

Planning in Smart Grids

Maurice Bosman

Members of the dissertation committee:

Prof. dr.	J.L. Hurink	University of Twente (promotor)
Prof. dr. ir.	G.J.M. Smit	University of Twente (promotor)
Dr. ir.	B. Claessens	VITO
Prof. dr. ir.	J.A. La Poutré	Utrecht University
Prof. dr.	A. Bagchi	University of Twente
Prof. dr.	J.C. van de Pol	University of Twente
Prof. dr.	M. Uetz	University of Twente
Prof. dr. ir.	A.J. Mouthaan	University of Twente (chairman and secretary)



This research has been funded by Essent, GasTerra and Technology Foundation STW, in the SFEER project (07937).

CTIT Ph.D. thesis series No. 11-226
Centre for Telematics and Information Technology
University of Twente, P.O.Box 217, NL-7500 AE Enschede

Copyright © 2012 by Maurice Bosman, Enschede, The Netherlands.

All rights reserved. No part of this book may be reproduced or transmitted, in any form or by any means, electronic or mechanical, including photocopying, microfilming, and recording, or by any information storage or retrieval system, without prior written permission of the author.

Typeset with \LaTeX .

This thesis was printed by Gildeprint, The Netherlands.

ISBN 978-90-365-3386-7
ISSN 1381-3617
DOI 10.3990/1.9789036533867

PLANNING IN SMART GRIDS

PROEFSCHRIFT

ter verkrijging van
de graad van doctor aan de Universiteit Twente,
op gezag van de rector magnificus,
prof. dr. H. Brinksma,
volgens besluit van het College voor Promoties
in het openbaar te verdedigen
op donderdag 5 juli 2012 om 14.45 uur

door

Maurice Gerardus Clemens Bosman

geboren op 2 november 1983
te Eindhoven

Dit proefschrift is goedgekeurd door:

Prof. dr. J.L. Hurink (promotor)

Prof. dr. ir. G.J.M. Smit (promotor)

ABSTRACT

The electricity supply chain is changing, due to increasing awareness for sustainability and an improved energy efficiency. The traditional infrastructure where demand is supplied by centralized generation is subject to a transition towards a Smart Grid. In this Smart Grid, sustainable generation from renewable sources is accompanied by controllable distributed generation, distributed storage and demand side load management for intelligent electricity consumption. The transmission and distribution grid have to deal with increasing fluctuations in demand and supply. Since realtime balance between demand and supply is crucial in the electricity network, this increasing variability is undesirable.

Monitoring and controlling/managing this infrastructure increasingly depends on the ability to control distributed appliances for generation, consumption and storage. In the development of control methodologies, mathematical support, which consists of predicting demand, solving planning problems and controlling the Smart Grid in realtime, is of importance. In this thesis we study planning problems which are related to the Unit Commitment Problem: for a set of generators it has to be decided when and how much electricity to produce to match a certain demand over a time horizon. The planning problems that we formulate are part of a control methodology for Smart Grids, called TRIANA, that is developed at the University of Twente.

In a first part, we introduce a planning problem (the microCHP planning problem), that considers a set of distributed electricity generators, combined into a Virtual Power Plant. A Virtual Power Plant uses many small electricity generating appliances to create one large, virtual and controllable power plant. In our setting, these distributed generators are microCHP appliances, generating Combined Heat and Power on a domestic scale. Combined with the use of a heat buffer, operational flexibility in supplying the local heat demand is created, which can be used in the planning process, to decide when to generate electricity (which is coupled to the generation of heat). The power output of a microCHP is completely determined by the decision to generate or not.

The microCHP planning problem combines operational dependencies in sequential discrete time intervals with dependencies between different generators in a single time interval, and searches for a combined electricity output that matches a desired form. To illustrate the complexity of this problem, we prove that the microCHP planning problem is \mathcal{NP} -complete in the strong sense. We model the

microCHP planning problem by an Integer Linear Programming formulation and a basic dynamic programming formulation. When we use these formulations to solve small problem instances, the computational times show that practical instance sizes cannot be solved to optimality. This, in combination with the complexity result, shows the need for developing heuristic solution approaches. Based on the dynamic programming formulation a local search method is given that uses dynamic programs for single microCHP appliances, and searches the state space of operational patterns for these individual appliances. Also, approximate dynamic programming is proposed as a solution to deal with the exponential state space. Finally, a column generation-like technique is introduced, that divides the problem in different subproblems for finding operational patterns for individual microCHPs and for combining individual patterns to solve the original problem. This technique shows the most promising results to solve a scalable Virtual Power Plant.

