Utrecht University

Master's thesis Computing Science (40 ECTS)

Forecast-based optimal operation of islanded microgrids

Author:
R.J.J. Brouwer

Supervisors:
dr. ir. J.M. van den Akker
W. Vermeiden MSc.

October 13, 2017

Abstract. In this thesis, we provide an introduction to microgrids and associated concepts. We look at the central challenges to their operational management, and provide a discussion on the classification of microgrids based on a number of characteristics and how this influences the planning process.
In the second part, we describe a forecast-based centralized planning approach for the (economically) optimal operational management of (islanded) microgrids. Using mathematical programming techniques, we formulate four models: (1) a simple MILP-formulation of our problem, (2) an extension adding safety constraints, (3) a scenario-based two-stage model and (4) a scenario-based multi-stage model. Additionally, we describe an algorithm for the generation of the scenario tree needed for the fourth model.
The performance of the models is compared on three test instances: two are based on an existing microgrid in the Netherlands, and one is artificially constructed. The economic advantage of the described planning approach is not unequivocally shown for all described instances, but it is shown to have its advantages.

Acknowledgments. I would like to thank my supervisor, Marjan van den Akker, for her help during this project. She provided me with a lot of useful feedback on previous drafts of this thesis and on ideas that I had while developing the approach presented in this thesis. Her questions forced me to explain what I had in mind, and I usually came out of a meeting with a much more clear idea of what to do next then I went in with.
Additionally, I would like to thank my supervisor at Zown, Wouter Vermeiden, for providing me with the opportunity of writing this thesis at Zown and supporting me throughout the process. He always made clear that writing my thesis should be my first priority, and provided me with a lot of freedom during the project. I would also like to thank everyone I worked with from Zown and iCarus. They are too many to list them all here, but I have enjoyed working with all of them over the past ten months. I have had a great time working on the Analytics project at Zown and, at the same time, on my own thesis.

Table of Contents

Part I.
Introduction & Theory

1. Introduction

We live in a time where it becomes increasingly clear that the climate is changing, and that we are to blame. In its most recent assessment report (AR5), the Intergovernmental Panel on Climate Change (IPCC) stated that: "it is extremely likely that human influence has been the dominant cause of the observed warming since the mid-20th century" [57]. Between 90% and 100% of publishing climate scientists agree that human activity is the cause of recent global warming [18].

The need to mitigate the potentially dramatic effect of global warming by reducing emissions is widely recognized. On the global level, the Paris agreement, negotiated at the 21st United Nations conference on climate change (COP21) in December 2015, commits the signatories to limit the emission of greenhouse gases and prevent the increase in global average temperature to be more than 2 ℃ compared to pre-industrial levels. It expresses the intention, however, of limiting the increase in temperature even more, to 1.5 ℃ [125].

The European Union has set targets for 2020 – 2050, focused on reducing emissions (by 20% each decade compared to 1990 levels), increasing the share of renewables (20% (2020) and 27% (2030) of the total production) and increasing energy efficiency/savings (saving 20% (2020) and 27% (2030) compared to a "business as usual"-scenario) [32,33].

The Netherlands used to be a front-runner in the early days of climate politics (the late 1980s). Over the years, however, Dutch climate policy and efforts have become less ambitious, possibly due to the changing political tide in the 2000s and the dominance of the traditional fossil fuel sector [72]. As of January 2017, the Netherlands has not yet ratified the Paris agreement. Current (short-term) targets result from an agreement between the government and forty organizations (including environmental groups, unions and financial institutions): the "Energieakkoord voor duurzame groei" [117]. The main aims of the agreement are to achieve energy savings of 1.5% per year and increase the share of renewables in energy production to 14% in 2020 and 16% in 2023. A recent report states that the targets for 2020 will probably not be met, although it is optimistic about the observed trends [27].

The agreements discussed above exemplify a (policy) move to change our energy supply chain (a *energy transition*): shifting production from fossil fuel-dependent generation to renewable energy producers, reducing energy use and increasing efficiency. A prime motivator for this move, aside from reducing the effects of global warming, is decreasing the dependence on fossil fuels. It becomes increasingly clear that the supply of these fuels is not infinite and that their extraction will become more and more expensive. Additionally, a large portion of the supply of these fuels originates from (politically) less stable regions. An energy supply that is less dependent on fossil fuels is therefore highly

desirable. The rapid development and improvement of renewable energy technologies support this move.

The electricity grid faces considerable challenges because of this change. The traditional system was of a hierarchical nature. Large fossil fuel-based generators 'simply' matched the demand of the consumers in the grid. Over the past few years, the share of distributed energy resources (DER) in the grid has increased significantly. The energy is no longer only produced in a few large power plants, but the production is distributed over the grid. This situation causes a variety of problems (e.g. frequency and voltage instability), some of which are investigated (along with potential benefits) in [83].

The balance between supply and demand should be maintained in the grid, but this becomes increasingly difficult as a large portion of the DER installed are of an intermittent nature due to the natural resource on which they depend (e.g. sun or wind). The flexibility that existed on the generation side because of the large fossil fuel-based generators is thereby decreased. On the other hand, technological developments allow us to (potentially) exploit more flexibility on the consumption side, through demand-side management (DSM), for example. The traditional approach is not capable of exploiting this new flexibility. This has led to a move towards a "next-generation power grid", the smart grid. This smart grid, while based on an entirely different philosophy, has to be built on top of the existing (old) infrastructure.

In this thesis, we will focus on microgrids. Microgrids are just one of the increasingly popular new technologies that can be considered part of the larger "smart grid"-movement. A more detailed description of these concepts will be provided in Section 2. We will develop and evaluate a number of algorithms for the operational management of microgrids. The main research questions of this thesis will be:

1. How can we manage the operation of microgrids with centralized control in a (near-)optimal way?
   a) What is the best way to deal with uncertainty in the input for the model?
   b) Is considering an abstract supply-demand balancing problem sufficient to arrive at good solutions?
2. What is the value of intelligent planning in microgrid operation?
   a) What is the potential reduction in operational cost and how does it depend on the composition of the microgrid?
   b) Does the potential reduction in operational cost outweigh the cost of constructing a planning (cost of the IT infrastructure, etc.)?

The first main question will be the focus of my thesis. This thesis was written during an internship at Zown (see Section 2.5), and the second question serves to provide Zown with an *indication* of the potential value of this work and possible future follow-ups for their system.

The remainder of this thesis will be organized as follows: the first part will provide a general background on microgrids and introduce some concepts that will be used in

the later parts of this thesis. In the second part, the problem is described and models are developed. The algorithms used are described here as well, and a few details on the implementation are discussed. Finally, in the third part of this thesis, we will apply the algorithms from the previous part to a number of test instances and evaluate the results. The latter parts use a number of concepts that are explained in the first part but can otherwise be read separately.

2. Microgrids

2.1. The electricity system

We will sketch a very brief overview of the Dutch electricity system. A more elaborate description can be found in [129].

As discussed above, the system is arranged to facilitate a centralized electricity supply: demand occurs on one end of the system, and large plants match this demand on the other end. The grid that connects both sides can be divided based on voltage level:

- The high voltage level (>50 kV) is primarily used to transport electricity over long distances. This part of the grid is usually referred to as the transmission system, which is managed by a transmission system operator (TenneT in the Netherlands). The transmission grid serves hardly any demand directly, except for a very limited number of high load consumers.
- The medium voltage level (3–30 kV) is connected to the high voltage transmission system and distributes the electricity further. This part of the grid is known as the distribution system. It satisfies the demand of some large consumers directly, but its primary function is to distribute the electricity to the low voltage nets. This part of the grid is managed by a distribution system operator.
- The low voltage level (230–400 V) is the level where most (smaller) consumers, like households and offices, connect to. This part of the grid is also the responsibility of the distribution system operator.

This system is designed to facilitate the transportation of electricity from high voltage levels downwards to consumers. It is not designed to deal with the developments we discussed briefly at the end of the previous section. Consumers are no longer exclusively consuming electricity. Peaks caused by local generation on the (traditional) consumer side, require upgrades of the grid to be able to transport the surplus. With an ever-increasing penetration of (renewable) distributed generation, this practice is not tenable in the long term. As mentioned above, the move towards a smart grid promises to solve these problems.

2.2. Smart Grid

The term “Smart Grid” knows many definitions. In the context of this thesis, the following definition by the European Commission is probably the most appropriate:

> "A Smart Grid is an electricity network that can cost efficiently integrate the behaviour and actions of all users connected to it generators, consumers and those that do both in order to ensure economically efficient, sustainable power system with low losses and high levels of quality and security of supply and safety." [31]

The Smart Grid is seen as the next-generation power grid, and relies heavily on *smart* digital technology to achieve a variety of goals: improving reliability, facilitating distributed generation and increasing energy efficiency, to name a few. It differs from the standard electrical grid mainly in its organization (hierarchical, unidirectional vs. distributed, bidirectional) and the extensive use of ICT. A more elaborate introduction to the "Smart Grid" concept and the role of ICT in its development is provided in [35]. A smart grid that integrates multiple commodities (electricity, heat, natural gas, hydrogen, biomass, etc.) is sometimes referred to as an Intelligent Energy Network [95].

An extensive survey of all smart grid-related work until 2011 can be found in [34]. A more recent survey in [16] provides an overview of the critical issues that are currently faced in smart grid technology.

2.3. Microgrid

The concept of microgrids, much like the smart grid, has been defined in many different ways. We present here the definition from [47], which is a good representation of the definitions used in most European projects:

> "Microgrids comprise Low Voltage distribution systems with distributed energy sources, such as micro-turbines, fuel cells, PVs, etc., together with storage devices, i.e. flywheels, energy capacitors and batteries, and controllable loads, offering considerable control capabilities over the network operation. These systems are interconnected to the Medium Voltage Distribution network, but they can be also operated isolated from the main grid, in case of faults in the upstream network. From the customer point of view, Microgrids provide both thermal and electricity needs, and in addition enhance local reliability, reduce emissions, improve power quality by supporting voltage and reducing voltage dips, and potentially lower costs of energy supply."

Microgrids, then, are relatively small energy distribution networks comprised of three keys types of components: distributed energy generators, storage devices and (controllable) loads. Microgrids are able to operate autonomously (islanded mode), but can be connected to the traditional electrical grid (macro grid) as well. Microgrids are not limited to electrical components but may handle multiple other commodities, such as heat and (bio-)gas. A visualization of an example microgrid is presented in Fig. 1.

Finally, microgrids are obviously contained in the definition of a smart grid. The properties of a smart grid apply to microgrids as well, on a small scale. The efficient management of microgrids requires intelligent ICT to facilitate reliable and efficient operation.

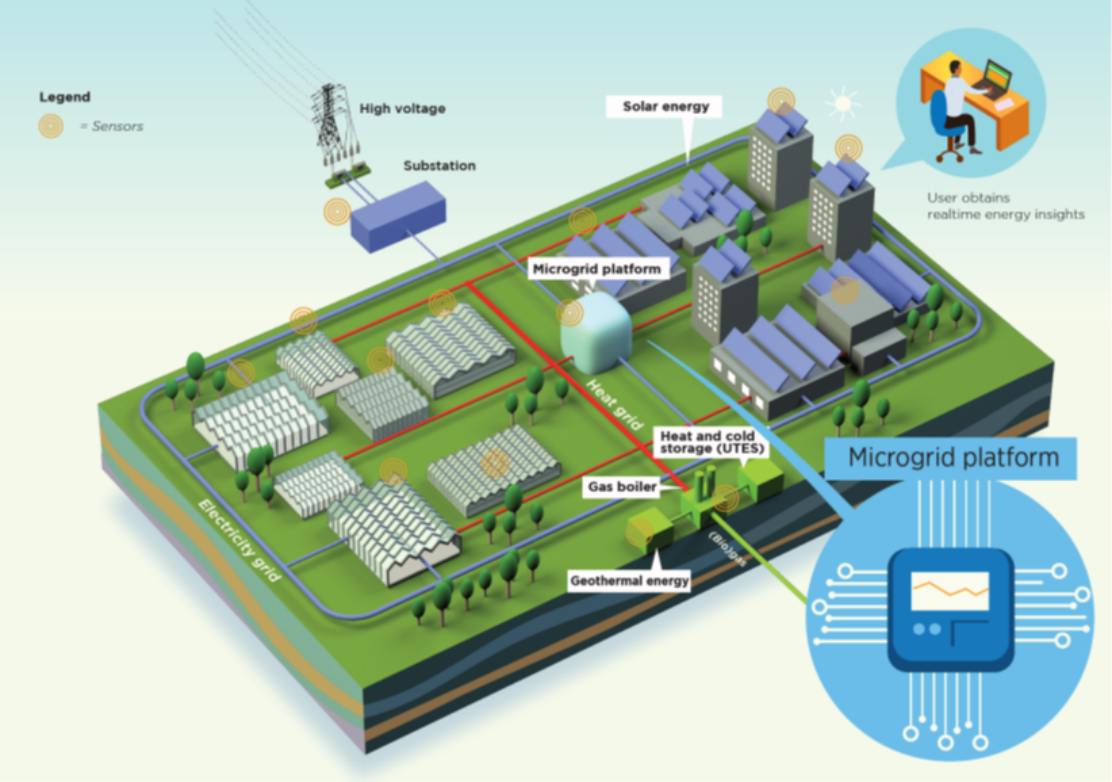

Figure 1: Example lay-out of a microgrid, including the Microgrid platform (Section 3.4) [131]

The central problem of microgrid management is the balancing of power demand and supply. This is complicated by several characteristics of microgrids:

- The large share of intermittent distributed energy resources (DER) in microgrids. The output of DERs such as Photo Voltaic (PV) and Wind Turbine (WT) systems is hard to predict and may fluctuate significantly depending on the local conditions (i.e. the availability of primary sources such as solar radiation and wind).
- The (relatively) small scale of microgrids. The small scale on which a microgrid balances supply and demand makes forecasting even harder. In large networks, prediction errors and demand/supply fluctuations will usually be smoothed out because of the number and geographical spread of producers and consumers. Within a microgrid, however, there is a high spatiotemporal correlation between all of them.
- The possibility of islanded operation. While a grid-connected network can, in principle, always counteract shortage and surplus using the main grid, a microgrid in islanded mode is required to get the balance of supply and demand exactly right by itself.

On the other hand, a microgrid (usually) contains a considerable number of controllable loads, storage elements and controllable energy sources. The problem translates to the optimal management of these assets taking the operational constraints of the microgrid (primarily supply/demand balance) into account. The types of planning problems encountered in microgrids will be further investigated in Section 3. Next, we will consider the typical components of a microgrid and their role in solving this problem.

For further reading, introductions to the microgrid and surrounding concepts can be found in [63, 110]. For a broad review of the literature about microgrids in general see [98, 126]. In [81, 103, 126] lists of microgrid locations are presented.

2.4. Microgrid components

As mentioned before, a microgrid consists mainly of distributed energy generators, storage devices and (controllable) loads. In this section, each of these categories will be discussed in terms of its relevance for the balancing of supply and demand. Note that these categories are not necessarily mutually exclusive. In multi-commodity microgrids, for example, a generator might need to be supplied with one commodity (e.g. electricity or natural gas) to produce another (e.g. heat).

2.4.1. Generation

A variety of distributed generators can be included in microgrids. They can, broadly, be divided along the lines of the amount of control that an operator has over the output of a generator.

On the one hand, intermittent Renewable Energy Sources (RES) can only be controlled to a very limited extent. The output depends on the availability of the primary energy resources (e.g. solar radiation in case of PV), and curtailing the output of these units is both inadvisable and undesirable, due to their low operating costs and environmental benefits. Although curtailment is possible, the output of RES should generally be considered as a given lower bound on energy generation in supply/demand balancing.

On the other hand, dispatchable generators such as microturbines and fuel cells enable the operator to exert much more control over the amount of generated energy. The level of control differs depending on the exact type of generators, and there are constraints that limit the possibilities somewhat (e.g. ramp up constraints). Combined Heat/Power (CHP) units are a special case. As they supply both heat and electrical power, the (lack of) demand of both commodities needs to be considered when making control decisions. Another challenge is to let the operation of the generators be as efficient as possible.

A review of distributed energy resources in microgrids can be found in [59].

2.4.2. Storage

Storage units provide a large amount of flexibility, as they can be used both as consumers and producers of energy, depending on the situation. Storage devices of all kinds (e.g. lead-acid batteries, flywheels) provide a buffer and play a large part in maintaining system stability, power quality and reliable operation.

Literature reviews on energy storage technologies can be found in [6, 76, 108, 122].

The line between buffers/storage on the one hand and loads on the other is not clear-cut: flexible loads such as boilers may serve as buffers, effectively storing one form of energy (heat), while consuming another (electricity or gas).

2.4.3. Loads

Finally, a microgrid usually has a number of loads, which can broadly be divided into two categories: fixed (uncontrollable) and flexible (controllable) loads [98]. Fixed loads are uncontrollable: their demand needs to be satisfied at any given time. Flexible loads provide flexibility: they can be curtailed (*load shedding*) or shifted in time (*load shifting*).

The uncontrollable loads are not that interesting from a control perspective, but flexible loads are. In microgrids, part of the flexibility on the generation side is lost because of the large portion of non-dispatchable generation (e.g. PV). Controllable loads offer an opportunity to regain some flexibility on the demand-side of the problem. This is commonly referred to as Demand Side Management (DSM). Demand response is the most important associated concept: it is based on changing the energy usage in response to changes in the availability (or price) of energy. This leads to an end-user changing its consumption of energy from its normal pattern. As mentioned above, there are two main ways of doing this: flat-out reducing power consumption (load shedding), or shifting the demand in time (load shifting). A cooling system, for example, should keep the temperature within a certain range. When the microgrid has an excess of energy (e.g. during the sunny hours of the day), the system can use extra energy to bring the temperature down. When, sometime later, the generation drops significantly, it can refrain from cooling for some time (as long as the temperature stays within the allowed range), effectively reducing the overall demand at that time.

2.5. Zown: your own energy zone

This thesis was written at Zown, an initiative of Alliander. Alliander consists of a group of companies, one of which, Liander, is the distribution system operator for a large part of the Netherlands [3]. They will have to face the challenges that the energy transition brings with it. Alliander does not just want to move along with the changes but aims to actively contribute to the transition. In this context, Alliander facilitates the development of innovative smart grid solutions. Zown is one of these so-called "new markets" [2].

Zown offers two types of services. On the one hand, Zown advises its clients on technical, financial and organizational aspects of meeting their energy needs. On the other, it provides the Microgrid Platform (MGP), an application that balances supply and demand and provides insight into the operation of the microgrid [139].

This thesis fits in the process of integrating automated planning in their system. The Microgrid Platform has recently proven to be capable of balancing and controlling a microgrid in real-time. The next step is to integrate forecasting and planning in the operation of the grid. This thesis focuses on the development of optimization models

for the planning step of this problem and aims to describe algorithms for its efficient implementation.

The mission of Zown in providing advise on microgrids and facilitating their operation can be approached from multiple perspectives:

- **The client** might be interested in creating a microgrid for various reasons. Motivations could be financial, but may also spring from a desire for a more autonomous, sustainable and/or reliable energy supply;
- **Society** benefits from the implementation of microgrids as they are part of the energy transition as a reaction to global warming and the development of technologies like it ensures the continuing existence of a reliable electricity grid;
- **The system operator** (during *operation*) profits from well-managed microgrids because of the localized balancing of supply and demand, solving some of the problems that distributed (renewable) generation might cause in the grid on the local level;
- **The system operator** (during *construction*) can save a considerable amount of investments in infrastructure as the connection of the (remote) microgrid to the main grid can be downsized or even left out altogether.

The last two justify the investment of Alliander in the microgrid concept. In the future, it may well be that the main grid "merely" exists as an insurance policy for a number of microgrids that aim to operate autonomously, as much as possible. It is unlikely that the entire grid will undergo such a transformation in the foreseeable future, in places where extensive electric infrastructure already exists, but microgrids will be integrated to a certain extent, bringing along all the benefits mentioned above.

This thesis aims to apply techniques from the scientific literature to the problems encountered in managing microgrids. In this way, we hope to contribute to the development of innovative solutions to present-day problems.

3. Microgrid operation

In this section, we will briefly investigate some trends in microgrid operational control. There is a number of planning problems related to microgrids, concerning different time scales (Table 1). In this thesis, the focus will be on mid- to longterm operational control. In reference to Table 1, unit commitment and economic dispatch are the central problems. The integration with electricity markets is included at this level. The detailed, reactive, control will not be treated, nor will overall system planning and maintenance problems. A few reviews on these issues can be found in [90, 98, 113] (real-time control) and [41] (system planning).

Problem	Time span	Description
System planning (generation mix and sizing)	>1 years	Dimensioning of the system, determining layout and components of the grid
System maintenance	1 week – 1 year	Determining maintenance schedules
Unit commitment	∼4 hours – 1 week	Determining on/off-state of generators
Economic dispatch	∼10 minutes – 4 hours	Determining actual output of generators and managing load based on short-term forecasts
Regulation, control and protection	<10 minutes	Voltage and frequency regulation, power quality control, fault protection

Table 1: Overview of planning problems on different time-scales in microgrids, based on Table 1 from [70].

3.1. Classical approaches

Energy balancing problems have been around for quite some time. The Unit Commitment problem is a classic optimization problem, which aims to optimally schedule the operation of electrical power generating units. A solution to a UC problem requires two types of decisions: which units to turn on/off (unit commitment) and the generation level of the committed units (economic dispatch). A recent overview of Unit Commitment formulations and solution approaches can be found in [53]. The traditional unit commitment problem applies to large-scale systems, where the demand is always met by a group of generators. Often, this includes planning for *spinning reserve*. The spinning reserve is a part of the capacity of the generators that are online that is not used to supply the expected load, but can be brought in fairly quickly in case of shortages. This yields a more robust solution.

An example MILP formulation of a typical UC problem looks like this (adapted from [51, 53]):

$$\text{minimize} \sum_{i,t} (z_{it}F_i + g_{it}C_i + y_{it}S_i) \text{ s.t.}$$

$$\sum_i g_{it} = D_t \qquad \forall t \qquad (1)$$

$$\sum_i r_{it} = SD_t \qquad \forall t \qquad (2)$$

$$g_{it} \geq z_{it}Cap_i^{min} \qquad \forall i,t \qquad (3)$$

$$g_{it} + r_{it} \leq z_{it}Cap_i^{max} \qquad \forall i,t \qquad (4)$$

$$r_{it} \leq z_{it}SP_i^{max} \qquad \forall i,t \qquad (5)$$

$$g_{it} + r_{it} \leq g_{it-1} + RR_i^{inc} \qquad \forall i,t \qquad (6)$$

$$g_{it} \geq g_{it-1} + r_{it-1} - RR_i^{dec} \qquad \forall i,t \qquad (7)$$

$$z_{it} \leq z_{it-1} + y_{it} \qquad \forall i,t \qquad (8)$$

$$z_{it} \geq z_{it-1} - x_{it} \qquad \forall i,t \qquad (9)$$

$$y_{it} \leq z_{i\tau} \qquad \forall i,t,\tau = t,...,\min\{t + Up_i^{min}, |T|\} \qquad (10)$$

$$z_{it-1} - z_{it} \leq 1 - z_{i\tau} \qquad \forall i,t,\tau = t,...,\min\{t + Down_i^{min}, |T|\} \qquad (11)$$

$$g_{it}, r_{it} \geq 0 \qquad \forall i,t$$

$$0 \leq x_{it}, y_{it}, \leq 1 \qquad \forall i,t$$

$$z_{it} \in \{0,1\} \qquad \forall i,t$$

where
(decision variables):

g_{it}	is the amount of energy produced by generator i at interval t
r_{it}	is the amount of spinning reserve available from generator i at interval t
x_{it}	is 1 if generator i shuts down at the start of interval t
y_{it}	is 1 if generator i is started at the start of interval t
z_{it}	is 1 if generator i is active during interval t

(constants):

D_t	(expected) total demand at interval t
SD_t	(desired) amount of spinning reserve at interval t
F_{it}	operating cost of generator i at interval t
C_{it}	generation cost of generator i at interval t (cost per generated unit of energy)
S_{it}	start-up costs of generator i at interval t
Cap_i^{min}	minimum capacity of generator i
Cap_i^{max}	maximum capacity of generator i
SP_i^{max}	maximum reserve contribution of generator i
RR_i^{inc}	ramp up rate of generator i for increasing output
RR_i^{dec}	ramp up rate of generator i for decreasing output
Up_i^{min}	minimum up time for generator i after start-up
$Down_i^{min}$	minimum down time for generator i after shut down

The objective of the UC problem is to minimize overall operational cost, while satisfying a number of constraints: the demand should be met (1); a sufficient amount of spinning reserve should be available (2); units should operate within their capacity limits (3)-(4); units should contribute at most their maximum to the spinning reserve (5); ramp rate limits should be respected (6)-(7). (8) and (9) ensure the start up and shut down variables get the correct value. Finally, (10) and (11) represent the minimum up- and downtime for generators.

While this example model shows some of the constraints that need to be satisfied in a typical Unit Commitment problem, there are more that could be included, emission constraints for example. A more in-depth discussion of these constraints and typical solution methods to both the deterministic and the stochastic two-stage version of the problem can be found in [53]. The authors of [36] provide an overview of the litera-

ture in relation to the Short-Term HydroThermal Coordination problem, a variant of UC dealing with hydro and thermal resources instead of electric ones. Similarly, [121] provides an overview of mathematical programming approaches to the Hydro Unit Commitment problem. In [109] an overview of the (standard) Unit Commitment problem and many solution approaches is provided, mainly focusing on so-called "non-conventional" approaches (i.e. approaches other than dynamic programming, mathematical programming and local search). A more recent survey can be found in [120]. This survey also treats a large number of approaches, but focuses on dealing with uncertainty in unit commitment problems.

Aside from the "missing" constraints mentioned above, there are some other factors that prevent the direct application of the traditional unit commitment problem to microgrids. We will discuss some of them below.

