquad \forall i \in \mathcal{I}_{PV}, t \in \mathcal{T}, \tag{17n}$$

$$p_{it} = \frac{1}{\eta_i} \cdot x_{it}^c + \eta_i \cdot x_{it}^d, \qquad \forall i \in \mathcal{I}_{bat}, t \in \mathcal{T}, \tag{17o}$$

$$s_{it} = s_{it_0} + \sum_{t'=t_0}^{t'=t} \left(x_{it'}^c + x_{it'}^d\right) \cdot \tau, \qquad \forall i \in \mathcal{I}_{bat}, t \in \mathcal{T}, \tag{17p}$$

$$s_{it} = s_{it_0} + \sum_{t'=t_0}^{t'=t} \left(x_{it'} - \beta_i\right) \cdot \tau, \qquad \forall i \in \mathcal{I}_{HP}, t \in \mathcal{T}, \tag{17q}$$

$$\boldsymbol{p}_v = \sum_{v \in D(j)} \boldsymbol{p}_v, \qquad \forall v \in \mathcal{J} \cup v_{mo}, \tag{17r}$$

Domain constraints

$$x_{it} \in \{0,1\}, \qquad \forall i \in \mathcal{I}_{PV}, t \in \mathcal{T}, \tag{17s}$$

$$y_{it}^c \in \{0,1\}, \qquad \forall i \in \mathcal{I}_{bat}, t \in \mathcal{T}, \tag{17t}$$

$$y_{it}^d \in \{0,1\}, \qquad \forall i \in \mathcal{I}_{bat}, t \in \mathcal{T}, \tag{17u}$$

$$x_{it}^d \leq 0, \qquad \forall i \in \mathcal{I}_{bat}, t \in \mathcal{T}, \tag{17v}$$

$$x_{it}^c \geq 0, \qquad \forall i \in \mathcal{I}_{bat}, t \in \mathcal{T}. \tag{17w}$$

In this model, the set of device constraints are given first. For the batteries, equation (4) had to be rewritten in a linear way. To this end, the "charging" decision variable x_{it} needs to be split into a negative discharging variable x_{it}^d and a positive charging variable x_{it}^c. To make sure that the battery cannot charge and discharge at the same time, the binary variables y_{it}^c and y_{it}^d are introduced for $i \in \mathcal{I}_{bat}$. These variables are united into the vector $\boldsymbol{y} = (y_{it}^c, y_{it}^d)$ to simplify notation. Constraint (17b) together with (17v) make sure that when discharging ($y_{it}^d = 1$), the battery limits are respected. When not discharging ($y_{it}^d = 0$), the discharge rate must be equal to zero. Constraints (17c) and (17w) achieve the same thing for the charging rate x_{it}^c. The constraint (17d) makes sure that the battery cannot charge and discharge at the same time, i.e. it must be either charging, discharging, or idle. The heat pump and storage constraints do not require such linearisation and are directly adapted into (17e)-(17h). The congestion constraints are specified in (17i) and (17j), and the market constraints in (17k) and (17l). Finally, the definitions and domain specifications are given in (17n)-(17w).

3.2.3 Decentralised optimisation model

In this subsection, a decentralised optimisation model is introduced that is part of "transactive control" cluster of [26] as described in Subsection 2.3.3. The model is a hierarchical agent-based optimisation model that aims to solve the dispatch problem described in Section 3.1 and obeys all general constraints and features that were outlined in Subsection 3.2.1. The decentralised approach described in this subsection forms the mathematical basis of the *Local pricing receding horizon* (LPRH) model. In this subsection, specific attention is given to the agent structure and "local pricing" part of LPRH, and the "receding horizon" concept is treated in Section 4.3.

In the decentralised optimization, dynamic pricing is used in a way that is very similar to the *profile steering* approach mentioned in Section 2.3.3. The resulting approach is referred to as *price steering* following the terminology of [45]. Price steering is a transactive control approach where prices are used to influence, or *steer*, agents that control devices. From a system perspective, it may also be referred to as a *market based control approach* because there is essentially a price negotiation happening between the agents. An important assumption is that the *device agents* react to price signals in a rational and predictable way. For example, heat pumps are used when prices are relatively low and batteries are discharged when prices are relatively high. This behaviour can be modelled through a device specific cost function that the agent tries to minimise.

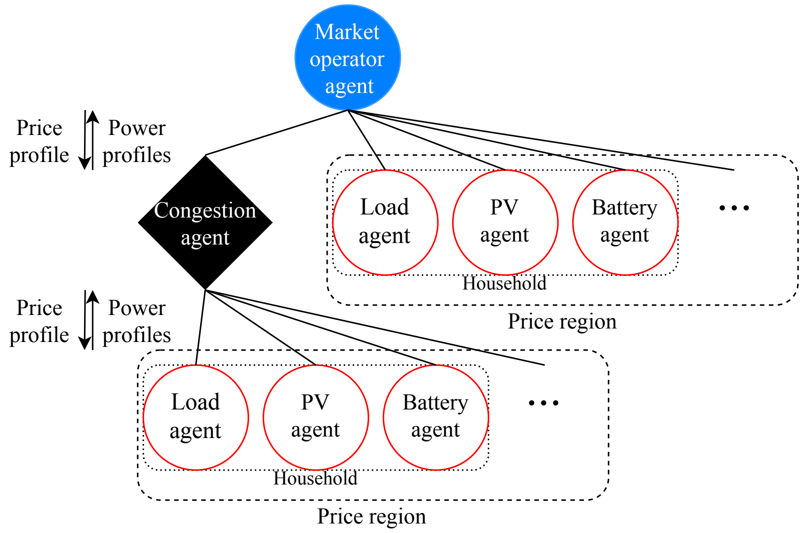

FIGURE 5: *Systematic diagram of the agent and communication structure in the decentralised optimisation. Two households are included, one with a battery installed and another with a heat pump installed. Arrows denote the flows between nodes along the network connections*

An obvious choice for this cost function is the expected monetary cost, but many other aspects (e.g. sustainability related) may also be factored into the cost function. The price signals are given in the form of *price profiles* $\boldsymbol{\lambda}_v \in \mathbb{R}^N$ that consist of electricity prices λ_{vt} for each PTU t in the time horizon. In the case of the decentralised model, these profiles may differ for each congestion agent or the market operator, i.e. $v \in v_{mo} \cup \mathcal{J}$. As such, this is an example of *local pricing*. As is explained in more detail in the next paragraphs, the market operator v_{mo} and congestion nodes $j \in \mathcal{J}$ are respectively represented by a *market operator agent* and *congestion agents*. These "parent agents" v have the ability to choose and transmit a price profile $\boldsymbol{\lambda}_v$ to their set of immediate children $D(j)$. In this way, distinct pricing zones may emerge, i.e. prices are identical within a zone and may differ between these zones. A schematic representation of this situation is included in Figure 5.

AGENT DESCRIPTIONS

Before going into the exact operation of the algorithm, it is helpful to first explain the role of each agent type. All node types that were introduced in Subsection 3.2.1, i.e. device nodes, congestion nodes, and the market operator node, are assigned a specific agent type (see also Figure 5 for a simple schematic representation). The following agent types are considered:

- **Device agents**: Device agents control the device nodes $i \in \mathcal{I}$ in the LV grid. Because devices are owned by users in the LV grid, the device agents are modelled to represent user incentives. In the current model, devices may include static loads, PV, batteries, and heat pumps. Device agents are denoted by the corresponding node indices $i \in \mathcal{I}$. They get as input a price signal $\boldsymbol{\lambda}_{P(i)}$ from their parent node $P(i)$. Based on this price signal, the device agent runs a local optimisation to minimise some device specific cost function $C_i(\boldsymbol{\lambda}_{P(i)}, \boldsymbol{x}_i)$ subject to the relevant device constraints from equations (2)-(11). The resulting *proposed* power profile $\boldsymbol{p}_i^*$ is subsequently sent back up to the parent node $P(i)$. Note that for the inflexible load devices the power profile is fixed (see definition (2)), i.e. the cost function is constant in $\boldsymbol{x}$ and the optimal power profile is always the same. Summing up:

  input $\boldsymbol{\lambda}_{P(i)}$

  goal To minimise $C_i(\boldsymbol{\lambda}_{P(i)}, \boldsymbol{x}_i)$ s.t. device constraints at node i

  output $\boldsymbol{p}_i^*$

- **Congestion agents**: Congestion agents control congestion nodes from the set $\mathcal{J}$. Congestion agents are also denoted by the corresponding node indices $j \in \mathcal{J}$. They represent the interest of the DSO by mitigating congestion at critical locations. Their input is a price signal $\boldsymbol{\lambda}_{P(i)}$ from their parent node $P(j)$ and they initially send it down unchanged as $\boldsymbol{\lambda}_{P_j}$ to their immediate children $D(j)$. In response, children send back *proposed* power profiles (or first propagate the price profile if they are congestion agents themselves; this situation is discussed later). Note that these power profiles must be device feasible as device agents only send up feasible profiles. When a congestion agent has received the power profiles of all children, it aggregates them through equation (12) and evaluates grid feasibility of this cumulative profile, i.e. constraint (13). If the grid constraints are not satisfied, the congestion agent modifies the price signal and sends down the updated price $\boldsymbol{\lambda}_{P_j}$ to all children. This process is repeated until grid feasibility is achieved. At this point, the congestion agent sends up its aggregated power profile. Summing up:

  input $\boldsymbol{\lambda}_{P(j)}, \{\boldsymbol{p}_v : v \in D(j)\}$

  goal To achieve a power profile $\boldsymbol{p}_j$ that satisfies congestion constraint (13).

  output $\boldsymbol{\lambda}_j, \boldsymbol{p}_j$

- **Market operator agent**: The market operator agent controls the market operator node v_{mo}. The market operator is referred to by its corresponding node index v_{mo}. The market operator agent strives to match the total power profile of the LV network to the target profile $\boldsymbol{\Theta}$ that has been agreed on at the transformer level. Similarly to congestion agents, the market operator agent sends down a price profile $\boldsymbol{\lambda}_{P_{v_{mo}}}$ and aggregates the received power profiles through equation (14). Note that all power profiles will only end up at the market operator agent when they constitute a dispatch that is free of congestion. This is a consequence of the congestion agents that only send up a power profile when it is grid feasible. When the aggregated power profile does not match the target profile, i.e. it does not satisfy constraint (16), the price $\boldsymbol{\lambda}_{P_{v_{mo}}}$ is modified and sent back down. This process is repeated until market feasibility is achieved. At this point, the dispatch is feasible at the device, grid, and market level and the proposed power profiles are scheduled. Summing up:

  input $\{\boldsymbol{p}_v : v \in D(v_{mo})\}$

  goal To achieve a power profile $\boldsymbol{p}_{v_{mo}}$ that satisfies market constraint (16).

  output $\boldsymbol{\lambda}_{v_{mo}}$

The overall communications structure is quite simple: Price signals are sent down and proposed power programmes up. It follows naturally that the market operator agent, who has only children, only receives proposed power programmes and sends down price signals. The device agents only have parents and therefore only receive price signals and send back proposed power programme. The congestion agents, i.e. the inner nodes of the rooted tree, do both. They receive price signals from above and send them down (potentially updating them), and they receive proposed power programmes from below and send them up (aggregating them).

ALGORITHM DESCRIPTION

The core of the decentralised optimisation approach lies in finding a dispatch that meets the goal mentioned in Section 3.1. Most of this optimisation task happens in the price/power update function "UPDATEPROFILES". This function is presented in pseudocode form in Algorithm 1. It takes as input a node specification v and the (local) price vector $\boldsymbol{\lambda}$. This node specification dictates the node from which the optimisation is carried out. The optimisation happens in a recursive fashion with a hierarchical structure, i.e. calling the UPDATEPROFILES function at an arbitrary node $v \in V$ means that it will automatically be called for all its children, their children, etc. To better understand this process, the function is explained line by line with reference to Algorithm 1.

Let $v \in V$ be an arbitrary agent in the network and let the initial price vector $\boldsymbol{\lambda}$ be given. If v happens to be a device agent, line 3 is executed. In this line, the local device specific optimisation

Algorithm 1 Price/power update function

1: **function** $\boldsymbol{p}_v, \boldsymbol{\lambda}_v = $ UPDATEPROFILES$(v, \boldsymbol{\lambda}_v)$
2:     **if** $v \in \mathcal{I}$ **then**
3:         $\boldsymbol{p}_v = \operatorname{argmin}_{\boldsymbol{p}_v \in \mathcal{F}} C_v(\boldsymbol{\lambda}_v, \boldsymbol{p}_v)$
4:     **else if** $v \in \mathcal{J} \cup v_{mo}$ **then**
5:         **repeat**
6:             **for** $u \in D(v)$ **do**
7:                 $\boldsymbol{p}_u, \boldsymbol{\lambda}_v = $ UPDATEPROFILES$(u, \boldsymbol{\lambda}_v)$
8:             **end for**
9:             $\boldsymbol{p}_v = \sum_{u \in D(v)} \boldsymbol{p}_u$
10:             error$_v$ = VIOLATION$_v(\boldsymbol{p}_v)$
11:             $\boldsymbol{\lambda} = $ ADJUSTPRICES$(v, \boldsymbol{\lambda}$, error$_v)$
12:         **until** error$_v < \epsilon$
13:     **end if**
14:     **return** $\boldsymbol{p}$
15: **end function**

is carried out by minimising the local cost function $C_v(\cdot)$ after which UPDATEPROFILES terminates. If v is a parent agent, i.e. either a congestion agent or market operator agent, lines 2-3 are skipped and the loop 5-12 is entered. In lines 6-8, UPDATEPROFILES is called within itself for each agent in the set of children $D(v)$. In this call lies the recursive nature of the function: all power profiles of underlying child nodes need to first be specified before the next part of the function can be reached.

Assume first that the congestion agent/market operator agent v has only device agents in his set of direct child agents $D(v)$. Then, UPDATEPROFILES is called for each of them, running the local device optimisations in line 3. The updated power profiles of all the child device agents are then summed into $\boldsymbol{p}_v$. The function VIOLATION$_v$ checks the constraints at v. If v is a market operator agent, constraint (16) is evaluated. If v is a congestion agent, constraint (13) is evaluated. The function VIOLATION$_v$ then returns the violation of the relevant constraint in the form of error$_v$. If this error is smaller than ϵ, the loop is finished and UPDATEPROFILES terminates and returns a feasible dispatch. If the error exceeds the tolerance ϵ, the price signal is adjusted through the "ADJUSTPRICES" function in line 11 using a gradient descent method. This process is repeated until a feasible dispatch is found. If the above assumption does not hold and the constraints of agent v are not completely satisfied, the above process repeats (recursively) for each of them. Because all leaf nodes are controlled by device agents, the function eventually converges (assuming a feasible dispatch can be found).

In a full run of the decentralised optimisation model, an initialisation of the tuple $(v, \boldsymbol{\lambda})$ needs to be specified before calling the function UPDATEPROFILES. For v, the market operator is always chosen, i.e. $v = v_{mo}$, such that the profiles are recursively updated for the complete network. For $\boldsymbol{\lambda}$, an identical price is first assigned to each device and time period. Using this "flat" price profile, the local device optimisations (line 3) are run leading to an initial dispatch $\boldsymbol{p}$.

3.3 Theoretical comparison

In the following chapters, performance of the centralised and decentralised approaches are compared on different criteria. Because centralised and decentralised optimisation differ fundamentally in their mathematical structure, it is important to have a critical look at this comparison.

In Section 3.2.1, the features that were common to both models are outlined. Most notably, the structure of the grid, the discretisation of time, and the feasible set were included in this general description. The description of the feasible set was based on allowed behaviour at the different device nodes, limits and a description of aggregation at the congestion nodes, and limits and a description of aggregation at the market operator node. In these aspects, the centralised and decentralised models are completely equivalent. However, a careful look at the individual models

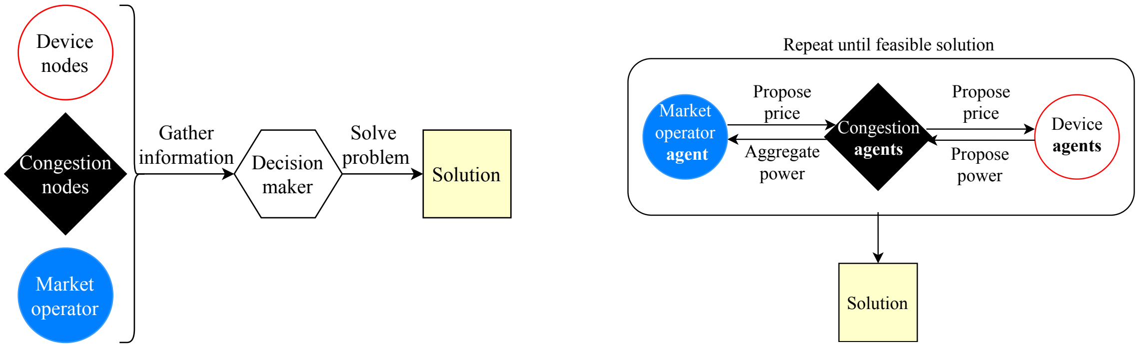

(A) *Centralised optimisation*

(B) *Decentralised optimisation*

FIGURE 6: *Simplified flow diagrams of the centralised and decentralised optimisation models. Note that although visualised identically, entities represent nodes in diagram 6a and agents in diagram 6b.*

reveals that they differ in some other key parts. To give a quick intuitive overview of this difference, Figure 6 provides a schematic visualisation of the control and communication structures of the centralised model (Figure 6a) and the decentralised model (Figure 6b). These differences are discussed in the following in a more detailed fashion and are summarised in Table 2:

- **Decision makers**: The centralised approach has a single *centralised* decision maker whereas the decentralised approach has multiple *decentralised* decision makers that are referred to as *agents* in Section 3.2.3.
- **Availability of information**: The centralised approach has access to "perfect information" about the grid and its participants over the time horizon. This information is represented by the parameters and constraints of programme (17) (Section 3.2.2). It is referred to as *global* information in Table 2 because information is known to the decision maker irrespective of location in the grid. In the decentralised case, only *local* information is available as can be seen in Figure 5. This information is limited to proposed price profiles from parent agents and proposed power profiles from child agents together with the local constraint at the current agent. Therefore, no single agent has global or perfect information about the system state.
- **Availability of control**: For the availability of control, a similar comparison can be made. The decision maker of the centralised optimisation approach can freely choose any dispatch within the feasible set $\mathcal{F}$ taking into account the full set of devices, i.e. the global system state. In the decentralised approach, agent decisions are only communicated locally: congestion agents and the market operator agent send (updated) price profiles to their child agents, and device agents send up power profiles to their parent agent. In other words, control is only exerted locally in the decentralised model.
- **Mode of control**: Another difference between the algorithms lies in the mode of control. In the centralised optimisation approach, the market operator sends a "control signal" to all devices simultaneously and this *directly* results in the desired behaviour of the devices. In the decentralised case, "control signals" consist of *proposed* price/power profiles that are *indirect*. They do not lead to direct implementation of behaviour because the local decision makers do not know for certain whether this behaviour will be optimal or even feasible globally.