To apply the microCHP planning problem in a realistic setting, the planned total output of the Virtual Power Plant is offered to an electricity market and controlled in realtime. For a day ahead electricity market, we propose stepwise bid functions, which the operator of a Virtual Power Plant can use in two different auction mechanisms. Based on the probability distribution of the market clearing price, we give lower bounds on the expected profit that a Virtual Power Plant can make. To control in realtime the operation of the Virtual Power Plant in the TRIANA approach, the planning is based on a heat demand prediction. It has been shown that deviations from this prediction can be ‘absorbed’ in realtime. In addition to that, we discuss the relation between operational freedom and reserve capacity in heat buffers, to be able to compensate for demand uncertainty.

As a second planning problem, we integrate the microCHP planning problem with distributed storage and demand side load management, in the classical framework of the Unit Commitment Problem. In this general energy planning problem we give a mathematical description of the main controllable appliances in the Smart Grid. The column generation technique is generalized to solve the general energy planning problem, using the real-world electricity infrastructure as building blocks in a hierarchical structure. Case studies show the practical applicability of the developed method towards an implementation in a real-world setting.

SAMENVATTING

De elektriciteitsvoorziening is aan verandering onderhevig door een toenemende bewustwording van duurzaamheid en een verhoging van de energie-efficiëntie. De traditionele infrastructuur die ingericht is om lokale vraag centraal te bedienen, ondergaat een transitie richting een Intelligent Net (Smart Grid). Dit Intelligente Net ondersteunt duurzame opwekking uit hernieuwbare bronnen en krijgt te maken met bestuurbare decentrale opwekking, decentrale opslag en decentrale consumptiemogelijkheden die slim beheerst kunnen worden. De transmissie- en distributienetwerken krijgen hierdoor te maken met toenemende fluctuaties in de vraag naar en het aanbod van elektriciteit. Deze toenemende variabiliteit is ongewenst, aangezien in de elektriciteitsvoorziening een continue balans tussen vraag en aanbod dient te worden behouden.

Het monitoren en beheersen van deze infrastructuur hangt in toenemende mate af van het vermogen om decentrale opwekking, opslag en consumptie te kunnen sturen. In de ontwikkeling van beheers- en regelmethodologieën speelt de wiskunde een belangrijke rol, in het voorspellen van vraag, het oplossen van planningsproblemen en het realtime aansturen van het Intelligente Net. Dit proefschrift behandelt planningsproblemen. In de context van het Intelligente Net zijn deze planningsproblemen verwant aan het Unit Commitment Problem, dat bestaat uit een verzameling generatoren waarvoor beslissingen voor iedere generator genomen dienen te worden: wanneer en hoeveel elektriciteit moet een generator opwekken zodat een zeker vraagprofiel over een tijdshorizon bediend kan worden. De planningsproblemen in dit proefschrift zijn onderdeel van een beheers- en regelmethodologie voor Intelligente Netten genaamd TRIANA, die is ontwikkeld aan de Universiteit Twente.

Allereerst wordt een planningsprobleem geïntroduceerd (het microWKK planningsprobleem) dat een verzameling elektriciteitsopwekkers beschouwt, die verenigd zijn in een Virtuele Elektricitieitscentrale. Een Virtuele Elektricitieitscentrale bestaat uit een grote groep kleinschalige elektriciteitsopwekkers, zodanig dat een grote virtuele en bestuurbare centrale wordt gevormd. De generatoren die wij bekijken zijn microWKK (Warmte Kracht Koppeling) installaties, die op een huishoudelijk niveau warmte en elektriciteit gecombineerd opwekken. Het niveau van warmte- en elektriciteitsgeneratie is volledig vastgelegd door de beslissing om te produceren of niet. Door toevoeging van een warmtebuffer wordt flexibiliteit gecreëerd in de planningsmogelijkheden om aan de lokale warmtevraag te voldoen, waar-

door er operationele vrijheid ontstaat voor de beslissing om - aan warmteproductie gekoppelde - elektriciteit te produceren.