The unit commitment problem focuses on scheduling generators to meet the demand. It does not consider storage (devices that can either consume or produce energy at different times) or demand side management (load shedding/shifting). In the large conventional grid, this is not really an issue: the vast majority of the energy produced is consumed instantly. The contribution of storage in this network is small, and the network operators typically do not have control over flexible demand. In a microgrid, this changes. Due to the small size of the grid and the large share of intermittent RES, storage devices and DSM are essential to balance supply and demand. They need to be integrated into the problem.

Another consequence of the small size of microgrids and their large share of RES is that it is difficult to forecast (uncontrollable) demand and supply. In a large grid, for example, individual deviances from typical use profiles hardly influence the total demand, as this is an aggregation of a large number of users. Because of the large size of the conventional grid, forecasts can be relatively reliable. In a microgrid, individual changes can have a comparatively large effect on the total supply/demand balance. The localized nature of microgrids enlarges the problem, as deviations from the forecasted generation can be relatively big. On an otherwise sunny day, at some point during the day a cloud might pass by, decreasing the output of all solar generation in the network at once. As significant deviations from the forecasted levels are to be expected, it is much more important that the solutions to the problem are robust. The problem needs to be modified to take this into account somehow.

Microgrids that operate in grid-connected mode are not (necessarily) entirely self-sufficient. They can either buy from or sell to the main grid. This might be seen as some type of buffer and modeled in a way similar to storage devices with time-varying cost vectors, but this is a serious simplification. This approach might work for the short-term imbalance market, because it does not require bidding, but only facilitates the sale of reserve capacity. But as the planning reaches further into the future, energy markets typically use auctioning/bidding strategies. Once the bidding is done, of course, the amount of energy the microgrid is supposed to consume or produce is "known", and the prices for deviating from this amount will be established. It then can be integrated in the problem in the aforementioned way. As long as the main goal is to balance supply

and demand in the grid, this simplification is acceptable: bidding strategies can then be adapted to follow the outcomes of the problem. To maximize profit (or minimize expenses) a different modeling that takes bidding strategies into account is required. In [21], an MIP-formulation of the problem of managing a battery for this purpose that takes into account the Dutch short-term imbalance markets is provided, showing that it can be done.

Modifying the unit commitment problem to suit microgrids results in an energy management problem, which will be discussed towards the end of Section 3.3. The problem we will be studying in this thesis can be seen as a generalization of Unit Commitment. First, a brief discussion on the control hierarchy in microgrids is provided.

3.2. Control hierarchy

Typically, the control of a (micro)grid is split into multiple levels. Higher levels of control work over larger time-scales, with the lowest levels of control providing reactive real-time control, and the highest levels responsible for long-term economic operation. This hierarchical approach is used in bulk power systems [124], but has often been applied to microgrids as well [44, 45, 62, 63, 85, 86, 96, 135].

Such a hierarchical control structure in microgrids typically has three types of elements [60, 63, 64, 86, 123, 135]:

- **Microsource Controller** (MC) and **Load Controller** (LC). These operate at the device level, and control the operation of controllable devices in the microgrid.
- **Microgrid Central Controller** (MGCC). The MGCC is the central controller of a single microgrid and is responsible for the energy management of the grid.
- **Distribution Management System** (DMS). A system responsible for the operation of medium and low voltage areas with (possibly) more than one microgrid connected to it.

The DMS is out of scope for this thesis. Depending on the control strategy (decentralized vs. centralized), more intelligence can be located at the MCs/LCs, but they are typically mainly responsible for low level control (e.g. voltage control), whereas the MGCC is responsible for (longer term) planning, sending set-points down to the MCs/LCs.

3.3. Centralized/distributed control

One of the central issues of microgrid control is whether it should be controlled using a centralized or a distributed approach. Both approaches have a number of benefits and drawbacks and it ultimately depends on the situation which should be preferred. This thesis, and thereby much of the discussion below, focuses on the centralized approach. In Section 5 we will explain why the centralized approach is the best fit for this work. A few examples of distributed approaches are: PowerMatcher [66], TRIANA [5, 9, 87, 88] and an associated profile steering approach [127] and others [17, 111].

	Advantages	Disadvantages
Centralized	• Easy to implement and maintain • Lower cost • Decisions can be based on knowledge of the whole system • Real-time observability of entire system	• Computationally heavy • Requires a lot of data-flow • Single point of failure • Not easy to expand (no 'plug-and-play' capability)
Distributed	• Better suited for 'plug-and-play' expansion • Lower computational cost • More robust (no single point of failure) • Better scalability: more suitable for large, complex systems	• Synchronization challenges • More time required to reach consensus • More challenging and expensive to implement • Narrow control over the entire system

Table 2: Advantages and disadvantages of centralized and distributed control strategies.

A short review of the benefits and drawbacks of both approaches, in part based on [81,85,111,118,135], can be found in Table 2. The choice for one or the other is considered to be "ultimately a philosophical question" [118], but in general the centralized approach is preferred for smaller, monopolized microgrids, while the distributed approach is more appropriate for large, complex, heterogeneous systems [85, 103].

The optimal management of the assets in the grid is the responsibility of the MGCC. This is commonly referred to as an energy management problem, and a number of energy management systems (EMS) has been proposed to take care of this. An EMS is ideally implemented in the MGCC [118]. The main functionalities of a microgrid EMS are forecasting, optimization, data analysis, and providing a human-machine interface [111, 118].

3.4. The Microgrid Platform

At Zown, the *Microgrid Platform* (MGP) fulfills the role of the EMS. An EMS by the same name has been proposed before [112], but this is unrelated to Zown's MGP.

The basic architecture of the MGP is summarized in Fig. 2. The LC/MC controllers are on the device level (the bottom square). The lowest (reactive) levels of control have there place at the device level, outside the MGP.

The role of MGCC, in this system, is shared between *myZown Local Controller* and *myZown Operator* (including *myZown Analytics*). OSGP (Open Smart Grid Platform), in the middle, provides an abstraction layer. We will not go into the details of the interaction between these components. The myZown Local Controller is programmed into a remote terminal unit (RTU) that is physically located inside the microgrid, and it controls the devices directly. On the other hand, myZown Operator is a cloud-based

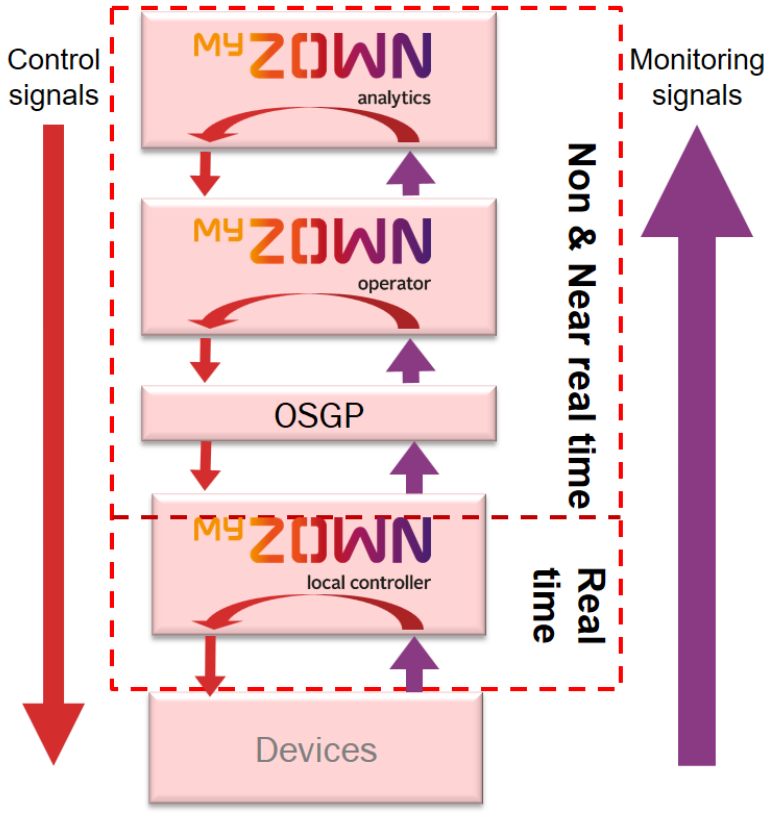

Figure 2: Basic architecture of the MGP [131]

application that receives/sends data from/to the local controller via a wireless connection. These components together (contained within the dashed red box in Fig. 2) are known as the Microgrid Platform.

In the system, two flows of information can be identified: monitoring signals are passed upwards and control signals are passed along downwards. The main control loop runs from the devices, all the way up through myZown Operator and back down again to the devices. This thesis is part of the effort to extend this loop using intelligent planning (shown as the longest loop through myZown Analytics). A shorter loop is also shown, that only runs through the local controller. The local controller is able to manage the microgrid on its own, through heuristic rules and reactive control. The intelligence (forecasting, planning and visualization) is located above myZown Operator, however. The implementation of the optimization models will be located in this part of the architecture, in a separate module called myZown Analytics. It comprises four steps:

- **Collect & Transform**. Data is collected from multiple sources: data from the microgrid itself (e.g. profiles, device data) and from external sources (e.g. market data). The purpose of this step is to transform the data in such a way that it can be

used by the system and store it in a uniform way. Outlier detection, interpolation and aggregation all have their place here.

- **Forecast**. The purpose of this step is to provide forecasts of e.g. load and solar generation that can be used in the planning. Where external services can be used for forecasting, this boils down to retrieving the forecast and possibly manipulating its format. Otherwise, it includes the development and usage of a predictive model.

- **Optimize**. This step solves an optimization problem to obtain a planning for the operation of the system. It takes the forecasts of the previous step as input, and produces a planning for the microgrid that ensures economical operation. This step is the main focus of this thesis.

- **Execute**. The execution of the planning is listed as a separate step in the process. The visualization of the planning, and the translation of the planning to device set points are the responsibility of this final step.

The architecture of the MGP is highly centralized and hierarchical. In the next section, we will provide an overview of the types of algorithms that are typically used to solve the energy management problem in similar setups.

3.5. Modeling and algorithms

At the level of the MGCC, unit commitment and economic dispatch-like problems are solved. A wide variety of models have been developed to capture these problems, and the full range of optimization algorithms have been applied to solve them. Table 3 shows an overview of approaches used in studies consulted for this thesis.

A number of studies use a heuristic approach to supply the highest priority loads using the most convenient generation first. This is referred to as "priority list method" [39, 123] or "priority load control" [37]. This is approach is very similar to the heuristic dispatch methods recently developed at Zown, and has been implemented in myZown Local Controller. In [123], this method is only used to solve the unit commitment problem, while the solution for economic dispatch is found using sequential quadratic programming. Other heuristic strategies, like load following and cycle charging, are also used [23], and in the case of [23] a genetic algorithm is used to select good dispatch strategies. A comparison between the performance of the priority list method and a genetic algorithm in [39] shows that these heuristic strategies can easily be improved upon.

In [116] dynamic programming is used as the basis for an EMS for optimal (all-electric) microgrid management, while [38] presents two dynamic programming approaches for supplying heat demand: one for minimizing peaks, and one for minimizing costs of heating systems. The authors of [92] propose a dynamic programming formulation of a microgrid that supplies both heat and electricity.

The problem is often simplified to an abstract supply-demand balancing problem, where non-linear constraints are linearly approximated. In many cases, this results in

Heuristic/priority-based	[23], [37], [39], [123]
Dynamic programming	[38], [61], [92], [116]
Mathematical programming	[8], [12], [15], [20], [48], [52], [67], [80], [84], [89], [94], [97], [99], [100], [104], [107], [114], [123], [136]
Evolutionary/genetic algorithms	[14], [23], [24], [39], [137]
Agent-based	[73], [68], [69], [106]
Other	[30], [78], [82]

Table 3: Overview of solution approaches

a mathematical programming formulation. The most accurate model would be a mixed integer non-linear programming (MINLP) formulation, but this is most often simplified to a mixed integer linear programming (MILP) formulation by approximating non-linear constraints by (piece-wise) linear functions (e.g. in [97]).

Some other examples include evolutionary algorithms (e.g. [24]), simulated annealing [30], fuzzy logic [78], multi-agent systems [68,69], particle swarm optimization [106], bee colony optimization [73] and many more.

The problem of determining the lay-out of a microgrid (dimensioning, selection and sizing) is often considered separately from the operational management (commitment and dispatch decisions) and is therefore not in scope for this thesis. Some studies, however, combine the two problems and solve them together (e.g. [12, 48]).

Surveys that (partly) treat this topic and give a broad overview of modeling techniques and algorithms used can be found in [1, 41, 85, 98].

4. Optimal control of microgrids

The problem of optimally managing a microgrid in the operational phase has already been touched upon in Section 3.1. In this section, it became clear that the traditional unit commitment problem has to be adapted to be successfully applied in the context of microgrids. Depending on the type of microgrid and the desired results a number of choices has to be made before constructing a model. The most important of these are discussed below.

For every section below a (non-exhaustive) table is provided with some examples of studies that represent different approaches in optimizing the operational aspect of microgrids.

4.1. Operation modes

A microgrid can operate in two different modes. It can either be islanded, meaning that it has to balance supply and demand by itself, without a connection to the main electricity grid, or grid-connected, in which case the microgrid *is* connected to the main grid. The two types of microgrid may have different operation objectives. Islanded microgrids simply want to minimize their operational costs, while ensuring the balance

	Islanded	Grid-connected	Both
Electricity	[20], [23], [24], [37], [46], [80], [82], [89], [94], [97], [107], [133], [137], [138]	[60], [99], [100], [104], [123], [134]	[14], [39], [52], [114], [111], [116], [130]
Electricity; Heat	[15]	[8], [10], [11], [12], [67], [136]	[48], [50], [92]
Various	[69], [84]		

Table 4: Overview of microgrid commodities and operation modes

of supply and demand within the system. Grid-connected microgrids, on the other hand, may use the external grid for balancing. It can even exploit the connection to the main grid by trading on energy markets and making a profit. Table 4 provides an overview of the operation modes considered in the studies consulted for this thesis.

Both operation modes come with their own challenges. Aside from the shared problems, the main challenge of grid-connected operation is the (possible) integration of electricity markets in the problem, while an islanded microgrid must deal with all aspects of grid operation on its own (e.g. voltage/frequency control, power quality and exact supply/demand balancing) [77]. In grid-connected microgrids, the main utility grid takes care of most of the lower level problems.

4.2. Commodities

While microgrids are primarily thought of as small electrical grids, they may serve other demands as well. The most common commodity to be include besides electrical power is thermal energy. Many microgrids include cooling/heating appliances, effectively resulting in a multi-commodity grid. Such microgrids may contain (micro)-CHP generators, heat pumps or cooling systems. Models of multi-commodity grids need to take the interaction between the different forms of energy into account. Fig. 3 displays a model of the energy flows in a multi-commodity microgrid, taking electricity, heat, diesel and gas into account. Table 4 provides an overview of the commodities modeled in the studies consulted for this thesis.

In [43], a review is presented of modeling, planning and energy management of combined cooling, heating and power microgrids. The accurate modeling of thermal generators, storage systems and loads is challenging, and as with electricity, there is a trade-off between the accuracy and the complexity of the model. The inclusion of multiple commodities further complicates the energy management of a microgrid due to the coupling of the different forms of energy.

An example of a very broad approach can be found in [84], which presents a general model for a “multigood microgrid”. A good can be anything in this context, ranging from electricity and heat to woodchips and ice cubes.

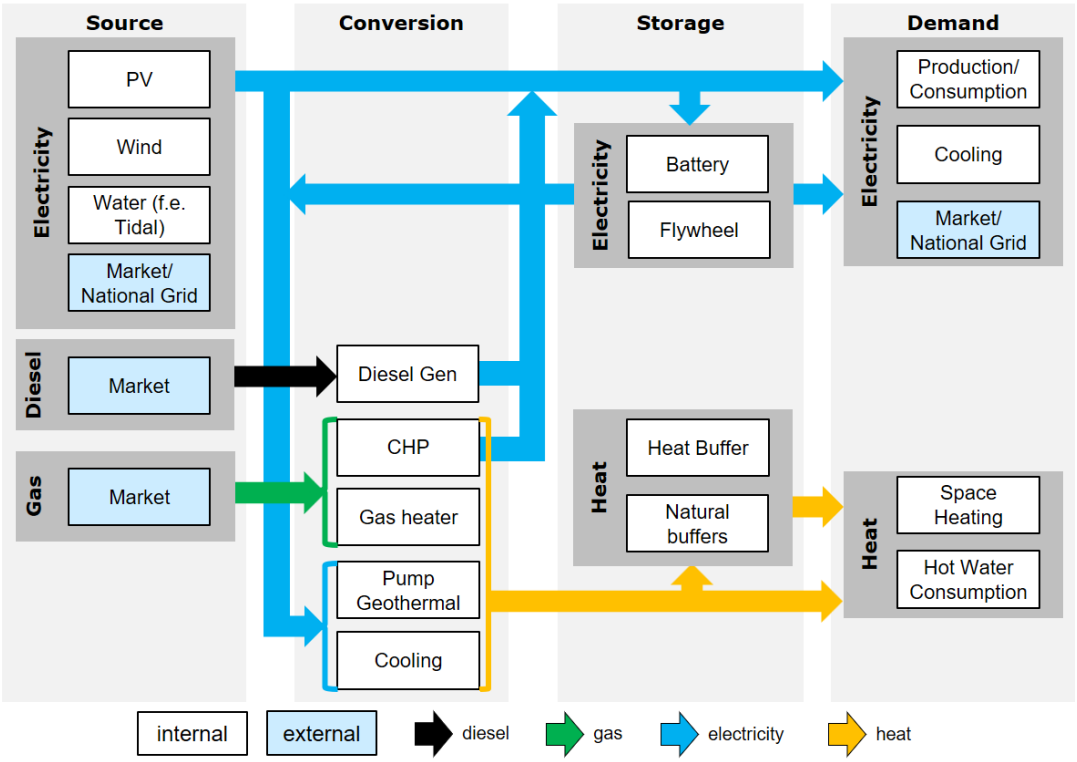

Figure 3: Energy flows in a multi-commodity microgrid [131]

4.3. Level of detail of the model

The bulk of the algorithms proposed so far treat the problem as an abstract supply-demand balancing problem, ignoring the underlying network and the associated constraints [111]. While this simplifies the problem considerably, and yields satisfactory results in many cases, the control decisions taken by such algorithms are not guaranteed to be feasible in the real world.

Opinions differ considerably on this subject. If we consider the entire problem at once to facilitate (near) real-time control, it is essential to have a very detailed model of the system. One of the main advantages of the hierarchical control structure, as described in Section 3, however, is that these problems can be (partially) separated. The lower levels of control are responsible for the (real-time) control of phase, voltage and frequency. These are responsible for ensuring the system follows power reference values and remains stable. The higher levels consider a more abstract balancing problem, ensuring *economically* optimized power dispatch [99].

Limitations that arise due to the topology of the network often need to be included in the model, but this does not mean that the full topology should be modeled. A carefully designed grid may not impose too many restrictions on the solution, which shows the dependency between the design and operational phases of microgrids.

4.4. Objective

There is a number of possible objectives in optimizing the energy management of a microgrid, which all might result in different management strategies. Fig. 1 in [1] identifies four main categories:

1. Environment cost: minimizing or penalizing carbon emissions;
2. Capital and operational costs: fuel, maintenance, energy import-costs, etc.;
3. Energy storage cost: battery efficiency, etc.;
4. Miscellaneous: e.g. power losses, load shedding cost.

A detailed description of the objective functions of a large number of studies, based on this classification, can be found in Tables 3–6 of [1].

An optimization problem may be formulated to maximize/minimize one or multiple of these objective functions. Most often, a single objective is selected, but multi-objective optimization is not uncommon. Not all formulations allow for the explicit incorporation of multiple objectives, but a linear combination of objectives into a single objective is always a possibility.

In most models, it is relatively easy to replace the objective function, while leaving the rest of the model unchanged. It is therefore not crucial to select a single objective in advance. Changing the objective from minimizing operational costs to minimizing CO_2 emissions simply requires the change of the coefficients in the function. In other words: the monetary cost per kWh generated has to be converted to the amount of CO_2 emitted per kWh generated.

Of course, this is not always as straight forward. It has to be possible to express the objective in terms of the units that the problem uses. In a simple unit commitment/economic dispatch problem, these are decisions on the on/off-state of generators and power in- or output levels of generators, storage units and loads. Any objective function needs to be expressed in terms of cost per generated kW and cost for being switched on or off. If different costs are associated with multiple efficiency levels of a generator, for example, the modeling of the problem *does* need to change to account for this in the objective function.

4.5. Uncertainty and robustness

A central challenge in the management of microgrids is dealing with the uncertainty that exists. As discussed in section 3.1, the uncertainty associated with the amount of renewable energy generated on the one hand and the demand on the other complicates the problem significantly. It is unlikely that any approach that does not take this into account in some way will be capable of managing microgrids efficiently.

There is number of ways that uncertainty can be dealt with. An overview is provided in Table 5, as a supplement to the discussion below.

One of the simplest ways of dealing with uncertainty is the planning of spinning reserve. In this way, there is always back-up capacity available that can compensate for

MPC/RH	[11], [46], [67], [80], [84], [94], [97], [99], [100], [104], [114], [133], [134], [136], [138]
Two/multi-stage stochastic programming	[58], [105], [107], [115], [119]
Chance-constrained programming	[74]
Robust optimization	[68]

Table 5: Overview of approaches to uncertainty

fluctuations in actual demand. An example of the usage of spinning reserve in microgrids is [39].

Many studies apply Model Predictive Control (MPC) or Rolling Horizon strategies, which boils down to solving an updated version of the same deterministic problem every time step, thereby adjusting the planning to the current situation (see e.g. [97]). These strategies deal with uncertainty without explicitly modeling it. It is countered by reevaluating the planning quite often based on updated information.

Other approaches, that do modify the model significantly to take uncertainty into account, include two-stage (scenario-based) stochastic programming (e.g. [115]), multi-stage stochastic programming [105], chance-constrained programming [74] and robust optimization [68]. These studies mainly consider all-electric microgrids. The number of studies considering multi-commodity microgrids is much more limited [43, 75, 102, 132].

The authors of [19] aim to give an overview of stochastic programming approaches to solve unit commitment problems. Stochastic unit commitment is also extensively discussed in [53]. The authors of [71] review stochastic modeling and optimization tools used for microgrid planning, operation and control.

Roughly, uncertainty can be dealt with in one of three ways:

- Increasing the model complexity to explicitly take the uncertainty into account (e.g. two-stage stochastic programming);
- Increasing the number of times a new planning is generated, re-evaluating the problem based on the current situation (e.g. rolling horizon strategies);
- Adding large safety margins in the constraints of the model, such that it can be expected that the planning will be feasible in most cases, but will typically be far from optimal (e.g. robust optimization).

Part II.
Models & Algorithms

In this part, we will introduce the problem of the optimal operational management of microgrids. We will present an outline of the solution process and describe it step-by-step.

5. Problem definition

The optimal operational management of a microgrid involves making decisions about the control of all devices in the grid. As explained in Section 3, we will not consider real-time control, but the generation of a planning for the microgrid for the next 24 hours. We aim to develop a flexible model and solution algorithm that is easily adapted to various microgrid configurations.

Our optimization problem is to find the cheapest solution (in terms of operational costs) that ensures the energy balance in every time step (i.e. supply and demand are balanced), given the constraints that the grid and the devices in it impose. This objective fits in the second category mentioned in Section 4.4.

As we explained in Section 3.1, the problem can be seen as an extension of the classical unit commitment problem, where adding the possibility of including buffers (like batteries) into the problem is the most substantial change with respect to the problem that we described in Section 3.1. It is, therefore, a more general version of unit commitment. We will show that the basic problem is already NP-hard by providing a polynomial time reduction from the knapsack problem. We have presented a typical unit commitment problem in Section 3.1. Stripping the possible extensions from this problem, a basic unit commitment problem can be formulated as follows:

$$\text{minimize } \sum_{i,t} (y_{it} f_i + g_{it} c_i + x_{it} s_i) \text{ s.t.}$$

$$\begin{aligned}
\sum_i g_{it} &= D_t && \forall t \\
y_{it} P_i^{min} \leq g_{it} &\leq y_{it} P_i^{max} && \forall i, t \\
y_{it} &\leq y_{it-1} + x_{it} && \forall i, t \\
x_{it}, y_{it} &\in \{0, 1\} && \forall i, t
\end{aligned}$$

where

(decision variables):

g_{it}	is the amount of energy produced by generator i at interval t
x_{it}	is 1 if generator i is started at the start of interval t
y_{it}	is 1 if generator i is active during interval t

(constants):

D_t	(expected) total demand at interval t
f_i	operating cost of generator i at interval t
c_i	generation cost of generator i at interval t (cost per generated unit of energy)
s_i	start-up costs of generator i at interval t
P_i^{min}	minimum production of generator i
P_i^{max}	maximum production of generator i

Note that this is a special case of our more general problem, obtained by omitting buffers and disregarding efficiency constants and reserve constraints. In other words, by reducing the knapsack problem to this problem we show that the inclusion of a balance constraints and generator capacity and start-up constraints suffice to make the problem NP-hard.

Now consider the knapsack problem:

$$\text{maximize } \sum_{i=1}^{n} v_i q_i \text{ s.t.}$$

$$\sum_{i=1}^{n} w_i q_i \leq W$$

$$q_i \in \{0,1\} \qquad \forall i$$

where
(decision variables):

q_i	is 1 if we select item i and 0 otherwise

(constants):

W	the size of the knapsack
w_i	the weight of item i
v_i	the value of item i

Any knapsack problem of this form can be reduced to a unit commitment problem as follows, assuming a formulation of the knapsack problem is given, where n is the number of items:

1. Set $s_i, c_i = 0 \ \forall i$
2. Set $T = 0$ (i.e. $t \in \{0\}$) and $G = \{0, 1, ..., n\}$ $(i \in G)$
3. Set $f_i = -v_i \ \forall i \in \{1, ..., n\}$
4. Set $P_i^{min} = P_i^{max} = w_i \ \forall i \in \{1, ..., n\}$
5. Set $D_0 = W$
6. Set $P_0^{min} = 0, P_0^{max} = W, f_0 = 0$

This reduces the knapsack problem to a single-time step unit commitment problem where every item is represented by a generator that can be turned on or off at only one output level corresponding to the weight of the item in the original problem. Turning a generator on has a negative cost corresponding to the value of the item in the original problem. The demand that has to be met equals the size of the knapsack and an additional generator 0 ensures that the equality of the demand constraint is satisfied. An optimal solution to the problem thus obtained can be translated to an optimal solution to the original knapsack problem by selecting all the items of which the corresponding generator was turned on in the transformed problem. The value of the objective function for this solution will equal the value of the objective function of the transformed problem, multiplied by minus one.