In the following Chapters, performance of the centralised and decentralised optimisations is compared by evaluating the global cost function f_{cost} for both models. The centralised optimisation uses this cost function as its global objective function as shown in (17). In contrast, the decentralised optimisation does not have an explicitly defined global objective function. Instead, local cost functions for device agents and constraint violations for "parent agents" may be seen as local objective functions that are minimised (see Algorithm 1). For the centralised approach, global availability of information further ensures that the decision maker knows $f_{cost}(\boldsymbol{x}, \boldsymbol{y})$ for any dis-

Model feature	Centralised optimisation model	Decentralised optimisation model
Decision makers	single	multi
Availability of information	global	local
Availability of control	global	local
Mode of control	direct	indirect

TABLE 2: *Differences between the centralised and the decentralised optimisation models*

patch $(\boldsymbol{x}, \boldsymbol{y})$. Together with directness and global availability of control this then ensures that, when given a sufficient amount of time, the centralised optimisation always finds a dispatch that is globally optimal, i.e. a pair $(\boldsymbol{x}^*, \boldsymbol{y}^*) \in \mathcal{F}$ with $(\boldsymbol{x}^*, \boldsymbol{y}^*) = \mathrm{argmin}_{(\boldsymbol{x},\boldsymbol{y})\in\mathcal{F}} f_{cost}(\boldsymbol{x}, \boldsymbol{y})$. Such an assumption cannot be made for the decentralised case for multiple reasons. First of all, the local nature of available information means that no single agent "knows" either the full dispatch $(\boldsymbol{x}, \boldsymbol{y})$ or the resulting total cost $f_{cost}(\boldsymbol{x}, \boldsymbol{y})$. Furthermore, local and indirect control may prevent some feasible solutions from being reached. Take for example two identical heat pumps i_1 and i_2 at time t that are connected under the same congestion node j_1. In the centralised case, a controller can freely choose any values x_{i_1t}, x_{i_2t} including pairs where $x_{i_1t} \neq x_{i_2t}$. However, in the decentralised case, j_1 can only send a single price to both heat pumps to which they react in the same way, i.e. we must have $x_{i_1t} = x_{i_2t}$. Apparently, there exist feasible solutions $(\boldsymbol{x}, \boldsymbol{y})$ that are not *reachable* by the decentralised optimisation model. Letting $\mathcal{R}$ represent the *reachable set* of solutions for the decentralised model, it then holds that $\mathcal{R} \subset \mathcal{F}$. This raises the following question:

> "*Do the centralised and decentralised optimisations solve the same problem and are they outcome equivalent?*"

It turns out that comparison of centralised and decentralised optimisation is not a new challenge, it has been described extensively in for instance [27, 51, 2]. In [51], the comparison is introduced in the context of climate control in an office building. The authors emphasise the difference by referring to the decentralised optimisation as a "market approach" and the centralised optimisation as a "control approach". In [2], the analysis from [51] is extended using a combination of microeconomic theory and control theory to give a more general mathematical comparison. The authors give a proof for outcome equivalence of centralised and decentralised models when a set of assumptions hold. Before extending this result to the present analysis, care should be taken to consider the differences between the two applications. Most importantly, the agent-based market approach analysed in [2] uses an auction mechanism where a limited, *one-dimensional*, resource is allocated amongst agents in a way that is shown to be globally optimal. For the decentralised approach considered in this thesis, the *N-dimensional* power profiles represent the equivalent resource. To find an optimal allocation, an N-dimensional price vector $\boldsymbol{\lambda}$ needs to be specified instead of a simple one-dimensional clearing price. Because of the resulting cost functions $C_i(\boldsymbol{\lambda}_{P(i)}, \boldsymbol{x}_i^*)$ and the differing control strategy, three out of the four assumptions from [2] are not, in general, satisfied for the model of Subsection 3.2.3.

Summing up, the theoretical proofs that show perfect equivalence between centralised and decentralised models may not be applicable to the present analysis. However, this result does not necessarily invalidate the present comparison. As long as all constraints are respected, the performance of the models may be compared based on the global cost function f_{cost}. In Section 5, the comparison between the centralised and decentralised models will be addressed through empirical methods.

4 Modelling uncertainty

Uncertainty plays a central role in the operation of smart grids. Therefore, this aspect is discussed in more detail in this chapter. In Section 4.1, sources of uncertainty in the design and optimisation of smart grids are described and characterised. Then, in Section 4.2 modelling of these uncertainties is described in the treatment of input data. Next, a receding horizon scheme is introduced in Section 4.3. With these elements in place, the centralised and decentralised optimisation models are updated to include the discussed uncertainty in Subsections 4.3.3 and 4.3.4 respectively.

4.1 Sources of uncertainty

In Subsection 2.3.2 it is reported that uncertainty embodies an important challenge to EM in smart grids. For the current analysis, two main sources of uncertainty can be classified: intermittent generation and human behaviour. These sources are first described in Subsections 4.1.1 and 4.1.2 respectively. The section ends with some general remarks about the uncertainty sources and their typical manifestations at different locations in the grid.

4.1.1 Intermittent generation

The energy transition is accompanied by a fast increasing volume of distributed, often intermittent, energy sources in the distribution grid. These energy sources are based on natural phenomena that are not only intermittent, but impossible to predict perfectly. At the LV grid level, the most influential intermittent energy source is PV (it is also the only source which is considered in this thesis). In the relevant day-ahead timescale, significant uncertainty in the predictions still exists for PV production as was reported in a recent review paper [8]. In the context of PV forecasting, this timescale is referred to as “mid-term” which consists of predictions that range somewhere between 30 minutes and several days ahead. For forecasts that extend a maximum of two hours, it was demonstrated in [6] that direct solar power observations represent the preferred source of forecasting input data. Beyond this horizon, numerical weather predictions are typically better suited for this task. These weather predictions are notorious for having large variability which is the main reason why PV forecasting remains challenging. In [39], historical data was compared to weather forecasts to get an idea of the potential performance of mid-term PV forecasting. In the paper, daily forecasts were classified in one of four weather types: cloudy, foggy, rainy, and sunny. To quantify forecasting performance, the mean relative error (MRE) was used, defined by

$$\text{MRE} = \frac{1}{N}\sum_{n=1}^{N}\frac{|m_n - r_n|}{|r_n|},$$

where N is the number of measurements and m_n and r_n represent the predicted and realised values respectively for time period $n \in [1, 2, \dots, N]$. For each weather type, the Mean Relative Error (MRE) was calculated using the above formula and letting m_n and r_n be the predicted and realised PV power production values. Sunny days were accompanied by the most accurate predictions with an average MRE of just 4.85% whereas cloudy days scored the worst with an MRE of 12.42%. The average MRE over all weather types was reported at 8.64%. From these results it was concluded that mid-term PV forecasting based on numerical weather predictions has highly varying success rates depending on the overall weather type. In any case, prediction inaccuracies remain an important challenge in the realm of smart-grids and in day-ahead markets especially. Therefore, an EM tool such as the one considered in this work should be robust with respect to these inaccuracies and it should account for them explicitly.

4.1.2 Human behaviour

The second main source of uncertainty, human behaviour, has many facets that are relevant here. Because of the ongoing energy transition and its accompanying trends (see e.g. Subsection

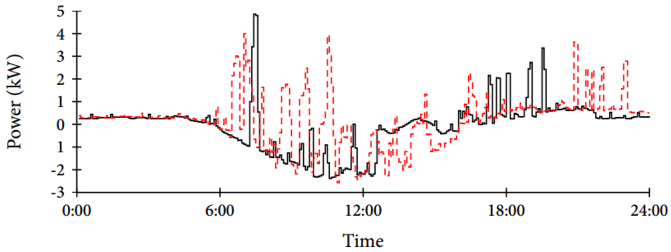

(A)

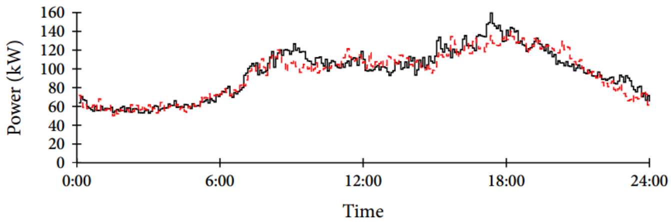

(B)

FIGURE 7: *Comparison of demand profiles for two separate days, exactly a week apart, at (a) household level and (b) transformer level. Measurements have a time resolution of 5 minutes. (figure taken from [45], data from Alliander)*

2.1.3), humans influence the energy market as consumers and producers simultaneously; they are therefore often referred to as *prosumers*. The producer part of this role consists mostly of PV production and some DERs running alongside it. In addition to PV yields, prosumers may own storage devices such as household battery units. By operating these devices prosumers interact in the energy market, leading to times where they have a net consumption but also to times with a net feed-in of power. Such interactions result in significant stochasticity, especially when users are given complete autonomy on the flexibility of DERs. In this thesis, a fixed, pre-specified, amount of flexibility is assumed to be available to decision makers independent of human behaviour, e.g. heat pumps having a known range of comfortable household temperatures. This choice therefore bypasses a significant portion of uncertainty by ignoring variations in available flexibility.

Uncertainty in the “conventional” consumer part of the prosumer role is considered in the models. This mainly consists of uncertainty in predicting the inflexible household loads. Here, the unpredictable variations in time and intensity of usage for electrical devices and resources such as lights, entertainment devices, and household appliances cause this uncertainty. The highly volatile nature of these loads combined with increasing levels of electrification makes them challenging to model. As seen in for instance [48], a large body of work exists on the topic of load forecasting. For instance, in [38], a state-of-the-art forecasting technique is introduced that aims to predict load profiles through a deep-learning algorithm. The authors claim to improve upon the existing literature by achieving an MRE that lies between $20 - 25\%$ for a typical single household load.

4.1.3 General remarks

In addition to the sources of uncertainty mentioned in the previous subsection, many other uncertain factors may be identified. Some of these sources may be classified as *disruptive* and could for instance be caused by physical tempering with the grid assets or other unexpected failures in the grid. Even cyber-attacks are in the realm of possibilities with the potentially vulnerable data communications between grid actors. Even though these uncertainties are not modelled explicitly in this work, it may be important to consider them in the process of developing an EM tool.

To get an idea of the implications of uncertainty for the overall energy dispatch and proper functioning of the grid it is crucial to consider different levels of aggregation. The *law of large numbers* states that variance of uncertain parameters goes down when summing up multiple independent instances. In the context of electricity grids, this phenomenon is called *load diversification* as was already touched upon in Subsection 2.3.2 [45]. A nice example of load diversification is presented in Figure 7 where load profiles for different days are compared at the household and transformer level. The aggregated profiles in Figure 7(b) are far more similar between the two days and relative peaks are a mere fraction of those in Figure 7(a). It turns out that household loads are uncorrelated enough to apply the law of large numbers (satisfying the independence requirement). However, when considering PV production, the profiles cannot be treated as independent and there will not be a significant amount of load diversification. Synchronised RES conditions therefore have the risk to lead to large prediction errors also at the transformer level (up to 20% in generation power) [12]. More on this is discussed in Section 4.2. Ignoring this problem of systematic uncertainties due to synchronisation of RES, it is tempting to only consider the transformer level and conclude that the situation without RES is relatively stable at this level. However, congestion can happen more locally, e.g. when cable tolerances are exceeded between household and transformer.

4.2 Realistic data input

In this section, the input data to the power profiles is considered in detail. In Subsection 4.2.1, the structure of the power profiles and correlation factors are discussed. In Subsection 4.2.2, data input techniques are introduced, explicitly taking into account the described structure and correlation factors.

4.2.1 Characterisation of input data

In order to use real-life data in a way that is realistic, it is important to first consider the structure of this data. In this thesis, we focus on a distribution network for a neighbourhood and the energy loads of the houses in this neighbourhood. The input data has two main elements that both result from processes that are uncontrollable for the decision maker:

- **Uncontrollable load profiles:** As was briefly described in Section 3.2.1, static load profiles represent the power consumption of the full set of *uncontrollable* appliances for a household (e.g. television, water cooker, lighting). By definition, these profiles cannot be influenced by the decision maker. For each household, the uncontrollable load profiles are modelled as a single load device. For each of these load devices $i \in \mathcal{I}_{load}$, a profile $\boldsymbol{\chi}_i$ is taken from a given data set which contains potential load profiles for his device. The power profile $\boldsymbol{p}_i$ of the load device node is directly specified by this data profile as described in equation (2).
- **Curtailable PV production profiles:** The base production profile of a PV unit results directly from the given solar intensity and is uncontrollable like the static load profile. These production profiles are also taken from a data set and are used in the model as $\boldsymbol{\xi}_i$ for all $i \in \mathcal{I}_{PV}$. The difference with static loads is that PV devices are *curtailable*, i.e. the decision maker can turn them off during any time slot $t \in \mathcal{T}$ through the decision variable x_{it}. The resulting power profile $\boldsymbol{p}_i$ is the product of the input data and a decision variable as described in equation (3).

In the preceding chapters, the input data is only described in a general way. To get a better idea of the characteristics of the power profiles contained in this data set, some concrete examples are considered here. First, three example PV power profiles are considered that originate from three real Dutch households situated in the same neighbourhood (numbered 1, 2, 3). In Figure 8, the daily PV production profiles of the houses are presented for a single month (April). Heat plots are used to visualise PV yields over the course of the month where each horizontal bar corresponds to a daily profile for a single household and measurements are presented with a resolution of five minutes. Based on e.g. the data of Household 1 in Figure 8a, it is easy to distinguish different weather types. Clear days have a relatively smooth profile that has a gradual increase of power production towards the middle of the day, see e.g. the 18th until the 21st of April. Overcast days have a relatively constant, low, power output, see e.g. the 28th until the 30th of April. Finally,

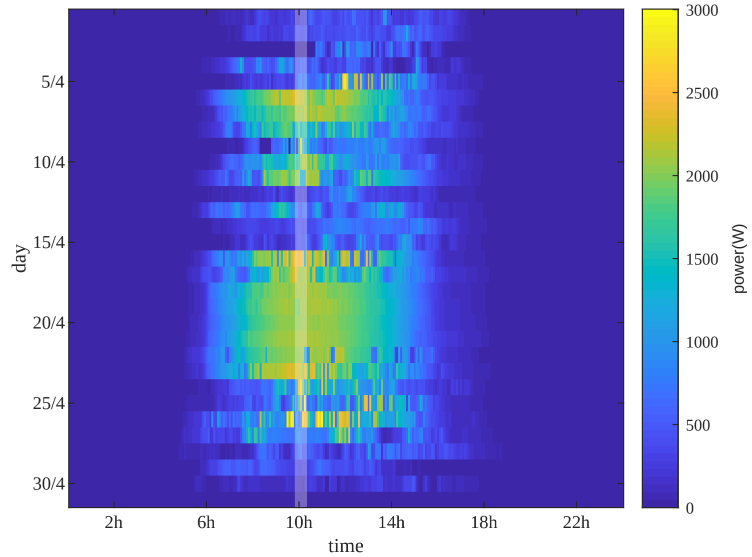

(A) *Household 1*

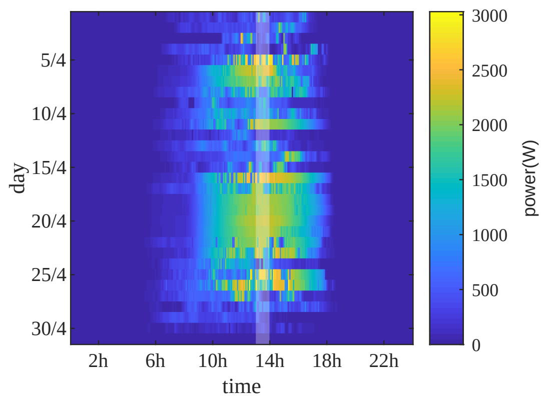

(B) *Household 2*

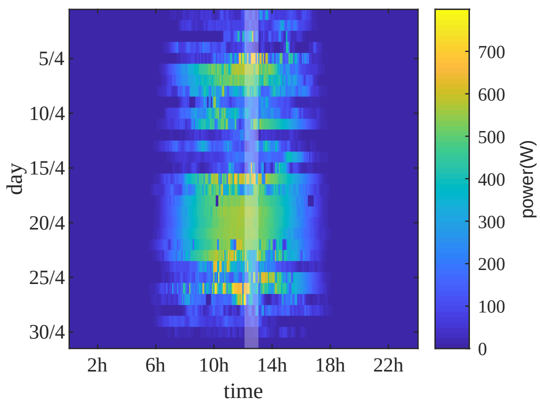

(C) *Household 3*

FIGURE 8: *Heat plots of PV power production of three households for the month of April. Power outputs are average values over the 5 minute intervals. The average peak production is indicated by a vertical bar.*

cloudy days have significant variability in power output due to the passing of clouds, see e.g. the 26th of April (verified through [24]). From this data it is directly clear that PV forecasting is significantly less accurate in the case of cloudy weather compared to clear or overcast days as was already reported in Subsection 4.1.1. In addition to weather conditions, the subtle increase in daylight hours over the course of April may be seen in the widening of the nonzero bands in the heat plot. Comparing now the plot for household 1 to that of the other two, the above mentioned features are seen as well. Because of their physical proximity, the relative power intensities are heavily correlated between the three households.