Het microWKK planningsprobleem combineert operationele afhankelijkheid voor individuele installaties in opeenvolgende discrete tijdsintervallen met afhankelijkheid tussen installaties in enkelvoudige tijdsintervallen, en vraagt naar een gecombineerde elektriciteitsopwekking die overeenkomt met een gewenst profiel. In het kader van complexiteitstheorie wordt \mathcal{NP} -volledigheid van dit probleem bewezen. Door het microWKK planningsprobleem te modelleren als geheeltallig lineair probleem of via een structuur die gebruik maakt van dynamisch programmeren, worden pogingen beschreven om praktijkvoorbeelden optimaal op te lossen. Naast het gevonden complexiteitsresultaat tonen de benodigde rekentijden voor het optimaal oplossen van deze praktijkinstanties aan dat een oplossing voor dit planningsprobleem gevonden moet worden in een heuristiek. Een eerste heuristiek is gebaseerd op de exacte aanpak die gebruik maakt van dynamisch programmeren. Deze methode lost de operationele planning op voor individuele microWKKs (in een relatief kleine toestandsruimte per microWKK) en doorzoekt de oorspronkelijke toestandsruimte door kunstmatige prijssignalen aan te passen voor deze individuele problemen. Een tweede methode benadert de bijdrage van de toestandsovergangen in de volledige toestandsruimte en stuurt deze toestandsovergangen bij naargelang de uitkomst van de planning. Ten slotte wordt een methode voorgesteld die ideeën overneemt uit kolomgeneratie, waarin het planningsprobleem wordt opgedeeld in verschillende deelproblemen voor het vinden van beslissingspatronen voor individuele microWKKs en voor het combineren van zulke patronen om het oorspronkelijke probleem op te lossen. Deze methode geeft veelbelovende resultaten om een schaalbare Virtuele Elektriciteitscentrale te kunnen plannen.

In de praktijk zal een Virtuele Elektriciteitscentrale ook moeten acteren op een elektriciteitsmarkt en is op basis van de gemaakte planning een continue aansturing vereist. Voor een elektriciteitsmarkt waarop een dag van tevoren elektriciteit wordt verhandeld, geven wij advies voor stapsgewijze biedingsfuncties, die de exploitant van de Virtuele Elektriciteitscentrale kan gebruiken in twee verschillende veilingmechanismen. Gebaseerd op de kansverdeling van de marktprijs geven we ondergrenzen voor de verwachte winst die een Virtuele Elektriciteitscentrale kan maken. De TRIANA aanpak kiest voor een samenwerking tussen voorspelling, planning en continue aansturing. Afwijking ten opzichte van de voorspelling kan grotendeels worden opgevangen in de continue aansturing. Daarnaast maken we onderscheid tussen het deel van de warmtebuffer dat gebruikt wordt in de planningsfase en de reservecapaciteit die gebruikt wordt om afwijkingen van de voorspelling op te vangen, zodat bijsturing in de praktijk vermeden kan worden.

Een tweede planningsprobleem integreert het microWKK planningsprobleem met andere vormen van decentrale opwekking, opslag en consumptie in het klassieke raamwerk van het Unit Commitment Problem. Dit algemene energie-planningsprobleem geeft een wiskundige beschrijving van de combinatie van de belangrijkste beheersbare decentrale elementen in het Intelligente Net. De kolomgeneratie methode wordt gegeneraliseerd naar het algemene energie-planningsprobleem, welke gebruik maakt van de hiërarchische infrastructuur van de elektriciteitsvoor-

ziening om een methode op te bouwen die schaalbaar is. Onderzoeksvoorbeelden tonen aan dat de ontwikkelde methode praktisch toepasbaar is richting een implementatie in het bestaande netwerk.

DANKWOORD

Normaal gesproken komt het toetje pas na het hoofdgerecht. In dit geval echter vind ik het gepast om met een dankwoord te beginnen, dat u in staat stelt om de - schitterende - context te bepalen waarin dit proefschrift tot stand is gekomen. Daarnaast bespaart het sommigen de moeite om het gehele boekwerk door te bladeren op zoek naar het dankwoord.

Allereerst wil ik uiteraard mijn promotoren bedanken, Johann Hurink en Gerard Smit. Johann weet als geen ander het onderzoek op een prettige manier in de juiste richting te sturen. Zijn commentaar, hoewel kalligrafisch niet erg hoogstaand, heeft mij erg geholpen om mijn tekst inhoudelijk te verbeteren. Door Gerard ben ik met de vakgroep CAES in aanraking gekomen. De onuitputtelijke stroom afstudeerders die binnen de vakgroep blijft om te promoveren toont aan dat zowel het onderzoek als de sfeer binnen de vakgroep uitstekend is. Iedereen binnen de vakgroepen CAES en DWMP ben ik dankbaar voor de afgelopen vier jaar; vakgroepuitjes, beachhandbal, competitieve EK-, WK- en Tourpools, potjes Go in de koffiepauze, zaalvoetbal en stukjes cabaret op promotiefeestjes, teveel om op te noemen.

U ziet het, een onderzoeker heeft een druk bestaan. Gelukkig is er ook nog tijd voor afwisseling in de vorm van onderwijs en afstudeerbegeleiding. Het is met name leuk om zowel Master- als Bachelorstudenten te mogen begeleiden bij hun eindopdrachten; het hele proces is vaak een feest der herkenning.