Knapsack is a weakly NP-hard problem. A reduction from the strongly NP-hard 3-partition problem was provided in [7]. However, this requires the addition of minimum downtime constraints or a restriction on the number of start-ups to the basic problem described above. Including fuel tank constraints allows for a similar reduction.

Every device in the microgrid either consumes or produces energy (or both). So, the essential variable defining the behavior of a device is its in- or output over time, otherwise called the device profile. While some devices may require additional variables to model them, all devices will have variables that represent this profile. The variable g_{it} in the problem above is an example of such a profile variable that defines the output of generator i at every time step t. A profile for a single device may be represented by two variables if the device can both produce and consume energy, as we will see exemplified by the batteries in the next sections. The profile variables should sum to zero for every time step, to ensure energy balance in the system.

The general problem thus is to find profiles for all devices in a microgrid such that the energy balance can be maintained in every time step.

To account for multiple commodities (such as heat, electricity and gas), one might include multiple balance constraints (one for every commodity) and multiple profile variables for each device. However, we will limit ourselves to a single commodity, electricity, in this thesis.

As it is unnecessary to model devices of which the profiles cannot be controlled (such as PV-panels, for example) they can be eliminated from the problem by including their contribution in the expected demand at every time step. However, uncertainty may exist about the output or consumption of these uncontrollable components (e.g. loads and PV). So, we may need to cover a range of possibilities, depending on the expected output of PV panels and the expected demand of the loads in the system. In general: we need to account for the uncertainty around the behavior of uncontrollable components in our grid.

We will use this basic understanding of the problem in the models in Section 9. Although we limit ourselves to one commodity and a limited number of devices, it should be clear from the discussion above how these models can be extended to include multiple commodities and new devices. Similarly, it is therefore straightforward to adapt

the model described in the following sections to use a different objective (e.g. CO_2 minimization) by changing cost constants and possibly introducing additional auxiliary variables.

In Section 3.3, we compared the (dis)advantages of using a centralized or distributed control architecture. We will assume centralized control of the microgrid, for the following reasons:

- The system architecture of the MGP, the layout of the existing microgrids, and the plans for the myZown Analytics are (currently) not suited for a distributed approach, thereby requiring a centralized approach;
- The microgrids that are currently managed by Zown and those expected to be added in the near future, are all relatively small microgrids that are managed by a single stakeholder. This is the kind of microgrid for which a centralized approach is the most suitable, as explained in Section 3.3;

We assume that the local microgrid controller (programmed into an RTU, see Section 3.4) implements a recovery policy that monitors the grid for extreme situations. Our model generates a planning (or more than one, in case of a scenario-based approach) that is feasible with a high probability. However, due to a number of reasons (e.g. a low-quality forecast or large fluctuations within a single time step), it may be found to be infeasible during operation, when the local controller is implementing it. When this occurs, the RTU steps in and overrules the decisions from the generated problem. This is done by monitoring certain bounds that the system should never exceed (e.g. a minimal state of charge for the batteries when all generators are off-line). When such a bound is exceeded, the system detects that an extreme situation is occurring, causing it to take over control to avoid a black-out.

We will use mathematical programming techniques to model the problem, which will be explained in more detail in Section 9.

To avoid unnecessary repetition, we will first provide a list of symbols that will be used in the formulations of the models below. Example values, based on the Island test case (Section 12), are provided for the constants listed.

(decision variables):		*stage*
g_{its}	is the amount of energy produced by generator i at interval t in scenario s	1
x_{its}	is 1 if generator i is started at the start of interval t in scenario s	1
y_{its}	is 1 if generator i is active during interval t in scenario s	1
b^{in}_{its}	is the amount of energy going into battery i at interval t in scenario s	2
b^{out}_{its}	is the amount of energy coming out of battery i at interval t in scenario s	2
ch_{its}	is the level of charge of battery i at interval t in scenario s	2
w^+_{ts}	unsatisfied demand at interval t in scenario s	2
w^-_{ts}	unused surplus at interval t in scenario s	2
r_{it}	is 1 if the reserve is above the maximum threshold for battery i at interval t	n.a.
(constants):		*value*
p_s	relative probability of scenario s occurring	
D^-_{ts}	Negative demand at interval t in scenario s	
D^+_{ts}	Positive demand at interval t in scenario s	
f_i	operating cost of generator i	*€0.40*
c^g_i	generation cost of generator i (cost per generated unit of energy)	*€0.30*
s_i	start-up costs of generator i	*€0.50*
P^{min}_i	minimum production of generator i	*40/32 kW*
P^{max}_i	maximum production of generator i	*8/6.4 kW*
c^b_i	cost per unit of energy discharged from battery i	*€0.00057*
E_i	maximum charge level of battery i	*20 kWh*
v^c_i	maximum charging speed of battery i	*12 kW*
v^d_i	maximum discharging speed of battery i	*12 kW*
R^{max}_i	maximum reserve threshold of battery i	*3 kWh*
R^{min}_i	minimum reserve threshold of battery i	*1 kWh*
η^b_i	efficiency of battery i	*0.93*
c^w	penalty for 1 kWh of unsatisfied demand or surplus	*€2*
η^g	Grid efficiency	*0.97*
(other symbols):		
T	Index set of time steps in the problem (i.e. $t \in T$)	
G	Index set of generators in the problem (e.g. $i \in G$)	
B	Index set of batteries in the problem (e.g. $i \in B$)	
S	Index set of scenarios in the problem (e.g. $s \in S$)	

Note that the subscripts mentioned may not always all be present. For example, in some models, we do not work with scenarios, which means that the subscript s is dropped from most constants and variables. However, if a variable is missing a subscript, it will always be the last one mentioned here, i.e. the subscripts of a variable may be interpreted in the order that they are introduced here.

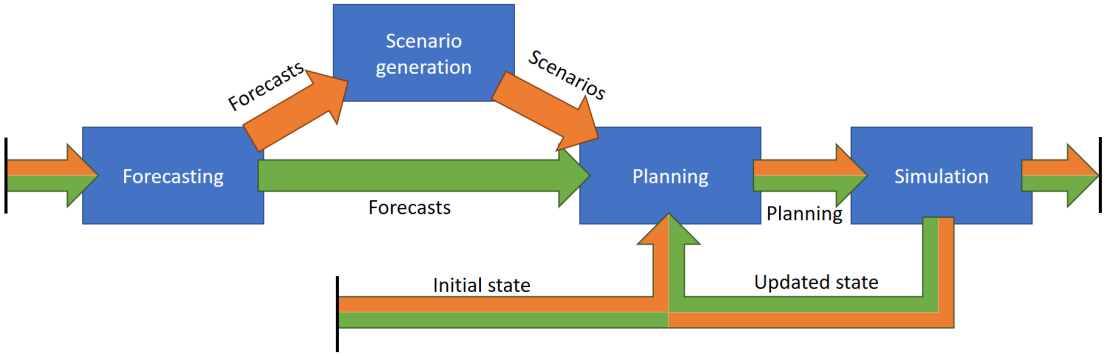

Figure 4: High level representation of the solution process. The orange arrows represent the flow for the stochastic models, while the green arrows represent the flow for the deterministic models. Note that "Simulation" is replaced by the application to the actual grid when the algorithms are operational.

The rest of this part will be organized as follows. It may be noted that the organization of the discussion in this part follows the flow of the solution approach (Fig. 4), treating each aspect as it occurs chronologically in the process.

In Section 6 we will briefly discuss the models that generate the forecasts used as input for the optimization models and some observations on what makes a forecast 'good' for our purpose. We will then discuss the generation of scenarios in Section 7. In Section 8 we discuss how the individual components of a microgrid are modeled. This is followed by a description of the optimization models in Section 9. All models in this section use the forecasts as described in Section 6 and (part of) the device models as described in Section 8. The stochastic models also use the scenarios generated by the approach described in Section 7. Following the optimization models, we will describe the simulation model in Section 10 that simulates the application of the planning generated by the optimization models in a microgrid and evaluates its performance. Finally, we will provide some details on the implementation that were not yet included in other sections in Section 11.

6. Forecasts

In our models, we use forecasts to predict the output of PV panels and the demand in the microgrid(s) under study. No prediction is made for the generation of wind energy, although a small number of WTs is installed on Island (see Section 12 for a description of this microgrid). The contribution of these generators, however, is very small. At the time of writing, they have even been temporarily removed, after being severely damaged during a storm. Their output was not measured separately, but is part of the load measurements: the WTs are connected to a small load in the network, and only

the total power output/usage of this group is measured. The contribution of the WTs is therefore not considered separately but is included in the load forecast. This may, theoretically, lead to a negative load. However, the contribution of the WTs is far too small to satisfy the base load in this microgrid, even when fully functional. In practice, the load will always be positive.

We will describe the models and algorithms used for both PV-output and load forecasting briefly below.

6.1. PV

The forecasting models and algorithms for the PV-forecast were developed by iCarus [56].

The algorithm uses a number of different models to predict the output of a group of solar panels based on historical data from this system and weather data. The forecasts of these models are combined and scaled by a time-specific weight. These weights determine how much the forecast of each of the models influences the final forecast. They are learned from minimizing the forecast-error on historical data and can vary based on the time that the forecast is created and how far ahead the predicted value is in time.

The algorithm also learns how well it performs from its own performance on past data. Based on this, a standard deviation is determined and passed along with each predicted value to indicate the distribution at that particular time, based on past experience. The underlying assumption that the forecast error is approximately normally distributed was made after analyzing the distribution of the forecast error on past predictions.

We used this existing algorithm to obtain forecasts for the test instances discussed in Section 12.

6.2. Load

The forecasting models and algorithms for the load-forecast were developed as a joint effort by iCarus and Zown during this project. The general principle of these algorithms is the same as for the PV-forecast: the outcomes of a number of models are combined to obtain a single forecast. The standard deviation is determined in the exact same way.

The models that were studied to be used for the prediction of the load profiles are: neural networks, (S)ARIMA(X) models, random forest and a number of "educated guess-models". These last models include the "profile predictor", which uses the exact profile of the most comparable historical day as a prediction, i.e. it tries to forecast the type of day rather than the individual load values, and the "basecases" which simply predict today to be the same as yesterday (basecase yesterday) or the same day last week (base case last week).

In the end, only the two basecases were included in the combined forecasting algorithm, which we use in our test instances.

6.3. Forecast quality

While developing the models described below (Section 9) we have constructed artificial forecasts and performed a simulated test of two weeks using the (preliminary) optimiza-

tion model that was available at that particular time. These tests were otherwise set up exactly as the tests we discuss in Section 13 and were aimed at determining the effect of specific features of the forecast on the cost of operation. We generated a new planning every 15 minutes and determined the costs using the same simulator described in Section 10. This provided some insight in the relevant features of a good forecast. We will briefly describe some of these investigations and draw some conclusions about the necessary features of a good forecast. Keep in mind that all tests were run for a single microgrid (Island, see Section 12), currently managed by Zown, and that the results can therefore by no means be seen as general principles. To elicit what aspects of a forecast are most important for a particular application, one should run their own set of artificial forecasts through a simulator and analyze the results.

The first set of artificial forecasts was constructed based on the data of two weeks of operation from 6 until 19 March 2017.

The PV forecasting algorithm used has proven itself in similar situations before, but the load forecasting algorithms described above are themselves currently being developed. To elicit which features of a load forecast contribute to a good result we constructed a number of artificial scenarios and fed them to the simulator (described in Section 10), which provided an indication of the performance of the entire solution approach using these forecasts as input. Note that we only constructed artificial load forecasts for this purpose: the PV forecast was not artificially constructed. The artificial scenarios include the following:

- The "perfects". These predict the *actual* values (perfect prediction) for a number of hours, and after that just predict the overall average at every time step. We tested scenarios that have a perfect prediction for the first 1, 2, 4, 6 and 12 hours.
  ⇒ The quality of the solution improves significantly as the length of the perfectly predicted part increases. However, this effect only holds for the first four hours. There is hardly any difference in the quality between the 6- and 12-hour versions. This indicates that the accuracy of the prediction for later time steps does not have that much impact.

- The "shifts". These predict the *actual* values, but shifted by a number of hours. We have tested scenarios with a shift of 1, 2, 4 and 8 time steps in both directions (early and late).
  ⇒ This has a very large impact on the performance. This indicates that the prediction of the timing of peaks in demand is very important.

- The "bias's". These predict the *actual* values, but with a bias of -10%, -5%, -2%, -1%, 1%, 2%, 5% or 10% (of the average value). These structurally predict either too low or too high values.
  ⇒ Bias decreases the performance on the short term, but not dramatically.

- The "blocks". These predictions are aggregated into blocks. During a block, the average actual value of that block is predicted. We constructed scenarios with blocks of 2, 4 and 8 time steps.

⇒ Aggregation results in solutions that are a little worse, but the effect is not that large.

- The "errors". These predictions are the *actual* values with a random error (either positive or negative) of up to 1%, 2%, 5%, 10% or 20% of the average value. This error is determined separately for every value.
  ⇒ For problems with a smaller planning horizon (e.g. 4 hours) these errors decrease the solution quality. On problems with a longer planning horizon, however, they result in a better solution. Keep in mind that these tests were run with a naive deterministic model. The artificial errors seem to emulate the effect of scenarios.

The main conclusions drawn from this are that it is important for the forecast to be accurate on the short term (first 4-6 hours). The accuracy of prediction on the long-term is less important. However, not considering the planning after six hours does reduce the solution quality significantly. One might say that it is relevant to consider that the future exists, but it is not particularly crucial to predict it accurately after six hours.

Furthermore, it is more important to predict the timing of the peaks correctly than getting their height exactly right.

To investigate the effect of improving the forecast on the overall quality of the planning resulting from the optimization, we constructed another set of artificial forecasts.

We based this set on the data for the first seven days of May 2017. Every artificial forecast samples its "prediction" at every time step from a normal distribution with *bias* times the realized value as its mean and *stdev* times the realized value as its standard deviation. The values considered for *bias* are: 0.8, 1.0, 1.2 and 1.4 and for *stdev*: 0.0, 0.1, 0.2, 0.4 and 0.6. The standard deviation linked with the resulting values is learned in the same way as before.

Unsurprisingly, decreasing the standard deviation results in significantly better solutions, and a bias closer to 1.0 yields better results, both in terms of overall costs of the planning and the number of times that the reactive control has to overrule the generated planning.

7. Scenario trees

For the stochastic models, we need to represent the uncertainty associated with the forecasts described in the previous section. We will do this by generating scenarios and constructing scenario trees based on these scenarios. In general, this process consists of two major steps: (1) sampling scenarios and/or constructing an initial scenario tree; (2) reducing the size of the initial scenario tree and adapting its structure to obtain a tree of the desired size and form that can be used in the optimization model.

For the two-stage model presented in Section 9.3, only the first step is relevant, while both are crucial for the multi-stage model presented in Section 9.4.

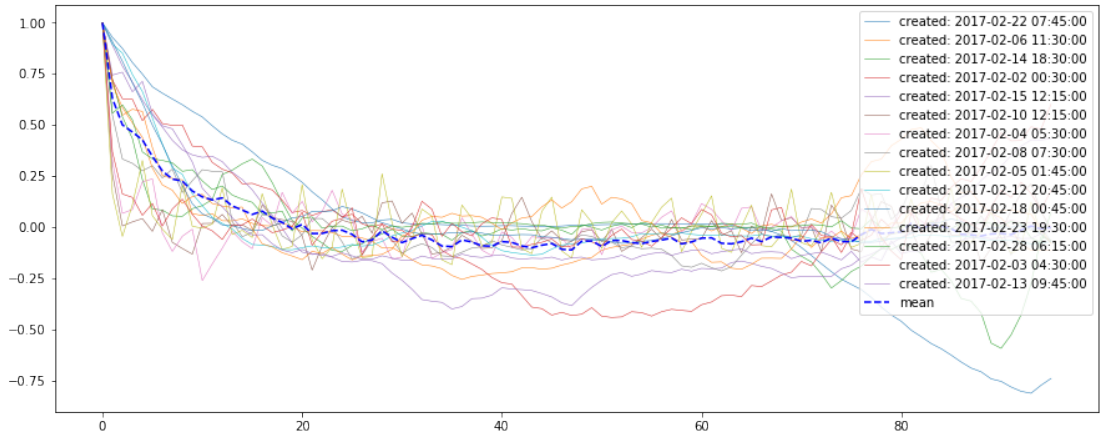

Figure 5: The correlation of the forecast-error (load) as a function of the time-lag (the number of time steps between two values) for 15 different 24-hour forecasts.

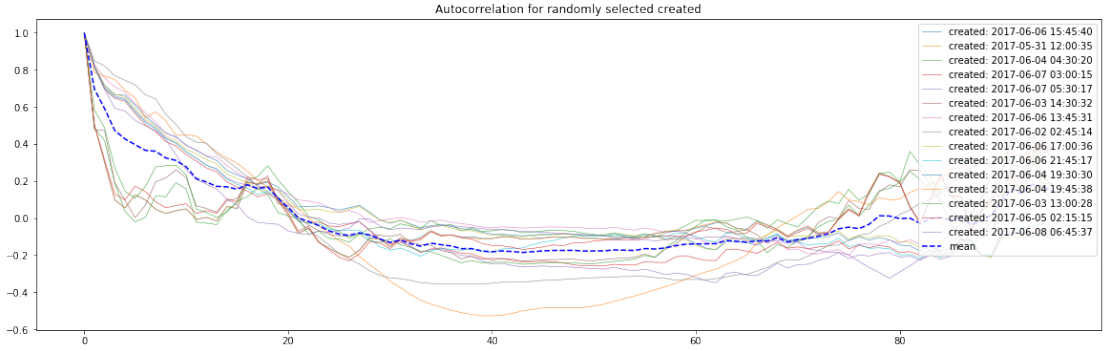

Figure 6: The correlation of the forecast-error (PV) as a function of the time-lag (the number of time steps between two values) for 15 different 24-hour forecasts.

7.1. Sampling and initial tree construction

First, we will sample a large number of random scenarios. In this thesis, we are only considering scenarios for load and PV-output, but this approach can be generalized to any number of uncertain factors.

We assume the solar and load scenarios to be independent. Additionally, we assume that the errors on the forecast at different time steps within the same scenario are correlated. In an analysis on the autocorrelation of the error of older forecasts for the Island system we found this to be the case in the load forecasts (see Fig. 5) and the PV forecasts (see Fig. 6). The mean correlation for values that are one time step apart was around 0.63 (load) and 0.74 (PV). We will simplify our modeling of the scenarios by only considering the correlation between values that directly follow each other, and we will use a correlation coefficient of 0.63 when constructing load scenarios and of 0.74

when constructing solar scenarios.

Note, however, that this is by no means a general conclusion. The correlation coefficient might change as the forecasting method changes, different algorithms are used for forecasting, a different period in time or microgrid is considered or another variable is forecasted. Before implementing this algorithm anywhere, the particular system and its forecasting algorithms should be studied to validate the choice for any correlation coefficient to be used. This choice should also be periodically re-evaluated.

Assuming that the correlation coefficient (ρ) is known, the sampling of a single scenario is done as described in Alg. 1. We sample a value for each time step based on the distribution defined by the predicted value and the associated standard deviation at that time step, taking the correlation of the errors in subsequent time steps into account. As we assume the solar and load scenarios to be independent, a single scenario is obtained by performing Alg. 1 twice: once to obtain a solar scenario and once to obtain a load scenario. Combining these independently sampled scenarios yields a single scenario to be plugged into the problem.

Input: planning horizon T, correlation ρ, forecasts $\mu = (\mu_1, ..., \mu_T)$, standard deviations $\sigma = (\sigma_1, ..., \sigma_T)$
$X_1 \sim \mathcal{N}(0, 1)$
$s_1 \leftarrow \mu_1 + \sigma_1 X_1$
for $i \leftarrow 2$ **to** T **do**
  $X_i \sim \mathcal{N}(0, 1)$
  $X_i \leftarrow \rho X_{i-1} + \sqrt{1 - \rho^2} X_i$
  $s_i \leftarrow \mu_i + \sigma_i X_i$
end
return $(s_1, ..., s_T)$

Algorithm 1: Sampling a single scenario

We repeat this process $|S|$ times to come up with the initial scenario fan.

There are alternative ways of constructing an initial scenario tree. In [101] a tree is constructed by determining the b points that best represent the probability distribution in each step. The value of b can differ in every stage and is called the *bushiness*. It is a level-wise branching factor. Picking $b = 3$ for all levels would result in a tree with 3^N leaves, where N is the number of time steps in the model. As our problem considers between 40 and 100 time steps, even a small bushiness of 2 would yield extremely large trees: the smallest possible tree would already have $2^{40} \approx 1.1 \cdot 10^{12}$ leaves, each representing a separate scenario. Reducing that to a reasonable amount for use in our model will take way too much computational time and effort, considering that the resulting tree might not represent the full range of possibilities as much as we would like: in the first time step just two scenarios are considered.

One might solve this problem by increasing the bushiness in the earlier time steps and keeping it low in the later levels. At this point, it is guesswork to determine what the best bushiness-values are to obtain a good initial tree. An algorithm is presented in [101] that generates such a tree with "flexible" bushiness. This increases the bushiness in each

node, until it has found a bushiness that is within some fixed bound on a certain distance measure, which is a way to ensure that the tree represents the underlying distribution well.

Our approach is to "trust" the random sampling of a large number of scenarios, as described above, to come up with a fairly representative set. Typically, such a sample will represent the distribution a lot better in the earlier time steps and will be a much less accurate representation in the later stages. We already touched upon this issue before: this is exactly in line with the importance of the earlier time steps compared to the later ones. It is important that the scenarios represent the first few time steps well. The further we get in time, the less important accuracy becomes. This is partly due to the fact that we are not using many of the decisions taken for the later time steps in most cases.

The idea is to join the generated scenarios at the root, and start the reduction process on that initial tree, grouping similar scenarios as we go. This approach follows the general scenario tree construction approach outlined in [25], where it is recommended to use cluster analysis to group similar scenarios after the initial sampling. We will follow an approach that is similar to the one taken in [42].

7.2. Scenario reduction

From the initially constructed large scenario fan, we want to construct a much smaller (in terms of the number of nodes) scenario tree (see Fig. 7). We want the probability distribution of this scenario tree to match the distribution of the original fan as close as possible.

We will be using some of the algorithms developed in [26, 49] and applied in [42] for this purpose. More specifically, we will be using the simultaneous backward reduction algorithm, using the Kantorovich distance as a distance measure. An overview of the algorithm is presented in Alg. 3.

Figure 7: The scenario reduction process reduces the scenario fan (left) to a scenario tree (right), by grouping scenarios at every level of the tree. This means that the tree has fewer leaves than the fan, and that it branches on all levels, instead of just the first.

7.2.1. Kantorovich Distance

We will be using the Kantorovich Distance as a measure to compare scenarios. The description in this section is based on [42].

The Kantorovich Distance D_K, for discrete probability distributions, is the optimal solution to the following linear programming problem:

$$\text{minimize} \sum_{i \in S} \sum_{j \in S'} \eta_{ij} c_{|T|}(\xi^i, \tilde{\xi}^j) \text{ s.t.}$$

$$\begin{aligned} \sum_{i \in S} \eta_{ij} &= q_j && \forall j \in S' \\ \sum_{j \in S'} \eta_{ij} &= p_i && \forall i \in S \\ \eta_{ij} &\geq 0 && \forall i \in S, \forall j \in S' \end{aligned}$$

Where S is the index set of scenarios from one distribution with scenarios ξ^i and corresponding probabilities p_i, and S' is the index set of scenarios from the other distribution with scenarios ξ^j and probabilities q_j.

The problem comes down to finding a mapping of minimal weight between the scenarios representing both distributions, where c_t is the distance measure between two scenarios defined as:

$$c_t(\xi^i, \xi^j) = \sum_{\tau=1}^{t} dist(\xi^i_\tau - \xi^j_\tau)$$

So, $c_{|T|}$ represents the distance of two scenarios over the entire planning horizon.

Our scenarios contain two values: a forecast for the PV-output, and a forecast for the load. To determine the distance between scenarios, however, we can reduce these to a single value. Only the *difference* between uncontrollable demand (load) and uncontrollable production (PV-output) matters in our models. So, we define $dist()$ on these composite scenarios as follows.

Let $\xi^i_{t,pv}$ be the forecast for the PV-output at time step t in scenario i and $\xi^i_{t,load}$ the load forecast:

$$dist(\xi^i_t - \xi^j_t) = |(\xi^i_{t,load} - \xi^i_{t,pv}) - (\xi^j_{t,load} - \xi^j_{t,pv})|$$

This distance measure defines the distance between two scenarios at time t as the difference in *reduced* or net demand (the difference between PV-output and load) between two scenarios. A scenario with a high PV-output and a high load may, therefore, be treated as very similar to a scenario with a low PV-output and a low load, as the net demand that has to be satisfied is nearly the same. While this is a good idea in the context of the models presented in Section 9, this may not always be the best choice for a distance measure. For example, when explicitly modeling PV curtailment, the two example scenarios mentioned before are actually very different. In this case one might consider the difference in both forecasts separately (i.e. $dist(\xi^i_t - \xi^j_t) =$

$|(\xi^i_{t,load} - \xi^j_{t,load})| + |(\xi^i_{t,pv} - \xi^j_{t,pv})|)$. But as we consider both PV-output and load to be uncontrollable in our model, the aforementioned distance measure is a good fit.