Despite this correlation, some notable differences are observed as well. Firstly, the total area of solar panels is clearly not equal between the different households as can be observed in the significantly lower outputs for Household 3. Judging from this data, it seems plausible that Households 1 and 2 have an area of solar panels that is around four times as large as that of Household 3. Another difference that is observed is the *shift* in intensity between the households. Compared to Household 1, the peak solar power output of Household 2 seems to be shifted by 3,5 hours, from 10:05h to 13:40h. This gives a strong indication that the solar panels of Household 1 are oriented more towards the east than those of Household 2 by a sizable angle. Household 3 seems to lie somewhere in between the other two with a peak production around 13:00h.

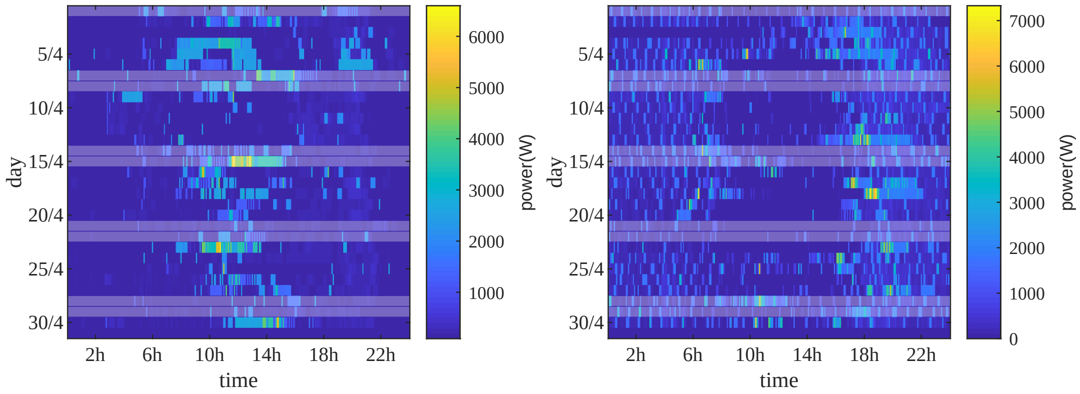

(A) *Household 1* (B) *Household 2*

FIGURE 9: *Heat plots of household power consumption of three households for the month of April. Power outputs are average values over the 5 minute intervals. The weekends are indicated by horizontal bars.*

To get some insight into typical static load profiles of households, we consider a similar data set as for the PV production on power *consumption* for two different example households (numbered 1 and 2). This power consumption data is presented in the form of two heat plots in Figure 9. The structure of this data differs fundamentally from that of the PV data set. There is less of a pattern between different days, and the load profiles seem to more closely resemble random noise. Referring first to the data for Household 1 in Figure 9a, higher loads mostly occur during the day, with more pronounced peaks during the weekends and on Mondays. Compared to the PV data, peaks on average have a shorter duration and higher relative intensities. This is to be expected with the relatively short running time and abrupt switching that is typical for household appliances such as washing machines, water cookers, microwaves etc. The data for Household 2 in Figure 9b has a structure that looks fundamentally different. The power demand seems generally to be lowest *during* the day with small equally spaced intervals of increased loads roughly concentrated between dusk and dawn. This pattern may very well be indicative for the presence of a heat pump in household 2. During the colder hours outside the daytime, the heat pump stabilises the room temperature in the house by periodically turning on and injecting heat. Compared to PV data, physical proximity does not seem to be a significant correlation factor for load data. However, other correlation factors such as day of the week may play a noticeable role.

4.2.2 Model data generation

In Subsection 4.2.1, some features of the source data sets are discussed. These data sets consist of specifications of the uncontrollable processes, i.e. the static load and (potential) PV production. In this subsection, we describe how the model data is retrieved from this source data. The goal hereby is to create realistic models, i.e. representative of a real neighbourhood distribution grid.

Before discussing the generation of model data, we first consider the scope of a single *simulation run*. In this thesis, we focus on the day-ahead market (see also Section 3.2.1). A simulation run therefore consists of a 24-hour period during which all decisions need to be specified and an objective function can be evaluated. Referring to the input power profile data discussed in Section 4.2.1, this means that each static load and each PV device is assigned a 24-hour power profile from the input data set as described in equations (2) and (3). Therefore, a total of $|\mathcal{I}_{load}|$ consumption profiles and $|\mathcal{I}_{PV}|$ production profiles need to be taken from the data set, where $|\mathcal{I}_{load}|$ denotes the number of static load devices and $|\mathcal{I}_{PV}|$ the number of PV devices in the model. One option would be to assign uniform probabilities to all profiles in the data set and randomly select a sample for each load and PV device. However, a realistic simulation run needs to account for all known correlation factors, and this naive approach ignores both spacial and temporal correlation. Such a naive approach typically leads to an increased level of load diversification (see Subsection

4.1.3), resulting in a simulation run that is far less challenging. To prevent this, we account for all correlation factors by gathering all load and PV profiles from the same date and from the same neighbourhood. This ensures that the given real-life spacial and temporal correlation factors are retained in the simulation run.

4.3 Receding horizon models

Based on the analysis of the uncertainty of the source data, the base models of Section 3.2 may be adapted to be more realistic. In the models of Chapter 3, a major simplification is made by assuming perfect knowledge of future predictions. In this section, an effort is made to introduce a modelling technique that does take into account the uncertainty described in Section 4.1. This technique, called *receding horizon* modelling, is introduced first in Subsection 4.3.1 followed by a description of the updated centralised model (Subsection 4.3.3) and decentralised model (Subsection 4.3.4).

4.3.1 Receding horizon definition

In Subsection 3.2.1 the discretisation of time was introduced. This discretisation allows the 24-hour time horizon to be modelled as a (finite) sequence of N discrete blocks, PTUs. As described in the previous subsection, information about future data inevitably has some uncertainty associated with it. As time passes, future time slots become present time slots and uncertainty vanishes as unknown information becomes realised. To take into account this inherent uncertainty, a real EM approach has to optimise for a dispatch based on only a prediction of future device states. In modelling the dispatch problem, a realistic approach should account for this natural flow of time and the evolution of uncertainty associated with it. A popular way to achieve this is through a receding horizon scheme.

The main idea of a receding horizon scheme is to split up the *original time horizon* into several smaller time windows and solve the model for each of them separately. These smaller time windows are referred to as *planning windows*. These planning windows may overlap and are typically of equal size. The planning window corresponding to the first sub-problem covers the very first time slots of the original horizon. After solving the sub-problem, the planning window moves a fixed amount of time slots further into the future. The number of slots that the planning window moves between sub-problems is referred to as the *step size* of the receding horizon model. In order to cover all time slots of the original time horizon, the step size needs to be smaller than (or equal to) the size of the planning window. After solving a sub-problem and before moving to the next iteration, the time slots covered by the step-size are fixed to their solution value, and are used to calculate the initial state for the next. This process iterates until all time slots of the original time horizon are fixed. Summing up, the receding horizon scheme is uniquely characterised by the length of the time window and the step size.

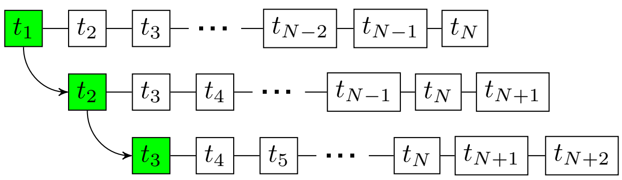

FIGURE 10: *The principle of a receding horizon market with step size 1 and planning window size N. Planned power schedules (white) are turned into contracts for the current time slot (green) in the next iteration.*

For the present dispatch problem, a receding horizon is implemented with a step size of one PTU, and a planning window, referred to as $\mathcal{H}$, of N PTUs. The original time horizon $\mathcal{T}$ is the one defined in Section 3.2.1, i.e. $\mathcal{T} = \{t_1, \ldots, t_N\}$. This scheme is visualised in Figure 10 where the first three iterations are displayed. The main objective of this receding horizon scheme is not to split the original problem into smaller problems, but instead to simulate the passing of time and accompanying changes in uncertainty. In the real world, an operator does not know exactly what will be the PV yield or the inflexible loads at the households (see Section 4.1). Generally, the further into the future, the more uncertain the predictions for such data are.

4.3.2 General model description

In this subsection, the general structure of the receding horizon models described in Subsection 4.3.3 and Subsection 4.3.4 is introduced. In the models from Chapter 3, future realisations of the PV yield and the household load are assumed to be known "perfectly" for each device in the form of the sample profiles $\boldsymbol{\chi}_i$ and $\boldsymbol{\xi}_i$. However, to model the real-world uncertainty also in its time-varying component, a different approach needs to be taken. To this end, we consider an arbitrary planning window $\mathcal{H}$ and let the present (or current) time period be defined as t_c. To model uncertainty in future PTUs, the perfect knowledge needs to be hidden for the optimisation model. To achieve this, we model predicted profiles as a convex combination of "historic data profiles" and "perfect information data profiles". The historic data profile can for instance be the average household consumption/PV production over the previous year. The perfect information data profile consists of the actual consumption/production levels that are eventually *realised* in the optimisation model. We first consider predictions of the load profiles defined by $\chi_{it}^{prediction}$. To model realistic predictions of the power profiles, we take a convex combination of the *perfect information* values χ_{it} and the *historic data* values $\overline{\chi}_{it}$. To this end, a non-decreasing sequence $(\mu_0^{load}, \mu_1^{load}, \ldots, \mu_{N-1}^{load})$ is introduced with $\mu_0^{load} = 0$. The predicted values are then calculated for the planning window $\mathcal{H}$ as

$$\chi_{it}^{prediction} = \mu_{t-t_c}^{load} \cdot \overline{\chi}_{it} + (1 - \mu_{t-t_c}^{load})\chi_{it}, \qquad \forall t \in \mathcal{H}. \tag{18}$$

Because $|\mathcal{H}| = N$, each PTU in the planning window $\mathcal{H}$ is matched to an element of the sequence. The first PTU $t = t_c$ gives $\mu_{t_c - t_c} = \mu_0 = 0$, i.e. perfect predictions at the present time slot are assumed. Time slots $t > t_c$ that are part of the planning window $\mathcal{H}$ are associated with "worse" predictions. This is because the factor μ_{t-t_c} is non-decreasing. We have chosen to let the sequence μ_k approach 1 for $k = N$ (and therefore $(1 - \mu_k)$ approaches 0) implying that predictions become less accurate (or stay the same) for PTUs that lie further into the future. The exact same principle can be applied to the PV profile predictions, defined as $\xi_{it}^{prediction}$. They are given by

$$\xi_{it}^{prediction} = \mu_{t-t_c}^{PV} \cdot \overline{\xi}_{it} + (1 - \mu_{t-t_c}^{PV})\xi_{it}, \qquad \forall t \in \mathcal{H}, \tag{19}$$

where μ_k^{PV} is also a non-decreasing sequence $(\mu_0^{PV}, \mu_1^{PV}, \ldots, \mu_{N-1}^{PV})$ is introduced with $\mu_0^{PV} = 0$.

Because of the receding horizon scheme, decisions applied to time periods towards the end of the planning window have a decreasing influence on the resulting dispatch. The market constraint (16) is therefore also adapted, replacing the fixed error tolerance α by a dynamic error tolerance vector $\boldsymbol{\alpha}^{RH} \in \mathbb{R}^{|\mathcal{H}|}$ to reduce the computational effort and to achieve faster convergence of the optimisation models. The error tolerance vector is defined as

$$\alpha_t^{RH} = \nu_{t-t_c} \cdot \alpha + (1 - \nu_{t-t_c}) \cdot \alpha_+ \qquad \forall t \in \mathcal{H}, \tag{20}$$

where α is the standard error tolerance from Chapter 3, and α_+ is an additional error tolerance. As in equations (18) and (19), a non-decreasing sequence $(\nu_0, \nu_1, \ldots, \nu_{N-1})$ is introduced with $\nu_0 = 0$. In this way, equation (20) ensures that, for a given planning window, the error tolerance starts at α for the first PTU t_c and ends at $\alpha + \alpha_+$ for the last PTU t_{c+N-1}. Consequently, the "dynamic" error tolerance vector relaxes the market constraints for PTUs towards the end of $\mathcal{H}$.

Note that we set the first element of each sequence at zero, i.e. $\mu_0^{load} = \mu_0^{PV} = \nu_0 = 0$. Because of this, at the present PTU t_c, equations (18), (19), and (20) simplify to their "Chapter 3 counterparts" (2), (3), and (16) respectively. Combined with the fact that a step size of 1 is used, this ensures that the feasible set $\mathcal{F}$ remains unchanged compared to Chapter 3. Note also that the original time horizon ends at N. This means that after a single receding horizon step, the planning window moves outside of the original horizon. To prevent this, $\mathcal{T}$ is cyclically updated for the loads, PVs, and the target profile, i.e. t_{N+1} is interpreted as t_1, t_{N+2} is interpreted as t_2, etc.

The receding horizon approach described above may straightforwardly be applied to both the centralised model (Section 3.2.2) and the decentralised model (Section 3.2.3). The updated receding horizon versions of the models are described for the centralised and decentralised models in Subsection 4.3.3 and Subsection 4.3.4 respectively.

4.3.3 Receding horizon centralised model

The centralised model that was described in the programme (17) is changed slightly as described in Subsection 4.3.2 to build the receding horizon sub-model. Instead of the total time range $\mathcal{T}$, the objective function, constraints, and definitions are evaluated over the planning window $\mathcal{H}$ for each receding horizon iteration. Deterministic load and PV profiles $\boldsymbol{\chi}_i$ and $\boldsymbol{\xi}_i$ are replaced by predictions $\boldsymbol{\chi}_i^{prediction}$ and $\boldsymbol{\xi}_i^{prediction}$. Furthermore, the market constraints are updated, replacing the fixed error tolerance α by the error tolerance vector $\boldsymbol{\alpha}^{RH}$.

Implementing the above changes to the model in programme (17), the resulting receding horizon sub-model programme is given in (21).

$$
\begin{aligned}
\min_{\boldsymbol{x},\boldsymbol{y}} \quad & f_{cost}(\boldsymbol{x},\boldsymbol{y}) && && \text{(21a)}\\
\text{s.t.} \quad & \textit{Device constraints} && && \\
& x_{it}^d \geq y_{it}^d \cdot X_i^{min}, && \forall i \in \mathcal{I}_{bat}, t \in \mathcal{H} && \text{(21b)}\\
& x_{it}^c \leq y_{it}^c \cdot X_i^{max}, && \forall i \in \mathcal{I}_{bat}, t \in \mathcal{H} && \text{(21c)}\\
& y_{it}^c + y_{it}^d \leq 1, && \forall i \in \mathcal{I}_{bat}, t \in \mathcal{H} && \text{(21d)}\\
& x_{it} \geq X_i^{min}, && \forall i \in \mathcal{I}_{HP}, t \in \mathcal{H} && \text{(21e)}\\
& x_{it} \leq X_i^{max}, && \forall i \in \mathcal{I}_{HP}, t \in \mathcal{H} && \text{(21f)}\\
& s_{it} \geq S_i^{min}, && \forall i \in \mathcal{I}_{bat} \cup \mathcal{I}_{HP}, t \in \mathcal{H} && \text{(21g)}\\
& s_{it} \leq S_i^{max}, && \forall i \in \mathcal{I}_{bat} \cup \mathcal{I}_{HP}, t \in \mathcal{H} && \text{(21h)}\\
& \textit{Congestion constraints} && && \\
& \boldsymbol{p}_{jt} \geq P_{jt}^{min}, && \forall j \in \mathcal{J}, t \in \mathcal{H} && \text{(21i)}\\
& \boldsymbol{p}_{jt} \leq P_{jt}^{max}, && \forall j \in \mathcal{J}, t \in \mathcal{H} && \text{(21j)}\\
& \textit{Market constraints} && && \\
& \boldsymbol{p}_{v_{mo}t} \geq \Theta_t - \boldsymbol{\alpha}^{RH}, && \forall t \in \mathcal{H} && \text{(21k)}\\
& \boldsymbol{p}_{v_{mo}t} \leq \Theta_t + \boldsymbol{\alpha}^{RH}, && \forall t \in \mathcal{H} && \text{(21l)}\\
& \textit{Further constraints} && && \\
& p_{it} := \mu_{t-t_c}^{load} \cdot \overline{\chi}_{it} + (1-\mu_{t-t_c}^{load})\chi_{it}, && \forall i \in \mathcal{I}_{load}, t \in \mathcal{H} && \text{(21m)}\\
& p_{it} := (\mu_{t-t_c}^{PV} \cdot \overline{\xi}_{it} + (1-\mu_{t-t_c}^{PV})\xi_{it}) \cdot x_{it}, && \forall i \in \mathcal{I}_{PV}, t \in \mathcal{H} && \text{(21n)}\\
& p_{it} := \frac{1}{\eta_i} \cdot x_{it}^c + \eta_i \cdot x_{it}^d, && \forall i \in \mathcal{I}_{bat}, t \in \mathcal{H} && \text{(21o)}\\
& s_{it} := s_{i0} + \sum_{t'=t_0}^{t'=t} \left(x_{it'}^c + x_{it'}^d\right) \cdot \tau, && \forall i \in \mathcal{I}_{bat}, t \in \mathcal{H} && \text{(21p)}
\end{aligned}
$$

$$s_{it} := s_{i0} + \sum_{t'=t_0}^{t'=t} \left(x_{it'} - \beta_i\right) \cdot \tau, \qquad \forall i \in \mathcal{I}_{HP}, t \in \mathcal{H} \tag{21q}$$

$$\boldsymbol{p}_v := \sum_{v \in D(j)} \boldsymbol{p}_v, \qquad \forall v \in \mathcal{J} \cup v_{mo} \tag{21r}$$

Domain constraints

$$x_{it} \in \{0, 1\}, \qquad \forall i \in \mathcal{I}_{PV}, t \in \mathcal{H} \tag{21s}$$

$$y^c_{it} \in \{0, 1\}, \qquad \forall i \in \mathcal{I}_{bat}, t \in \mathcal{H} \tag{21t}$$

$$y^d_{it} \in \{0, 1\}, \qquad \forall i \in \mathcal{I}_{bat}, t \in \mathcal{H} \tag{21u}$$

$$x^d_{it} \leq 0, \qquad \forall i \in \mathcal{I}_{bat}, t \in \mathcal{H} \tag{21v}$$

$$x^c_{it} \geq 0, \qquad \forall i \in \mathcal{I}_{bat}, t \in \mathcal{H} \tag{21w}$$

Compared to (17), all references of $\mathcal{T}$ are replaced by $\mathcal{H}$, and definitions (17m) and (17n) are replaced by (21m) and (21n) respectively.