Het onderzoek zelf valt binnen een relatief nieuw onderzoeksgebied - voor promovendi en hun begeleiders. Relatief nieuw, want mijn voorgangers/collega's Albert en Vincent hebben in een korte tijd het kennisniveau op gebied van Smart Grids binnen de Universiteit Twente enorm opgeschroefd. Zonder hun harde werk was dit proefschrift niet zo uitgebreid geworden als het nu is, waarvoor dank.

De ceremoniële taak van paranimf wordt uitgevoerd door mijn broers Rob en Matthieu. Het is altijd gezellig om weer onder elkaar te zijn. Datzelfde geldt voor mijn ouders, die mij ook altijd geweldig gesteund hebben. Tenslotte, dit boekwerk had er niet heel veel anders uitgezien zonder Marinke, maar in de rest van mijn leven heeft ze al heel veel toegevoegd.

Voordat ik te sentimenteel begin te worden, wordt het tijd voor een lichte afsluiter van dit toetje, aangezien er nog genoeg zware kost zal volgen in de komende pagina's: laten we hopen dat PSV maar weer eens kampioen mag worden.

Maurice

CONTENTS

Abstract	v
Samenvatting	vii
Dankwoord	xi
Contents	xiii
List of Figures	xvi
1 Introduction	1
1.1 The electricity supply chain	2
1.1.1 The basic electricity supply chain	3
1.1.2 Electricity markets	6
1.1.3 Developments in the electricity supply chain: the emergence of the Smart Grid	10
1.2 Flexible and controllable energy infrastructure	14
1.2.1 Virtual Power Plant	14
1.3 Problem statement	15
1.4 Outline of the thesis	16
2 Contextual framework	17
2.1 Unit Commitment	18
2.1.1 Traditional Unit Commitment	18
2.1.2 Recent developments in Unit Commitment	21
2.2 A Virtual Power Plant of microCHP appliances	23
2.2.1 Existing approaches	23
2.2.2 Business case	24
2.3 A three step control methodology for decentralized energy management	25
2.3.1 Management possibilities	26
2.4 Energy flow model	30
2.5 Conclusion	34

3	The microCHP planning problem	35
3.1	Problem formulation	36
3.1.1	MicroCHP as an electricity producer	36
3.1.2	Requirements	37
3.1.3	Optimization objectives	40
3.2	Complexity	40
3.2.1	Complexity classes	41
3.2.2	3-PARTITION	51
3.2.3	Complexity of the microCHP planning problem	52
3.2.4	Optimization problems related to the microCHP planning problem	57
3.3	An Integer Linear Programming formulation	58
3.3.1	ILP formulation	59
3.3.2	Benchmark instances	64
3.3.3	ILP Results	67
3.3.4	Conclusion	68
3.4	Dynamic Programming	69
3.4.1	Basic dynamic programming	70
3.4.2	Results	74
3.4.3	Conclusion	74
3.5	Dynamic programming based local search	75
3.5.1	Separation of dimensions	76
3.5.2	Idea of the heuristic	77
3.5.3	Dynamic programming based local search method	78
3.5.4	Results	79
3.5.5	Conclusion	85
3.6	Approximate Dynamic Programming	85
3.6.1	General idea	85
3.6.2	Approximate Dynamic Programming based heuristic to solve the microCHP planning problem	89
3.6.3	Conclusion	90
3.7	Column generation	91
3.7.1	General idea	93
3.7.2	Problem formulation	94
3.7.3	Results	100
3.7.4	Lower bounds for a special type of instances of the microCHP planning problem	104
3.7.5	Conclusion	111
3.8	Conclusion	112
4	Evaluation of the microCHP planning through realtime control	113
4.1	Realtime control based on planning and prediction	114
4.2	Prediction	115
4.3	Realtime control	116
4.4	Evaluation of heat capacity reservation	118

4.5	Conclusion	120
5	Auction strategies for the day ahead electricity market	121
5.1	Auction mechanisms on the day ahead electricity market	122
5.2	A Virtual Power Plant acting on a day ahead electricity market . .	124
5.2.1	The bid vector	125
5.2.2	Price taking	126
5.2.3	Quantity outcome of the auction	127
5.2.4	Market clearing price distribution	127
5.3	Bidding strategies for uniform pricing	128
5.4	Bidding strategies for pricing as bid	130
5.4.1	Natural behaviour of the market clearing price distribution	131
5.4.2	Lower bounds on optimizing for pricing as bid	132
5.4.3	Computational results	134
5.5	Conclusion	140
6	The general energy planning problem	141
6.1	Application domain	143
6.1.1	Distributed generation	143
6.1.2	Distributed storage	145
6.1.3	Load management	147
6.2	The general energy planning problem	148
6.2.1	The Unit Commitment Problem	148
6.2.2	The general energy planning problem	150
6.3	Solution method	152
6.3.1	Hierarchical structure	152
6.3.2	Sub levels and sub problems	153
6.3.3	Phases and iterations	154
6.4	Results	156
6.4.1	Case study 1	157
6.4.2	Case study 2	167
6.5	Conclusion	176
7	Conclusion	179
7.1	Contribution of this thesis	179
7.2	Possibilities for future research	182
A	Creation of heat demand data	183
	Bibliography	185
	List of publications	195
	Refereed	195