We want to use the Kantorovich distance to express to which extent the distribution represented by a subset of scenarios resembles the distribution represented by the original set. Let J be an index set of the scenarios that are deleted from the original set to obtain the reduced set. If the original set contained $|S|$ scenarios, we have that $J \subset \{1, ..., |S|\}$. Our reduced set of scenarios contains all ξ^j such that $j \in \{1, ..., |S|\} \setminus J$. Now the Kantorovich distance of this subset to the original set equals $\sum_{i \in J} p_i \min_{j \notin J} c_{|T|}(\xi^i, \xi^j)$. This follows directly from our definition of the Kantorovich distance above if we consider that all non-deleted scenarios will map onto themselves, and we therefore only have to consider the distance of the deleted scenarios from the original set to one of the remaining scenarios in the subset.

It follows that the probability q_j of the any preserved scenario ξ^j should equal its original probability, plus the probability of all scenarios mapped onto it: $q_j = p_j + \sum_{i \in J(j)} p_i$, where $J(j) = \{i \in J | \arg\min_{j' \notin J} \{c_T(\xi^i, \xi^{j'})\} = j\}$.

Minimizing the Kantorovich distance over all possible subsets of a given size can now be formulated as a MILP, and presents an NP-hard problem in itself.

As there is no efficient way to find the optimal solution for this problem, we will use a heuristic algorithm called *simultaneous backward reduction* that has been shown to provide good results in [49].

7.2.2. Simultaneous backward reduction

The idea of the *simultaneous backward reduction* algorithm is to repeatedly delete a single scenario: the one that increases the Kantorovich distance of the reduced set to the original one the least. It is fairly straight-forward to determine which scenario this is:

$$\min_{j \in \{1,...,|S|\}} \{p_i \min_{j \neq i} c_T(\xi^i, \xi^j)\}$$

This is used in Alg. 2. It uses a stopping condition θ, that controls how many scenarios remain after the reduction. In this way, the complexity of the resulting tree can be controlled. Alternative stopping conditions include a tolerance ϵ. This tolerance expresses a maximum distance increase that the deletion of a single scenario may result in. Alternatively, this tolerance could apply to the overall distance increase (over all deleted scenarios).

The idea is to apply this algorithm on every level of the tree and fixing the child/parent pointers of every scenario based on which scenario it collapses into. This idea is implemented in Alg. 3, which uses a slightly modified version of Alg. 2 as a subroutine.

For this algorithm, consider every scenario to be a list of T entries, each with one value for every forecasted variable and an associated probability at that time step. Additionally, imagine a $T \times S$ matrix M defining the parent of every scenario. E.g. $M_{ts} = s'$ means that scenario s is joined with scenario s' at time step $t - 1$. If a scenario is not

Input: set of scenarios $\{\xi^i\}, i \in \{1, ..., |S|\}$ and associated probabilities p_i, target θ

`/* Compute the distances of all scenario pairs */`

for $i \leftarrow 1$ **to** $|S|$ **do**
  for $j \leftarrow 1$ **to** $|S|$ **do**
    $c_{ij} \leftarrow c_T(\xi^i, \xi^j)$
  end
end

`/* Iteratively remove a single scenario */`

$J \leftarrow \emptyset$

repeat
  foreach $l \in \{1, ..., |S|\} \setminus J$ **do**
    foreach $k \in J \cup \{l\}$ **do**
      $c'_{kl} \leftarrow \min_{j \notin J \cup l} c_{kj}$
    end
    $z_l \leftarrow \sum_{k \in J \cup \{l\}} p_k c'_{kl}$
  end
  $l_i = \arg\min_{l \notin J} z_l$
  $J \leftarrow J \cup \{l_i\}$
until $|S| - |J| = \theta$;

`/* Recompute the probabilities of the remaining scenarios */`

foreach $i \in \{1, ..., |S|\} \setminus J$ **do**
  $J' \leftarrow \{j \in J | \arg\min_{j' \in J} c_{ij'} = i\}$
  $q_i \leftarrow p_i + \sum_{j \in J'} p_j$
end

return $\{q_1, ..., q_{|S|}\}$

Algorithm 2: Simultaneous backward reduction

removed at t we have that $s = s'$. These fully represent the tree structure of the scenario tree. Note that the probability of a scenario may be computed by summing the probabilities of its children, but for ease of notation, we store the result of this sum at every node in the tree.

The algorithm will start at $t = |T|$ by reducing the number of scenarios at this level to $\theta_{|T|}$ by iteratively removing the scenario that causes the smallest increase in distance of the reduced set to the original set of scenarios. This is repeated for every level in the tree, where deleted scenarios will contribute their probability mass to the closest remaining scenario. A scenario that is deleted at time step t is only deleted in all levels above t in the tree, representing the time steps before t (i.e. $0, 1, ..., t-1, t$), but is still present in the lower part of the tree (time steps $t+1, ..., |T|$). The node representing the deleted scenario one level deeper (at $t+1$) becomes a child of the node representing the

Input: set of scenarios $\{\xi^i\}, i \in \{1, ..., |S|\}$ and associated probabilities $p_{i(|T|+1)}$, $|T| \times |S|$-matrix of parent pointers M, targets $\theta_t, t \in \{1, ...|T|\}$

for $t \leftarrow |T|$ **to** 1 **do**

  `/* Compute the distances of all scenario pairs */`

  for $i \leftarrow 1$ **to** $|S|$ **do**

    for $j \leftarrow 1$ **to** $|S|$ **do**

      $c_{ij} \leftarrow c_t(\xi^i, \xi^j)$

    end

  end

  `/* Iteratively remove a single scenario */`

  $J \leftarrow \emptyset$

  repeat

    foreach $l \in \{1, ..., |S|\} \setminus J$ **do**

      foreach $k \in J \cup \{l\}$ **do**

        $c'_{kl} \leftarrow \min_{j \notin J \cup l} c_{kj}$

      end

      $z_l \leftarrow \sum_{k \in J \cup \{l\}} p_k c'_{kl}$

    end

    $l_i = \arg\min_{l \notin J} z_l$

    $J \leftarrow J \cup \{l_i\}$

  until $|S| - |J| = \theta_i$;

  `/* Recompute the probabilities of the remaining scenarios */`

  foreach $i \in \{1, ..., |S|\} \setminus J$ **do**

    $J' \leftarrow \{j \in J | \arg\min_{j' \in J} c_{ij'} = i\}$

    $p_{it} \leftarrow p_{i(t+1)} + \sum_{j \in J'} p_{j(t+1)}$

    foreach $j \in J'$ **do**

      $m_{(t+1)j} \leftarrow i$

    end

  end

  $S \leftarrow S - J$

end

Algorithm 3: Scenario reduction algorithm

$\|S\|$	l_x	time (s)
300	l_1	60.44
	l_2	60.98
	l_3	60.75
350	l_1	111.50
	l_2	111.08
	l_3	110.60

$\|S\|$	l_x	time (s)
400	l_1	190.33
	l_2	190.63
	l_3	190.31
450	l_1	>300.00
	l_2	>300.00
	l_3	>300.00

$\|S\|$	l_x	time (s)
500	l_1	>300.00
	l_2	>300.00
	l_3	>300.00

Table 6: Timings of the scenario selection method for three different "branching models". The computation time in the last column is an average over ten runs of the algorithm.

closest remaining scenario at time step t, effectively merging the deleted scenario into the closest remaining scenario at this point.

We can use the level-wise target θ_t to control the complexity of the resulting tree. Note that we only control the number of nodes (or scenarios) present at every level of the tree, and thereby define the *average* branching factor at every level. We do not control the maximum or minimum branching factor. We only control the number of scenarios that is removed, into which other branches the deleted parts are then merged is determined based on the distance measure, as explained above.

We tried three ways of determining our targets θ_t:

l_1 has a linear increase in the number of scenarios. So, a tree of depth 42 (the depth of the scenario tree for problems where $\tau = 24$) and 300 scenarios will have a linear increase of 300/42 at every level. Each time, we round the resulting number to the nearest integer: the first level will have 7 scenarios, the second will have 14, and so on until the last level, which contains 300 leafs;

l_2 has a constant branching factor. For $\tau = 24$ and $|S| = 300$, this branching factor equals $\sqrt[42]{300}$. The first level will have round($\sqrt[42]{300}$) scenarios, the second round(($\sqrt[42]{300}$)^2), and so on until the last level, which again contains 300 leafs;

l_3 is a variation on l_1, as the target number of scenarios also increases linearly. However, it increases linearly over only the first 24 levels, and after that remains constant. So, again for $\tau = 24$ and $|S| = 300$, we have an increase of 12.5 per level. We round the resulting number to the nearest integer at every level: the first level has 13 scenarios, the second 25, the third 38 and so on until the 24th, which has 300 scenarios. All levels after that will also have 300 nodes.

First, we ran the scenario reduction algorithm with these three ways of determining the target number θ_t for a few different values of $|S|$. The results of this are presented in Table 6. There is virtually no difference in the time it takes to create the scenario tree using l_1, l_2 or l_3. However, the time needed increases rapidly when increasing the initial number of scenarios $|S|$. We allow a total of 10 minutes for the entire planning process. This means that the generation of the scenario tree takes time away from the

solver. Based on this, we can discard any combination of $|S|$ and l_x that takes more than five minutes immediately.

This leaves the values of $|S|$ between 300 and 400. Initial tests with these values showed, however, that using 350 or 400 initial scenarios led to problems for which a solution that was reasonably close to the optimum could often not be found within the available time. Using 300 scenarios, a solution that was proven to be reasonably close to the optimum *could* be found in most cases. Therefore, we chose to construct the scenario tree for the multi-stage problems with 300 initials scenarios.

We have evaluated the performance of each of these "branching models". The results are presented in Section 13.5.

8. Grid components

In this section, we will describe the modeling of the various components in the micro-grids under study. This includes controllable and uncontrollable generation, uncontrollable load and storage technologies. Controllable loads are not included. Everything is modeled in terms of energy (kWh) produced or consumed.

8.1. Generators

The generators are represented by three variables every time step: y_{it} describes whether generator i is on or not, x_{it} is 1 if it is started up at time step t and g_{it} represents the power output of the generator. The variables x_{it} and y_{it} depend on each other as follows:

$$y_{it} - y_{it-1} \leq x_{it}$$

The three generators in the Island scenario run on diesel. Two of them are Olympian GEP55-1 diesel generators. They have a maximum production of 40 kW [13]. The other one is a Kohler SDMO K44C3, with a maximum production of 32 kW [65]. In the Residential scenario, we add a CAT C4.4-110, with a maximum production of 80 kW [4]. For all three generators we use 20% of their maximum production as an estimate of the minimum possible production (i.e. 8 kW, 6.4 kW and 16 kW). This is modeled in the following constraint:

$$y_{it} P_i^{min} \leq g_{it} \leq y_{it} P_i^{max}$$

where y_{it} is the on/off-state of generator i at time step t, g_{it} is its power output and P_i^{min} and P_i^{max} are its production limits.

As both generators can be brought on-line fairly quickly, no ramp-up rate constraints are included in the model. Additionally, fuel consumption and power output seem to have an almost linear relationship, as shown in Fig. 8. This is based on information about the fuel consumption of all three generators running at 100%, 75% and 50% of their maximum production [4, 13, 65]. All types of diesel engines use approximately 0.3 liters of diesel per kWh produced. The diesel price is estimated at around €1 per liter, which means that the production of one kWh by a diesel generator costs €0.30.

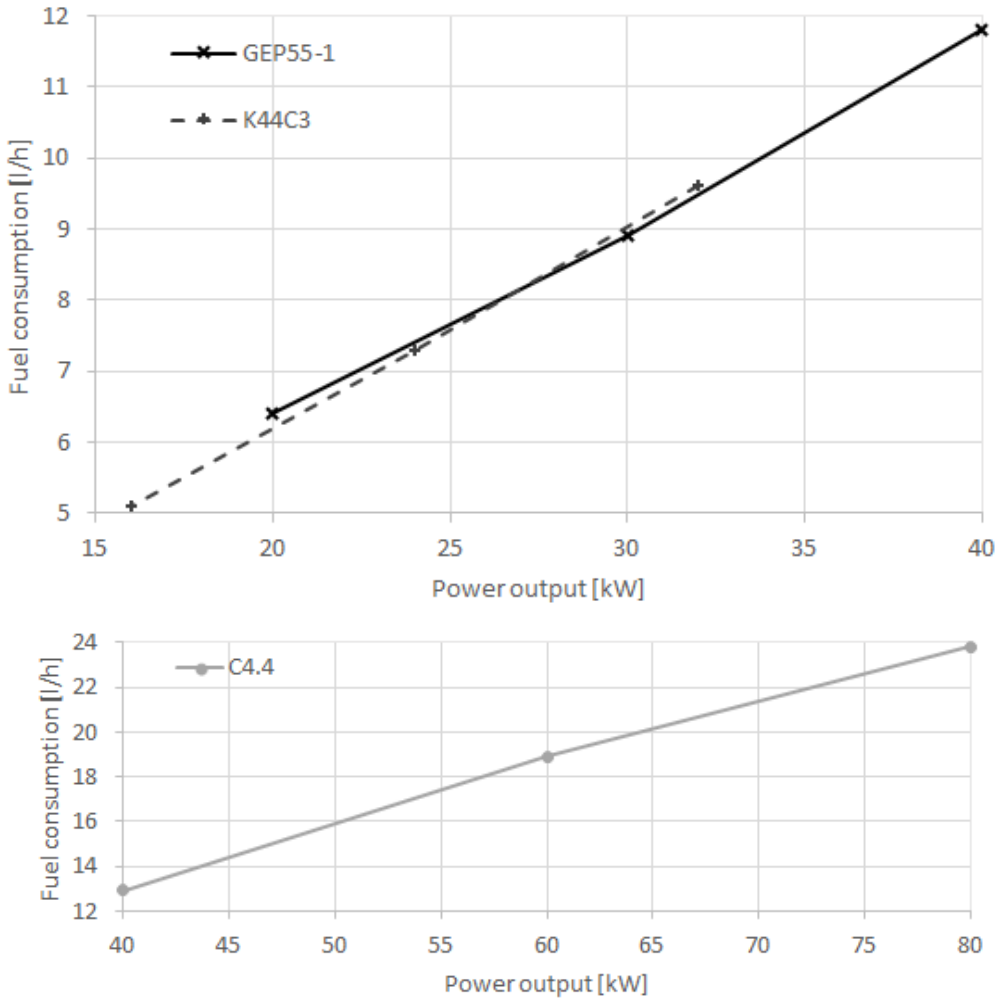

Figure 8: Fuel consumption / power output of the diesel generators

We validated these values by looking at the measurements over the first six months of 2017. Here we found that all generators reached somewhat higher maximum outputs (42 kW, 43 kW and 33 kW, respectively). The production rarely drops below the minimum that we determined, but all generators seem to have a real minimum at around 4 kW. Although the generators can produce more than their theoretical maximum, their efficiency presumably goes down. We chose to only model the part of the range of the generator where we know the cost-curve to be approximately linear (between 20% and 100% of its capacity).

Every generator has three associated cost coefficients. One (c_i^g) expresses the (fuel) cost per generated kWh of energy. Another (f_i) models basic (fixed) operational cost, including device wear, for every time unit that the generator is on-line. The last one (s_i) expresses the additional cost associated with bringing a generator on-line.

We have seen that the costs of the production of a single kWh of energy are estimated to be €0.30, so $c_i^g = 0.3$ for all i.

It is hard to determine and validate the operational and start-up costs for the diesel engines. Most comparable studies do not explicitly list the values used for these constants, and it is hard to determine them empirically. In a previous model of the same generators, start-up costs were determined to be €0.10 and operational costs €4 per hour. However, one would expect the start-up cost to be larger than the operational cost - as the reverse will cause the state of the generator to switch quite often, which is the type of behavior we would like to avoid.

In [14], microturbines and fuel cells that are similar in terms of their maximum power output are modeled. Their start-up costs are determined to be between $0.14 and $0.24, and no additional operational costs are used. As the initially assumed cost constants caused the generators to behave in a way that was determined to be undesirable by field experts, we decided to assume different values (start-up cost of €0.50 and operational costs of €0.40 per hour). However, these values cannot be validated, although they do result in a more intuitive planning.

8.2. Batteries

A relatively simple model has been assumed for the batteries. Self-discharge is not considered in this model. As the batteries are used intensively in the Island scenario, we expect this to be of little consequence.

We do take the inverter efficiency of the battery into account. To some extent, losses due to self-discharge etc. are covered by this as well. The efficiency applies both when charging and when discharging the battery. The efficiency of the battery (η_i^b) is multiplied with the energy to model the losses. The state of charge of the battery is therefore related to the energy charged to/discharged from the battery as follows:

$$ch_{it-1} - ch_{it} = \frac{b_{it}^{out}}{\eta_i^b} - b_{it}^{in}\eta_i^b$$

where b_{it}^{out} is the output of battery i at time step t, b_{it}^{in} its input and ch_{it} its state of charge (all in kWh).

The model does include costs per kWh discharged to model battery degradation, the wear of the device by (dis)charging it. This is based on the decrease in value of the battery before it is installed, and after it has out-lived its usefulness.

In the Island scenario, we have two batteries: one lithium-ion and one lead-acid battery. The lithium-ion battery has a maximum charge level of 20 kWh. We assume that it will degrade to 75% of this (15 kWh) after 8000 cycles. This will decrease its value by 25%. It was originally purchased for €350. This means it will decrease in value by €87,50 over 8000 cycles of 20 kWh. This comes down to €0,00056875 per kWh discharged. The lead-acid battery has no remaining value, which would lead to a cost of €0 per kWh discharged. But, as it is an old, unreliable battery, the use of the lithium-ion battery is preferred. It is currently rarely used (only in cases of emergency). We therefore decided to exclude it from the model altogether.

We validated the capacity of the batteries by looking at the measurements over the first six months of 2017. The measurements of the output of the batteries is only reliable when the generators are off, so we determined the amount of energy drawn from the lithium-ion battery over a period of time during the night, when no generators are active. Between 11-06-2017 03:48 and 11-06-2017 07:53, the state of charge of the lithium-ion battery goes down from 100% to 28%. During this time, 16.9 kWh is extracted from the battery. This would indicate that the storage capacity of the battery is actually higher than specified, as this would mean than approximately 23.5 kWh can be drawn from the battery in total.

Similarly, the lead-acid battery goes from 100% to 62% between 05:33 and 08:20 on 24-06-2017. During this time, approximately 8.5 kWh is discharged from the battery. This means that the capacity of the lead-acid battery is more than twice the specified amount, at 22.5 kWh.

To make the model more robust, the batteries can be operated using minimum reserve thresholds. The thresholds have the following purpose:

- Minimum reserve threshold (R_i^{min}): The minimum reserve threshold. If the state of charge starts below this threshold (before the first time step), it has to be charged to at least this threshold in the first time-step. The state of charge of the battery can never be below this threshold after the first time step;
- Maximum reserve threshold (R_i^{max}): If the state of charge is below this threshold, the battery cannot be discharged in the planning, but it is not required to be charged either. The battery can only be used as a producer of energy if its state of charge stays above this threshold (in the planning).

In essence, the minimum threshold guarantees a minimal reserve that will always be available to compensate any mismatch in the realization. By requiring a higher state of charge for discharging the battery, we aim to avoid continuous fluctuations around the minimum threshold. Without this threshold, the planning keeps the state of charge almost continuously at its minimum threshold, or very near to it, causing it to dive below the threshold with every mismatch that occurs. This would force the starting of a generator in the next time step to get the battery state of charge back up. This is not a desirable situation, and we hope to avoid it by limiting the use of the battery to higher charge levels. This translates to the following bounds and constraints:

$$ch_{it} \geq R_i^{min}$$

This bound makes sure that the state of charge of the battery is above the minimum threshold at the end of every time step. The following two constraints ensure that the battery is only discharged when the state of charge is above the maximum reserve threshold:

$$R_i^{max} r_{it} \leq ch_{it}$$

$$b_{it}^{out} \leq r_{it} v_i^d$$

Where r_{it} is an auxiliary variable to force the described behavior.

8.3. Other

In practice, it is possible to curtail the output of a PV-park, but this is not modeled. The same applies to load shedding, which is even less desirable. This allows for the treatment of all uncontrollable components as a single "reduced demand".

The total demand D_t equals (the forecast of) the uncontrollable load minus the (expected) uncontrollable generation (e.g. PV-output). This means that the demand can assume both positive and negative values. It is therefore modeled using two constants: D_t^+ and D_t^- for the positive and negative part of the demand, respectively. This results in the following relationship:

$$D_t = D_t^+ - D_t^-$$

Determining the value of these constants is a matter of preprocessing.

The main aim of the model is to balance supply and demand. This is expressed in the following constraint:

$$\frac{D_t^+}{\eta^g} - D_t^- \eta^g = \sum_{i \in G} g_{it} + \sum_{j \in B} \left(b_{jt}^{out} - b_{jt}^{in} \right)$$

where $0 \leq \eta^g \leq 1$ is the grid-efficiency, modeling grid losses.

9. Models

In this section we will propose a number of optimization models, in order of increasing complexity. First we will present a simple deterministic model in Section 9.1, which we will extend with safety constraints in Section 9.2. Then, we will describe our first stochastic (two-stage) model in Section 9.3, followed by a more complex multi-stage model in Section 9.4. Finally, we will discuss a model improvement that can be applied to all of these models in Section 9.5.

The models use the device models as described in Section 8. They solve the energy balancing problem by balancing the sum of the energy produced and consumed in every time step. This is equivalent to balancing the average power output (input) of the devices in the grid for every time step: average power output in kilowatts over one time step is convertible to energy in kWh by dividing by the number of time steps in an hour. Input that has a different unit (other than kWh, like watts or kilowatts) is converted before being plugged into the model.

The models will be formulated as MILPs. This technique is most often used in practice (c.f. Table 3 in Section 3), and seems suitable for the problem at hand. The characteristics of the appliances in the microgrids can be approximated reasonably well with

(piecewise) linear functions. When these characteristics are modeled accurately, the resulting model will be able to arrive at solutions that are very close to the optimum. Finally, as we consider (monopolized) microgrids, we expect the problem size to remain manageable. Most (commercial) solvers will be able to deal with MILP problems of the size that we expect to encounter reasonably fast. These considerations justify the selection of a mathematical programming approach to solve the energy management problem for these microgrids.

9.1. Naive model

This model takes a similar approach to [97], re-solving a simple MILP problem over a number of time steps every time step. As we will see in the section on implementation, we have chosen to solve the problem for the next 24 hours every 15 minutes. The length of one time step is 15 minutes, our planning horizon is 24 hours (or 96 time steps) and we solve the problem every 15 minutes. Only the decisions for the first time step are then used to actually control the micro grid. In short, we use a *rolling horizon* strategy.

This initial model does not take the uncertainty in the forecast into account: it assumes that the forecasted values will be realized exactly as predicted.

The objective is to minimize operational costs. This includes fuel costs but also costs associated with device wear. The cost function sums over the cost of operation of all devices. The batteries only have a discharge cost, as we have seen in Section 8.

The total contribution of the batteries to the cost function is:

$$\sum_{t \in T} \sum_{j \in B} c_j^b b_{jt}^{out}$$

Using the cost constants discussed in Section 8, the total contribution of the diesel generators to the cost function becomes:

$$\sum_{t \in T} \sum_{i \in G} (f_i y_{it} + c_i^g g_{it} + s_i x_{it})$$

Let G be the set of all traditional generators, and let B be the set of all batteries. Finally, let T be the set of time steps within the planning horizon. The unit of all constants and variables is converted to kWh. The full model then becomes:

$$\text{minimize} \sum_{t \in T} \left(\sum_{i \in G} (f_i y_{it} + c_i^g g_{it} + s_i x_{it}) + \sum_{j \in B} \left(c_j^b b_{jt}^{out} \right) \right) \text{ s.t.}$$

$$\frac{D_t^+}{\eta^g} - D_t^- \eta^g = \sum_{i \in G} g_{it} + \sum_{j \in B} \left(b_{jt}^{out} - b_{jt}^{in} \right) \qquad \forall t \in T \qquad (12)$$

$$y_{it} P_i^{min} \leq g_{it} \leq y_{it} P_i^{max} \qquad \forall t \in T, i \in G \qquad (13)$$

$$y_{it} - y_{it-1} \leq x_{it} \qquad \forall t \in T, i \in G \qquad (14)$$

$$ch_{it-1} - ch_{it} = \frac{b_{it}^{out}}{\eta_i^b} - b_{it}^{in}\eta_i^b \qquad \forall t \in T, i \in B \qquad (15)$$

$$y_{it} \in \{0,1\} \qquad \forall t \in T, i \in G$$

$$0 \leq x_{it} \leq 1 \qquad \forall t \in T, i \in G$$

$$0 \leq ch_{it} \leq E_i \qquad \forall t \in T, i \in B$$

$$0 \leq b_{it}^{in} \leq v_i^c \qquad \forall t \in T, i \in B$$

$$0 \leq b_{it}^{out} \leq v_i^d \qquad \forall t \in T, i \in B$$

Refer to Section 5 for a list of variables and constants used in this model.

Constraint (14) does not work for $t = 0$. To deal with this special case, we define that x_{it} expresses the initial state of generator i if $t = -1$. Similarly, in constraint (15), ch_{it} expresses the initial charge level of battery i for $t = -1$.

9.2. Adding safety constraints

To make sure the planning always has a small buffer to (cheaply) recover from a mismatch, the model is extended by adding some safety constraints that ensure a reserve is present in the batteries (as explained in Section 8). Note that, when for all batteries $R_i^{min} = R_i^{max} = 0$, the model is equivalent to the naive model.

This addition is intended to make the planning a bit more robust against fluctuations of the realized demand. All other parts of the model are kept exactly the same.