Algorithm 2 Receding horizon centralised model

```
1: t_c ← t_1
2: repeat
3:     Solve (21) for H
4:     Fix all variables for t_c
5:     p ← p + 1
6: until p > N
7: Evaluate f_cost for T
```

To run the full receding horizon model, the sub-model needs to be embedded in an algorithm. This algorithm is presented in pseudocode in Algorithm 2. In line 1, the present time period is initialised to the first period of the complete time horizon. Then, in the loop between lines 2-6, in every iteration the different sub-problems are solved (line 3). The solutions for the first time period in the planning window are fixed (line 4), and the present time period is updated (line 5). The loop is terminated when all sub-problems have been solved, i.e. the total time interval has been realised. At this point, the value of the realised cost function is evaluated (line 7). The complete optimisation model represented by the combination of sub-model (21) and Algorithm 2 is referred to as the receding horizon centralised optimisation (RHCO).

4.3.4 Receding horizon decentralised model

In this subsection, the decentralised model from Subsection 3.2.3 is adapted following Subsection 4.3.2 to create a decentralised receding horizon model. In accordance with the receding horizon scheme explained in Section 4.3, a sub-model is solved multiple times for a set of planning windows. Analogously to the approach for the RHCO model, the constraints and functions are all evaluated for the current planning window $\mathcal{H}$. Predicted profiles $\boldsymbol{\chi}_i^{prediction}$ and $\boldsymbol{\xi}_i^{prediction}$ are used together with the dynamic error tolerance $\boldsymbol{\alpha}^{RH}$ in accordance to Subsection 4.3.2.

Algorithm 3 LP-RH Algorithm

```
λ ← λ_flat
t_c ← t_1
repeat
    x, λ = UPDATEPROFILES(v_mo, x, λ)
    p ← p + 1
until p > N
Evaluate f_cost for T
```

The above updates are applied to the optimisation model described in Subsection 3.2.3 and presented in Algorithm 3. Note that not all details described above are explicitly visible in the pseudocode representations. The initialisation of the flat price profile is given in line 1. In line 2, the first planning horizon $\mathcal{H}$ is initialised by specifying the present time slot t_c. Then, the receding horizon execution is modelled in lines 3-6 where the sub-problem is solved in line 4 and the rolling of the horizon in line 5. Finally, the full set of realisations is used to calculate the value of the cost function f_{cost} over the total time interval in line 7. The resulting model is called the *receding horizon local pricing* model and denoted by the abbreviation RHLP.

5 Numerical experiments

Using the theoretical background from Chapter 2 and the models from Chapters 3 and 4, a number of numerical experiments are carried out in this chapter. An experiment setup is first given in Section 5.1 including the considered models and their role in the analysis, the relevant performance criteria, and a detailed (general) case description. Specific experiments are introduced, including results and a brief discussion for each of them, in the Sections 5.2, 5.3, and 5.4.

5.1 Experiment setup

In this section, the setup of the numerical experiments is given. The used models are briefly summarized in Subsection 5.1.1. In Subsection 5.1.2, the different measures are introduced for comparing the models and assessing overall performance. Finally, in Subsection 5.1.3, a general case description is presented where all the settings for the models are described that are specific to the analysis in this chapter.

5.1.1 Models

In Chapters 3 and 4, four optimisation models have been introduced. In the following numerical experiments, three out of these four models are compared based on the performance criteria described in the next subsection. The three models and the accompanying role in the analysis are as follows:

- **Perfect information centralised optimisation (PICO)**: This model corresponds to the centralised optimisation model from Section 3.2.2 with global information. The perfect information refers to the fact that no uncertainty in future predictions is simulated and the complete system state and future states are known to the central decision maker at all times. This model has a benchmark role as the resulting objective function value represents a lower bound to that of any other candidate schedule.
- **Receding horizon centralised optimisation (RHCO)**: This model is an adaptation of PICO where uncertainty is incorporated through a receding horizon scheme. This centralised optimisation model is described in detail in Section 4.3.3. Due to the uncertainty following from the receding horizon, its overall performance cannot surpass that of its perfect information counterpart.
- **Receding horizon local pricing (RHLP)**: This is the decentralised optimisation model that is described in Section 4.3.4 and is based on local information. In addition to the uncertainty associated with a receding horizon scheme, the decentralised nature of the RHLP model limits the control structure compared to centralised models (see also Section 3.3 for a detailed discussion on this comparison). Because the RHLP model uses a novel transactive control approach, its behaviour is analysed in more detail than that of the two centralised models in this chapter.

5.1.2 Performance criteria

In Chapters 3, the dispatch problem is described in detail and two optimisation models are introduced. In Section 3.1, the goal of these models was formulated as follows:

> "*To find a dispatch that follows a target demand profile in a grid feasible way while simultaneously being economic and sustainable*"

To address this goal in a way that is as realistic as possible, in Chapter 4, the models are adapted to include real-life uncertainty. In the numerical experiments, the following performance criteria are chosen to analyse and compare the different models:

- **Feasibility**: For an EM approach to be effective, it should function in a broad range of scenarios. Feasibility is addressed by setting a maximum time for the models to run (stopping time), and checking whether a solution is found that respects all constraints. This is done for each model and a representative set of cases in Section 5.2.

- **Computation time** Computation time of an EM approach may be decisive in the process of ensuring grid stability. The allowed computation time is often limited in real-life EM, and models can account for this by having proper stopping criteria. The resulting trade-off between maximum computation time and quality of the solution is considered in Section 5.3.
- **Energy loss**: *Energy loss* was selected as the main quantitative criteria for measuring to what degree a dispatch is "*economic and sustainable*". The energy loss for the different models is analysed in detail in Section 5.4.

5.1.3 Case description

To go from the general models described in Chapters 3 and 4 to the specific case considered in this chapter, specific parts of the model need to be addressed. This includes the numerical values for the parameters and specification of the grid topology. Next, the global and local cost functions are considered for the centralised and decentralised models respectively. The subsection ends with a description of the data set that is used for this case study.

NUMERICAL SPECIFICATIONS

For the numerical analysis, the "IEEE low voltage test feeder" is used, which is a well-known benchmark LV grid (see Figure 11) [35]. The grid contains a total of 54 household connections, visualised by red circles in the figure. A set of six congestion nodes is added for the present analysis which are indicated in the figure as black diamonds. The market operator is represented by a solid blue circle at the root of the network. Note that Figure 11 only shows the households (red circles), but that each household may contain multiple *independent* devices. As a consequence, only device nodes are considered as leaves and households are not explicitly described in the model.

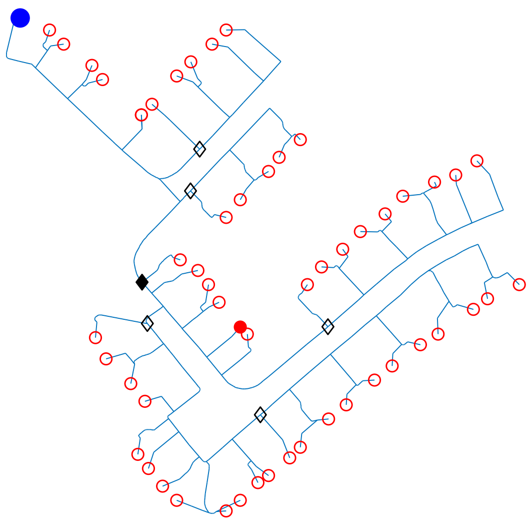

FIGURE 11: *The distribution network topology. Schematic visualisation adapted from IEEE low voltage test feeder. [35].*

The set of available devices depends on the allocation sampling and is considered in more detail in Section 5.4. All devices are dimensioned in a way that reflects realistic current market standards. The batteries $i \in \mathcal{I}_{bat}$ are all assigned a storage capacity of 10.8 kWh, meaning that in terms of the bounds of the SoC model we get $S_i^{min} = 0$ and $S_i^{max} = 10.8$ kWh. The maximum battery charge and discharge rates are given by $X_i^{min} = -4000$ W and $X_i^{max} = 4000$ W respectively and battery charge/discharge efficiency were all set at $\eta_i = 0.90$. This corresponds to a round-trip efficiency of 81% which is realistic for modern day household batteries [42]. The heat pumps $i \in \mathcal{I}_{HP}$ are assumed to have an identical working energy capacity *range* of 2.0 kWh which, after multiplying by the coefficient of energy performance, corresponds to the heat capacity difference between the minimum and maximum comfortable user temperatures. The resulting SoC bounds are then given by $S_i^{min} = 0$ and $S_i^{max} = 2.0$ kWh. The heat pumps can only be charged (not discharged) and are estimated to have $X_i^{min} = 0$ and $X_i^{max} = 1600$ W. A constant leakage rate was assumed and set at $\beta_i = 360$ W for all heat pumps.

The congestion nodes are dimensioned in a way that reflects realistic grid limitations. Upper and lower bounds P_j^{min} and P_j^{max} range somewhere between ± 15 kW and ± 150 kW depending on their location in the grid. Note that the bounds are symmetric, i.e. $P_j^{min} = -P_j^{max}$. In this thesis, the target profile $\boldsymbol{\Theta}$ at the market operator represents a day-ahead prediction of the total load in the distribution grid. It is assumed here that the average of the total profiles for the given month can reasonably model such a prediction. The target is therefore calculated by summing the monthly average of all devices in the network, i.e. PVs, loads, and heat pump leakage. Error margins are provided to the target constraints through α in PICO, and $\boldsymbol{\alpha}^{RH}$ in RHCO and RHLP. The value of α is defaulted at 4000 W for both parameters and $\boldsymbol{\alpha}^{RH}$ is described by equation (20) with $\alpha^+ = 4000$ W. This means that for the receding horizon models, the error margin at the present time slot t_c is given by $\alpha = 4000$ W and the error margin at the last time slot of the planning window $t_c + N - 1$ is given by $\alpha + \alpha^+ = 8000$ W (see also Subsection 4.1.3). The sequence ν_k is chosen in such a way that the error margins increase linearly with increasing k, i.e. $\nu_t = \frac{t-t_c}{N-1}$.

All simulations were run on an i7 processor with a clockspeed of 2.50GHz (8 CPUs).

COST FUNCTIONS

In this paragraph the specific choices of cost functions are specified. This includes the global cost function f_{cost}, which is the objective function of the PICO and RHCO models. It is set to represent the total loss of the dispatch as mentioned in Subsection 5.1.2. Furthermore, the local cost functions $C_i(\boldsymbol{\lambda}_i, \boldsymbol{p}_i)$ for each agent $i \in \mathcal{I}$ in the RHLP model need to be described. In the following, both the device specific contributions to f_{cost} for PICO and RHCO and the local cost functions for RHLP are given for each device category:

- **Loads**: As loads are uncontrollable, they do not have any loss associated with them. The device specific loss, written as f_{cost}^{load}, is therefore assumed to be exactly zero and not contributing to the total loss f_{cost}.

  No local cost function is used at the inflexible load agents because there is no flexibility and therefore no local optimisation process.

- **PVs**: For PVs, the energy loss f_{cost}^{PV} is defined as the amount of lost energy that is incurred from curtailing the PV units. This loss is defined as

$$f_{cost}^{PV} = \sum_{i \in \mathcal{I}_{PV}} \sum_{t \in \mathcal{T}} (x_{it} - 1) \cdot \chi_{it} \cdot \tau, \tag{22}$$

  where we note that the PV yields χ_{it} are negative by definition. The local cost function for the PV units is based on a constant operational cost c_i^{PV}. This cost is incurred for each time slot in which the PV panels are turned on (i.e. not curtailed). To prevent synchronised behaviour between the different PV-panels, this cost is randomly generated for each PV agent from an interval $[0, 0.2]$ with a uniform probability distribution. The local cost functions are then defined as

$$C_i(\boldsymbol{\lambda}_i, \boldsymbol{x}_i) = x_{it} \cdot (c_i^{PV} + \chi_{it}\lambda_{it}). \tag{23}$$

  The decision rule for arbitrary x_{it} is then given by

$$\underset{x_{it} \in \mathcal{F}}{\operatorname{argmin}}\, C_i(\boldsymbol{\lambda}_i, \boldsymbol{x}_i) = \begin{cases} 1, & \text{for } \lambda_{it} < -\frac{c_i^{PV}}{\chi_{it}}, \\ 0, & \text{for } \lambda_{it} \geq -\frac{c_i^{PV}}{\chi_{it}}, \end{cases} \qquad \forall i \in \mathcal{I}_{PV}, t \in \mathcal{T}. \tag{24}$$

  The decision is made on economic incentive such that it will curtail if the expected benefits of selling the produced energy does not outweigh the incurred costs.

- **Batteries**: The batteries have an efficiency given by $\eta_i = 0.90$. The power *transmitted* to the grid is described by p_{it} in constraint (4) for $i \in \mathcal{I}_{bat}$ and $t \in \mathcal{T}$ or $t \in \mathcal{H}$ depending on the model. The power *absorbed* by the battery is described by the decision variable x_{it} in the same equation, and their difference $|p_{it} - x_{it}|$ represents the loss for battery i and time slot t. To keep the centralised models in a linear form, the power absorbed by the battery x_{it} was split into a charging decision variable $x^c_{it} \in [0, X^{max}_i]$ and a discharging decision variable $x^d_{it} \in [X^{min}_i, 0]$ (see Subsection 3.2.2). For the total set of batteries, the loss f^{bat}_{cost} is defined as

$$f^{bat}_{cost} = \sum_{i \in \mathcal{I}_{bat}} \sum_{t \in \mathcal{T}} ((\frac{1}{\eta} - 1) \cdot x^c_{it} - (1 - \eta) \cdot x^d_{it})\tau, \tag{25}$$

for the centralised optimisation models (PICO and RHCO).

For the battery agents in RHLP, a decision rule needs to be specified. To reflect the economic aspect of realistic user behaviour, batteries should charge below and discharge above some price threshold. With only this rule, batteries would either charge at the maximum rate below the threshold or discharge at the maximum rate above the threshold. To help prevent this extreme behaviour, two transition intervals $[c^{bat}_1, c^{bat}_2]$ and $[c^{bat}_3, c^{bat}_4]$ are instated where the rate of charge/discharge depends linearly on the price. In the interval $[c^{bat}_2, c^{bat}_3]$, the battery neither charges nor discharges. The battery agent decision rule is now given by

$$\operatorname*{argmin}_{x_{it} \in \mathcal{F}} C_i(\boldsymbol{\lambda}_i, \boldsymbol{x}_i) = \begin{cases} X^{max}_i, & \text{for } \lambda_{it} \in [0, c^{bat}_1], \\ X^{max}_i \cdot \frac{c^{bat}_2 - \lambda_{it}}{c^{bat}_2 - c^{bat}_1}, & \text{for } \lambda it \in (c^{bat}_1, c^{bat}_2], \\ 0, & \text{for } \lambda it \in (c^{bat}_2, c^{bat}_3], \\ X^{min}_i \cdot \frac{c^{bat}_3 - \lambda_{it}}{c^{bat}_4 - c^{bat}_3}, & \text{for } \lambda it \in (c^{bat}_3, c^{bat}_4], \\ X^{min}_i, & \text{for } \lambda it \in (c^{bat}_4, 1], \end{cases} \quad \forall i \in \mathcal{I}_{bat}, t \in \mathcal{T} \tag{26}$$

where it is noted that $\lambda_{it} \in [0, 1]$ because LPRH uses normalised price profiles (see Subsection 3.2.3). A graphical representation of this decision rule is included in Figure 12. For the current case study, the cost values are calculated from the efficiency parameter η to model the desired behaviour: $c^{bat}_1 = 0.26, c^{bat}_2 = 0.36, c^{bat}_3 = 0.64, c^{bat}_4 = 0.74$.

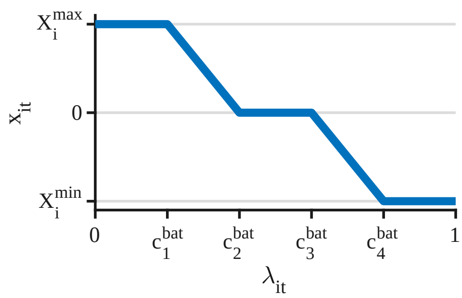

FIGURE 12: *Decision rule showing the battery charge decision x_{it} as a function of the local price λ_{it}.*

- **Heat pumps**: The heat pumps are assumed to have a perfect charge efficiency. The leakage rate β_i is not included in the total loss because it is fixed and cannot be optimised in the model. Therefore, no loss is associated with the operation of heat pumps.

For the heat pump agents in the RHLP model, the local cost function is essentially a special case of the battery cost function. Decision rule (26) holds for heat pumps with $X^{min}_i = 0$, which results in

$$\operatorname*{argmin}_{x_{it}\in\mathcal{F}} C_i(\boldsymbol{\lambda}_i, \boldsymbol{x}_i) = \begin{cases} X_i^{max}, & \text{for } \lambda_{it} \in [0, c_1^{HP}], \\ X_i^{max} \cdot \frac{c_2^{HP} - \lambda_{it}}{c_2^{HP} - c_1^{HP}}, & \text{for } \in (c_1^{HP}, c_2^{HP}], \\ 0, & \text{for } \lambda_{it} \in (c_2^{HP}, 1]. \end{cases} \qquad \forall i \in \mathcal{I}_{HP}, t \in \mathcal{T} \tag{27}$$

In the present case study we use the following parameters for heat pumps: $c_1^{HP} = 0.40$, $c_2^{HP} = 0.60$. A graph of the decision rule is included in Figure 13.