LIST OF FIGURES

1.1	The development of the Dutch electricity production	3
1.2	The transmission grid of The Netherlands	4
1.3	The market clearing prices of the APX day ahead market for the period 22/11/2006 - 9/11/2010	8
1.4	The traded volumes of the APX day ahead market for the period 22/11/2006 - 9/11/2010	9
2.1	Classical Unit Commitment for a generation company	19
2.2	The three step approach	26
2.3	The hierarchical structure of the domestic Smart Grid	27
2.4	An energy flow model of the example of the generation company . . .	32
2.5	A model of the smart grid infrastructure	33
3.1	Electricity output of a microCHP run	38
3.2	Solution space for the microCHP planning problem	40
3.3	A feasible and an infeasible 2-opt move	46
3.4	Feasible 3-opt moves	46
3.5	Sequential construction of k -opt moves	47
3.6	Example: the capital cities of the 12 provinces of The Netherlands . . .	48
3.7	Comparison of runtimes for TSP instances	50
3.8	Two instances of 3-PARTITION	51
3.9	One of 16 feasible partitions in the given 3-PARTITION example . . .	52
3.10	An example of the output of the microCHP planning problem	53
3.11	The cluster C_a , consisting of $m(B - s(a) + 1)$ production patterns for the house corresponding to the element a of length $s(a)$	55
3.12	Production patterns in a more realistic example	56
3.13	The structure of dynamic programming by example of the Held-Karp algorithm	70
3.14	Two possible representations of decision paths until interval j	71
3.15	State changes from $(3, 13, 2)$ with corresponding costs	72
3.16	The detailed planning of a case with a different number of intervals .	84
3.17	A (partial) transition graph of a DP formulation and a sample path through this structure	87

3.18	The idea of the column generation technique applied to the microCHP planning problem	95
3.19	The calculation of the lower bound of the group planning problem . .	105
3.20	An example of a desired production pattern; a sine with amplitude 30 and period 18	107
3.21	Calculated lower bounds and solutions derived from the column generation technique, for sines with varying amplitude and period	109
3.22	Computation times related to the number of iterations for the column generation technique	110
3.23	A counterexample for the natural fleet bounds	111
4.1	The necessary buffer reserve capacity for different values of MAPE and MPE for a planning using 24 intervals	119
4.2	The necessary buffer reserve capacity for different values of MAPE and MPE for a planning using 48 intervals (hourly prediction!)	120
5.1	An example of supply/demand curves	123
5.2	A price/supply curve for one hour on the day ahead market	126
5.3	The acceptance rate of single bids whose hourly price is based on the hourly price of the previous day	128
5.4	Graphical representations of the difference between uniform pricing and pricing as bid	131
5.5	The behaviour of a_t for different values of γ	135
5.5	The behaviour of a_t for different values of γ (continued)	136
5.5	The behaviour of a_t for different values of γ (continued)	137
5.6	The lower bound for different values of γ and T_{\max}	138
5.7	Evaluation of constructed bids for different history lengths	139
6.1	The hierarchical structure of the general energy planning problem . .	153
6.2	The general energy planning problem	153
6.3	The division into master and sub problems	155
6.4	The operational cost functions of the power plants	158
6.5	The solution of the UCP	162
6.6	The mismatch during the column generation for the four use cases . .	165
6.7	The second use case in more detail (for a legend, see Figure 6.5)	166
6.8	Comparison of rough planning and final found solution for the planning of the local generators in the second use case	166
6.9	Operational costs related to additional electricity consumption	174
6.10	Comparison of rough planning and final found solution for the planning using 10 small power plants, 3000 microCHPs, 2000 heat pumps, 1000 electrical cars, 5000 freezers and 5000 batteries	177

INTRODUCTION

It is hard to imagine a world without electricity. In the current organization structure of our society electricity plays a key role in communication, lifestyle, security, transportation, industry, health care, food production; in fact almost any aspect of society makes use of electricity. In this context we may state that the availability of electricity enabled the world population to grow towards the current size. Moreover, it is not unrealistic to state that the high standard of living cannot be kept when the electricity system collapses. Reliability of electricity supply is therefore extremely important.