Again, let G be the set of all traditional generators, and let B be the set of all batteries. The unit of all constants and variables is converted to kWh. The full model then becomes:

$$\text{minimize} \sum_{t \in T} \left(\sum_{i \in G} (f_i y_{it} + c_i^g g_{it} + s_i x_{it}) + \sum_{j \in B} \left(c_j^b b_{jt}^{out} \right) \right) \text{ s.t.}$$

$$\frac{D_t^+}{\eta^g} - D_t^- \eta^g = \sum_{i \in G} g_{it} + \sum_{j \in B} \left(b_{jt}^{out} - b_{jt}^{in} \right) \qquad \forall t \in T \qquad (16)$$

$$y_{it} P_i^{min} \leq g_{it} \leq y_{it} P_i^{max} \qquad \forall t \in T, i \in G \qquad (17)$$

$$y_{it} - y_{it-1} \leq x_{it} \qquad \forall t \in T, i \in G \qquad (18)$$

$$ch_{it-1} - ch_{it} = \frac{b_{it}^{out}}{\eta_i^b} - b_{it}^{in}\eta_i^b \qquad \forall t \in T, i \in B \qquad (19)$$

$$R_i^{max} r_{it} \leq ch_{it} \qquad \forall t \in T, i \in B \qquad (20)$$

$$b_{it}^{out} \leq r_{it} v_i^d \qquad \forall t \in T, i \in B \qquad (21)$$

$$y_{it} \in \{0,1\} \qquad \forall t \in T, i \in G$$

$$r_{it} \in \{0,1\} \qquad \forall t \in T, i \in B$$

$$0 \leq x_{it} \leq 1 \qquad \forall t \in T, i \in G$$

$$
\begin{aligned}
R_i^{min} \leq ch_{it} \leq E_i & \qquad \forall t \in T, i \in B \\
0 \leq b_{it}^{in} \leq v_i^c & \qquad \forall t \in T, i \in B \\
0 \leq b_{it}^{out} \leq v_i^d & \qquad \forall t \in T, i \in B
\end{aligned}
$$

where the definitions of all decision variables and constants are the same as before, with the addition of decision variables r_{it} and constants R_i^{max} and R_i^{min}. See Section 5 for a brief description and overview of variables and constants used in the model.

The additional provisions for x_{it} and ch_{it} at $t = -1$ in (18) and (19) are the same as before.

9.3. Two-stage model

We want to take the uncertainty in the forecast explicitly into account in the planning. In order to do so, we need to include a way to represent that in the model.

We will do this using scenarios in a two-stage model. As stages we consider two timescales: planning and realization, where the first stage (planning) variables are those that we can control and decide on, and the second stage variables are those that are (reactively) adjusted to the scenario that is materializing during operation. In our example below, the battery decisions would form the second stage and everything else (generator state and output) would be the first stage decisions. In this example, the two-stage structure of the problem is relatively simple, as the output of the batteries is the only thing that is decided in the second stage. One can, however, imagine extending the microgrid with other devices or a grid-connection, and changing which decisions are contained in the second stage.

The constraints of the resulting model are very similar to those of the naive model but the constraints concerning energy balance and battery output are repeated for every scenario. The objective function is slightly different as well: the contribution of the generators remains unchanged, but the contribution of the battery is now the expected cost over all scenarios. In other words: the contribution of the battery is a weighted sum of the battery costs over all scenarios, where the cost of every scenario is weighted by its relative probability.

It may happen that two scenarios are too different (either in a single time step or over multiple time steps), meaning that the difference between them cannot be compensated by the batteries. This means that there is no feasible solution to the entire problem. Even when this is the case, we still want to generate a planning that is feasible for *most* scenarios. To make this possible, we will add two variables for every time step and scenario that represent the unsatisfied demand (or surplus). A high penalty ensures that it is still optimal to satisfy the demand in all scenarios and match the demand as closely as possible in those where it is not feasible to satisfy it completely, without causing infeasibilities.

We obtain the scenarios using the sampling method discussed in Section 7. However, we do not execute Alg. 2, and we keep the original scenario fan.

Let S be the full (index) set of scenarios. Also, let p_s be the (relative) probability

assigned to scenario $s \in S$, such that $\sum_{1}^{|S|} p_s = 1$. Then, we get the following deterministic equivalent of the stochastic problem:

$$\text{minimize } \sum_{t\in T}\sum_{i\in G}\left(f_i y_{it} + c_i^g g_{it} + s_i x_{it}\right) + \sum_{s\in S} p_s \left(\sum_{t\in T}\sum_{j\in B} c_j^b b_{jts}^{out} + c^w\left(w_{ts}^+ + w_{ts}^-\right)\right) \text{ s.t.}$$

$$\frac{D_{ts}^+}{\eta^g} - D_{ts}^-\eta^g = \sum_{i\in G} g_{it} + \sum_{j\in B}\left(b_{jts}^{out} - b_{jts}^{in}\right) + w_{ts}^+ - w_{ts}^- \qquad \forall t \in T, s \in S \quad (22)$$

$$y_{it}P_i^{min} \le g_{it} \le y_{it}P_i^{max} \qquad \forall t \in T, i \in G \quad (23)$$

$$ch_{i(t-1)s} - ch_{its} = \frac{b_{its}^{out}}{\eta_i^b} - b_{its}^{in}\eta_i^b \qquad \forall t \in T, s \in S, i \in B \quad (24)$$

$$y_{it} - y_{it-1} \le x_{it} \qquad \forall t \in T, i \in G \quad (25)$$

$$0 \le ch_{its} \le E_i \qquad \forall t \in T, s \in S, i \in B$$

$$0 \le b_{its}^{in} \le v_i^c \qquad \forall t \in T, s \in S, i \in B$$

$$0 \le b_{its}^{out} \le v_i^d \qquad \forall t \in T, s \in S, i \in B$$

$$y_{it} \in \{0, 1\} \qquad \forall t \in T, i \in G$$

$$0 \le x_{it} \le 1 \qquad \forall t \in T, i \in G$$

See Section 5 for a brief description and overview of variables and constants used in the model.

The additional provisions for x_{it} and ch_{its} at $t = -1$ in (24) and (25) are the same as before.

9.4. Multi-stage model

From the discussion above, it has become clear that the problem as we have defined it displays a stage-wise structure in two dimensions: on the one hand we have the decision stages (as made explicit in the model from Section 9.3) and on the other we have the time steps, dividing the problem in $|T|$ intervals, each with its own associated decisions.

The temporal structure of the problem is best modeled by a multi-stage model. Each time step becomes a stage in the problem.

The general form of a multistage stochastic program is stated in [101] as follows:

$$\text{minimize } \{F(x) := \mathcal{R}[Q(x, \xi)] : x \in \mathbb{X} \text{ and } x \triangleleft \mathfrak{F}\}$$

Where x are the decision variables, ξ are the random values, $\mathbb{X}$ is the set of constraints, $Q(x, \xi)$ is the objective function, $\mathcal{R}$ is a risk functional of the cost function and $x \triangleleft \mathfrak{F}$ represents the non-anticipativity constraint. For some background on multi-stage stochastic models, refer to [101].

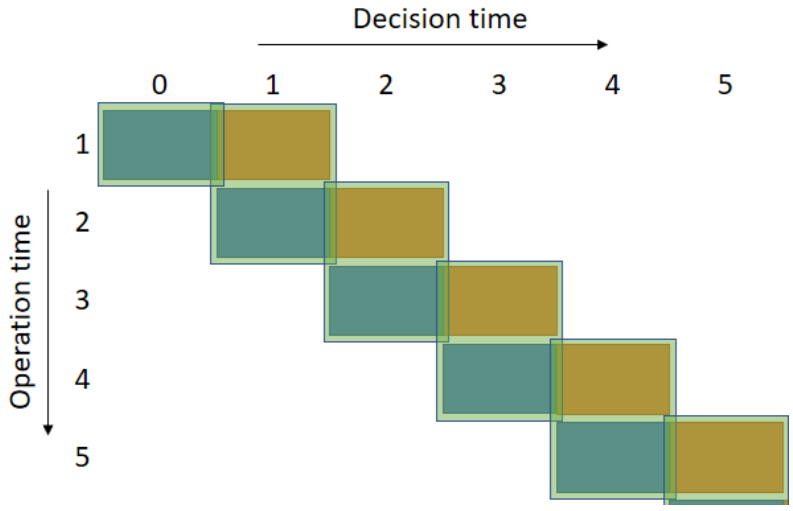

Figure 9: The model captures the stage-wise structure of the model both in the time that decisions are meant to be implemented (*operation time*) and the time that the decision is fixed (*decision time*). This image gives an impression of the model structure for the first five time steps of a problem. The blue boxes represent the "first-stage" decisions and the orange boxes represent the "second-stage" decisions. Adjacent boxes are meant to be implemented in the same time step. The transparent green boxes represent the actual stages of the problem, grouping the decisions together that are made in the same time step, with the same available information.

The non-anticipativity constraint requires that decisions for stage t can only be based on information that is available at t. In our model, because of the use of a scenario tree, the non-anticipativity is implicit.

By defining our model this way, however, we lose the capability of differentiating between decision stages. Each time step is represented by a stage in the model, meaning that all decisions for that time step are made on the basis of the exact same information. In the model from section 9.3, the first-stage decisions were made *before* the actual values of the demand and PV-output would become known, while the second-stage decisions were used to adjust the solution *after* these were known. In the simplest multi-stage model, however, all these decisions are made at the same time.

To resolve this problem, we need to realize that the information used for the "first-stage" decisions of time step $t + 1$ is exactly the same as the information available in the "second-stage" of time step t. By shifting the decisions that were previously in the first stage of period t to period $t - 1$, we are again able to differentiate between decision stages for a single time step. The basic idea is illustrated in Fig. 9.

Our objective function becomes a recursive expression, recursively adding the expected cost of all decisions further ahead in time. Every node in the scenario tree has its own set of constraints concerning balance and batteries and a set of constraints modeling the

behavior of the generators in its children.

Assuming a scenario tree to be given (refer to Section 7 for the algorithms used for its construction), we can once again express this as a MILP.

Let S_t be the index set of all scenarios present at time step t. In other words: S_t is a list of the scenario indices of all nodes at depth t in the scenario tree. Let $par(s), s \in S_t$ be (the scenario index of) the direct parent of s in the tree, so $par(s) \in S_{t-1}$. Additionally, let $ch(s)$ be the set of (scenario indices of) children of s, i.e. all nodes s' for which $par(s') = s$. Remember that p_s is the *relative* probability assigned to scenario s at every step (i.e. the conditional probability of s occurring at t given that its parent scenario did occur at $t-1$, etc.). This means the following. Let $s' = par(s)$. Now, p_s is the probability of scenario s occurring relative to its siblings. In other words, $\sum_{s'' \in ch(s')} p_{s''} = 1$. Note that we have that $S_0 \subseteq S_1 \subseteq ... \subseteq S_{|T|-1} \subseteq S$. The combination of the s and t indices uniquely identify a node in the tree.

$$\text{minimize } \sum_{i \in G} \left(f_i y_{i00} + c_i^g g_{i00} + s_i x_{i00}\right) + \sum_{s \in S_0} p_s \left(\sum_{j \in B} c_j^b b_{j0s}^{out} + \right.$$

$$c^w \left(w_{0s}^+ + w_{0s}^-\right) + \sum_{i \in G} \left(f_i y_{i1s} + c_i^g g_{i1s} + s_i x_{i1s}\right) + \sum_{s' \in ch(s)} p_{s'} \left(... + \right.$$

$$\left. \left. \left. \sum_{s''' \in ch(s'')} p_{s'''} \left(\sum_{j \in B} c_j^b b_{j(|T|-1)s'''}^{out} + c^w \left(w_{(|T|-1)s'''}^+ + w_{(|T|-1)s'''}^-\right) \right) \right) \right) \right)$$

$$\begin{aligned}
\text{s.t.} \quad \frac{D_{ts}^+}{\eta^g} - D_{ts}^- \eta^g &= \sum_{i \in G} g_{it(par(s))} + \sum_{j \in B} \left(b_{jts}^{out} - b_{jts}^{in}\right) && \forall t \in T, s \in S_t \\
&\quad + w_{ts}^+ - w_{ts}^- \\
y_{its} P_i^{min} &\leq g_{its} \leq y_{its} P_i^{max} && \forall t \in T, i \in G, s \in S_{t-1} \\
ch_{i(t-1)(par(s))} - ch_{its} &= \frac{b_{its}^{out}}{\eta_i^b} - b_{its}^{in} \eta_i^b && \forall t \in T, s \in S_t, i \in B \\
y_{its} - y_{i(t-1)(par(s))} &\leq x_{its} && \forall t \in T, i \in G, s \in S_{t-1} \\
0 \leq ch_{its} &\leq E_i && \forall t \in T, s \in S_t, i \in B \\
0 \leq b_{its}^{in} &\leq v_i^c && \forall t \in T, s \in S_t, i \in B \\
0 \leq b_{its}^{out} &\leq v_i^d && \forall t \in T, s \in S_t, i \in B \\
y_{its} &\in \{0, 1\} && \forall t \in T, i \in G, s \in S_{t-1} \\
0 \leq x_{its} &\leq 1 && \forall t \in T, i \in G, s \in S_{t-1}
\end{aligned}$$

See Section 5 for a brief description and overview of variables and constants used in the model.

The additional provisions for x_{it} and ch_{its} at $t = -1$ are the same as before.

Figure 10: Graphical representation of the change in time resolution. The time steps up to $t = \tau$ will be 15 minutes long, while all time steps after that will be considered in groups of four, in other words: the time resolution changes to 1 hour per time step.

9.5. Model-independent improvements

As we have seen in Section 6.3, the accuracy of the forecast is not as crucial for the later time steps in the planning. This inspired a modification of the models that is graphically represented in Fig. 10. The idea is to change the time resolution from 15 minutes to one hour after a number of time steps (τ). This leads to a significant reduction in the size of the problem: fewer scenarios will be needed to approximately express the same range of possibilities as the depth of the scenario tree significantly decreases, and a problem with the same amount of scenarios will become easier from a computational point of view, allowing for the consideration of larger problems in the same computation time.

This time shift is fairly easily modeled. It requires the aggregation of the forecasts after time step τ. This can be achieved by taking the average over the predicted values for each of the quarters of the aggregated hour.

The resulting pooled standard deviation can be computed by taking the square root of the average variance (i.e. the squared standard deviation) and the variance of the mean over all periods:

$$\sigma_{agg} = \sqrt{\sum_{i=t}^{t+3} \left(\frac{\sigma_i^2}{4} + \frac{\sum_{j=i+1}^{t+3} (\mu_i - \mu_j)^2}{4^2} \right)}$$

where t is the index of the first time step being included in the aggregation and μ_t is the predicted value (i.e. the mean of the distribution) at time t. We have chosen this approach to determine the aggregated standard deviation, based on how it is computed in the first place. As explained in Section 6, the standard deviation is determined by looking at the distribution of the error of the forecast in the past. The standard deviations are determined separately for forecasts of different time steps created at a different time of day. Effectively, the four standard deviations that we want to aggregate express the spread of the error for non-overlapping groups of predictions in the past. To obtain a standard deviation that expresses the spread of the error for the combined group of predictions, we use the formula above to determine the aggregate standard deviation of the four time steps together.

Aside from this, a number of constants (f_i, P_i^{min}, P_i^{max}, v_i^c, v_i^d) need to be multiplied by four relative to the quarter-long time periods. Otherwise, the model remains the same. This can be easily implemented by including a parameter τ in the generation of the model files and modifying the constraints over variables where $t \geq \tau$ to work with the multiplied constants. The value of τ can be anything between 0 and $|T|$, the only limitation on τ is that it has to be a multiple of four.

10. Simulation model

To evaluate the effects of the planning constructed in the previous step, a simple simulator was implemented. The simulator determines the effect of the planning by following it as closely as possible and solving any mismatch of prediction and realization by applying simple rules-of-thumb. The cost of operation is then computed, which results in an estimation of the cost of operation when following the generated planning. More details will be provided below.

The simulation takes over the role of the microgrid controller, which initiates the planning process and implements the decisions during operation. A rolling horizon strategy is followed: every simulated time step, the loop as depicted in Fig. 4 is executed. New forecasts are created, from which scenarios are (possibly) constructed. These are in turn fed to the optimization model, from which the relevant decisions are retrieved. The simulation then determines the effect of the decisions for one time step, considering the realized values of the uncertain factors in the model, and updates the state of the system accordingly (as described below). This entire process is then repeated for the next time step.

We have, however, also implemented variants of this process where a new planning is only generated after the old one is not valid anymore.

In the *-variants, this means that re-planning occurs when the predicted value had a distance larger than δ to the realization in the previous time step (for the naive and safety models) or there was no scenario that had a distance smaller than δ (for the scenario-based models). In other words, if a sufficiently similar scenario was included, we can re-use the current planning. Otherwise, we have to re-plan. The values of δ for the different scenarios will be listed in Section 12.

The **-variants only re-plan after the planning had to be adjusted, meaning that the batteries were unable to compensate the mismatch on their own in that time step.

The simulator computes the *realized cost* of the planning. For this, it uses the same cost coefficients as the optimization model (f_{it}, c_{it}^g, s_{it} and c_{it}^b above). The only difference is that the actual (realized) values of the PV production and the demand have become known. These values are almost never exactly the same as the forecast, and this difference is accounted for by simple rules that can be summarized in the following schematic:

A shortage occurs when the reduced demand (demand minus PV output) is higher than the forecast, and a surplus occurs when it is lower. First, the batteries are used to resolve the mismatch (within their operational limits). If the batteries cannot solve

Shortage	Surplus
Compensate with battery	
Adjust generator output	
Bring additional generator on-line	PV-curtailment

the mismatch on their own, we turn to the generators. If one or more generators are on-line, their output will be adjusted (again, within their operational limits) to compensate the remainder. Finally, if the batteries and on-line generators together are unable to resolve the situation, an additional generator is brought on-line (in case of shortage) or PV-output is curtailed (in case of surplus).

The first two steps of this process (adjusting the batteries and adjusting the generator output) can be inverted. In Section 13.1 we will evaluate both strategies.

The operational limits that are taken into account are the same as in the optimization models. For batteries, this includes the maximum charge level and charging/discharging speeds. For generators, this includes maximum and minimum production.

11. Implementation details

11.1. Forecasting

The forecasting models are implemented as Python scripts. The models produce a forecast for each quarter of an hour in the next 24 hours every 15 minutes, which comes down to 96 forecasted values every 15 minutes, or 9216 values a day.

For PV forecasting, the output of all PV-installations is aggregated and predicted as one. This results in a single prediction in the Island scenario for example, even though it has three PV-fields.

Similarly, the load forecasting predicts just one aggregated value for the expected load.

11.2. Optimization

As stated above, the model uses time steps of 15 minutes, which means that some of the constants from the model presented above are divided by 4 (all constants expressed in kW, plus the operating cost f_i).

The model is written in AMPL [40], saved in `.mod` and `.dat` files. The solvers used are GLPK (v4.60) [79] and CPLEX (v12.7.1) [54] (after using GLPK to convert the model to the LP-format). A cloud-based version of CPLEX (DOcplexcloud [55]) was also used, but as this does not add anything to the analysis, we will exclude it from our results. In addition to the timelimit of 10 minutes (discussed below), we use the standard optimality gap of 0.01% as a stopping condition when using the models from Section 9.1 or 9.2. When using the models from Section 9.3 or 9.4 we use an optimality gap of 1%

as a stopping condition. Additionally, we use the CPLEX parameter `objdifference` to instruct the solver to only look for (new) solutions with an objective value that differs by at least €0.01 from the current solution. This helps to speed up the proof of optimality. However, it may lead to skipping of an interval in which the true integer optimum may lie. This weakens the guarantee of optimality somewhat, but by keeping the allowed difference as low as €0.01 we do profit from a speed-up, without skipping large intervals.

The models are a direct translation of the ones presented above.

11.3. Simulation

The basic simulator is implemented as a Python script. The script gathers the input from the forecasts, the current system state and the actual realized values for every time step. It performs some preprocessing and provides the input for the optimization. From the solution that the solver comes up with, it extracts the control actions. These control actions are adjusted using the rules described in section 10 to match the realized reduced demand. The simulation updates the system state and the (cumulative) costs afterward and provides the input for the optimization problem for the next time step.

We chose to put a time limit of 10 minutes on the optimization process. Most models can be solved to optimality within this time. In some cases, however, especially for problems with long planning horizons, it may take multiple hours to find the optimal solution. This is, obviously, impossible to work with in practice when a new problem has to be solved every 15 minutes. Considering the other work that has to be done in this 15 minutes (retrieving forecasts, implementing decisions), it seems reasonable to set a maximum of 10 minutes for the optimization to use. If no feasible solution is found within this time, the problem remains unsolved and the decisions from a previous time step will be used.

Part III.
Results & Discussion

12. Test instances

The models and algorithms described in Part II were applied to three test cases. Each of these test cases consists of:

1. A set of system parameters defining the devices in the grid and their associated constants (discussed in detail in Section 8);
2. A forecast of the PV-output of the solar panels in the grid for the period of 1-7 June (Section 6);
3. A load forecast for the period of 1-7 June (also Section 6);
4. Realizations of the PV-output and load for the same period.

These four elements make up a test case, to which the process described in Part II can be applied. This is done by running a simulation of either one day (June 1st) or one week (1-7 June). June is a month in which both load and PV-output are comparatively high in the Island case, which makes for an interesting test case. June 1st is a fairly typical day in this month, while the week of 1-7 June features a mix of interesting days: typical days like June 1st but also a day that the site was open for longer than usual and a Monday, when it is usually closed, that it was open. This means that the load forecast will be much too low for these days, which will help in testing the robustness of the different approaches.

The simulation will be set up as described in Section 10: based on the provided forecasts and the system parameters a model will be created and solved every time step to obtain a planning for the next time step, which will then be corrected by the simulator based on the realizations using the rules-of-thumb described in Section 10.

In this section, we describe the three test cases that will be used in the simulations. In Section 13, we will discuss the simulation runs in more details and present the results.

12.1. Island

The Island case is based on a microgrid that is currently managed by Zown. The PV forecasts and load forecasts used in this case are the actual forecasts created for the period of 1-7 June 2017 for this grid. Similarly, the realized values used in the simulation to determine the actual demand and PV output are the actual realized values of this same period.

On Island, three generators are installed. One Kohler SDMO K44C3 and two Olympian GEP55-1 diesel generators, as described in Section 8. There are two batteries on Island, one lead-acid battery and one lithium-ion battery. Only the latter is included in the model, as was explained in Section 8, along with its specifications. A number of solar

panels is installed, that together have a maximum output of 12 kW. A visitor center and a restaurant on Island are responsible for the majority of the electrical load.

The values used for the battery thresholds were determined by experience. These values were found to work well for the Island case: they reduced the number of times that the batteries were unable to solve the mismatch, but were not too limiting for the planning. As a result, the overall costs were often brought down.

In summary:

- Devices:
  - 2x Olympian GEP55-1 diesel generator: $P^{min} = 8$ kW, $P^{max} = 40$ kW, $s = 0.50$, $f = 0.40$, $c^g = 0.30$;
  - 1x Kohler SDMO K44C3 diesel generator: $P^{min} = 6.4$ kW, $P^{max} = 32$ kW, $s = 0.50$, $f = 0.40$, $c^g = 0.30$;
  - 1x Lithium-ion battery: $E = 20$ kWh, $v^c = v^d = 12$ kW, $R^{max} = 3$ kWh, $R^{min} = 1$ kWh, $\eta^b = 0.93$, $c^b = 0.00057$.
- Other constants: $c^w = 2$, $\eta^g = 0.97$, $\delta = 0.25$;
- Average load: 9.6 kW; Maximum PV-output: 12 kW.

12.2. Island+

We tested a variation on the Island case, slightly altering the relative magnitude of system components to observe the effects on the planning. The PV-output and the battery size (both in terms of storage capacity and maximum power output) are currently relatively small compared to the generators and peak load on Island. Indeed, plans have been developed to increase both the number of PV panels on Island and the size of the batteries. This test case investigates the potential benefits of such a change, while also evaluating if the planning algorithms are still valuable in this scenario. The devices in this case are the same as those used in the Island case, with the exception of the battery, which differs in its storage capacity and maximum output, as described below.

The average quarterly load on Island is around 2-2.5 kWh (min. ∼0.75 kWh, max. ∼7 kWh). The total load to bridge the gap overnight is around 50 kWh in total. During this period, the Island is closed and the demand for electricity is at its lowest and most predictable. We want the battery to be able to satisfy this nightly demand, in principle. The total daily load is around 200 - 250 kWh on average. In addition to satisfying the nightly demand, we want the capacity of the battery to match the load of around 4-6 typical hours of use. The storage capacity of the battery in Island+ will therefore be 60 kWh, or three times that of the Island battery. Similarly, we want the battery to be able to provide the peak load (i.e. the highest load that occurs over fifteen minutes). The peak load is around 30 kW, so this will be the maximum output of our battery as well. This means that the battery can charge/discharge about 2.5 times as fast as in the Island case.

Because the load does not increase, we do not need to scale the reserve thresholds for the battery with the same factor. We double the minimum reserve threshold (to 2

kWh), and keep the maximum reserve threshold at 2 kWh above the minimum reserve threshold (at 4 kWh).

On Island, all solar panels together have a maximum output of 12 kW, which means that even when they provide their maximum output, they are hardly able to match the average load. Averaged over the whole day, they produce about 0.8 kWh every quarter (in June, the sunniest month of the year). To make things a little more interesting, we multiply their capacity by two, reducing the need for fossil back-up. This means that we scale the PV forecast (both the predicted value and the standard deviation) and the PV realization. The load forecast and realization remains the same as in the Island case.

In summary:

- Devices:
  - 2x Olympian GEP55-1 diesel generator: $P^{min} = 8$ kW, $P^{max} = 40$ kW, $s = 0.50$, $f = 0.40$, $c^g = 0.30$;
  - 1x Kohler SDMO K44C3 diesel generator: $P^{min} = 6.4$ kW, $P^{max} = 32$ kW, $s = 0.50$, $f = 0.40$, $c^g = 0.30$;
  - 1x Lithium-ion battery: $E = 60$ kWh, $v^c = v^d = 30$ kW, $R^{max} = 4$ kWh, $R^{min} = 2$ kWh, $\eta^b = 0.93$, $c^b = 0.00057$.
- Other constants: $c^w = 2$, $\eta^g = 0.97$, $\delta = 0.5$;
- Average load: 9.6 kW; Maximum PV-output: 24 kW

12.3. Residential

A final test case was constructed to investigate the performance of the developed approach in situation where all components are larger (more PV-panels, more loads, a larger battery and larger generators), and where the demand is of a different type: households. The load profile in particular is much more stable, and much easier to predict.