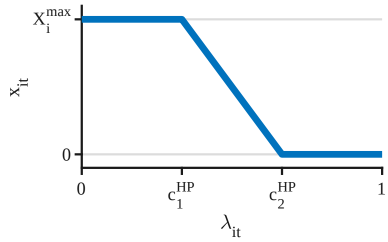

FIGURE 13: *Decision rule showing the heat pump charge decision x_{it} as a function of the local price λ_{it}.*

Using the device specific energy losses, the combined total loss function, denoted f_{cost}, is given as

$$f_{cost} = f_{cost}^{PV} + f_{cost}^{bat}. \tag{28}$$

Considering the local cost functions, we note that they do not directly describe losses. However, they are designed in a way that should reflect realistic user behaviour while simultaneously leading to a system wide minimisation of the total loss given by f_{cost}. The values of the cost parameters for the heat pump, battery, and PV agents are calculated in a way that prioritises operation of more efficient devices. In the present case study, this means that price profiles are used to first exhaust the flexibility of heat pumps (100% efficiency), then that of the batteries (90% efficiency), and finally consider PV curtailment (0% efficiency). The order of prioritisation will stay this way as long as we make sure to choose $c^{PV} < c_2^{bat} < c_1^{HP} < c_2^{HP} < c_3^{bat}$.

MODEL DATA GENERATION

The data set that is used for the present analysis was collected during a pilot study in [32] under the name of "*energy front runners*" (*Energie koplopers* in Dutch) that took place in the Dutch town of Heerhugowaard. In the study, a real-life prototype smart grid was developed including a physical grid that connected more than 200 households. Households were equipped with solar panels, heat pumps, electric boilers and other smart devices. The study was carried out in two phases spanning over the period between 2013 and 2018 during which a large amount of data was gathered. The data that was made available for the present study represents a subset that includes power production data of *PV units* and power consumption data of *total household loads* for around 90 households. A total of 11 months have been made available (February until December 2018) and the power outputs are (generally) given with a resolution of five minutes. Unfortunately, a significant fraction of the data is missing. After pre-processing of the data, a selection of 27 households was made for all of which 93 days had a sufficiently complete and non-erroneous set of data. These dates are unevenly spread out over the months as can be seen in Table 3.

To model uncertainty in load and PV predictions, the results from Section 4.1 are used. For loads, predictions were reported to be typically associated with a mean relative error (MRE) between 20% and 25%. For PVs, it was reported that depending on weather type, prediction MREs varied

Month (of 2018)	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Number of usable days	0	0	8	4	18	0	27	16	20	0	0

TABLE 3: *Table showing the amount of days for each month for which all of the households had a usable dataset.*

between 4.86% and 12.42%. In this case description, the upper limits of these ranges are used. This means that MREs of 25% and 12.42% are applied to the loads and PVs respectively. In order to simulate these uncertainties, equation (18) is used for the loads, and equation (19) for the PVs.

We assume that for each model run, the date (day and month) is known specifying the available data sets $\boldsymbol{\chi}$ and $\boldsymbol{\xi}$. The realisations of the power outputs corresponding to this date are given per time slot $t \in \mathcal{H}$ by χ_{it} for static loads $i \in \mathcal{I}_{load}$ and ξ_{it} for PVs $i \in \mathcal{I}_{PV}$ for each household. These are the "perfect information" power values that are used in the PICO model. To model uncertainty in the RHCO and RHLP models, equations (18) and (19) calculate the predictions $\chi_{it}^{prediction}$ and $\xi_{it}^{prediction}$ by combining these perfect information values with "general data". As was explained in Subsection 4.1.3, monthly averages for a given user from the data set are used to represent this general data. This means that the average of all χ_{it} and ξ_{it} values for a given month and fixed i and t is calculated from the data set. Now, $\overline{\chi}_{it}$ and $\overline{\xi}_{it}$ represent these average power outputs for static loads $i \in \mathcal{I}_{load}$ and PVs $i \in \mathcal{I}_{PV}$ respectively during a time slot $t \in \mathcal{H}$. The increase in uncertainty as a function of the number of time slots into the future was hypothesised to be best modeled by a concave function, i.e. the level of uncertainty increases quickly at first and more slowly as the considered horizon grows larger. This behaviour has also been observed in other studies, e.g. in [20, 41] for wind- and in [6] for solar power generation. A square root dependency on time was chosen as a simple concave function for the two sequences μ_k^{load} and μ_k^{PV}. It is described by

$$\mu_t^{load} = q^{load} \cdot \sqrt{\frac{t - t_c}{N - 1}}, \qquad \forall t \in \mathcal{H}, \tag{29}$$

for static loads, and by

$$\mu_t^{PV} = q^{PV} \cdot \sqrt{\frac{t - t_c}{N - 1}}, \qquad \forall t \in \mathcal{H}, \tag{30}$$

for PVs where in both cases N is the number of time slots in the planning window $\mathcal{H}$. It has been verified empirically that to achieve the mentioned MREs over the complete set of possible days, we should have $q^{load} = 0.35$ and $q^{PV} = 0.60$.

5.2 Feasibility

The goal of the feasibility assessment is to provide a clear overview of how the set of feasible solutions depends on the chosen constraints and parameters within the models. Furthermore, the results give information about how well the different models are equipped to function in a range of scenarios. The effects of the different categories of constraints are addressed by varying their strictness using additional artificial parameters. These experiments are described in Subsections 5.2.1 - 5.2.3. The three models are run for each of the 93 available data profile sets (days). As a stopping criterion for the RHCO and RHLP models we use a maximum time of 5 seconds per receding horizon iteration. Because the PICO model is used as a benchmark, it is given a more generous time limit of 60 seconds. Upon completion, the percentage of runs that result in a feasible solution is compared for each of the three models and analysed in Subsection 5.2.4.

5.2.1 Device constraints

A general description of the device behaviour was given in Subsection 3.2.1 and specifically in the equations (2)-(11). Out of these equations, most are either definitions or auxiliaries to other equations, and only a subset of these equations may strictly be interpreted as *device constraints*, namely: (5), (7), (10), (11). These four constraints describe the allowed charge rates and SoC states of the batteries and heat pumps. They can be combined for both of these device classes into the following two constraints:

$$x_{it} \in [X_i^{min}, X_i^{max}], \qquad \forall i \in \mathcal{I}_{bat} \cup \mathcal{I}_{HP}, t \in \mathcal{T}$$
$$s_{it} \in [S_i^{min}, S_i^{max}], \qquad \forall i \in \mathcal{I}_{bat} \cup \mathcal{I}_{HP}, t \in \mathcal{T}$$

The bounding parameters are the lower and upper bounds on the charge rate, described by X_i^{min}, X_i^{max} respectively, and the lower and upper bounds on the SoC described by S_i^{min}, S_i^{max} where $i \in \mathcal{I}_{bat} \cup \mathcal{I}_{HP}$. To get an idea of the effect of the above constraints on model feasibility, a parameter $\gamma_1 \in [0, 1]$ is introduced and used to modify the bounds. The modified constraints are described by

$$x_{it} \in [\gamma_1 \cdot X_i^{min}, \gamma_1 \cdot X_i^{max}], \qquad \forall i \in \mathcal{I}_{bat} \cup \mathcal{I}_{HP}, t \in \mathcal{T} \tag{31}$$
$$s_{it} \in [\gamma_1 \cdot S_i^{min}, \gamma_1 \cdot S_i^{max}], \qquad \forall i \in \mathcal{I}_{bat} \cup \mathcal{I}_{HP}, t \in \mathcal{T} \tag{32}$$

The value of γ_1 can then be modified to reflect constraints with varying tightness where lower values of γ_1 correspond to tighter constraints. The effect of varying γ_1 was investigated by running all three models on a total of 93 sample instances for $\gamma_1 \in \{0.4, 0.5, \ldots, 1\}$, i.e. we consider only cases where the bounds are further tightened compared to the base case. The achieved results are found in Figure 14, where the percentage of instances for which a feasible solution was found within the computation time is given. These results are discussed in Subsection 5.2.4.

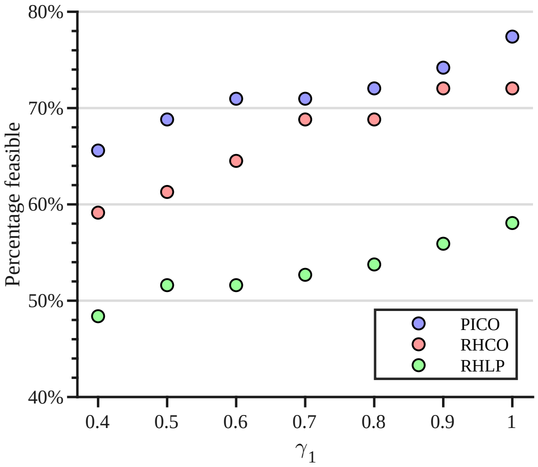

FIGURE 14: *Results from feasibility experiment for device constraints.*

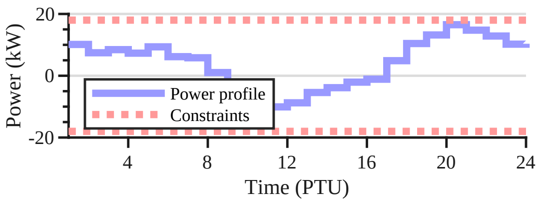

FIGURE 15: *Congestion feasibility.*

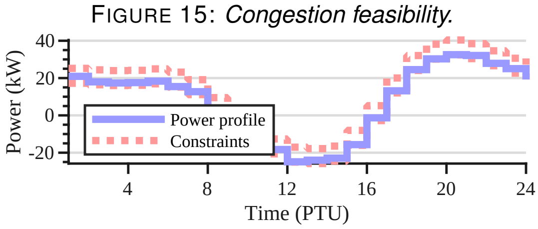

FIGURE 16: *Market feasibility.*

5.2.2 Congestion constraints

In a next step, we consider the congestion constraints, described in Subsection 3.2.1 in equation (13) for each congestion node. Similarly to the device constraints, the tightness depends on the bounds P_{jt}^{min} and P_{jt}^{max}. In the case of the congestion constraints, these bounds correspond to the minimum and maximum power level that each congestion node $j \in \mathcal{J}$ can withstand. To give a visual aid to the constraints, the power profile and bounds for a given congestion node are plotted for an example run in Figure 15. Note that this example constitutes a feasible run, the constraints are both respected. Similarly to the device constraints, a parameter $\gamma_2 \in [0, 1]$ is introduced to modify the tightness of these constraints. The modified constraint are given by

$$p_{jt} \in [\gamma_2 \cdot P_{jt}^{min}, \gamma_2 \cdot P_{jt}^{max}], \qquad \forall j \in \mathcal{J}, t \in \mathcal{T}. \tag{33}$$

Again, the modified constraint is applied to the three models for the values $\gamma_2 \in \{0.4, 0.5, \ldots, 1\}$. The percentage of instances with a feasible solution are again documented and compared for the different model runs and γ_2 values. The results are shown in Figure 17 and discussed in Subsection 5.2.4.

5.2.3 Market constraints

Finally, the market constraints are considered. An example of a typical power profile at the market operator node is given in Figure 16 together with the bounds corresponding to equation (16) from Subsection 3.2.1. These bounds are adapted slightly in the receding horizon models through the dynamic error tolerance vector $\boldsymbol{\alpha}^{RH}$ that was defined in equation (20) in Subsection 4.3.3. Rewriting the general market constraints from (16) using the dynamic tolerance, we get

$$p_{v_{mo}t} \in [\Theta_t - \alpha_t^{RH}, \Theta_t + \alpha_t^{RH}], \qquad \forall t \in \mathcal{T}. \tag{34}$$

Now similarly to the previous constraint classes, a parameter $\gamma_3 \in [0, 1]$ may be defined that modifies the tightness of (34). This modification can be described by the following general constraint given by

$$p_{v_{mo}t} \in [\Theta_t - \gamma_3\alpha_t^{RH}, \Theta_t + \gamma_3\alpha_t^{RH}], \qquad t \in \mathcal{T}. \tag{35}$$

For the PICO model, the resulting adapted market constraint is given by

$$p_{v_{mo}t} \in [\Theta_t - \gamma_3\alpha, \Theta_t + \gamma_3\alpha], \qquad t \in \mathcal{T}. \tag{36}$$

To investigate the impact of the tightness of these constraints, the models are run for the full set of profiles and values $\gamma_3 \in \{0.25, 0.5, 0.75, 1\}$. The results of this experiment are shown in Figure 18 and discussed in Subsection 5.2.4.

5.2.4 Discussion

The results of the feasibility study are presented in Figures 14, 17, and 18. For all three constraint types and settings, the ranking of the three models is the same: the PICO model always has the highest percentage of feasible runs, followed by the RHCO model, and finally the RHLP model. For the full set of runs considered in this section, the percentage of feasible runs for the three models is given by: 74% for PICO, 70% for RHCO, and 54% for RHLP. The difference between the centralised models and the decentralised model can, at least in part, be attributed to the distinction between the feasible and reachable set of solutions that was discussed in Section 3.3. This leads to an effectively smaller feasible set for the decentralised model compared to the centralised models. For the PICO and RHCO models, the difference in percentage of feasible runs can be explained by two main reasons. Firstly, the additional uncertainty in future predictions for the RHCO model can lead to infeasible runs whereas the PICO model always has perfect predictions to optimise over. A decision that was optimal for an (uncertain) set of future predictions is not

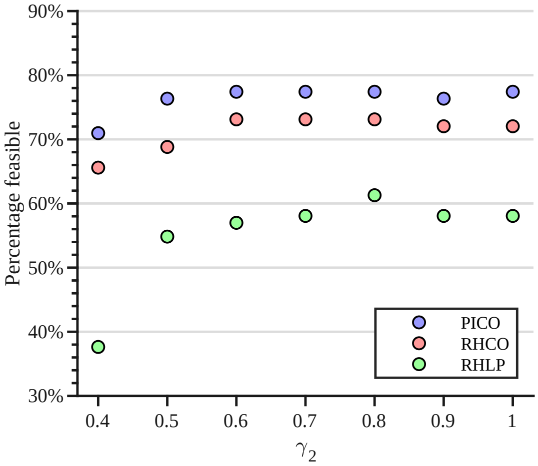

FIGURE 17: *Results from feasibility experiment for congestion constraints.*

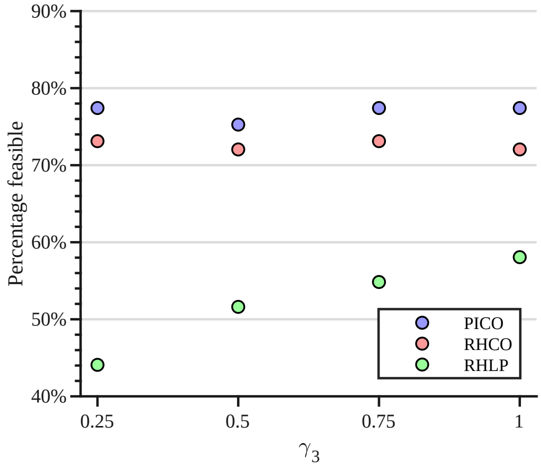

FIGURE 18: *Results from feasibility experiment for market constraints.*

necessarily optimal when the realisations become known, and may even become infeasible. Secondly, the PICO model was given a benchmark role and assigned a stopping time of 60 seconds whereas each RHCO iteration was only given 5 seconds to finish, implications of these stopping criteria are considered in detail in Subsection 5.3.2. Referring now to the specific constraint sets, the following results are observed:

- **Device constraints**: For all three models, a monotonically increasing dependency on γ_1 is observed. Increasing γ_1 increases the total "flexibility capacity" by modifying the SoC ranges and at the same time increases the flexibility potential per time slot by modifying the charge/discharge limits of the storage devices. In this way, the value of γ_1 directly affects peak shaving and overall balancing potential of the models. It is therefore not surprising that the results for the device constraint run as seen in Figure 14 show the steepest increase over the range of constraint settings compared to the results from Figures 17 and 18.
- **Congestion constraints**: For the congestion constraint runs, none of the models has a monotonically increasing dependency on γ_2. The percentage of feasible model runs seems to increase between $\gamma_2 = 0.4$ and $\gamma_2 = 0.6$ and to then stabilise with one small outlier at $\gamma_2 = 0.9$. From this behaviour we can conclude that there is a percentage of runs (20% for PICO, 25% for RHCO, and 40% for RHLP) where infeasibility seems to come from market or device constraints instead of congestion constraints. The outlier might be explained by the stopping criteria: because a time limit is imposed (instead of e.g. an iteration limit), the model runs are not completely reproducible.
- **Market constraints**: For the market constraint runs, a large variation in model behaviour is observed. For the centralised models, no significant change in feasibility percentages is seen for the different values of γ_3. In contrast, a monotonic increase in percentage of feasible runs is observed for the LPRH model, from 44% to 58% over the γ_3 domain. For the RHLP model, larger error margins on the market constraints seem to significantly increase the reachable set for each iteration. For the centralised models, it seems that market constraints stop leading to infeasibilities beyond γ_3 values of around 0.25.

5.3 Computation time

In this section, two main concepts related to computation time of the different models are considered. The goal of the section is to find out what happens between start and termination of the optimisation. This may help to determine proper stopping criteria and to address differences in behaviour between the models. In Subsection 5.3.1, we analyse the evolution of different model values for the LPRH model specifically. Then, in Subsection 5.3.2, an analysis is included for addressing the effect of increasing the stopping time of the different models. A discussion is included in Subsection 5.3.3. We note that the computation time is heavily dependent on factors out of scope of the analysis, such as the specifics of the implementation and the hardware used.