To offer a reliable and stable electricity supply an enormous infrastructure is used, which takes care of the transmission and distribution of electricity from the production side to the consumption side. Different measures are taken to secure this infrastructure from local disruptions in the system. These measures include technical equipment to disconnect failing parts of the electricity grid and control mechanisms that can adapt to changing demand with respect to these kinds of disruptions in the system. To this end it is necessary to have backup (spinning reserve) capacity at hand at all times. Furthermore, different electricity markets exist and offer organizational structures for supply and demand matching, including spinning reserve capacity. This emphasizes the realtime nature of electricity supply: electricity demand needs to be supplied instantly.

The electrical energy originates from different energy resources. These energy resources are divided into two groups: depletable energy sources and renewable energy sources. Examples of depletable sources are fossil fuels (e.g. gas, coal and oil), where wind, sun and water are examples of renewable energy sources. Due to the ongoing global debate on sustainability and climate, a trend can be identified in the electricity supply, that shows a move from depletable energy sources towards renewable energy sources.

Next to this shift towards sustainability, the energy efficiency of the electricity production and consumption is continuously improved. The primary usage of en-

ergy resources can be decreased by improving the energy efficiency, which together with the sustainable shift helps reducing greenhouse gas emissions.

Both the sustainability shift and the search for improving energy efficiency lead to a decentralization of the electricity supply chain: an increasing amount of electricity is produced (on a smaller scale) distributed at the consumption side of the supply chain. This decentralization leads to increasing challenges for the electricity grid; as opposed to the previously occurring one-way traffic of electricity, now electricity may flow bidirectionally through the grid and comes from more dispersed sources. Also, due to the increasing amount of renewable energy sources the electricity production is subject to increasing uncertainty; renewable energy sources are not ideally suited for use as controllable production units in the electricity supply.

The above mentioned electricity generation, consumption, transmission, distribution, storage, and the management and control of these elements play an essential role in the *electricity supply chain*. This electricity supply chain is subject to many changes, that lead to the idea for an improved/adapted infrastructure: the concept of Smart Grids. It is an interesting field for developing new control and management methodologies. A control methodology that especially takes the partial decentralization of the electricity supply into account is developed at the University of Twente. This methodology is called TRIANA. The work in this thesis is part of the TRIANA methodology and especially focuses on mathematical planning problems involving decentralized generators, consumption and storage. We focus on combinatorial problems where generators cooperate in a so-called Virtual Power Plant, and on extensions of the well studied Unit Commitment Problem. In the case of the Virtual Power Plant we use the outcome of the planning problems to act on an electricity market.

In the following sections we give an extended introduction to the background of Smart Grids that underlies this thesis. We discuss the electricity supply chain in Section 1.1. Section 1.1.2 introduces the different electricity markets. The developments in the electricity supply chain are given in Section 1.1.3. Then we give the organizational structure of a Virtual Power Plant in Section 1.2. We conclude with a description of the problem statement in Section 1.3 and an outline for the rest of the thesis in Section 1.4.

1.1 THE ELECTRICITY SUPPLY CHAIN

The electricity supply chain deals with the challenge of continuously matching electricity demand with supply. In the electricity supply chain five main areas of interest can be identified:

- production (we also use the terms generation or supply)
- consumption (demand)
- transmission and distribution
- storage

- management and control.

Technological, economical and political developments lead to an interesting evolution of the classical infrastructure towards the so-called Smart Grid. In this section we sketch the basic behaviour of the electricity supply chain, and show the developments that lead to the Smart Grid.

1.1.1 THE BASIC ELECTRICITY SUPPLY CHAIN

We start by giving a general overview of the basic principles by which electricity is produced and delivered to the customer. The actors in the different areas are identified and the interaction between them is sketched.




Figure 1.1: The development of the Dutch electricity production

The growth of the electricity production in The Netherlands is depicted in Figure 1.1. This data is derived from [3]. In the classical infrastructure, this production is mostly given by the electricity generation from centrally organized power plants that are connected directly to the transmission grid. Examples of traditional power plants are gas-, coal- or oil-fired power plants or nuclear power plants. These generation plants differ in size: their capacity ranges from tens/hundreds of MW for the largest part of these generators, up to more than 1 GW for some very large plants. An increasing amount of electricity production is directly connected to the distribution grid, as Figure 1.1 shows. Opposite to most common supply chains, electricity has to be instantly supplied, whenever demand occurs. An important

are used to change the voltage level for the different connections. The distribution grid is connected to the transmission grid and is operated by Distribution System Operators (DSOs). In The Netherlands there are 9 DSOs. Where a TSO is responsible for large-scale electricity transmission, a DSO is responsible for the final part in the electricity supply chain, i.e. the delivery towards the customer. It uses medium voltage lines of 50 kV and 10 kV, and transforms the voltage level eventually to the 230 V that is currently used in The Netherlands at the consumption side. TSOs and DSOs are monopolists in their respective areas. Therefore, they are bounded by regulations set by governmental authorities (e.g. *Energiekamer* in The Netherlands) with respect to price setting for transporting electricity.