We model this case as an isolated village: a residential area of 200 households without a connection to the main utility grid. The demand needs to be satisfied with locally generated energy, either produced by the PV-panels or by one of the two generators connected to the grid. The storage of energy is facilitated by a large battery. We look at the same period as the Island and Island+ cases: 1-7 June 2017.

The load profiles and predictions were constructed in the following way. Using the (fractional) electricity profiles [128] provided for 2017 by NEDU [91] and the average yearly electricity consumption of Dutch households (2910 kWh) [93], we constructed the average profile of a household for the period of 1-7 June. Given this average profile, we sample values for each household individually, from a normal distribution with the average demand as its mean and 50% of this demand as its standard deviation, much like the sampling of a scenario in Alg. 1. We assume a correlation of 0.63 between adjacent time steps. To come up with a realization, we simply sum 200 of such samples, while we use the expected profile (i.e. 200 times the average profile of a household) as a prediction. As these profiles are independently sampled, the variance of the prediction is equal to the sum of the variances of the summed distributions, which means that the

standard deviation of the prediction equals $\sqrt{200}$ times the standard deviation of the distribution from which the household profiles were sampled.

For the PV profiles, we use a scaled version of a set of predictions and realizations. The original data set we used is an average constructed for a number of PV systems in the Enschede area in the Netherlands. For each of the systems, the realizations and forecasts were collected. The data set provides an average realization and prediction for a PV-system with a maximum output of 2.7 kW in this area. When scaling up, we simply multiply the realization and mean of the forecast by a scaling factor, which means that the standard deviation can also simply be multiplied by this factor.

The average expected demand is just over 14 kWh per quarter. We want the PV-panels to satisfy a significant portion of the daily demand. During the day, the load is around 60-80 kW. As we have no option to get rid of excess energy (only the battery, to a limited extent), but we want to use as much solar energy as possible, we will scale the PV-panels such that they satisfy the demand of an average day at maximum production. Multiplying the forecast for a system with a maximum output of 2.7 kW by 30 yields a forecast for a system with a maximum output of 81 kW, close to the demand of an average day.

The composition of the grids in terms of the devices they contain will be similar to Island, meaning that they may contain PV panels, loads, diesel generators and batteries. As argued before, the model can easily be extended to include other types of devices, but we will limit ourselves to these for the time being.

The battery used in this instance is based on the battery-cluster that was used in the "Smart Storage" project, running from 2010 till 2013 in Etten-Leur, The Netherlands and will remain operational until 2017 [22, 28, 29]. The batteries used here are lithium-ion batteries, similar to the one in the Island scenario. However, they are significantly larger. The total capacity is 230 kWh, the maximum power output is 400 kW, and the maximum charging power is 100 kW.

In addition we assume two diesel generators to be present, while no grid connection exists. One small generator (the Kohler SDMO K44C3 also present in the Island scenario) and one larger generator (a CAT C4.4-110, also discussed in Section 8).

We use comparable battery reserve thresholds of 10 kWh and 30 kWh. At the start of the simulation, the state of charge of the battery is at 40% (92 kWh) and all generators are turned off.

In summary:

- Devices:
  - 1x CAT C4.4-110: $P^{min} = 16$ kW, $P^{max} = 80$ kW, $s = 0.50$, $f = 0.40$, $c^g = 0.30$;
  - 1x Kohler SDMO K44C3 diesel generator: $P^{min} = 6.4$ kW, $P^{max} = 32$ kW, $s = 0.50$, $f = 0.40$, $c^g = 0.30$;
  - 1x Lithium-ion battery: $E = 230$ kWh, $v^c = 100$ kW, $v^d = 400$ kW, $R^{max} = 30$ kWh, $R^{min} = 10$ kWh, $\eta^b = 0.93$, $c^b = 0.00057$.
- Other constants: $c^w = 2$, $\eta^g = 0.97$, $\delta = 1$.
- Average load: 58 kW; Maximum PV-output: 81 kW

13. Results

All simulations described in this section have been run on a PC with an Intel Core i5-6600K CPU and 8 GB RAM.

In the tables below we report eight metrics on the different simulation runs. The column title associated with each metric is put in brackets:

- Average computation time, in seconds: the average time that the solver needed to generate the solution that was implemented in every time step of the simulation (time (s));
- Average computation time, in ticks: same as above, using the deterministic time-measurement that is a built-in feature of CPLEX (time (ticks));
- Average optimality gap: the average optimality gap, as a percentage of the value of the objective function, of the implemented solution (gap);
- Actual total cost: a summation of the cost of all individual time steps. These are the costs of operating the grid, following the generated planning as closely as possible, as determined by the simulator (cost, real);
- Expected total cost: a summation of the expected cost of all individual time steps. These are the costs of implementing the planning as-is before adjustments are determined in the simulation. In other words: this is what you expect to pay for implementing the generated planning (cost, exp.);
- Corrected total cost: a correction of the real cost, based on the change in the amount of energy stored in the batteries. This is computed by subtracting the minimal cost of the generation of this difference in energy, $c^g \cdot \Delta ch$, from the actual total cost (cost, cor.);
- Change in the amount of energy stored in the batteries. This represents the difference between the amount of energy (in kWh) stored in the system at the beginning of the simulation, and the amount stored at the end (Δch);
- Total number of times intervention was required: the number of time steps in which more than just the battery was used to adjust the planning based on the realizations. This means that the output of a generator was adjusted, the state of a generator was changed, energy generated by the PV-panels was wasted or the demand could not be satisfied in this time step (#adj.);
- Total number of times the planning failed: the number of time steps no feasible solution could be found, and no new planning was generated or no new planning was necessary (*- and **-variants, see Section 10) (#fail).

13.1. Simulation mode

The first set of experiments was done to determine which simulation mode works best. As described in Section 10, we first adjust the output of the batteries, and only if that is not sufficient, we adjust the generator output. However, this could just as well be the other way around. Note that we can only adjust the output of a generator if it is already on. In both situations, adjusting the output of the batteries takes precedence over starting an off-line generator.

		time			cost					
scenario	mode	(s)	(ticks)	gap	real	exp.	cor.	Δch	#adj.	#fail
Island	gen	0.57	578.65	0.008%	€58.25	€59.80	€54.18	+13.56	53	0
	bat	0.66	597.94	0.008%	€58.00	€61.84	€53.91	+13.62	16	0
Island+	gen	0.59	304.40	0.003%	€17.40	€19.29	€16.59	+2.71	8	0
	bat	0.48	278.04	0.003%	€17.24	€17.93	€16.25	+3.29	5	0
Residential	gen	0.07	65.80	0.006%	€230.01	€230.52	€240.79	-35.93	63	0
	bat	0.07	63.13	0.006%	€230.19	€230.19	€241.77	-38.60	0	0

Table 7: Results of simulation runs comparing adjustment strategies generator-first (gen) vs. battery-first (bat). For all runs: $|T| = 96$, $\tau = 96$, model: naive (see Section 9.1), solved with CPLEX (see Section 11).

		time			cost					
scenario	mode	(s)	(ticks)	gap	real	exp.	cor.	Δch	#adj.	#fail
Island	gen	5.02	5996.85	0.010%	€56.27	€58.80	€54.55	+5.75	46	0
	bat	5.55	6711.12	0.010%	€51.68	€56.05	€52.25	-1.91	13	0
Island+	gen	0.51	588.17	0.004%	€18.82	€21.82	€16.81	+6.71	8	0
	bat	0.46	548.57	0.004%	€17.85	€19.58	€15.70	+7.16	5	0
Residential	gen	101.12	133274.90	0.033%	€236.60	€236.15	€243.16	-21.85	67	0
	bat	101.10	128782.62	0.025%	€239.86	€239.86	€244.39	-15.09	0	0

Table 8: Results of simulation runs comparing adjustment strategies generator-first (gen) vs. battery-first (bat). For all runs: $|T| = 96$, $\tau = 96$, model: safety (see Section 9.2), solved with CPLEX (see Section 11).

We have run a number of single day-simulations to determine which one yields better results: adjusting the generator first or adjusting the battery first. The results are listed in Table 7 for the naive model (Section 9.1) and in Table 8 for the naive model extended with safety constraints (Section 9.2).

For the naive model we observe only minimal differences in the total cost but in the safety version, the differences are more pronounced. In two of the three test cases, the battery-first variant consistently performs better, while it is consistently beaten in the Residential case. At the very least this justifies the conclusion that using the battery first is no worse than using the generator first.

It is much easier to formulate a two-stage model if we use the battery output as the second stage: the battery is always available to adjust to an emerging shortage or surplus. A generator, however, has to be running to be able to contribute. Additionally, a battery can also actively store excess energy, while a generator can only decrease its contribution, and is, therefore, more limited when solving a surplus.

We will be using the battery-output as a second stage variable in our two-stage model 9.3. Given this, it is sensible to mirror this when solving a mismatch of planning and realization in the simulator by using the battery first. By using the same strategy in our simulations with the naive and safety models we are better able to compare the results.

The results of these simulations show that the battery-first strategy is not inferior to

		time			cost					
scenario	τ	(s)	(ticks)	gap	real	exp.	cor.	Δch	#adj.	#fail
Island	96	1.76	1704.29	0.009%	€421.56	€455.33	€419.09	+8.23	88	0
	24	0.29	196.78	0.007%	€422.73	€454.72	€419.49	+10.81	96	0
Island+	96	0.34	192.63	0.002%	€255.61	€263.85	€251.58	+13.43	60	0
	24	0.15	55.03	0.001%	€251.12	€255.98	€246.92	+14.01	61	0
Residential	96	0.05	44.73	0.003%	€2230.41	€2148.04	€2258.01	-92.00	197	0
	24	0.02	13.33	0.002%	€2229.08	€2154.35	€2256.68	-92.00	173	0

Table 9: Results of simulation runs comparing different values for τ. For all runs: $|T| = 672$, model: naive (see Section 9.1), solved with CPLEX (see Section 11).

		time			cost					
scenario	τ	(s)	(ticks)	gap	real	exp.	cor.	Δch	#adj.	#fail
Island	96	20.04	25538.45	0.009%	€421.72	€466.79	€421.46	+0.86	83	0
	24	0.63	610.73	0.008%	€419.82	€462.22	€419.82	+0.01	88	0
Island+	96	1.38	1773.56	0.004%	€251.34	€265.17	€246.11	+17.45	47	0
	24	0.11	97.29	0.002%	€246.24	€258.64	€241.52	+15.74	50	0
Residential	96	40.41	53531.66	0.012%	€2237.72	€2237.58	€2262.56	-82.80	0	0
	24	0.50	536.65	0.005%	€2237.78	€2237.66	€2260.00	-74.05	0	0

Table 10: Results of simulation runs comparing different values for τ. For all runs: $|T| = 672$, model: safety (see Section 9.2), solved with CPLEX (see Section 11).

the generator-first strategy. All simulations from this point onwards use the battery-first strategy.

13.2. Shrinking the problem: using τ

The next set of experiments were performed to check the validity of the model improvement described in Section 9.5: shifting the time resolution of the problem from 15 minutes to an hour after a number of time periods. We have simulated the results of running the microgrids for a week using $\tau = 96$ (i.e. no change in resolution) and $\tau = 24$ (shift to one-hour-resolution after six hours) for all our test grids. The results of doing this using the naive model are in Table 9, and those using the safety model are in Table 10.

We can see a large reduction in the average computation time in these results. With the naive and safety models, a solution proven to be within 0.01% of the optimum is found almost always, for both values of τ, but this shows that a lower value for τ allows for solving the problems much faster. This is not surprising, considering that a value of $\tau = 24$ cuts the amount of variables and constraints in the problem almost in half.

It is, however, interesting to see how this smaller version of the problem performs on other aspects. First, we see that the cost are generally comparable, and even lower for the $\tau = 24$ case, when a difference can be noted at all. The two approaches show

		time			cost					
scenario	solver	(s)	(ticks)	gap	real	exp.	cor.	Δch	#adj.	#fail
Island	CPLEX	0.22	144.82	0.005%	€53.36	€56.65	€53.36	-0.00	13	0
	GLPK	589.33	-	2.38 %	€56.36	€61.77	€52.98	+11.28	17	0
Island+	CPLEX	0.22	85.64	0.001%	€18.55	€18.48	€16.04	+8.36	5	0
	GLPK	399.17	-	1.15 %	€17.47	€18.20	€15.77	+5.68	5	0
Residential	CPLEX	0.03	18.46	0.002%	€242.10	€242.10	€241.66	+1.47	0	0
	GLPK	3.09	-	0.00 %	€234.21	€234.20	€242.22	-26.69	0	0

Table 11: Results of simulation runs using different solvers. For all runs: $|T| = 96$, $\tau = 24$, model: naive (see Section 9.1).

no consistent or large difference in the amount of energy they add to or take from the system (Δch). Finally, judging by the number of adjustments that have to be made, the $\tau = 96$ variant is the clear winner.

In all, the performance of both the full $\tau = 96$ case and the reduced $\tau = 24$ case is similar. It does seem that the $\tau = 96$ cases are more conservative, as they plan using a more detailed forecast of the later time steps. However, as the reduced version provides a large speed-up, we will be using a value of $\tau = 24$ for all upcoming tests. Especially for the more difficult two-stage and multi-stage models, where a solution within 1% of the optimum is often not found before the time limit is exceeded, this improvement allows us to find better solutions within the available time.

13.3. Solver comparison

As was mentioned in Section 11, we implemented our process in such a way that multiple solvers could be used to solve the resulting problems. In this section we look at the relative performance of two of these solvers: an open-source solver (GLPK [79]) and a commercial solver (CPLEX [54]). The results are summarized in Table 11.

Purely looking at the quality of the resulting solutions, both solvers perform approximately equally well. The difference in the cost can mainly be explained by the difference in Δch: the more expensive solutions end up with more energy stored in the battery at the end of the day.

The number of adjustments that have to be made is also equal in the Island+ and Residential scenarios. In both of these cases, both solvers were able to find a solution that was proven to be (close to) the optimal one in (almost) all time steps. In the case of the Island scenario, however, the run using GLPK almost never arrived at a near-optimal solution and were generally cut short by the time limit of 10 minutes. This may be one of the reasons why this run had to make more adjustment than the CPLEX run.

The problems presented to the solvers in this case were the easiest ones available. The fact that $\tau = 24$ means that the problem size was cut in half compared to runs with $\tau = 96$. Also, the naive model has the smallest number of constraints and variables, considering that the safety model adds a number of reserve variables and constraint and that the two- and multi-stage approaches multiply the number of variables and constraints by

		time			cost					
scenario	$\|S\|$	(s)	(ticks)	gap	real	exp.	cor.	Δch	#adj.	#fail
Island	100	17.10	15360.92	0.930%	€56.65	€56.91	€53.54	+10.36	1	0
	200	58.41	41822.19	0.903%	€56.69	€55.72	€54.07	+8.72	2	0
	300	106.63	81862.66	0.917%	€56.32	€56.60	€54.05	+7.58	1	0
Island+	100	19.31	18580.18	0.962%	€27.09	€27.06	€17.65	+31.46	0	0
	200	78.83	59976.25	0.957%	€25.76	€25.74	€16.68	+30.27	0	0
	300	206.46	148203.58	1.017%	€24.49	€24.47	€16.40	+26.96	0	0
Residential	100	1.89	1985.70	0.375%	€244.41	€244.37	€245.90	-4.96	0	0
	200	6.24	5318.35	0.391%	€239.82	€239.78	€244.79	-16.57	0	0
	300	11.88	11443.50	0.214%	€243.88	€243.85	€244.68	-2.68	0	0

Table 12: Results of simulation runs with different numbers of scenarios. For all runs: $|T| = 96$, $\tau = 24$, model: two-stage (see Section 9.3), solved with CPLEX (see Section 11).

a factor of $O(|S|)$. Considering that the GLPK solver already struggles with the Island scenario, it cannot be expected to perform well with the more complex models. Indeed, we found it was impossible to (consistently) find reasonably good solutions to two- or multi-stage problems with the GLPK solver. We therefore decided to use the CPLEX solver for all other test presented in this section.

The CPLEX solver, however, solves these problems reasonably fast. Most problems are solved almost instantly by CPLEX, which is impressive when compared to the time GLPK needs. To be fair, CPLEX is able to exploit the fact that the tests were run on a quad-core processor, as it uses multi-threading, while GLPK does not. But, even after correcting for this, the difference is still huge. There can be no doubt that CPLEX has the faster solver.

13.4. Two-stage: Determining the number of scenarios

The model described in Section 9.3, uses a number of randomly sampled scenario (see Alg. 1 in Section 7). We have run (one-day) simulations with different amounts of scenarios to evaluate the effect of increasing the number of scenarios on the solution quality.

The results in Table 12 show that the gains of increasing the number of scenarios from 100 to 300 are minor. However, the problem remains manageable and a solution within 1% of the optimum is usually found before the time limit is exceeded.

The small difference in cost is probably due to the way this model is set-up and the relative importance of the first time steps when using the rolling horizon approach. With 100 scenarios, we already have a quite accurate image of what to expect in the first few time steps. This does not improve much by increasing the number of scenarios to 300. However, in the later time steps, this does increase the quality of the representation a lot more, but this is not affecting the overall performance as much, as we have seen before.

It would, therefore, suffice to sample only 100 scenarios, or even less. However, we

		time			cost					
scenario	l_x	(s)	(ticks)	gap	real	exp.	cor.	Δch	#adj.	#fail
Island	l_1	416.02	230588.94	0.987%	€57.96	€57.01	€54.89	+10.22	2	0
	l_2	164.08	90404.24	0.971%	€62.80	€64.16	€58.41	+14.62	12	0
	l_3	500.91	263430.98	1.994%	€54.92	€55.45	€53.68	+4.12	1	0
Island+	l_1	499.52	271007.69	1.318%	€17.26	€16.04	€16.02	+4.15	1	0
	l_2	453.56	259248.22	1.326%	€24.03	€25.10	€18.43	+18.67	2	0
	l_3	487.54	245878.79	1.987%	€17.78	€16.56	€16.15	+5.43	1	0
Residential	l_1	8.91	6504.60	0.292%	€244.50	€244.45	€242.96	+5.14	0	0
	l_2	1.94	1423.97	0.436%	€250.73	€250.69	€247.21	+11.74	0	0
	l_3	21.02	13597.85	0.226%	€243.28	€243.23	€245.63	-7.84	0	0

Table 13: Results of simulation runs with different "branching models". For all runs: $|T| = 96$, $|S| = 300$, $\tau = 24$, model: multi-stage (see Section 9.4), solved with CPLEX (see Section 11).

chose to sample 300 scenarios for the comparison with the other algorithms. This is mainly due to two factors: (1) The problem is still manageable with 300 scenarios, so it is not absolutely necessary to use less, and, more importantly, (2) the multi-stage model starts out with 300 scenarios as well (as we will see in the next section). Choosing the same number of initial scenarios for both approaches allows for a better comparison.

13.5. Multi-stage: "branching model"

Comparing the different "branching models" defined in Section 7 in a single-day simulation led to the results presented in Table 13. We can see at a single glance that l_2 performs the worst. This is not surprising, considering that the first few time steps only have a very small number of scenarios, and they do not come close to representing the full range of possibilities.

The other two, l_1 and l_3, perform approximately equally well. They both have a large number of scenarios in the first few time steps, which gives them an edge over l_2. However, the increase in scenarios from l_3 when compared to l_1 does not seem to make much of a difference. This is probably due to the fact that we apply a rolling horizon strategy: only the decisions from the first time step are actually implemented. We do expect l_3 to generate scenario trees that produce solutions that are more often reusable when running the *-variant with the multi-stage model. It has more scenarios representing the earlier time step and therefore is more likely to contain one that is close enough to the realization to continue. Therefore, we chose to use l_3 from now on, even though l_1 performs similarly.

13.6. Re-planning strategies

We implemented the *-variant (described in Section 10) for the naive, safety and multi-stage models and the **-variant (also described in Section 10) for the naive and safety models. We will compare how well the decisions taken in the different cases are able to

be re-used, and what the effects of skipping the generation of a new planning until the old one is invalidated (according to the criteria described in Section 10) are. The results are shown in Table 14 (one day) and Table 15 (one week).

Surprisingly, it appears that the runs using *- and **-variants are not consistently beaten by the same run using the original variant. This indicates that re-planning every time step is not necessary, and may sometimes even worsen the result.

First, looking at the naive model, we see that (taking the difference in Δch into account) the three variants (original, * and **) perform almost equally well in the one-day simulations. The *- and **-variations even seem to be slightly better in the week-long simulations, with the exception of naive** in the Residential case. Now, if we include the number of adjustments necessary, we see that the *-variation usually increases the number of adjustments, while the **-variation needs fewer in almost all cases. Considering the number of time steps that the generation of a new planning was skipped, the performance of the *- and **-variations is impressive.

When we look at the safety model, the three variants again show similar results. We see that, although they are all fairly close together, the original model is slightly better than the *- and **-variations. Neither one of the *- and **-variations performs consistently better than the other. In the one-week simulations the original model performs worse than the other variations in the Island case but still wins in the other two test cases. Looking at the number of adjustments, the differences are small, although the original usually performs the best on this metric. The number of time steps that the planning is re-used in the variations of the safety model is comparable to the *- and **-variations of the naive model, and the results are equally impressive.

In the Island case, the multi-stage run has a somewhat better result than multi-stage*. In the Island+ and Residential cases, the differences are minimal, but multi-stage is the cheaper one in the Island+ case, while multi-stage* wins in the Residential case. The number of time steps that can be skipped using this approach is generally more than for the *-variations of the naive and safety models, but less than for the **-variations.

In summary, the approaches that skip the generation of a new planning, in an intelligent way, in a large number of time steps have a comparable performance to the original approach, that generates a new planning every time step.

13.7. Overall comparison

In Table 14 (one day) and Table 15 (one week) we present a comparison of the different approaches discussed in this thesis, using the settings that were in part determined based on the results presented in the previous sections.

For comparison, we also computed the cost of operating the grid based solely on the rules-of-thumb that are otherwise used to adjust the generated planning. These results are labeled as “none” in Table 14 and 15. In almost all cases, the approaches that do use any form of planning outperform the run without planning, although the differences in cost are not huge. In the number of adjustments needed, however, all planning approaches soundly beat this no-planning run.

		time			cost					
scenario	model	(s)	(ticks)	gap	real	exp.	cor.	Δch	#adj.	#fail
Island	none	-	-	-	€56.59	-	€54.51	+6.94	29	-
	naive	0.22	144.82	0.005%	€53.36	€56.65	€53.36	-0.00	13	0
	naive*	0.25	176.24	0.004%	€57.63	€60.25	€52.76	+16.22	15	57
	naive**	0.16	103.16	0.006%	€56.68	€56.55	€52.39	+14.29	7	88
	safety	0.57	402.59	0.008%	€51.56	€54.54	€52.18	-2.06	12	0
	safety*	0.64	524.85	0.008%	€53.11	€55.07	€52.87	+0.79	9	57
	safety**	0.72	612.57	0.009%	€55.46	€57.10	€52.46	+10.00	12	82
	two-stage	106.63	81862.66	0.917%	€56.32	€56.60	€54.05	+7.58	1	0
	multi-stage	500.91	263430.98	1.994%	€54.92	€55.45	€53.68	+4.12	1	0
	multi-stage*	531.18	274058.15	2.367%	€59.17	€58.64	€55.53	+12.12	4	50
Island+	none	-	-	-	€20.13	-	€16.74	+11.29	2	-
	naive	0.22	85.64	0.001%	€18.55	€18.48	€16.04	+8.36	5	0
	naive*	0.12	90.79	0.004%	€18.47	€19.49	€16.01	+8.20	5	65
	naive**	0.11	74.37	0.003%	€20.15	€21.03	€15.94	+14.03	7	86
	safety	0.12	102.03	0.002%	€17.88	€18.70	€15.00	+9.60	4	0
	safety*	0.15	100.82	0.002%	€18.93	€19.98	€15.71	+10.75	4	65
	safety**	0.16	112.90	0.002%	€21.61	€23.14	€16.87	+15.79	6	87
	two-stage	206.46	148203.58	1.017%	€24.49	€24.47	€16.40	+26.96	0	0
	multi-stage	487.54	245878.79	1.987%	€17.78	€16.56	€16.15	+5.43	1	0
	multi-stage*	501.16	254372.72	1.861%	€19.20	€20.11	€16.64	+8.52	2	65
Residential	none	-	-	-	€233.44	-	€261.04	-92.00	39	-
	naive	0.03	18.46	0.002%	€242.10	€242.10	€241.66	+1.47	0	0
	naive*	0.04	17.55	0.003%	€226.78	€226.78	€241.98	-50.68	0	60
	naive**	0.02	13.13	0.002%	€222.72	€222.75	€240.72	-60.01	0	92
	safety	0.63	703.31	0.008%	€239.43	€239.43	€241.68	-7.49	0	0
	safety*	1.08	796.36	0.008%	€233.34	€233.33	€241.76	-28.06	0	60
	safety**	0.52	465.23	0.008%	€231.74	€231.76	€240.78	-30.12	0	92
	two-stage	11.88	11443.50	0.214%	€243.88	€243.85	€244.68	-2.68	0	0
	multi-stage	21.02	13597.85	0.226%	€243.28	€243.23	€245.63	-7.84	0	0
	multi-stage*	20.09	13782.87	0.163%	€234.58	€234.52	€245.51	-36.62	0	74

Table 14: Results of simulation runs with different models. For all runs (if applicable): $|T| = 96$, $|S| = 300$, $\tau = 24$, $l_x = l_3$, solved with CPLEX (see Section 11).

		time			cost					
scenario	model	(s)	(ticks)	gap	real	exp.	cor.	Δch	#adj.	#fail
Island	none	-	-	-	€416.09	-	€417.28	-3.97	226	-
	naive	0.29	196.78	0.007%	€422.73	€454.72	€419.49	+10.81	96	0
	naive*	0.31	226.23	0.007%	€414.95	€453.57	€411.99	+9.88	109	312
	naive**	0.40	240.91	0.008%	€416.80	€443.69	€414.01	+9.31	90	570
	safety	0.63	610.73	0.008%	€419.82	€462.22	€419.82	+0.01	88	0
	safety*	0.99	816.69	0.009%	€415.46	€465.09	€414.63	+2.77	88	312
	safety**	1.08	969.34	0.009%	€417.33	€456.91	€416.01	+4.41	98	563
	two-stage	163.29	114539.78	0.951%	€428.94	€438.39	€427.68	+4.19	73	0
	multi-stage*	526.87	275213.52	7.852%	€438.00	€452.18	€434.41	+11.97	71	391
Island+	none	-	-	-	€265.64	-	€254.39	+37.51	61	-
	naive	0.15	55.03	0.001%	€251.12	€255.98	€246.92	+14.01	61	0
	naive*	0.09	61.93	0.002%	€249.89	€251.50	€245.68	+14.02	58	381
	naive**	0.11	66.56	0.002%	€250.26	€246.77	€244.51	+19.18	53	598
	safety	0.11	97.29	0.002%	€246.24	€258.64	€241.52	+15.74	50	0
	safety*	0.20	126.61	0.003%	€245.63	€257.89	€240.80	+16.10	49	381
	safety**	0.19	104.36	0.003%	€248.15	€252.87	€243.71	+14.79	47	606
	two-stage	154.79	119895.71	0.957%	€262.12	€260.54	€250.68	+38.14	24	0
	multi-stage*	418.74	216911.97	2.002%	€255.28	€256.78	€252.88	+7.99	30	473
Residential	none	-	-	-	€2299.01	-	€2326.61	-92.00	522	-
	naive	0.02	13.33	0.002%	€2229.08	€2154.35	€2256.68	-92.00	173	0
	naive*	0.03	12.89	0.003%	€2230.43	€2149.11	€2255.41	-83.25	192	382
	naive**	0.03	12.59	0.001%	€2237.91	€2220.65	€2262.89	-83.25	43	605
	safety	0.50	536.65	0.005%	€2237.78	€2237.66	€2260.00	-74.05	0	0
	safety*	0.80	593.89	0.006%	€2239.15	€2237.19	€2261.34	-73.96	13	382
	safety**	0.67	585.69	0.008%	€2246.30	€2233.48	€2267.68	-71.25	6	640
	two-stage	8.60	9257.26	0.360%	€2247.84	€2247.59	€2263.57	-52.42	0	0
	multi-stage*	17.61	12030.12	0.114%	€2248.62	€2248.28	€2265.88	-57.52	0	517

Table 15: Results of simulation runs with different models. For all runs (if applicable): $|T| = 672$, $|S| = 300$, $\tau = 24$, $l_x = l_3$, solved with CPLEX (see Section 11).