5.3.1 Optimisation paths

Because the (decentralised) LPRH model represents a novel EM approach, it is important to carefully consider its behaviour. In this subsection, we use so-called "optimisation paths" to carry out this analysis. These paths show the evolution of the objective function value and constraint violations as a function of the iteration number for a given model run. The shape of the optimisation paths may give us information about the model behaviour and help determine proper stopping criteria. The optimisation paths of a set of model runs consisting of three separate days are considered. These days are chosen in a way that is representative of the complete data set:

- **Normal day**: The term "normal day" is used to denote typical days where all three optimisation models are able to find a feasible solution within the default stopping time of 5 seconds. The 12th of August was selected for this experiment to represent a normal day.
- **Challenging day**: The set of "challenging day" refers to days where a feasible dispatch may be found by all of the optimisation models, but not within the default stopping time of 5 seconds. For the centralised models, optimal solutions are typically not obtained. The 26th of April was selected to represent a challenging day.
- **Impossible day**: The term "impossible day" refers to days where none of the optimisation models is able to find a feasible dispatch at any stopping time. In this case, the uncontrollable devices (PVs, inflexible loads) have profiles that lead to constraint violations that cannot be mitigated by the set of flexible devices. The 9th of May was selected to represent an impossible day.

For the three selected days, an experiment is run using the LPRH model to visualise the optimisation paths for the total loss (see Figure 19), the congestion constraints (see Figure 20), and the market constraints (see Figure 21). The device constraints are not considered as they are not allowed to be violated during the optimisation process. In the following analyses, we often apply the tightest congestion constraint settings, i.e. $(\gamma_1 = 1, \gamma_2 = 0.4, \gamma_3 = 1)$, to make sure that feasibility aspects occur in the model runs. In the following, this setting is referred to simply by *Tight congestion case* for clarity. The tight congestion case is also applied here, running each experiment for a single planning window $\mathcal{H}$, i.e. one iteration of the receding horizon scheme, with starting time t_c = 00:00h. The results are discussed in Subsection 5.3.3.

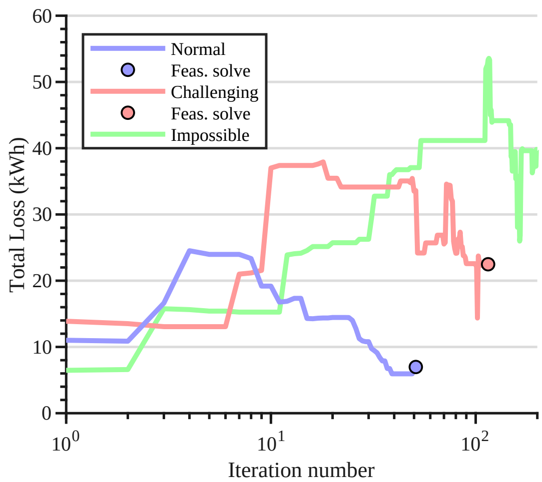

FIGURE 19: *Evolution of the total loss plotted as a function of the iteration count for three sample days.*

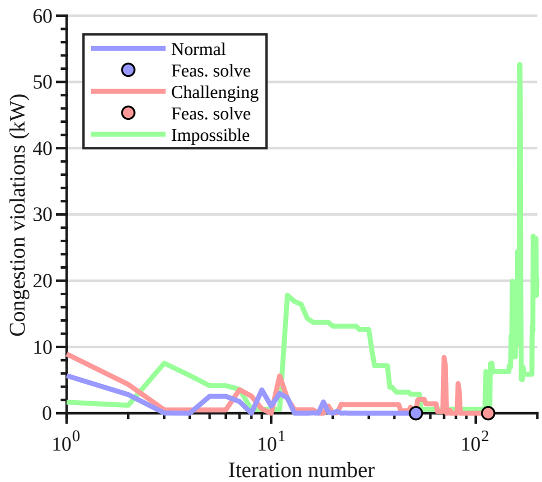

FIGURE 20: *Evolution of the congestion constraint violations plotted as a function of the iteration count for three sample days.*

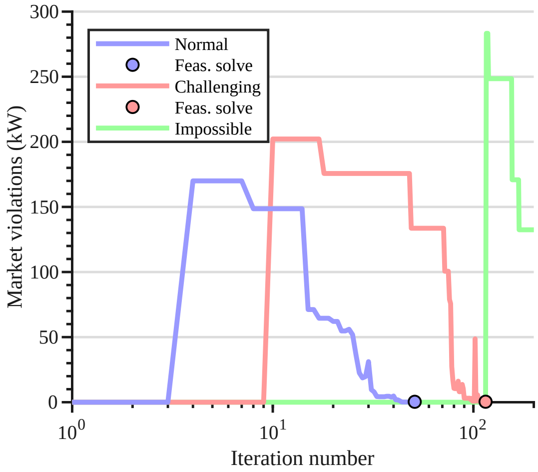

FIGURE 21: *Evolution of the market constraint violations plotted as a function of the iteration count for three sample days.*

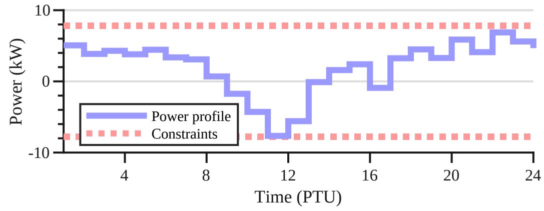

FIGURE 22: *Congestion node power.*

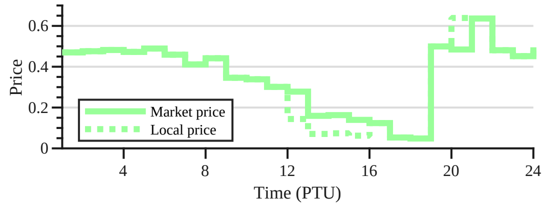

FIGURE 23: *Congestion node price.*

Model	Constraint settings	Percentage feasible (stopping time of 5s)	Percentage feasible (stopping time of 10s)	Difference
RHCO	$(\gamma_1 = 0.4, \gamma_2 = 1, \gamma_3 = 1)$	59%	61%	+2% (2 runs)
	$(\gamma_1 = 1, \gamma_2 = 0.4, \gamma_3 = 1)$	65%	65%	–
	$(\gamma_1 = 1, \gamma_2 = 1, \gamma_3 = 0.25)$	73%	74%	+1% (1 run)
RHLP	$(\gamma_1 = 0.4, \gamma_2 = 1, \gamma_3 = 1)$	48%	52%	+4% (3 runs)
	$(\gamma_1 = 1, \gamma_2 = 0.4, \gamma_3 = 1)$	38%	38%	–
	$(\gamma_1 = 1, \gamma_2 = 1, \gamma_3 = 0.25)$	44%	45%	+1% (1 run)

TABLE 4: *The effect of increasing the stopping time is reported for the two receding horizon models. A model run is executed for the full set of available days for the three strictest constraint settings.*

5.3.2 Stopping criteria

Another important aspect of optimisation is the stopping criterion. Especially when solving NP-hard problems such as the dispatch problem considered here, an optimal solution may not be obtained within a reasonable amount of time. However, limiting the allowed computation time may have a negative impact on the quality of the solution. This leads to a trade-off that needs to be considered when selecting a stopping criterion. Possible examples of stopping criteria are: (relative) gap (for MIPs), maximum number of iterations, and computation time. In the context of EM, a maximum computation time seemed to be the most realistic option. A small disadvantage of this criteria is that it is not entirely reproducible, as the number of iterations that can be completed before a given stopping time may vary due to differences in available processor capacity.

In Section 5.2, stopping times of 5 seconds were used for the receding horizon models, and a more generous stopping time of 60 seconds was used for the "benchmark" PICO model. In the following experiment, the goal is to see whether an increase in stopping time, we choose a doubling from 5 seconds to 10 seconds for the RHCO and RHLP models, leads to significant improvements in the percentage of feasible runs. For each constraint type (device, congestion, and market), the tightest setting from Section 5.2 was selected to run the stopping time experiments, i.e. $(\gamma_1 = 0.4, \gamma_2 = 1, \gamma_3 = 1)$, $(\gamma_1 = 1, \gamma_2 = 0.4, \gamma_3 = 1)$, and $(\gamma_1 = 1, \gamma_2 = 1, \gamma_3 = 0.25)$. This choice was made because the larger percentage of infeasible runs for these experiments enables us to more easily see improvements compared to the other runs. The results are presented in Table 4 and discussed in Subsection 5.3.3.

5.3.3 Discussion

In Subsections 5.3.1 and 5.3.2, experiments were described to analyse behaviour related to computation time for the models. The optimisation paths are plotted in Figures 19, 20, and 21, and the stopping time results are presented in Table 4. Figures 22 and 23 give example power and price profiles to better understand the process of congestion mitigation. A brief discussion about these results is included below:

- **Optimisation paths**: The three optimisation path runs all show a clear distinction between the "day types". For the total loss experiment (see Figure 19), the normal day has an overall lower total loss value and converges fastest out of the three days. The challenging day reaches a total loss that is slightly higher with some oscillating behaviour towards the end of the run. The total loss value for the impossible day keeps increasing over the plotted domain and shows oscillating behaviour that is even more pronounced. The shape of the optimisation paths for the total loss is quite similar between the three models. In the first iterations, the total loss remains stable, it then increases and finds a relatively stable value again around its maximum, after which it decreases again (during which some oscillation may occur) and reaches its final value.

  Turning to the market constraints (see Figure 21), the optimisation paths show similar behaviour: first a stable region, followed by increasing violations, a stable region again, and finally the decrease to its final (feasible) solution. For the normal and challenging days, this shape seems to coincide fairly well. However, the impossible instance only has its increase in market violations around iteration number 100 whereas the total loss starts increasing steadily from iteration 2. It seems that, at least for the normal and challenging runs, resolving market constraint violations coincides with decreasing of the total loss value.

  The congestion constraints (see Figure 20) start out with a nonzero violation for all three days. The process of resolving these violations is probably responsible for the regions of increased total loss and market constraint violations seen in Figure 19 and Figure 21 respectively. The optimisation paths in the congestion constraint plot show a chaotic behaviour consisting of various small peaks, contrasting the behaviour for the total loss and the market constraints. Each of these small peaks corresponds to the onset and subsequent mitigation of congestion somewhere in the grid. This process is shown in Figures 22 and 23 where one of the congestion agents is considered. Congestion occurs at two time intervals (PTUs 12-16 and 20-21), and the model solves it by *locally* changing the price as shown by the dashed line.

- **Stopping criteria**: Overall, doubling the stopping time showed relatively minor improvements in feasibility compared to changing the constraint settings $(\gamma_1, \gamma_2, \gamma_3)$.

  For the case with $(\gamma_1 = 1, \gamma_2 = 1, \gamma_3 = 0.25)$ an increase in feasible runs is observed for both RHCO and RHLP. This increase is not surprising when considering the large amount of iterations required as shown in Figure 21. The 1% improvement for both optimisation models corresponds to a run where the region of increased market constraint violations ends somewhere between the 5 and 10 second marks. The most significant increases are seen for the experiment with $(\gamma_1 = 0.4, \gamma_2 = 1, \gamma_3 = 1)$, i.e. the tightest device constraint setting. As discussed in Section 5.2, γ_1 affects both the target matching and congestion mitigation properties of the optimisation models, so this is not entirely surprising. For $(\gamma_1 = 1, \gamma_2 = 0.4, \gamma_3 = 1)$, no such improvements are observed. Apparently, when an infeasible run occurs due to congestion, increasing the allowed computation time does not help (for the cases considered here).

  Note that the few runs where increasing the stopping time resulted in a feasible solution correspond to the challenging days described in Subsection 5.3.1. For these days a feasible solution could be found, but it takes more than the default stopping time of 5 seconds. On the CPU that ran the optimisation path experiments, a stopping time of 5 seconds corresponds roughly to 70 iterations of the LPRH model, and a stopping time of 10 seconds corresponds to around 140 iterations, confirming the above assertion.

5.4 Total loss

The total (energy) loss is used to evaluate to what extent a solution is “economic and sustainable”. The centralised models (PICO and RHCO) have the total loss as their objective function and directly minimise it. The decentralised model (RHLP) minimises the local cost functions, and evaluates the total loss only after the optimisation. A theoretical discussion of the comparison between the models is given in Section 3.3. The numerical counterpart to this comparison is presented here. We compare for a set of runs the total loss values between the three different models in Subsection 5.4.1. In Subsection 5.4.2, the behaviour of the different device classes is considered for the centralised versus the decentralised models. Lastly in Subsection 5.4.3, the effect of varying device allocation on total loss values is considered for a selection of days.

5.4.1 Quality of solution

In this subsection, a general comparison of the total loss performance for the three models is considered to understand how their solution qualities differ. For this, the full set of runs for all constraint settings from Section 5.2, i.e. all different $(\gamma_1, \gamma_2, \gamma_3)$ instances, are used leading to a total of 1488 runs for each model type. However, for the total loss analysis, we only consider runs where all three models resulted in a feasible solution, leaving a total of 790 runs per model type. For each model type and the full set of (feasible) total loss results, a comprehensive characterization is given in the form of standard boxplots in Figure 24. The boxplots summarise the distribution of results through the five measures outlined below:

- **Median** The median total loss corresponds to the run that lies in the middle of the *ordered* set of total loss values for a given model. It is shown by the coloured vertical midline in each box.
- **First quartile (Q1)** The first quartile is the value that lies in the middle of the ordered set of values between the lowest total loss result and the median total loss result. It is essentially the median of the first half of the ordered set of results. It is shown by the leftmost edge of the box.
- **Third quartile (Q3)** The third quartile is the value that lies in the middle of the ordered set of values between the highest total loss result and the median total loss result. It is shown by the rightmost edge of the box. Furthermore, the inter-quartile range (IQR) is defined as the difference between the first and third quartile, i.e. IQR $=$ Q3 $-$ Q1. The IQR is used as a measure of spread of the data, and is also used to define outliers as described below.
- **Minimum** The minimum is the result with the lowest total loss value that is within a 1.5IQR total loss difference from the median. It is shown by the leftmost vertical line, also called whisker, of the boxplot.
- **Maximum** Similarly, the maximum is the result with the highest total loss value that is within a 1.5IQR total loss difference from the median. It is shown by the rightmost vertical line, also called whisker, of the boxplot.

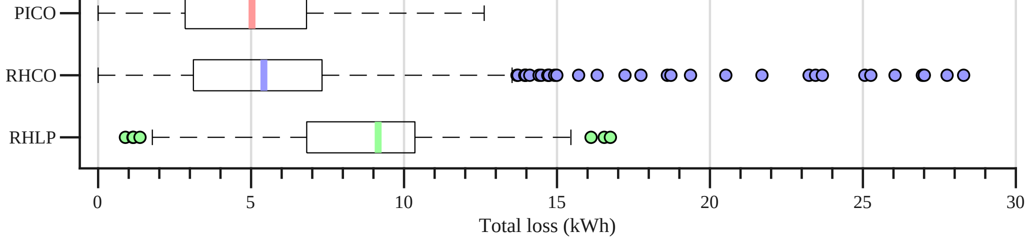

FIGURE 24: *Boxplots of the total loss results of all 790 runs for each of the three model types.*

Note that the minimum and maximum as defined above are not the actual lowest and highest values of the set of results. The 1.5IQR cut-off is chosen to limit the effect of extreme outliers on the boxplot representation. Any total loss values that are more than 1.5 times the IQR above or below the median are considered outliers and shown as coloured circles outside of the boxplots in Figure 24. In addition to the five measures shown in the boxplots, the mean of the distributions are given by 4.8 kWh for PICO, 5.7 kWh for RHCO, and 8.6 kWh for RHLP. In Subsection 5.4.4, the above statistical measures are compared and discussed for the three models.

5.4.2 Device activation

In this subsection, experiments are introduced with the goal of better understanding the behaviour of the different device classes in the centralised versus the decentralised models. We analyse in what way available flexibility is used by considering the level of “device activation”. In this context, device activation is assessed by considering power output of the flexible devices, comparing the usage between different device classes and between different models. From these results, we hope to achieve a better understanding of what causes differences in total loss results between the models.

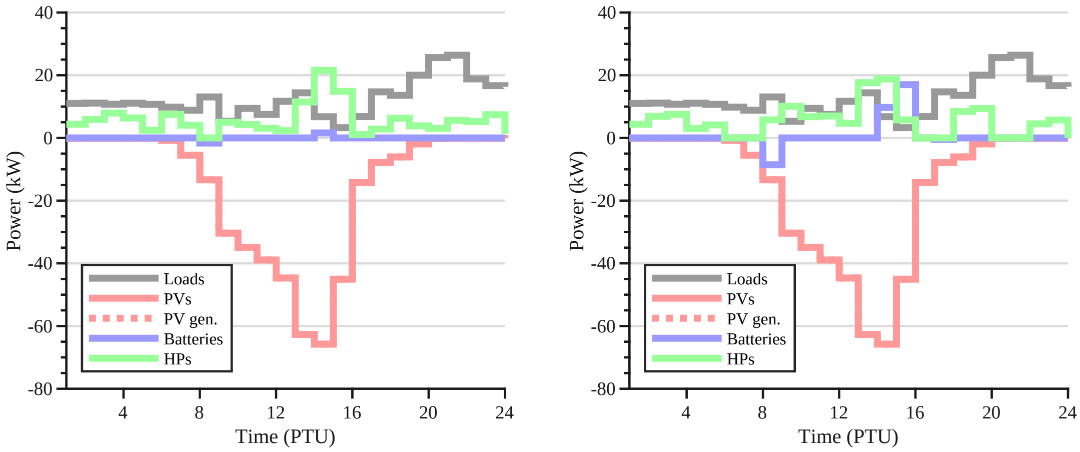

(A) *RHCO* (B) *RHLP*

FIGURE 25: *Comparison of device activation during the normal day (12th of August) for the (centralised) RHCO model in (A) and (decentralised) RHLP in (B).*

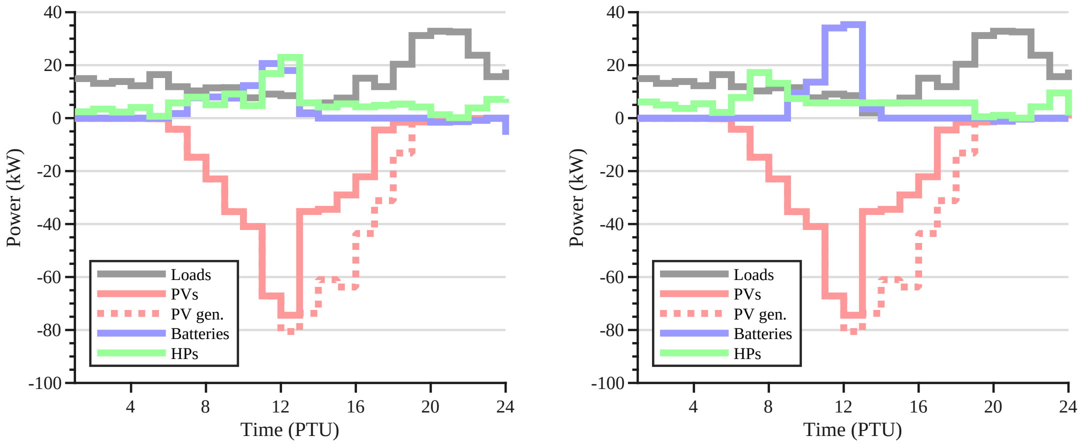

(A) *RHCO* (B) *RHLP*

FIGURE 26: *Comparison of device activation during the challenging day (26th of April) for the (centralised) RHCO model in (A) and (decentralised) RHLP in (B).*

To address device activation, the receding horizon models (RHCO and RHLP) are run for both the “normal day” and the “challenging day”, i.e. the 12th of August and the 26th of April. We focus on these two days in particular because they are assumed to be representative of the full set of feasible solutions, as was explained in Subsection 5.3.1. Also, because the aim of this subsection is to see the difference between centralised and decentralised decision-making, the benchmark role of the PICO model is not needed here. For each of the four model runs described above, the sum of all individual device power profiles is calculated for each of the device types, i.e. $\sum_{i \in \mathcal{I}_{load}} \boldsymbol{p}_i, \sum_{i \in \mathcal{I}_{PV}} \boldsymbol{p}_i$, etc., resulting in a set of cumulative power profiles. Cumulative power profiles are used instead of individual devices to reduce the effect of random noise, and to get a more complete picture for the specific model runs.