At the consumption side, stability and reliability are essential elements in the electricity supply chain. Reliability deals with the availability of the connection to the grid. Since the society depends heavily on electricity, the reliability of the grid should be large. In The Netherlands the reliability is very large; on average there is an interruption in the electricity supply of half an hour per year per household connection (23 minutes in 2011 [10]), which comes down to a reliability of 99.996%. This reliability is higher than in Germany (40 minutes), France (70 minutes) and the UK (90 minutes). Next to reliability, stability of the electricity supply is also important. Stability is the ability to keep the electricity supply at 230 V and 50 Hz (from a household perspective). Deviations from these values may lead to severe reductions in the lifetime of electronic equipment or even to defective equipment. Consumers, with the focus on domestic consumption in particular, pay for their consumption as well as for their connection, via contracts with an electricity retailer. Currently the electricity prices are determined by the retailer for a given time period (in the order of months), either at a constant rate or based on the time of use (e.g. a day/night tariff).

Storage of electricity is not applied at a large scale, due to efficiency losses and economical costs of storage systems. Therefore, the challenge in the electricity supply chain is to continuously find a match between consumption and production. To find this match the different actors within the electricity supply chain need to exchange information. However, their acting is driven by their own objectives. Electricity retailers can make fairly good predictions of the consumption of their consumers. Before the liberalization of the electricity market, which was finalized in the year 2004, these retailers were often also active at the production side by owning power plants. Currently a strict separation between retailing and producing actors is demanded, such that the market is more transparent. Production companies want to optimize their energy production, considering fuel costs, maintenance costs, revenue, etcetera. This leads to the situation that demand (in the form of electricity retailers) and supply (generation companies) are settled on an electricity market and cleared for a certain price. Note, that there are many forms of electricity markets, resulting in a dispersed settlement with a possible range of prices for each moment in time. TSOs/DSOs have the responsibility to secure a stable grid all the time. When the market actors operate exactly as they have settled by using the available market mechanisms on beforehand, demand and supply are balanced and stability measures by the TSO/DSOs are not required. However, the process of

electricity production and consumption is subject to uncertainty, which often leads to an imbalance in the supply chain. If such an imbalance occurs, a TSO has the ability to correct this imbalance by coordinating the increase/decrease of electricity generation. To this end, a reserve capacity is always standby. Moreover, the actors that are causing the imbalance are penalized.

1.1.2 ELECTRICITY MARKETS

As of July 1, 2004 the energy market was completely liberalized and consumers were able to choose their electricity and gas suppliers. From a supplier point of view this means that the supplier needs to offer a high quality of service to the consumer. In an ideal world this would mean that there is a full competition between energy suppliers (retailers). In practice, the liberalization led to an increase of the number of retailers. However, it is concluded in [26] that market entry is still difficult for small entities, since governmental regulations limit the way the electricity retailers may act. For that matter, these governmental regulations are intended to protect the consumer and recover/keep the confidence in the market. [106] shows that in practice consumers do not switch between retailers easily; [111] reports on increasing switches between retailers, but simultaneously reports that the three largest retailers in The Netherlands (i.e. Essent (RWE), Eneco and Nuon (Vattenfall)) still have a market share of 80.6% in July 2010.

Electricity retailers have contracts with their consumers to deliver electricity against a prescribed pricing system. To be able to really deliver the electricity, these retailers predict the consumption of their consumers and buy the corresponding amounts on the electricity market. In that way, their performance on the market determines to a large extent the profit they can make.

From the production point of view, generators are more and more subject to market competition. Generation companies need to actively bid their production capabilities on electricity markets. This enlarges the importance of minimizing operational costs, due to the fact that profit margins are under pressure.

Production and consumption meet at the electricity market [39, 118, 125]. The electricity market consists of different markets, based on the duration of the contract and the way in which the contract is realized. We differentiate between long/medium term markets and short term markets.

Long and medium term markets

Since energy balance is crucial in electricity markets, a good prediction of demand versus the available supply is necessary. A large share of the energy demand is very predictable, which implies that a large part of electricity can be traded at long term markets. For these amounts, electricity contracts are signed between electricity producers and retailers, up to three years in advance. These long term contracts are often agreed in a bilateral way [74], meaning that a single producer (power plant) and a single consumer (retailer) close a deal between each other. Standardized contracts are also available, to a smaller extent.

As the day of delivery comes closer, more electricity is traded in medium term contracts (months in advance), as the prediction of the demand gets more accurate. Again, most of these contracts are bilateral.