We can see that the safety model yields the best results overall (in terms of the cost of operation), both in the one-day and one-week simulations

Looking at the original intent of each of the models, they all seem to do what they were designed to do: the safety model improves over the naive model by bringing the number of adjustments down and thereby (usually) reducing the cost. The two-stage model is the most conservative of them all, bringing the number of adjustments down even further when compared with the other approaches. It does, however, pay a price for this: the costs are also the highest of all approaches. The multi-stage model, in turn, was designed to be less conservative, without bringing the number of adjustments up by too much, and this is exactly what it does: at a cost comparable to that of the safety model and a number of adjustments comparable to the two-stage model, this approach combines the best of both worlds. The downside of the multi-stage approach, however, is that it is computationally heavy.

13.8. Discussion

In the previous section, we briefly discussed the results of the different approaches. In this section, we aim to provide an explanation for some of the observed results and try to put them into perspective.

The first thing to note is that the difference in cost between most approaches is small. This is not surprising. In all scenarios, the output of the PV panels is not sufficient to satisfy the full load. The difference between the production of the PV panels and the consumption has to be produced by the generators at some point. This amount is the same for all approaches. How much has to be produced can only be slightly increased depending on how much energy is stored in the battery in between generation and consumption, as some of the energy stored there will be lost, due to the efficiency of the battery. Compared to the minimal production of the generators, their operational cost is small and the only other significant contribution comes from the start-up cost of the generators.

Some advantage can be gained by timing the production right, so less energy has to be routed through the batteries. This can further be improved by reducing the number of start-ups or reducing the operation time of the generators. However, the latter potentially counter-acts the timing of your production, and the difference in the number of start-ups will never be large. This means that any improvement should not be expected to be huge.

We approximated the cost of the *actual* operation of the Island scenario in the test period by looking at the reported values and compute an estimation of the actual cost using the same cost coefficients used in our models. This results in a cost of €50.76 for June 1st, and €405.26 for the entire week following June 1st. This seems a lot cheaper than even the cheapest of our planning approaches. However, a closer analysis of these results shows that the simulated approaches produce a larger amount of energy. Additionally, the battery costs are lower, indicating that less energy is routed through the batteries in the simulated cases. The difference in required production can be explained

by our choices for the grid- and battery-efficiencies to be on the conservative side. In reality, the losses turn out to be smaller. If we keep this in mind, most approaches do seem to improve the current situation, albeit only slightly.

In an earlier set of simulations, the difference between our approach and the current mode of operation was much larger. The cost of the simulation of the naive model was compared to the cost of real operation. There were two major differences when comparing it to our current tests: (1) the cost coefficients of the generators were different and (2) the test period was different. The cost coefficients used were those mentioned in Section 8: €1 operational cost per quarter and start-up cost of €0.10.

These cost coefficients reward the reduction of the time that a generator is operational far more and punish start-ups way less than the current coefficients. This makes the difference between approaches much more pronounced and allows for much more improvement on the current mode of operation, as the generators are operational for relatively long periods of time. However, this does result in a drastically different planning, where generators are switched on and off much more often.

In the current problem, the cost of operating the generators dominate. The usage of batteries costs very little, and the other sources of production (PV, and possibly wind) do not contribute anything to the cost. Determining the right values for the cost coefficients of the generator is therefore essential. We have found that changing these coefficients may result in a drastically different planning being generated, which leads to a similarly drastic difference in cost when compared to the (estimated) current operational cost. In this thesis, we have determined our cost coefficients mainly by adjusting them in such a way that they resulted in a planning that seemed sensible. This means that we cannot draw any strong conclusions about the economic advantage of using the described planning process.

However, we have mentioned before that the described approach can be extended relatively easily to include multiple commodities and more device types. For the problem used to obtain the results in this section, the main challenge is getting the planning of the generators right. This is relatively easy to grasp, and therefore it is possible to come up with reasonably good heuristic strategies for operation. When the grid becomes more complex it is harder to oversee the consequences of the interactions between different devices and/or commodities. The approach described here has the potential to provide good solutions in these cases as well, when devising good strategies becomes increasingly challenging.

The test period used to evaluate our approach also contributes to the small difference when compared to strategies that do not use planning. The production of the PV panels is at its peak during the month of June. For most of the year, the number of time steps where any PV output is predicted at all is smaller, and the peak is smaller as well. The variation is at its highest in June.

Additionally, a number of days were included in the test period of Island and Island+ where the load forecast was completely wrong, as the site was unexpectedly open during hours that it would usually have been closed. During these days, a planning generated

based on the load forecast is not very useful – it will probably only increase the cost. This shows in the results: in the Island and Island+ cases, the runs that use planning are only slightly better than the run that does not use any planning. If we look at the Residential case, where this unexpected deviation from the typical load did not occur, we observe that the planning approaches do perform better than the no-planning run.

In summary, the test period that we used is relatively challenging. So, even though the benefits of using any of the developed approaches seem small, it does show that using the described approach has the potential to save cost. We also have to consider that we are dealing with a relatively small microgrid, which makes it harder to predict the load accurately. This also leads to (in our case) a standard deviation that is relatively large compared to the overall load. For larger systems, this situation will improve, which will potentially allow the planning approaches to perform better as well.

One aspect that we have not discussed so far is the number of adjustments needed. All of the planning approaches need fewer adjustments than the no-planning approach. This means that only the output of the batteries has to be adjusted to keep the energy balance. The process using a two-stage or multi-stage model, in particular, produces a planning that will almost always be valid. This indicates that generators with more complex production constraints or efficiency functions may be effectively incorporated in the planning. Another advantage is the predictability of the system behavior. When a grid operator looks at the planning, it is actually likely to be an accurate representation of what will happen in the grid.

Finally, we have to consider the set-up of our experiments. It is hard to compare the results of different approaches, as there are multiple metrics that all define different aspects of a "good" solution: the cost is not the only thing that matters. At the very least, we need to take the change in the state of the batteries into account. One might have counteracted this by fixing the state of charge that the battery should be at when the simulation ends. However, this may also confuse the results, as this limits the freedom in choosing a battery management strategy, which may contribute to the quality of the overall solution. We have computed a corrected cost taking the state of charge of the batteries into account, but the correction is only an approximation. It is also unclear how to incorporate the number of adjustments in the cost. One could add a penalty for every adjustment necessary. While this would help in showing the difference in performance of the approaches, it has its problems. The resulting cost is no longer an indication of the actual cost of operation using a particular approach, and this makes it harder to evaluate *why* one approach is better than another.

14. Conclusion and future work

The main research questions of this thesis were the following:

1. How can we manage the operation of microgrids with centralized control in a (near-)optimal way?
   a) What is the best way to deal with uncertainty in the input for the model?
   b) Is considering an abstract supply-demand balancing problem sufficient to arrive at good solutions?
2. What is the value of intelligent planning in microgrid operation?
   a) What is the potential reduction in operational cost and how does it depend on the composition of the microgrid?
   b) Does the potential reduction in operational cost outweigh the cost of constructing a planning (cost of the IT infrastructure, etc.)?

With regard to the first question, we have shown that the safety model generally comes up with the cheapest solution, while the planning generated by the more conservative two-stage model is almost always valid. The multi-stage model combines these two advantages, resulting in solutions that are only slightly more expensive than the safety-solutions and need to be adjusted only a little more often than the two-stage-solutions.

Additionally, we found that it is not necessary to generate a new planning every time step to achieve a good result. In fact, in terms of computational cost and economic value, it may even be advisable to put off the generation of a new planning until it is absolutely necessary. This may, however, increase the number of adjustments necessary to keep the energy balance in the grid.

From what we have seen so far, it seems that repeatedly solving an abstract supply-demand balancing problem yields good solutions. We will only know for sure, however, when the algorithms using one of the developed models is actually put in charge of a real microgrid. At this time, Zown is taking steps to make this happen in the near future.

The second main question is a little harder to answer. The reduction in operational cost is not large, and it is hard to say whether it will outweigh the cost of developing and implementing the planning process in the long term, in particular for smaller microgrids. It does show potential to do so, however. For larger systems with more predictable energy profiles, it certainly provides an advantage.

There are multiple ways in which work on this topic could be continued:

- Implementing the approach described in thesis and evaluating the results. The results in this thesis are all based on simulations, with a number of underlying assumptions. An experimental study determining the performance of this approach in practice would be valuable;

- Extending the current model with multiple commodities or other device types. We expect the same approach (or a similar one) to be effective in these cases, but it would be interesting to put this expectation to the test;
- Repeating the test on different (types of) microgrids. This may be an extension of the previous point, but this one is aimed at strengthening the claim of general applicability of the current model;
- Refining the scenario selection and scenario tree construction approaches. The approach described in this thesis is only one of a large number of possible approaches, and it is not unthinkable that a considerable improvement on the current approach can be made. Multiple sampling methods, distance definitions or even alternative selection algorithms could be studied;
- Comparing the current approach to approaches using other optimization techniques. We listed our reasons for choosing this approach in Section 5 and Section 9, but other approaches are possible. Genetic algorithms have been tried in the past, for example, but different approaches have rarely been compared.

Aside from these, rather large, topics additional work may be done on determining values for the parameters of the current approach. The approach has a lot of parameters, ranging from the constants used in the models (the cost coefficients, efficiency, etc.) to parameters like τ, $|S|$ and $|T|$. While at least some analysis was performed to find good settings for all of these, there is room for improvement.

The operational management of microgrids remains a challenging problem that will require continued study. We have demonstrated the potential of the application of forecasting and planning techniques to this problem. We are confident that the implementation of this approach in practice and its continued development will contribute to the efficient operation of microgrids in the near future.

References

[1] A. Ahmad Khan, M. Naeem, M. Iqbal, S. Qaisar, and A. Anpalagan. A compendium of optimization objectives, constraints, tools and algorithms for energy management in microgrids. *Renewable and Sustainable Energy Reviews*, 58:1664–1683, 2016.

[2] Alliander N.V. Alliander looks ahead, 2016. Available at: `https://www.alliander.com/en/our-activities/new-markets`.

[3] Alliander N.V. Company Profile, 2016. Available at: `https://www.alliander.com/en/about-alliander/company-profile`.

[4] Avesco AG. Technical data Diesel Generator Set CAT C4.4-110 (DE110E3), 2015.

[5] V. Bakker. *TRIANA: a control strategy for Smart Grids*. PhD thesis, University of Twente, 2012.

[6] M. Beaudin, H. Zareipour, A. Schellenberglabe, and W. Rosehart. Energy storage for mitigating the variability of renewable electricity sources: An updated review. *Energy for Sustainable Development*, 14(4):302–314, 2010.

[7] P. Bendotti, P. Fouilhoux, and C. Rottner. Complexity of the min-up/min-down unit commitment problem. 2017.

[8] A. Bischi, L. Taccari, E. Martelli, E. Amaldi, G. Manzolini, P. Silva, S. Campanari, and E. Macchi. A detailed MILP optimization model for combined cooling, heat and power system operation planning. *Energy*, 74(C):12–26, 2014.

[9] M. G. C. Bosman. *Planning in smart grids*. PhD thesis, University of Twente, 2012.

[10] S. Bracco, F. Delfino, F. Pampararo, M. Robba, and M. Rossi. A mathematical model for the optimal operation of the University of Genoa Smart Polygeneration Microgrid: Evaluation of technical, economic and environmental performance indicators. *Energy*, 64(April):912–922, 2014.

[11] S. Bracco, F. Delfino, F. Pampararo, M. Robba, and M. Rossi. A dynamic optimization-based architecture for polygeneration microgrids with tri-generation, renewables, storage systems and electrical vehicles. *Energy Conversion and Management*, 96:511–520, 2015.

[12] S. Bracco, G. Dentici, and S. Siri. Economic and environmental optimization model for the design and the operation of a combined heat and power distributed generation system in an urban area. *Energy*, 55:1014–1024, 2013.

[13] Caterpillar. Olympian GEP55-1, 2008.

[14] C. Chen, S. Duan, T. Cai, B. Liu, and G. Hu. Smart energy management system for optimal microgrid economic operation. *IET Renewable Power Generation*, 5(3):258, 2011.

[15] Y.-H. Chen, S.-Y. Lu, Y.-R. Chang, T.-T. Lee, and M.-C. Hu. Economic analysis and optimal energy management models for microgrid systems: A case study in Taiwan. *Applied Energy*, 103(August):145–154, 2012.

[16] I. Colak, S. Sagiroglu, G. Fulli, M. Yesilbudak, and C. F. Covrig. A survey on the critical issues in smart grid technologies. *Renewable and Sustainable Energy Reviews*, 54:396–405, 2016.

[17] C. M. Colson. *Towards real-time power management of microgrids for power system integration: a decentralized multi-agent based approach.* PhD thesis, Montana State University, 2012.

[18] J. Cook, N. Oreskes, P. T. Doran, W. R. L. Anderegg, B. Verheggen, E. W. Maibach, J. S. Carlton, S. Lewandowsky, A. G. Skuce, and A. Sarah. Consensus on consensus: a synthesis of consensus estimates on human caused global warming. *Environmental Research Letters*, 11(2016):1–24, 2016.

[19] H. Dai, N. Zhang, and W. Su. A Literature Review of Stochastic Programming and Unit Commitment. *Journal of Power and Energy Engineering*, 03(04):206–214, 2015.

[20] R. Dai and M. Mesbahi. Optimal power generation and load management for off-grid hybrid power systems with renewable sources via mixed-integer programming. *Energy Conversion and Management*, 73(1-4):234–244, 2013.

[21] D. de Rie. Distributed Energy Storage in the Dutch Electricity Network. Master's thesis, Utrecht University, 2015.

[22] Department of Energy, Office of Electricity Delivery & Energy Reliability, and United States Government. Global Energy Stroage Database. Smart Storage., 2015. Available at: `http://www.energystorageexchange.org/projects/1752`.

[23] R. Dufo-López and J. L. Bernal-Agustín. Design and control strategies of PV-diesel systems using genetic algorithms. *Solar Energy*, 79(1):33–46, 2005.

[24] R. Dufo-López, J. L. Bernal-Agustín, J. M. Yusta-Loyo, J. A. Domínguez-Navarro, I. J. Ramírez-Rosado, J. Lujano, and I. Aso. Multi-objective optimization minimizing cost and life cycle emissions of stand-alone PV-wind-diesel systems with batteries storage. *Applied Energy*, 88(11):4033–4041, 2011.

[25] J. Dupačová, G. Consigli, and S. W. Wallace. Scenarios for Multistage Stochastic Programs. *Annals of operations research*, 100:25–53, 2000.

[26] J. Dupačová, N. Gröwe-Kuska, and W. Römisch. Scenario reduction in stochastic programming: An approach using probability metrics. *Mathematical Programming*, 95(3):493–511, 2003.

[27] Energieonderzoek Centrum Nederland. Nationale Energieverkenning 2016, 2016.

[28] Enexis B.V. Smart Storage in Etten-Leur, 2012.

[29] Enexis B.V., Alliander N.V., Energieonderzoek Centrum Nederland, and TNO. Openbaar Eindrapport Smart Storage. De eerste buurtbatterij van Nederland., 2013.

[30] R. Enrich, P. Skovron, M. Tolós, and M. Torrent-Moreno. Microgrid management based on economic and technical criteria. In *Energy Conference and Exhibition (ENERGYCON), 2012 IEEE International*, pages 551–556, 2012.

[31] European Commission. Definition, expected services, functionalities and benefits of smart grids. *EUR-Lex*, 202:1–10, 2011.

[32] European Commission. The roadmap for transforming the EU into a competitive, low-carbon economy by 2050, 2011.

[33] European Commission. 2020 climate & energy package, 2017. Available at: `http://ec.europa.eu/clima/policies/strategies/2020{\_}en`.

[34] X. Fang, S. Misra, G. Xue, and D. Yang. Smart Grid The New and Improved Power Grid: A Survey. *IEEE Communications Surveys & Tutorials*, 14(4):944–980, 2012.

[35] H. Farhangi. Smart Grid and ICT's Role in Its Evolution. In *Green Communications: Theoretical Fundamentals, Algorithms and Applications*, pages 29–50. CRC Press, 2012.

[36] I. A. Farhat and M. E. El-Hawary. Optimization methods applied for solving the short-term hydrothermal coordination problem. *Electric Power Systems Research*, 79(9):1308–1320, 2009.

[37] J. Faxas-Guzmán, R. García-Valverde, L. Serrano-Luján, and A. Urbina. Priority load control algorithm for optimal energy management in stand-alone photovoltaic systems. *Renewable Energy*, 68:156–162, 2014.

[38] J. Fink and J. L. Hurink. Minimizing costs is easier than minimizing peaks when supplying the heat demand of a group of houses. *European Journal of Operational Research*, 242(2):644–650, 2015.

[39] J. Fossati. Unit Commitment and economic dispatch in micro grids. *Memoria de Trabajos de Difusión Científica y Técnica*, 10:83–96, 2012.

[40] R. Fourer, D. M. Gay, and B. W. Kernighan. AMPL: A Mathematical Programming Language. *Management Science*, 36:519–554, 1990.

[41] C. Gamarra and J. M. Guerrero. Computational optimization techniques applied to microgrids planning: A review. *Renewable and Sustainable Energy Reviews*, 48:413–424, 2015.

[42] N. Gröwe-Kuska, H. Heitsch, and W. Römisch. Scenario Reduction and Scenario Tree Construction for Power Management Problems. In *Power tech conference proceedings, 2003 IEEE Bologna*, 2003.

[43] W. Gu, Z. Wu, R. Bo, W. Liu, G. Zhou, W. Chen, and Z. Wu. Modeling, planning and optimal energy management of combined cooling, heating and power microgrid: A review. *International Journal of Electrical Power and Energy Systems*, 54(May 2016):26–37, 2014.

[44] J. M. Guerrero, M. Chandorkar, T. L. Lee, and P. C. Loh. Advanced control architectures for intelligent microgrids - Part i: Decentralized and hierarchical control. *IEEE Transactions on Industrial Electronics*, 60(4):1254–1262, 2013.

[45] J. M. Guerrero, J. C. Vásquez, and R. Teodorescu. Hierarchical control of droop-controlled DC and AC microgrids - A general approach towards standardization. In *Industrial Electronics Conference Proceedings*, pages 4305–4310, 2009.

[46] L. Guo, W. Liu, X. Li, Y. Liu, B. Jiao, W. Wang, C. Wang, and F. Li. Energy Management System for Stand-Alone Wind-Powered-Desalination Microgrid. *IEEE Transactions on Smart Grid*, 7(2):1079–1087, 2016.

[47] N. D. Hatziargyriou, A. Engler, N. Jenkins, G. Strbac, J. A. Pecas Lopes, J. Ruela, G. Kariniotakis, J. Oyarzabal, and A. Amorim. Microgrids - Large Scale Integration of Microgeneration to Low Voltage Grids. In *Cigré 2006*, 2006.

[48] A. D. Hawkes and M. A. Leach. Modelling high level system design and unit commitment for a microgrid. *Applied Energy*, 86(7-8):1253–1265, 2009.

[49] H. Heitsch and W. Römisch. Scenario Reduction Algorithms in Stochastic Programming. *Computational Optimization and Applications*, 24:187–206, 2003.

[50] C. A. Hernandez-Aramburo, T. C. Green, and N. Mugniot. Fuel consumption minimization of a microgrid. *IEEE Transactions on Industry Applications*, 41(3):673–681, 2005.

[51] B. F. Hobbs, M. H. Rothkopf, R. P. O'Neill, and H.-P. Chao. Example of new capabilities: solving Unit Commitment problems using MIP. In *The Next Generation of Electric Power Unit Commitment Models*, chapter 2. Springer Science & Business Media, 2006.

[52] A. Hoke, A. Brissette, S. Chandler, A. Pratt, and D. Maksimović. Look-ahead economic dispatch of microgrids with energy storage, using linear programming. In *2013 1st IEEE Conference on Technologies for Sustainability (SusTech)*, pages 154–161. IEEE, 2013.

[53] Y. Huang, P. M. Pardalos, and Q. P. Zheng. *Electrical Power Unit Commitment. Deterministic and Two-Stage Stochastic Programming Models and Algorithms.* Springer US, 2017.

[54] IBM. CPLEX Optimizer. Available at: `https://www-01.ibm.com/software/commerce/optimization/cplex-optimizer/index.html`.

[55] IBM. IBM Decision Optimization - Onboarding. Available at: `https://onboarding-oaas.docloud.ibmcloud.com/`.

[56] ICarus. iCarus. Voorspeller energiestromen. Available at: `http://icarus.energy/`.

[57] IPCC. 2013: Summary for Policymakers. In T. F. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S. K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P. M. Midgley, editors, *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge and New York, 2013.

[58] L. Ji, D. X. Niu, and G. H. Huang. An inexact two-stage stochastic robust programming for residential micro-grid management-based on random demand. *Energy*, 67:186–199, 2014.

[59] H. Jiayi, J. Chuanwen, and X. Rong. A review on distributed energy resources and MicroGrid. *Renewable and Sustainable Energy Reviews*, 12(9):2465–2476, 2008.

[60] M. M. Josephine. Centralized Design and Control for Optimizing Microgrid Systems Using Distributed Generators. In *Proceedings of the World Congress on Engineering and Computer Science 2016*, volume I, 2016.

[61] H. Kanchev, B. Francois, and V. Lazarov. Unit commitment by dynamic programming for microgrid operational planning optimization and emission reduction. In *2011 International Aegean Conference on Electrical Machines and Power Electronics and 2011 Electromotion Joint Conference (ACEMP)*, pages 502–507. IEEE, 2011.

[62] H. Kanchev, D. Lu, F. Colas, V. Lazarov, and B. Francois. Energy management and operational planning of a microgrid with a PV-based active generator for smart grid applications. *IEEE Transactions on Industrial Electronics*, 58(10):4583–4592, 2011.

[63] F. Katiraei, R. Iravani, N. D. Hatziargyriou, and A. Dimeas. Microgrids management. *IEEE Power and Energy Magazine*, 6(3):54–65, 2008.

[64] A. Kaur, J. Kaushal, and P. Basak. A review on microgrid central controller. *Renewable and Sustainable Energy Reviews*, 55:338–345, 2016.

[65] Kohler SDMO. K44C3, 2017.

[66] J. K. A. H. P. J. Kok. *The PowerMatcher: Smart Coordination for the Smart Electricity Grid*. PhD thesis, Vrije Universiteit Amsterdam, 2013.

[67] P. O. Kriett and M. Salani. Optimal control of a residential microgrid. *Energy*, 42(1):321–330, 2012.

[68] E. Kuznetsova. *Microgrid agent-based modelling and optimization under uncertainty*. PhD thesis, Université de Versailles Saint-Quentin-en-Yvelines, 2014.

[69] G. Kyriakarakos, D. D. Piromalis, A. I. Dounis, K. G. Arvanitis, and G. Papadakis. Intelligent demand side energy management system for autonomous polygeneration microgrids. *Applied Energy*, 103(March 2013):39–51, 2013.

[70] H. Liang, A. K. Tamang, W. Zhuang, and X. S. Shen. Stochastic Information Management in Smart Grid. *IEEE Communications Surveys & Tutorials*, 16(3):1746–1770, 2014.

[71] H. Liang and W. Zhuang. Stochastic modeling and optimization in a microgrid: A survey. *Energies*, 7(4):2027–2050, 2014.