The tight congestion case is used, i.e. $(\gamma_1 = 1, \gamma_2 = 0.4, \gamma_3 = 1)$, for all four runs. An overview of the cumulative power profiles is given in Figure 25 for both models during the normal day and in Figure 26 for both models during the challenging day. Note that use of flexibility in the PV panels (i.e. curtailment) is displayed by the difference between the generation profile and the resulting power profile after potential curtailment of the decision maker(s). Furthermore, note that the uncontrollable load "activations" do not represent the use of flexibility (there is none for this class of devices), but are included for completeness and comparison. In Figure 27, we take a closer look at the behaviour of the battery and heat pumps. This is done by "zooming in" on their cumulative power profiles specifically, and comparing the two models in a single plot for the normal day in Figures 27a and 27c and for the challenging day in Figures 27b and 27d. Lastly, the total loss results for each of the four runs are summarised in Table 5, including PICO for completeness. A discussion of these results is included in Subsection 5.4.4.

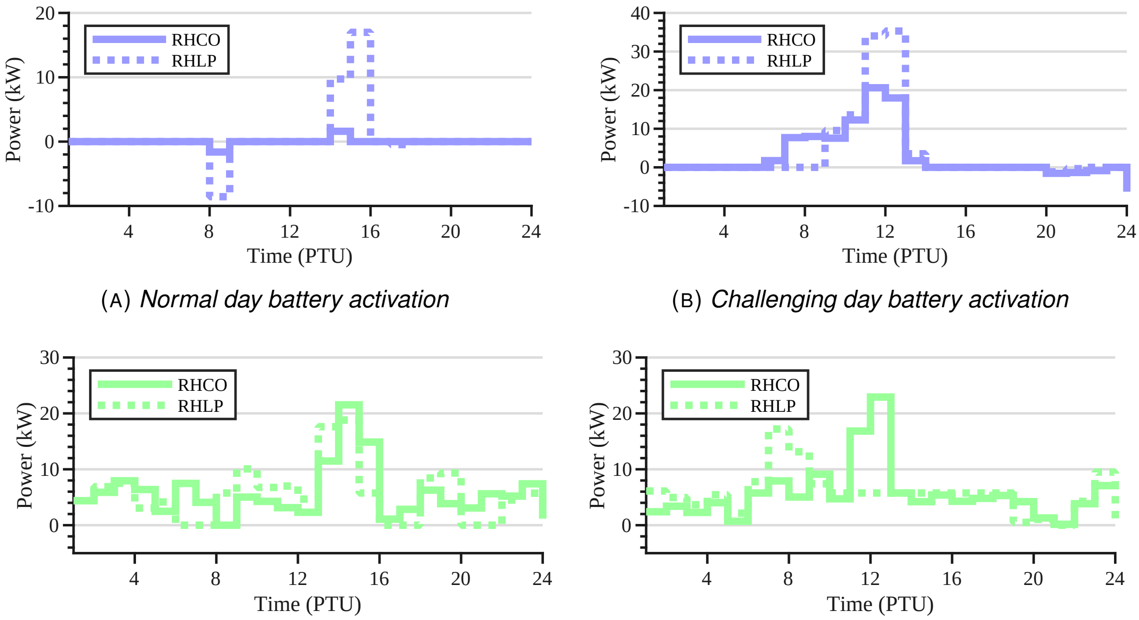

(A) *Normal day battery activation*

(B) *Challenging day battery activation*

(C) *Normal day heat pump activation*

(D) *Challenging day heat pump activation*

FIGURE 27: *Zoomed in plots of specific device activations from the overview plots shown in Figures 25 and 26. Each subplot shows a comparison of activations for the RHCO and RHLP models for a device type and day type.*

5.4.3 Allocation

In the previous sections, the location of the different devices in the model has been kept fixed, i.e. a household connection in the topology (see Figure 11 in Subsection 5.1.3) is associated with the same set of devices for each experiment. Furthermore, historic power profiles $\boldsymbol{\chi}_i$ and $\boldsymbol{\xi}_i$ were assigned to devices $i \in \mathcal{I}_{load}$ and $i \in \mathcal{I}_{PV}$ in a fixed way as well. In this subsection, we consider the effect of varying the device locations and model-data profile pairings described above. This aspect of the model, referred to here by the term *allocation* of devices and profiles, is analysed to better understand its (uncontrollable) effect on model performance.

Day	Model	Total loss value (kWh)
Normal day	PICO	0.34
	RHCO	0.34
	RHLP	3.7
Challenging day	PICO	7.8
	RHCO	8.9
	RHLP	9.8

TABLE 5: *Overview of total loss results for the device activation experiments.*

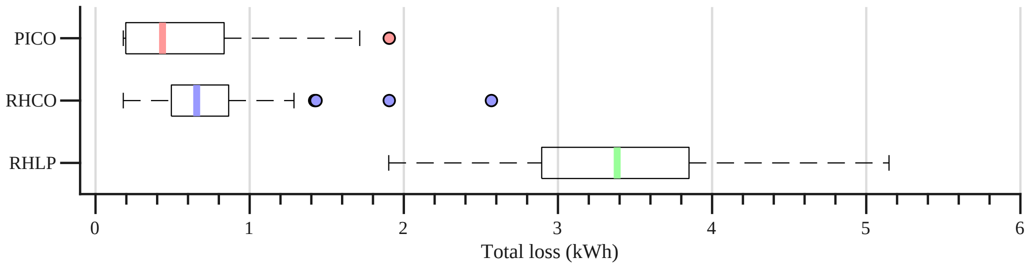

(A) *Normal day*

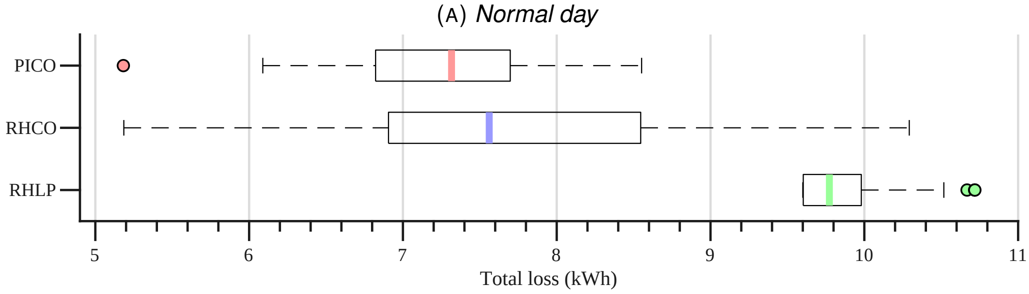

(B) *Challenging day*

FIGURE 28: *Boxplots of total loss results using data from the "normal day" in (A) and "challenging day" in (B). Each boxplot consists of all feasible results for a set of 100 random allocation and a specified model.*

Considering again the topology in Figure 11, we see that not every part of the distribution grid is the same. Some congestion nodes have more household connections than others and some are nested under another (or two other) congestion node(s). Furthermore, as was noted above, each household may have a different set of devices associated with it as was also described in Subsection 5.1.3. Because the amount of considered households is fairly small, the law of large numbers (see Subsection 4.1.3) does not hold at the congestion node levels, and different total loss results are expected for different allocation settings. To address the effect of allocation of devices and profiles, the "normal" and "challenging" days from Subsection 5.3.1 are again used in combination with the tight constraint setting. As in the previous subsection, we assume that these days represent the full set of feasible results. A set of 100 randomly assigned allocations is initiated for both days and the mentioned constraint settings. The three models are all run for these allocation settings, and the total loss results are documented in the form of boxplots in Figure 28. The normal day plot in Figure 28a consists of a total of 99 (out of 100) feasible runs and the challenging day plot in Figure 28b consists of just 78 (out of 100) runs where all three models found a feasible dispatch. A brief discussion is included in Subsection 5.4.4.

5.4.4 Discussion

A general total loss analysis was considered in Subsection 5.4.1. In Subsection 5.4.2, experiments were described to address activation levels of device types for the different models. In Subsection 5.4.3, the locations of devices and the pairing between model and data power profiles was varied in order to investigate the effect of changing allocation. The results of these experiments are outlined below:

- **Quality of solution** To discuss the results from Subsection 5.4.1, we introduce the concept of "*domination*" between two boxplots A and B. The results in A are said to dominate those in B, written as $A > B$, if A scores better than B on each of the five measures described in Subsection 5.4.1. The results of Figure 24 clearly show that PICO>RHCO>RHLP. The median total loss for the set of 790 runs are relatively close for the centralised PICO and RHCO models at values of 5.1 kWh and 5.4 kWh respectively, with the decentralised RHLP model scoring worse at a median of 9.2 kWh. The mean values presented in Subsection 5.4.1 give total loss values for the decentralised model that represents an 81% increase compared to the perfect information model and an 51% increase compared to the RHCO model. The overall worse performance of the decentralised RHLP model can be ascribed to the model differences that were discussed in Section 3.3, e.g. limited availability of information and control, leading to an effective decrease in the feasible set of solutions for the decentralised model. Furthermore, local optimisation by the storage agents in RHLP uses the decision rule given in Figures 12 and 13 separately for each PTU, whereas the centralised models account for the full planning window $\mathcal{H}$.

  The relative spread of the total loss values is fairly similar between the three models based on the similar IQRs. Furthermore, all three model results show a slight negative skew, most pronounced for the RHLP model, as evidenced by the slightly off-centre placement of the median in the boxes. The negative skew is almost certainly caused by the fact that total loss values are bounded by a minimum of 0kWh (must be positive).

  In addition to the statistical measures considered above, a significant difference is observed in the outliers between the three models. For PICO, no outliers are observed whereas RHCO and RHLP have a number of outliers that represents 3.7% and 0.7% of the total number of runs respectively. Furthermore, the outliers for RHCO model extend as far as 5 IQR lengths away from the median in the positive direction. From these outliers we may conclude that, even though we have RHCO>RHLP by the boxplot results, there exists a significant number of cases where the centralised model results in a sizable *increase* of total loss compared to the decentralised model. Under some conditions, RHCO may therefore proof to be a less reliable EM solution compared to the RHLP model.

- **Device activation** The results from Subsection 5.4.2 are discussed here to better understand the differences in device activations between models and their implications on the total loss results. For the normal day, relatively small losses were reported (see Table 5). This is also reflected in the figures: there is no PV curtailment of either model (see Figures 25a and 25b) and the battery usage is much smaller than the (lossless) use of HPs (see Figures 27a and 27c). The ten-fold increase in total loss for RHLP compared to RHCO can be seen in the increase in battery activation of LPRH compared to RHCO in Figure 27a. The centralised model seems to be more adapt at using the full flex potential of the heat pumps. This is not entirely surprising, since the battery and heat pump agents do not "proactively" charge to avoid having to charge later, instead they are directly controlled through the decision rule shown in Figures 12 and 13. Because the centralised model does have this capability, we expect it to perform better in this respect.

  For the challenging day, the total loss values of the RHCO and RHLP models are higher but relatively closer together. As seen in Figures 26a and 26b, flexibility is used from all (controllable) devices with battery activation centered around midday and similar levels of PV curtailment starting at midday and continuing until sunset around 20h. Referring to Figure 27b, the centralised model spreads out the use of battery flexibility over a larger time interval compared to the decentralised model where higher usage peaks are observed.

The centralised model effectively anticipates future PTUs and is able to spread out the flexibility capacity of the heat pumps so that less battery power is needed compared to the decentralised model. From Figure 27d, it is seen clearly that the heat pump in the RHLP model already uses up its flexibility around PTU 8, whereas the heat pump in the RHCO model retains it until the PV peak around PTU 12. The above observations clearly show the difference in (battery and HP agent) decision strategies between the centralised and decentralised models.

- **Device allocation** Based on the results from Section 5.4.3, i.e. Figure 28, some observations are discussed here. The goal is to better understand the effect of device allocation on total loss results for the different models.

  For the normal day, the RHCO model is able to find solutions that are feasible (100% of days) and even optimal (38% of days). However, for the challenging day, many solutions are infeasible (22% of days) and none of the feasible solutions are optimal. In contrast to the results of Subsection 5.4.1, PICO$>$RHCO$>$RHLP does not strictly hold. However, the medians and quartiles are ordered in this way and we do have PICO$>$RHLP and RHCO$>$RHLP, i.e. the centralised models still dominate the decentralised model for both days. This ordering can be explained by the same arguments that were presented for the total loss results from Subsection 5.4.1.

  We first consider the spread of the data by looking at the difference between the minimum and maximum, denoted here by the *reaches* of the boxplots, and the IQRs. The IQRs range between 0.5 and 1.6 kWh and the reaches between 0.9 and 5.1 kWh. Compared to the median values, these spreads are quite significant, indicating that the allocation has a sizable effect on total loss results. Between the normal day (see Figure 28a) and the challenging day (see Figure 28b), the spreads of the RHCO and RHLP models differ significantly. For the RHCO model, the spread is significantly larger for the challenging day compared to the normal day. Because the challenging day is associated with tighter constraints, more extreme total loss results are expected explaining the increased spread. For the RHLP model, the spread decreases significantly during the challenging day compared to the normal day. This apparently random difference in spreads may originate from the used stopping times or reasons specific to the chosen days. A follow up study may consider a larger number of different days for the allocation experiment to see if there is a more systematic nature to these results.

6 Conclusion and recommendations

In this final chapter, the implications of the results in this thesis are summarised. The two research questions are considered in Section 6.1, where the corresponding conclusions are summarised. Then, in Section 6.2, some recommendations are outlined for future research and for potential implementation of the findings.

6.1 Conclusion

In this thesis, a distribution level energy management problem has been considered. The problem was considered using several centralised and decentralised optimisation models. In Subsection 1.1, the following research questions have been formulated regarding this problem:

RQ1 How does the mathematical structure differ between the centralised and decentralised models and can we compare them?

RQ2 How effective are the centralised and decentralised models in operating the smart distribution grid?

The two research questions have been addressed in detail in the previous chapters. An outline of these findings is summarised in Subsection 6.1.1 for the first research question and in Subsection 6.1.2 for the second research question.

6.1.1 Research question 1

In Chapter 3, the theoretical comparison between the centralised and decentralised approaches was considered. In Section 3.1, the subclass of EM cases considered in thesis, i.e. smart distribution networks, was first introduced. The goal of the developed EM models is to solve the "dispatch problem" that was formulated as follows:

> "*To find a dispatch that follows a target demand profile in a grid feasible way while simultaneously being economic and sustainable*"

A general mathematical framework for the models was then considered in Subsection 3.2.1. Using this general framework, a centralised optimisation model was introduced in Subsection 3.2.2 in the form of a mixed integer program. A decentralised, agent-based model was introduced in Subsection 3.2.3. In Section 3.3, some key characteristics of the models were compared: number of decision makers, availability of information, availability of control, and mode of control. Even though the centralised and decentralised models are applied to the same problem, they differ fundamentally in these aspects. The centralised model has a single decision maker with global information availability, global and direct control (within flexibility limits) whereas the decentralised model has multiple decision makers (agents), local availability of information, and only local and indirect control. Because of this, the centralised model is able to eventually find a global optimum whereas this is not guaranteed for the decentralised model. The effective feasible set, or *reachable* set, is typically smaller for the decentralised model than the feasible set of the centralised models due to these differences. It is concluded that perfect equivalence between centralised and decentralised models holds only under very specific circumstances. It is also concluded that this does not necessarily invalidate the (numerical) comparison between the models.

6.1.2 Research question 2

To properly answer **RQ2**, we first specified what constitutes an effective EM approach. A good starting point is found in Subsection 2.3.2, where a number of EM challenges were outlined. The different models have varying ways of responding to these challenges, and the characterisation of these different responses forms the basis of this thesis. The more quantitative part of this analysis is contained in Section 5, focusing mostly on the challenges of dealing with *uncertainty* and *scalability of optimization*. The remaining challenges, i.e. (availability of) *flexibility*, *physical robustness*, *privacy and cyber-security* have a more qualitative nature. These "qualitative challenges" have been touched upon briefly in Subsection 2.3.3 and implicitly in the treatment of **RQ1** in Subsection 3.3. It was noted that availability of flexibility might be a concern for centralised approaches as they typically lack strong user incentives. Physical robustness was not considered in detail in this thesis, since neither model simulates the potential failing of grid assets. When considering privacy and cyber-security aspects, decentralised approaches were reported to have notable advantages because of the more localised nature of data sharing and grid optimization. Overall, the decentralised model scores much better than the centralised models on these qualitative aspects.