[72] D. Liefferink, D. Boezeman, and H. de Coninck. The Netherlands: a case of fading leadership. In R. K. W. Wurzel, J. Connelly, and D. Liefferink, editors, *The European Union in International Climate Change Politics: Still Taking a Lead?* Taylor & Francis, 2017.

[73] W.-M. Lin, C.-S. Tu, and M.-T. Tsai. Energy Management Strategy for Microgrids by Using Enhanced Bee Colony Optimization. *Energies*, 9(1):5, 2015.

[74] G. Liu, M. Starke, B. Xiao, X. Zhang, and K. Tomsovic. Microgrid optimal scheduling with chance-constrained islanding capability. *Electric Power Systems Research*, 145:197–206, 2017.

[75] X. Liu. Optimization of a combined heat and power system with wind turbines. *International Journal of Electrical Power & Energy Systems*, 43(1):1421–1426, 2012.

[76] X. Luo, J. Wang, M. Dooner, and J. Clarke. Overview of current development in electrical energy storage technologies and the application potential in power system operation. *Applied Energy*, 137:511–536, 2015.

[77] M. S. Mahmoud, S. Azher Hussain, and M. A. Abido. Modeling and control of microgrid: An overview. *Journal of the Franklin Institute*, 351(5):2822–2859, 2014.

[78] T. S. Mahmoud, D. Habibi, and O. Bass. Fuzzy logic for smart utilisation of Storage Devices in a typical microgrid. *2012 International Conference on Renewable Energy Research and Applications, ICRERA 2012*, 2012.

[79] A. Makhorin. GLPK (GNU Linear Programming Kit). Available at: `https://www.gnu.org/software/glpk/`.

[80] M. P. Marietta, M. Graells, and J. M. Guerrero. A rolling horizon rescheduling strategy for flexible energy in a microgrid. In *ENERGYCON 2014 - IEEE International Energy Conference*, pages 1297–1303, 2014.

[81] F. Martin-Martínez, A. Sánchez-Miralles, and M. Rivier. A literature review of Microgrids: A functional layer based classification. *Renewable and Sustainable Energy Reviews*, 62:1133–1153, 2016.

[82] M. Marzband, M. Ghadimi, A. Sumper, and J. L. Dom??nguez-Garc??a. Experimental validation of a real-time energy management system using multi-period gravitational search algorithm for microgrids in islanded mode. *Applied Energy*, 128(0):164–174, 2014.

[83] P. Mazidi and G. N. Sreenivas. Reliability Assessment of A Distributed Generation Connected Distribution System. *International Journal of Electrical and Electronics Engineering*, 3(2):82–88, 2013.

[84] S. Mazzola, M. Astolfi, and E. Macchi. A detailed model for the optimal management of a multigood microgrid. *Applied Energy*, 154:862–873, 2015.

[85] L. Meng, E. R. Sanseverino, A. C. Luna, T. Dragicevic, J. C. Vasquez, and J. M. Guerrero. Microgrid supervisory controllers and energy management systems: A literature review. *Renewable and Sustainable Energy Reviews*, 60:1263–1273, 2016.

[86] L. I. Minchala-Avila, L. E. Garza-Castañón, A. Vargas-Martínez, and Y. Zhang. A review of optimal control techniques applied to the energy management and control of microgrids. *Procedia Computer Science*, 52(1):780–787, 2015.

[87] A. Molderink. *On the three-step control methodology for Smart Grids*. PhD thesis, University of Twente, 2011.

[88] A. Molderink, V. Bakker, M. G. C. Bosman, J. L. Hurink, and G. J. M. Smit. A three-step methodology to improve domestic energy efficiency. *Innovative Smart Grid Technologies Conference, ISGT 2010*, pages 1–8, 2010.

[89] H. Morais, P. Kádár, P. Faria, Z. A. Vale, and H. M. Khodr. Optimal scheduling of a renewable micro-grid in an isolated load area using mixed-integer linear programming. *Renewable Energy*, 35(1):151–156, 2010.

[90] C. Natesan, S. K. Ajithan, S. Chozhavendhan, and A. Devendiran. Power Management Strategies In Microgrid:A Survey. *International Journal of Renewable Energy Research (IJRER)*, 5(2):334–340, 2015.

[91] NEDU. Verbruiksprofielen. Available at: http://www.nedu.nl/portfolio/verbruiksprofielen/.

[92] M. Y. Nguyen, N. H. Choi, and Y. T. Yoon. DP formulation of microgrid operation with heat and electricity constraints. *Journal of Power Electronics*, 9(6):919–928, 2009.

[93] NIBUD. Energie en water, 2017. Available at: https://www.nibud.nl/consumenten/energie-en-water/.

[94] D. E. Olivares, C. A. Cañizares, and M. Kazerani. A Centralized Energy Management System for Isolated Microgrids. *IEEE Transactions on Smart Grid*, 2013.

[95] F. Orecchini and A. Santiangeli. Beyond smart grids - The need of intelligent energy networks for a higher global efficiency through energy vectors integration. *International Journal of Hydrogen Energy*, 36(13):8126–8133, 2011.

[96] O. Palizban, K. Kauhaniemi, and J. M. Guerrero. Microgrids in active network management - Part I: Hierarchical control, energy storage, virtual power plants, and market participation. *Renewable and Sustainable Energy Reviews*, 36(July 2015):428–439, 2014.

[97] R. Palma-Behnke, C. Benavides, F. Lanas, B. Severino, L. Reyes, J. Llanos, and D. Saez. A microgrid energy management system based on the rolling horizon strategy. *IEEE Transactions on Smart Grid*, 4(2):996–1006, 2013.

[98] S. Parhizi, H. Lotfi, A. Khodaei, and S. Bahramirad. State of the art in research on microgrids: A review. *IEEE Access*, 3:890–925, 2015.

[99] A. Parisio, E. Rikos, and L. Glielmo. A model predictive control approach to microgrid operation optimization. *IEEE Transactions on Control Systems Technology*, 22(5):1813–1827, 2014.

[100] A. Parisio, E. Rikos, G. Tzamalis, and L. Glielmo. Use of model predictive control for experimental microgrid optimization. *Applied Energy*, 115(September 2016):37–46, 2014.

[101] G. C. Pflug and A. Pichler. *Multistage Stochastic Optimization*. Springer International Publishing, 2014.

[102] G. S. Piperagkas, A. G. Anastasiadis, and N. D. Hatziargyriou. Stochastic PSO-based heat and power dispatch under environmental constraints incorporating CHP and wind power units. *Electric Power Systems Research*, 81(1):209–218, 2011.

[103] E. Planas, A. Gil-De-Muro, J. Andreu, I. Kortabarria, and I. Martínez De Alegría. General aspects, hierarchical controls and droop methods in microgrids: A review. *Renewable and Sustainable Energy Reviews*, 17:147–159, 2013.

[104] I. Prodan and E. Zio. A model predictive control framework for reliable microgrid energy management. *International Journal of Electrical Power and Energy Systems*, 61:399–409, 2014.

[105] M. Qadrdan, J. Wu, N. Jenkins, and J. Ekanayake. Operating strategies for a gb integrated gas and electricity network considering the uncertainty in wind power forecasts. *IEEE Transactions on Sustainable Energy*, 5(1):128–138, 2014.

[106] J. Radosavljević, M. Jevtić, and D. Klimenta. Energy and operation management of a microgrid using particle swarm optimization. *Engineering Optimization*, 0273(May):1–20, 2015.

[107] N. Rezaei and M. Kalantar. Stochastic frequency-security constrained energy and reserve management of an inverter interfaced islanded microgrid considering demand response programs. *International Journal of Electrical Power and Energy Systems*, 69:273–286, 2015.

[108] P. F. Ribeiro, B. K. Johnson, M. L. Crow, A. Arsoy, and Y. Liu. Energy Storage Systems for Advanced Power Applications. *Proceedings of the IEEE*, 89(12):1744–1756, 2001.

[109] B. Saravanan, S. Das, S. Sikri, and D. P. Kothari. A solution to the unit commitment problem-a review. *Frontiers in Energy*, 7(2):223–236, 2013.

[110] C. Schwaegerl and L. Tao. The Microgrids Concept. In N. Hatziargyriou, editor, *Microgrid: architectures and control*, pages 1–24. John Wiley & Sons, Ltd., 2014.

[111] W. Shi. *Energy Management in Microgrids: Algorithms and System*. PhD thesis, University of California Los Angeles, 2015.

[112] W. Shi, E. K. Lee, D. Yao, R. Huang, C. C. Chu, and R. Gadh. Evaluating microgrid management and control with an implementable energy management system. *2014 IEEE International Conference on Smart Grid Communications, SmartGridComm 2014*, pages 272–277, 2015.

[113] Z. Shuai, Y. Sun, Z. J. Shen, W. Tian, C. Tu, Y. Li, and X. Yin. Microgrid stability: Classification and a review. *Renewable and Sustainable Energy Reviews*, 58:167–179, 2016.

[114] J. Silvente, G. M. Kopanos, E. N. Pistikopoulos, and A. Espuña. A rolling horizon optimization framework for the simultaneous energy supply and demand planning in microgrids. *Applied Energy*, 155:485–501, 2015.

[115] J. Soares, M. A. Fotouhi Ghazvini, N. Borges, and Z. Vale. A stochastic model for energy resources management considering demand response in smart grids. *Electric Power Systems Research*, 143:599–610, 2017.

[116] A. Sobu and G. Wu. Dynamic optimal schedule management method for microgrid system considering forecast errors of renewable power generations. *2012 IEEE International Conference on Power System Technology, POWERCON 2012*, 2012.

[117] Sociaal-Economische Raad. Energieakkoord voor duurzame groei, 2013.

[118] W. Su and J. Wang. Energy Management Systems in Microgrid Operations. *Electricity Journal*, 25(8):45–60, 2012.

[119] W. Su, J. Wang, and J. Roh. Stochastic Energy Scheduling in Microgrids With Intermittent Renewable Energy Resources Stochastic Energy Scheduling in Microgrids With Intermittent Renewable Energy Resources. *IEEE Transactions on Smart Grid*, 5(4):1876—-1883, 2014.

[120] M. Tahanan, W. van Ackooij, A. Frangioni, and F. Lacalandra. Large-scale Unit Commitment under uncertainty. *4OR*, 13(2):115–171, jun 2015.

[121] R. Taktak and C. D. Ambrosio. An Overview on Mathematical Programming Approaches for the Deterministic Unit Commitment Problem in Hydro Valleys. *Energy Systems*, pages 1–25, 2016.

[122] X. Tan, Q. Li, and H. Wang. Advances and trends of energy storage technology in Microgrid. *International Journal of Electrical Power and Energy Systems*, 44(1):179–191, 2013.

[123] A. G. Tsikalakis and N. D. Hatziargyriou. Centralized control for optimizing microgrids operation. *IEEE Transactions on Energy Conversion*, 23(1):241–248, 2008.

[124] UCTE. Policy 1: Load-Frequency Control and Performance [C], 2009.

[125] United Nations Framework Convention on Climate Change. Adoption of the Paris Agreement, 2015.

[126] T. S. Ustun, C. Ozansoy, and A. Zayegh. Recent developments in microgrids and example cases around the world - A review. *Renewable and Sustainable Energy Reviews*, 15(8):4030–4041, 2011.

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

[128] S. Van Langen, P. Van Tol, T. Quak, and M. Van Bruggen. Profielen elektriciteit 2017, 2017.

[129] P. van Oirsouw and J. F. G. Cobben. Netten. In *Netten voor distributie van elektriciteit*. Phase to Phase, 2011.

[130] A. Varela Souto. Optimization & Energy Management of a Microgrid Based on Frequency Communications. Master's thesis, TU Delft, 2016.

[131] W. Vermeiden. Micro Grid Platform Overview, 2017.

[132] L. Wang and C. Singh. Stochastic combined heat and power dispatch based on multi-objective particle swarm optimization. *International Journal of Electrical Power & Energy Systems*, 30(3):226–234, 2008.

[133] X. Wang, A. Palazoglu, and N. H. El-Farra. Operational optimization and demand response of hybrid renewable energy systems. *Applied Energy*, 143:324–335, 2015.

[134] Z. Wu, H. Tazvinga, and X. Xia. Demand side management of photovoltaic-battery hybrid system. *Applied Energy*, 148:294–304, 2015.

[135] R. Zamora and A. K. Srivastava. Controls for microgrids with storage: Review, challenges, and research needs. *Renewable and Sustainable Energy Reviews*, 14(7):2009–2018, 2010.

[136] Y. Zhang, T. Zhang, R. Wang, Y. Liu, and B. Guo. Optimal operation of a smart residential microgrid based on model predictive control by considering uncertainties and storage impacts. *Solar Energy*, 122(June):1052–1065, 2015.

[137] B. Zhao, X. Zhang, J. Chen, C. Wang, and L. Guo. Operation optimization of standalone microgrids considering lifetime characteristics of battery energy storage system. *IEEE Transactions on Sustainable Energy*, 4(4):934–943, 2013.

[138] B. Zhu, H. Tazvinga, and X. Xia. Model Predictive Control for Energy Dispatch of a Photovoltaic-Diesel-Battery Hybrid Power System. In *Preprints of the 19th World Congress of The International Federation of Automatic Control*, pages 11135–11140, 2014.

[139] Zown. Zown.eu, 2016. Available at: `http://www.zown.eu/`.

Appendices

A. List of Abbreviations

CHP	Combined Heat and Power
DER	Distributed Energy Resource
DMS	Distribution Management System
DSM	Demand Side Management
EMS	Energy Management System
GA	Genetic Algorithm
ICT	Information and Communication Technologies
LC	Load Controller
LP	Linear Programming
MC	Microsource Controller
MGCC	MicroGrid Central Controller
MGP	MicroGrid Platform
MILP	Mixed Integer Linear Programming
MINLP	Mixed Integer Non-Linear Programming
MPC	Model Predictive Control
OSGP	Open Smart Grid Platform
PV	PhotoVoltaic
RES	Renewable Energy Resources
RTU	Remote Terminal Unit
UC	Unit Commitment
WT	Wind Turbine

B. Results summary

A summary of the results is presented on the following pages. For a discussion on these results, refer to Section 13.

model	solver	scenario	$\|T\|$	τ	mode	time (s)	time (ticks)	gap	cost	exp. cost	Δch	#adj.	#fail
none	-	Island	96	-	bat	-	-	-	€56.59	-	+6.94	29	-
none	-	Island	672	-	bat	-	-	-	€416.09	-	-3.97	226	-
none	-	Island+	96	-	bat	-	-	-	€20.13	-	+11.29	2	-
none	-	Island+	672	-	bat	-	-	-	€265.64	-	+37.51	61	-
none	-	Residential	96	-	bat	-	-	-	€233.44	-	-92.00	39	-
none	-	Residential	672	-	bat	-	-	-	€2299.01	-	-92.00	522	-
naive	GLPK	Island	96	24	bat	589.33	-	2.38 %	€56.36	€61.77	+11.28	17	0
naive	GLPK	Island+	96	24	bat	399.17	-	1.15 %	€17.47	€18.20	+5.68	5	0
naive	GLPK	Residential	96	24	bat	3.09	-	0.00 %	€234.21	€234.20	-26.69	0	0
naive	CPLEX	Island	96	96	gen	0.57	578.65	0.008%	€58.25	€59.80	+13.56	53	0
naive	CPLEX	Island	96	96	bat	0.66	597.94	0.008%	€58.00	€61.84	+13.62	16	0
naive	CPLEX	Island	96	24	bat	0.22	144.82	0.005%	€53.36	€56.65	-0.00	13	0
naive	CPLEX	Island	672	96	bat	1.76	1704.29	0.009%	€421.56	€455.33	+8.23	88	0
naive	CPLEX	Island	672	24	bat	0.29	196.78	0.007%	€422.73	€454.72	+10.81	96	0
naive	CPLEX	Island+	96	96	gen	0.59	304.40	0.003%	€17.40	€19.29	+2.71	8	0
naive	CPLEX	Island+	96	96	bat	0.48	278.04	0.003%	€17.24	€17.93	+3.29	5	0
naive	CPLEX	Island+	96	24	bat	0.22	85.64	0.001%	€18.55	€18.48	+8.36	5	0
naive	CPLEX	Island+	672	96	bat	0.34	192.63	0.002%	€255.61	€263.85	+13.43	60	0
naive	CPLEX	Island+	672	24	bat	0.15	55.03	0.001%	€251.12	€255.98	+14.01	61	0
naive	CPLEX	Residential	96	96	gen	0.07	65.80	0.006%	€230.01	€230.52	-35.93	63	0
naive	CPLEX	Residential	96	96	bat	0.07	63.13	0.006%	€230.19	€230.19	-38.60	0	0
naive	CPLEX	Residential	96	24	bat	0.03	18.46	0.002%	€242.10	€242.10	+1.47	0	0
naive	CPLEX	Residential	672	96	bat	0.05	44.73	0.003%	€2230.41	€2148.04	-92.00	197	0
naive	CPLEX	Residential	672	24	bat	0.02	13.33	0.002%	€2229.08	€2154.35	-92.00	173	0
safety	CPLEX	Island	96	96	gen	5.02	5996.85	0.010%	€56.27	€58.80	+5.75	46	0
safety	CPLEX	Island	96	96	bat	5.55	6711.12	0.010%	€51.68	€56.05	-1.91	13	0
safety	CPLEX	Island	96	24	bat	0.57	402.59	0.008%	€51.56	€54.54	-2.06	12	0
safety	CPLEX	Island	672	96	bat	20.04	25538.45	0.009%	€421.72	€466.79	+0.86	83	0
safety	CPLEX	Island	672	24	bat	0.63	610.73	0.008%	€419.82	€462.22	+0.01	88	0
safety	CPLEX	Island+	96	96	gen	0.51	588.17	0.004%	€18.82	€21.82	+6.71	8	0
safety	CPLEX	Island+	96	96	bat	0.46	548.57	0.004%	€17.85	€19.58	+7.16	5	0
safety	CPLEX	Island+	96	24	bat	0.12	102.03	0.002%	€17.88	€18.70	+9.60	4	0
safety	CPLEX	Island+	672	96	bat	1.38	1773.56	0.004%	€251.34	€265.17	+17.45	47	0
safety	CPLEX	Island+	672	24	bat	0.11	97.29	0.002%	€246.24	€258.64	+15.74	50	0
safety	CPLEX	Residential	96	96	gen	101.12	133274.90	0.033%	€236.60	€236.15	-21.85	67	0
safety	CPLEX	Residential	96	96	bat	101.10	128782.62	0.025%	€239.86	€239.86	-15.09	0	0
safety	CPLEX	Residential	96	24	bat	0.63	703.31	0.008%	€239.43	€239.43	-7.49	0	0
safety	CPLEX	Residential	672	96	bat	40.41	53531.66	0.012%	€2237.72	€2237.58	-82.80	0	0
safety	CPLEX	Residential	672	24	bat	0.50	536.65	0.005%	€2237.78	€2237.66	-74.05	0	0

model	solver	scenario	$\|T\|$	$\|S\|$	τ	l_x	time (s)	time (ticks)	gap	cost	exp. cost	Δch	#adj.	#fail
two-stage	CPLEX	Island	96	100	24	-	17.10	15360.92	0.930%	€56.65	€56.91	+10.36	1	0
two-stage	CPLEX	Island	96	200	24	-	58.41	41822.19	0.903%	€56.69	€55.72	+8.72	2	0
two-stage	CPLEX	Island	96	300	24	-	106.63	81862.66	0.917%	€56.32	€56.60	+7.58	1	0
two-stage	CPLEX	Island	672	300	24	-	163.29	114539.78	0.951%	€428.94	€438.39	+4.19	73	0
two-stage	CPLEX	Island+	96	100	24	-	19.31	18580.18	0.962%	€27.09	€27.06	+31.46	0	0
two-stage	CPLEX	Island+	96	200	24	-	78.83	59976.25	0.957%	€25.76	€25.74	+30.27	0	0
two-stage	CPLEX	Island+	96	300	24	-	206.46	148203.58	1.017%	€24.49	€24.47	+26.96	0	0
two-stage	CPLEX	Island+	672	300	24	-	154.79	119895.71	0.957%	€262.12	€260.54	+38.14	24	0
two-stage	CPLEX	Residential	96	100	24	-	1.89	1985.70	0.375%	€244.41	€244.37	-4.96	0	0
two-stage	CPLEX	Residential	96	200	24	-	6.24	5318.35	0.391%	€239.82	€239.78	-16.57	0	0
two-stage	CPLEX	Residential	96	300	24	-	11.88	11443.50	0.214%	€243.88	€243.85	-2.68	0	0
two-stage	CPLEX	Residential	672	300	24	-	8.60	9257.26	0.360%	€2247.84	€2247.59	-52.42	0	0
multi-stage	CPLEX	Island	96	300	24	l_1	416.02	230588.94	0.987%	€57.96	€57.01	+10.22	2	0
multi-stage	CPLEX	Island	96	300	24	l_2	164.08	90404.24	0.971%	€62.80	€64.16	+14.62	12	0
multi-stage	CPLEX	Island	96	300	24	l_3	500.91	263430.98	1.994%	€54.92	€55.45	+4.12	1	0
multi-stage	CPLEX	Island+	96	300	24	l_1	499.52	271007.69	1.318%	€17.26	€16.04	+4.15	1	0
multi-stage	CPLEX	Island+	96	300	24	l_2	453.56	259248.22	1.326%	€24.03	€25.10	+18.67	2	0
multi-stage	CPLEX	Island+	96	300	24	l_3	487.54	245878.79	1.987%	€17.78	€16.56	+5.43	1	0
multi-stage	CPLEX	Residential	96	300	24	l_1	8.91	6504.60	0.292%	€244.50	€244.45	+5.14	0	0
multi-stage	CPLEX	Residential	96	300	24	l_2	1.94	1423.97	0.436%	€250.73	€250.69	+11.74	0	0
multi-stage	CPLEX	Residential	96	300	24	l_3	21.02	13597.85	0.226%	€243.28	€243.23	-7.84	0	0

model	solver	scenario	$\|T\|$	$\|S\|$	τ	l_x	time (s)	time (ticks)	gap	cost	exp. cost	Δch	#adj.	#fail
naive*	CPLEX	Island	96	-	24	-	0.25	176.24	0.004%	€57.63	€60.25	+16.22	15	57
naive*	CPLEX	Island	672	-	24	-	0.31	226.23	0.007%	€414.95	€453.57	+9.88	109	312
naive*	CPLEX	Island+	96	-	24	-	0.12	90.79	0.004%	€18.47	€19.49	+8.20	5	65
naive*	CPLEX	Island+	672	-	24	-	0.09	61.93	0.002%	€249.89	€251.50	+14.02	58	381
naive*	CPLEX	Residential	96	-	24	-	0.04	17.55	0.003%	€226.78	€226.78	-50.68	0	60
naive*	CPLEX	Residential	672	-	24	-	0.03	12.89	0.003%	€2230.43	€2149.11	-83.25	192	382
naive**	CPLEX	Island	96	-	24	-	0.16	103.16	0.006%	€56.68	€56.55	+14.29	7	88
naive**	CPLEX	Island	672	-	24	-	0.40	240.91	0.008%	€416.80	€443.69	+9.31	90	570
naive**	CPLEX	Island+	96	-	24	-	0.11	74.37	0.003%	€20.15	€21.03	+14.03	7	86
naive**	CPLEX	Island+	672	-	24	-	0.11	66.56	0.002%	€250.26	€246.77	+19.18	53	598
naive**	CPLEX	Residential	96	-	24	-	0.02	13.13	0.002%	€222.72	€222.75	-60.01	0	92
naive**	CPLEX	Residential	672	-	24	-	0.03	12.59	0.001%	€2237.91	€2220.65	-83.25	43	605
safety*	CPLEX	Island	96	-	24	-	0.64	524.85	0.008%	€53.11	€55.07	+0.79	9	57
safety*	CPLEX	Island	672	-	24	-	0.99	816.69	0.009%	€415.46	€465.09	+2.77	88	312
safety*	CPLEX	Island+	96	-	24	-	0.15	100.82	0.002%	€18.93	€19.98	+10.75	4	65
safety*	CPLEX	Island+	672	-	24	-	0.20	126.61	0.003%	€245.63	€257.89	+16.10	49	381
safety*	CPLEX	Residential	96	-	24	-	1.08	796.36	0.008%	€233.34	€233.33	-28.06	0	60
safety*	CPLEX	Residential	672	-	24	-	0.80	593.89	0.006%	€2239.15	€2237.19	-73.96	13	382
safety**	CPLEX	Island	96	-	24	-	0.72	612.57	0.009%	€55.46	€57.10	+10.00	12	82
safety**	CPLEX	Island	672	-	24	-	1.08	969.34	0.009%	€417.33	€456.91	+4.41	98	563
safety**	CPLEX	Island+	96	-	24	-	0.16	112.90	0.002%	€21.61	€23.14	+15.79	6	87
safety**	CPLEX	Island+	672	-	24	-	0.19	104.36	0.003%	€248.15	€252.87	+14.79	47	606
safety**	CPLEX	Residential	96	-	24	-	0.52	465.23	0.008%	€231.74	€231.76	-30.12	0	92
safety**	CPLEX	Residential	672	-	24	-	0.67	585.69	0.008%	€2246.30	€2233.48	-71.25	6	640
multi-stage*	CPLEX	Island	96	300	24	l_3	531.18	274058.15	2.367%	€59.17	€58.64	+12.12	4	50
multi-stage*	CPLEX	Island	672	300	24	l_3	526.87	275213.52	7.852%	€438.00	€452.18	+11.97	71	391
multi-stage*	CPLEX	Island+	96	300	24	l_3	501.16	254372.72	1.861%	€19.20	€20.11	+8.52	2	65
multi-stage*	CPLEX	Island+	672	300	24	l_3	418.74	216911.97	2.002%	€255.28	€256.78	+7.99	30	473
multi-stage*	CPLEX	Residential	96	300	24	l_3	20.09	13782.87	0.163%	€234.58	€234.52	-36.62	0	74
multi-stage*	CPLEX	Residential	672	300	24	l_3	17.61	12030.12	0.114%	€2248.62	€2248.28	-57.52	0	517