To investigate the quantitative challenges, a numerical analysis was included in Section 5. By introducing uncertainty to the centralised and decentralised models presented in Chapter 4 and by describing a realistic case study in Section 5.1, the basis for this numerical analysis was established. The general dispatch problem that was described in Chapter 3 is used to describe a specific EM problem for a reference topology using real power consumption and production profiles as input data. As a global cost function, total energy loss was chosen to embody the economic and sustainable dispatch goal. The considered models are the "Perfect information centralised optimization" model (PICO) as described in Subsection 3.2.2, the "Receding horizon centralised optimization" model (RHCO) as described in 4.3.3, and the decentralised "Receding horizon local pricing" model described in Subsection 3.2.3. These three models are compared and judged based on a set of three performance criteria. The criteria and the accompanying results and discussions from Chapter 5 are summarized below:

- **Feasibility:** A historic data set is used to simulate a number of days with varying weather conditions and consumption profiles. For each day in this set, an instance is run for all three models. The percentage of runs that lead to a feasible solution is then compared (see Section 5.2). This experiment is repeated for a number of different constraint settings, varying the tightness of the different constraint categories, i.e. device constraints, congestion constraints, market constraints. The results show a clear ranking amongst the models: PICO consistently has the highest percentage of feasible runs, followed by RHCO, and finally RHLP. Referring back to the conclusions from Subsection 6.1.1, this is not surprising as this ranking follows the degree of information and control availability. However, note that PICO served as a benchmark for the two other models allowing a larger computation time. The feasibility percentages reacted most clearly to tightening of the device constraints. This was explained by the fact that device constraints are the only category that affects both the instantaneous flexibility and total flexibility capacity.

  For the full set of considered cases, 72% of the cases from the benchmark set by the PICO model were successfully solved by the decentralised RHLP model. In other words, the perfect information model outperforms the decentralised model in 28% of days. When comparing the RHLP to the RHCO model, this percentage is just 23%.

- **Computation time:** Aspects related to computation time were considered for each model in Section 5.3. Special attention was given to the "optimization paths" of the decentralised RHLP model in Subsection 5.3.1. This consisted of analysing the evolution of total loss values and constraint violations as a function of the iteration number for three historic days. Results from this analysis clearly show the difference between congestion mitigation on the one hand, and resolving constraint violations at the market level on the other. When congestion occurs during a run, the decentralised model is able to either resolve it rather quickly, or not at all. Market constraint violations are slowly resolved with an increasing number of iterations. Furthermore, the effect of considering different historic days was investigated showing characteristics of days where the model easily finds a feasible solution and days where this is more challenging or even impossible.

  In Subsection 5.3.2, the effect of the stopping time was investigated, comparing the full set of model runs from Section 5.2 for a stopping time of 5 seconds and 10 seconds. The resulting improvements in feasibility percentages were relatively small for all three models, indicating that stopping times of 5 seconds were not a significant limiting factor in achieving feasible solutions. Especially for the RHCO model, this is not surprising since the feasibility percentages reported in Section 5.2 nearly approximate benchmark results of the PICO model, bounding the possible improvements seen in Subsection 5.3.

- **Total loss:** In Section 5.4, a number of experiments were carried out to analyse the behaviour of total loss results for the different models. First, the full set of runs from Section 5.2 was used, this time giving the resulting total loss values for all instances where each of the three models achieved a feasible solution. The distributions of the results were presented in the form of boxplots for each model, in order to consider some statistical characteristics. The three models were ranked PICO>RHCO>RHLP for each of the five measures contained in the boxplot (see Subsection 5.4.1). Although the spread of the distributions was fairly similar for the three models, RHCO has significantly more outliers than the other two models. It was therefore concluded that even though RHCO scores better on the statistical measures than RHLP, the decentralised model may still prove to be more reliable due to the relative absence of outliers in its results. The average total loss for the RHLP model was reported to be 81% higher then that of the PICO model and 51% higher than that of the RHCO model. The difference in total loss between PICO and RHLP is therefore 3.8kWh for the neighbourhood during one day. Using this value and a typical day-tariff of 0.22 euros, this total loss difference represents a potential cost-savings of just 5.65 euros per household per year when comparing the two models.

  In Subsection 5.4.2, the flexibility usage of the device classes was considered for two representative historic days. The results showed that the RHCO model was able to use flexibility in the storage devices more efficiently, requiring less battery power to be used in the dispatch for both the simulated days. Furthermore, it was shown that for a more challenging instance, the solution quality of the centralised and decentralised models are much closer than for a normal instance. In Subsection 5.4.3, the effect of varying the allocation of devices and data profiles was considered. The total loss values showed a significant spread, indicating that device allocation is an important factor when considering EM at the distribution level.

6.2 Recommendations

Based on the conclusions from Section 6.1, this section aims to give recommendations for future work, outlining aspects of the analysis that may be improved and suggesting potential extensions.

In answering the research questions, qualitative and quantitative aspects of the centralised and decentralised models have been compared and outlined. It was concluded that, for the specific case study, the novel decentralised model (RHLP) scored worse than the centralised model (RHCO) on percentage of feasible runs and total loss values by 23% and 51% respectively. The difference in percentage of feasible runs is quite significant, since a real-life EM approach should strive for near-perfect grid stability. Assuming the model and data inputs to be sufficiently realistic, this means that significant improvements are required before the RHLP model can be used for a real-life implementation. However, it may be the case that the constraints have been modeled unrealistically tight compared to a real-life distribution grid, resulting in unrealistic modelling instances and subsequent results. Furthermore, computation time was limited without discussing in detail what might constitute realistic values. Extensions to this work might consist of further investigating the effect of changing these constraints and settings and to subsequently test a case study that is more realistic than the one considered here. This may include a more detailed analysis of computation time, and the effect of varying stopping criteria, which may give insight into the scalability of the different models.

For the total loss results, it was determined that the percentage difference between RHLP and RHCO only constitutes a yearly cost difference in the order of 5.65 euros per household. It seems reasonable to assume that such a small quantitative advantage in favor of the centralised model does not outweigh the qualitative advantages in favor of the decentralised model (e.g. improved privacy and autonomy). To properly research this, an extension to this analysis may consist of calculating the total yearly savings that each of the models could achieve for a typical household, and to investigate the effectiveness of the resulting monetary incentives. Furthermore, future work may consider the potential cost savings for other stakeholders such as DSOs and BRPs.

Extensions to the models may include the introduction of time-shiftable devices such as smart washing machines, dishwashers or electric vehicles. Especially electric vehicles represent a class of devices that are already causing balancing problems in the distribution grid and modelling them may provide useful insights. Furthermore, the used data set has a lack of data on power profiles during the winter months. Since typical power profiles differ fundamentally during these months, e.g. higher loads, shorter day times and cloudy weather, it may be interesting to include them in a future study. Considering the case study, heat pumps may be described in a way that is more realistic, i.e. by letting the leakage rate be dependent on the time of year. Also, in the present case study, only one set of global and local cost functions was considered, i.e. total energy loss. Future studies could investigate the performance associated with alternative choices of cost functions.

As a final note, it may be interesting to consider the integration of the medium timescale models considered in this thesis with a real-time balancing algorithm such as [25]. Such an analysis might help pave the way for a practical implementation of the considered models, solidifying the contribution of this thesis to the field of smart grids.

References

[1] Paris Agreement. United nations treaty collection, 2016.

[2] Hans Akkermans, Jos Schreinemakers, and Koen Kok. Microeconomic distributed control: Theory and application of multi-agent electronic markets. In *Proceedings of International Conference on Critical Infrastructures*, 2004.

[3] Muhammad Raisul Alam, Marc St-Hilaire, and T. Kunz. Computational methods for residential energy cost optimization in smart grids. *ACM Computing Surveys*, 49:1–34, 2016.

[4] Trevor Alvey, Doug Goodwin, Xingwang Ma, Dan Streiffert, and David Sun. A security-constrained bid-clearing system for the new zealand wholesale electricity market. *IEEE Transactions on Power Systems*, 13(2):340–346, 1998.

[5] Ailin Asadinejad, Alireza Rahimpour, Kevin Tomsovic, Hairong Qi, and Chien-fei Chen. Evaluation of residential customer elasticity for incentive based demand response programs. *Electric Power Systems Research*, 158:26–36, 2018.

[6] Peder Bacher, Henrik Madsen, and Henrik Aalborg Nielsen. Online short-term solar power forecasting. *Solar Energy*, 83(10):1772–1783, 2009.

[7] Antimo Barbato and Antonio Capone. Optimization models and methods for demand-side management of residential users: A survey. *Energies*, 7(9):5787–5824, 2014.

[8] Florian Barbieri, Sumedha Rajakaruna, and Arindam Ghosh. Very short-term photovoltaic power forecasting with cloud modeling: A review. *Renewable and Sustainable Energy Reviews*, 75:242–263, 2017.

[9] Giulio Binetti, Ali Davoudi, Frank L Lewis, David Naso, and Biagio Turchiano. Distributed consensus-based economic dispatch with transmission losses. *IEEE Transactions on Power Systems*, 29(4):1711–1720, 2014.

[10] Christopher M Colson and M Hashem Nehrir. Comprehensive real-time microgrid power management and control with distributed agents. *IEEE Transactions on Smart Grid*, 4(1):617–627, 2013.

[11] European Commission. Progress report, energy from renewable sources in the netherlands, 2016.

[12] Bert den Ouden, A Moser, and E van der Hoofd. Market review 2014 h1. `https://www.tennet.eu/fileadmin/user_upload/Company/Publications/Technical_Publications/Dutch/2014_H1_TenneT_Market_Review.pdf`, 2014. Accessed: 2019-10-23.

[13] Ruilong Deng, Zaiyue Yang, Mo-Yuen Chow, and Jiming Chen. A survey on demand response in smart grids: Mathematical models and approaches. *IEEE Transactions on Industrial Informatics*, 11(3):570–582, 2015.

[14] Omar Ellabban, Haitham Abu-Rub, and Frede Blaabjerg. Renewable energy resources: Current status, future prospects and their enabling technology. *Renewable and Sustainable Energy Reviews*, 39:748–764, 2014.

[15] Rachel Emas. The concept of sustainable development: definition and defining principles. *Brief for GSDR*, 2015, 2015.

[16] Torgeir Ericson. Direct load control of residential water heaters. *Energy Policy*, 37(9):3502–3512, 2009.

[17] J Doyne Farmer and François Lafond. How predictable is technological progress? *Research Policy*, 45(3):647–665, 2016.

[18] Ahmad Faruqui, Dan Harris, and Ryan Hledik. Unlocking the €53 billion savings from smart meters in the EU: How increasing the adoption of dynamic tariffs could make or break the EU's smart grid investment. *Energy Policy*, 38(10):6222–6231, 2010.

[19] Ferdinando Fioretto, William Yeoh, Enrico Pontelli, Ye Ma, and Satishkumar J Ranade. A distributed constraint optimization (DCOP) approach to the economic dispatch with demand response. In *Proceedings of the Conference on Autonomous Agents and MultiAgent Systems*, pages 999–1007, 2017.

[20] Bri-Mathias Hodge and Michael Milligan. Wind power forecasting error distributions over multiple timescales. In *IEEE power and energy society general meeting*, pages 1–8. IEEE, 2011.

[21] Gerwin Hoogsteen. *A cyber-physical systems perspective on decentralized energy management*. PhD thesis, University of Twente, 2017.

[22] International Energy Agency (IEA). Technology roadmap - smart grids. Technical report, 2011.

[23] A Rezaee Jordehi. Optimisation of demand response in electric power systems, a review. *Renewable and Sustainable Energy Reviews*, 103:308–319, 2019.

[24] KNMI. Daggegevens van het weer in nederland. `https://projects.knmi.nl/klimatologie/daggegevens/index.cgi`, 2020. Accessed: 2020-04-16.

[25] Koen Kok and Arun Subramanian. Fast locational marginal pricing for congestion management in a distribution network with multiple aggregators. In *Proceedings of International Conference and Exhibition on Electricity Distribution*. AIM, 2019.

[26] Koen Kok and Steve Widergren. A society of devices: Integrating intelligent distributed resources with transactive energy. *IEEE Power and Energy Magazine*, 14(3):34–45, 2016.

[27] Lorenzo Kristov, Paul De Martini, and Jeffrey D Taft. A tale of two visions: Designing a decentralized transactive electric system. *IEEE Power and Energy Magazine*, 14(3):63–69, 2016.

[28] Claude Lemaréchal. Lagrangian relaxation. In *Computational Combinatorial Optimization*, pages 112–156. Springer, 2001.

[29] Ltd. MetalsTech. Market outlook 2016. `http://www.metalstech.net/market-outlook/`, 2016. Accessed: 2019-10-24.

[30] James A Momoh, Rambabu Adapa, and ME El-Hawary. A review of selected optimal power flow literature to 1993. I. nonlinear and quadratic programming approaches. *IEEE Transactions on Power Systems*, 14(1):96–104, 1999.

[31] Manisa Pipattanasomporn, Hassan Feroze, and Saifur Rahman. Multi-agent systems in a distributed smart grid: Design and implementation. In *Power Systems Conference and Exposition*, pages 1–8. IEEE, 2009.

[32] Abo Rassa, Coen van Leeuwen, Robertjan Spaans, and Koen Kok. Developing local energy markets: a holistic system approach. *IEEE Power and Energy Magazine*, 17(5):59–70, 2019.

[33] Serné Ronald. "nulwoning". `http://www.nulwoning.nl/`, 2016. Accessed: 2019-10-21.

[34] B Saravanan, Siddharth Das, Surbhi Sikri, and DP Kothari. A solution to the unit commitment problem—a review. *Frontiers in Energy*, 7(2):223–236, 2013.

[35] K Schneider, B Mather, B Pal, C Ten, G Shirek, H Zhu, J Fuller, J Pereira, L Ochoa, L De Araujo, et al. Analytic considerations and design basis for the IEEE distribution test feeders. *IEEE Transactions on Power Systems*, 33(3):3181–3188, 2017.

[36] Fred C Schweppe, Richard D Tabors, James L Kirtley, Hugh R Outhred, Frederick H Pickel, and Alan J Cox. Homeostatic utility control. *IEEE Transactions on Power Apparatus and Systems*, 99(3):1151–1163, 1980.

[37] Farhad Shahnia, Michael T Wishart, Arindam Ghosh, Gerard Ledwich, and Firuz Zare. Smart demand side management of low-voltage distribution networks using multi-objective decision making. *IET Generation, Transmission & Distribution*, 6(10):986–1000, 2012.

[38] Heng Shi, Minghao Xu, and Ran Li. Deep learning for household load forecasting—a novel pooling deep RNN. *IEEE Transactions on Smart Grid*, 9(5):5271–5280, 2017.

[39] Jie Shi, Wei-Jen Lee, Yongqian Liu, Yongping Yang, and Peng Wang. Forecasting power output of photovoltaic systems based on weather classification and support vector machines. *IEEE Transactions on Industry Applications*, 48(3):1064–1069, 2012.

[40] Soliman Abdel-Hady Soliman and Abdel-Aal Hassan Mantawy. Economic dispatch (ED) and unit commitment problems (UCP): formulation and solution algorithms. In *Modern Optimization Techniques with Applications in Electric Power Systems*, pages 185–279. Springer, 2012.

[41] Saurabh S Soman, Hamidreza Zareipour, Om Malik, and Paras Mandal. A review of wind power and wind speed forecasting methods with different time horizons. In *North American Power Symposium 2010*, pages 1–8. IEEE, 2010.

[42] Mike Steilen and Ludwig Jörissen. Hydrogen conversion into electricity and thermal energy by fuel cells: Use of H2-systems and batteries. In *Electrochemical Energy Storage for Renewable Sources and Grid Balancing*, pages 143–158. Elsevier, 2015.

[43] Enrico Telaretti, Mariano Ippolito, and Luigi Dusonchet. A simple operating strategy of small-scale battery energy storages for energy arbitrage under dynamic pricing tariffs. *Energies*, 9(1):12, 2016.

[44] Hermen Toersche. *Effective and efficient coordination of flexibility in smart grids*. PhD thesis, University of Twente, 2016.

[45] Thijs van der Klauw. *Decentralized energy management with profile steering: Resource allocation problems in energy management*. PhD thesis, University of Twente, 2017.

[46] Bram van der Waaij, Wilco Wijbrandi, and Mente Konsman. Whitepaper energy flexibility platform and interface EF-Pi. `http://flexiblepower.github.io/EF-Pi-Whitepaper/`, 2015. Accessed: 2020-04-15.

[47] J Vardakas, N Zorba, and C Verikoukis. A survey on demand response in smart grids. *IEEE Transactions on Industrial Informatics*, 11:1–8, 2015.

[48] Andreas Veit, Christoph Goebel, Rohit Tidke, Christoph Doblander, and Hans-Arno Jacobsen. Household electricity demand forecasting: benchmarking state-of-the-art methods. In *Proceedings of the International Conference on Future Energy Systems*, pages 233–234, 2014.

[49] Joseph Warrington, Sébastien Mariéthoz, and Manfred Morari. Negotiated predictive dispatch: Receding horizon nodal electricity pricing for wind integration. In *2011 8th International Conference on the European Energy Market*, pages 407–412. IEEE, 2011.

[50] Peng Yang, Gongguo Tang, and Arye Nehorai. A game-theoretic approach for optimal time-of-use electricity pricing. *IEEE Transactions on Power Systems*, 28(2):884–892, 2012.

[51] Fredrik Ygge and Hans Akkermans. Decentralized markets versus central control: A comparative study. *Journal of Artificial Intelligence Research*, 11:301–333, 1999.