Short term markets

To smoothen the rough profile of demand/supply amounts that are already settled via long term contracts, short term markets are used to exchange the final amounts of electricity via standardized trading blocks, day ahead markets, intraday markets and balancing markets.

In general, the prices on the day ahead market and balancing market are higher than on the long term market, due to relative inelastic demand. On the day ahead market electricity is traded in 24 hourly blocks, which are cleared a day in advance. Based on the latest predictions [17, 44], the last portion of the electricity profile is traded. This market is open for many demand and supply participants and is cleared by the market operator.

On the day of delivery, electricity can be traded on the intraday market. On this market, recently developments related to disturbances in demand or supply are settled by retailers and generation firms. The intraday market is organized by bilateral contracts (e.g. the APX intraday market) or standardized blocks (e.g. the APX strip market) [20]. The balancing market is a realtime market, in which realtime deviations from agreed long and short term contracts are settled by the TSO. In case demand differs from the predictions, or in case settled generation cannot be delivered, an imbalance occurs in the electricity network. This imbalance needs to be repaired to guarantee stability and reliability in the grid. Therefore the balancing market is a place where ancillary services as spinning reserve and congestion management are offered. Spinning reserve consists of the ability of generators to generate additional amounts of electricity when the TSO asks for it, to match balance disturbances. A generator gets paid for offering this ability, even if it does not have to produce electricity at all. Congestion management consists of a means for the TSO to manage loads that are exceeding the capacity of the network, which attracts more and more attention, due to recent developments towards the decentralization of the electricity supply chain. In this case the TSO can ask some generators to produce less electricity, and ask generators in a different part of the network to overtake this load, such that balance is preserved or network constraints are met.

Note, that in the literature often the term spot market is mentioned. However, it is used for both the day ahead market as well as for the balancing market. To avoid confusion between these terms, we stick to the terms day ahead market and balancing market.

The electricity market of The Netherlands

Since the day ahead market is a market that gets centrally cleared and is open to competition between different demand and supply participants, it is an interesting

market to study in more detail. In this thesis we focus on the electricity markets of The Netherlands [2]. The Amsterdam Power Exchange (APX) is a central market where electricity and gas is traded between market participants in The Netherlands and surrounding countries. The APX is established in 1999 as part of the liberalization of the electricity market. Currently the Dutch market is coupled to the markets of surrounding countries, which enables an interaction between the different markets.

To get some feeling for the prices on the day ahead market on the APX, we collected data from November 22, 2006 until November 9, 2010. Figure 1.3 shows

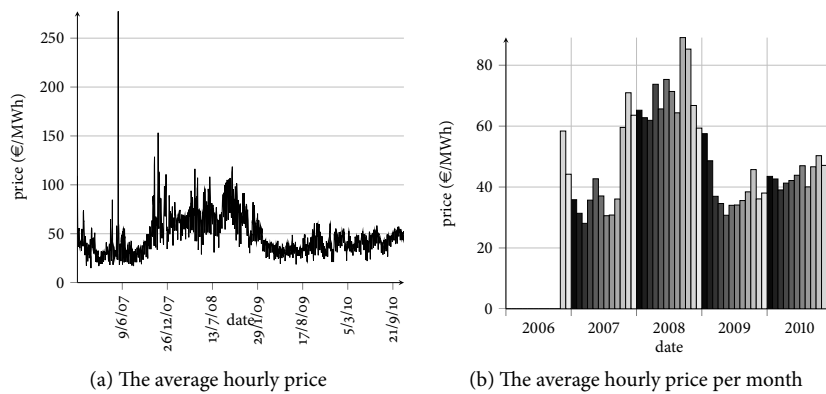


Figure 1.3: The market clearing prices of the APX day ahead market for the period 22/11/2006 - 9/11/2010

the development of the market clearing price on the day ahead market. In Figure 1.3a the average hourly price is depicted for each day. The average price is 48.87 €/MWh for the complete time horizon, with a minimum daily average of 14.83 €/MWh and a maximum of 277.41 €/MWh. In general no real trend in the development of the electricity prices can be found, other than that prices stabilize after a temporary peak in 2008. The average hourly price per month in Figure 1.3b filters daily peaks and shows the high prices in 2008 more clearly.

Figure 1.4 shows the development of the traded volumes during the same time horizon. Over the complete horizon, the average hourly volume is 3012.86 MWh, with a minimum of 1039.0 MWh and a maximum of 6744.8 MWh. The figure shows that an increasing amount of electricity is traded on the day ahead market in The Netherlands. In 2007 the market share of the (short term) day ahead market was 19.7% of the total generated electricity; in 2010 already 28.1% was traded on a daily basis.

Market power

This increase in market share of short term electricity markets is also reflected in the extensive literature that is available on market participation and market power.

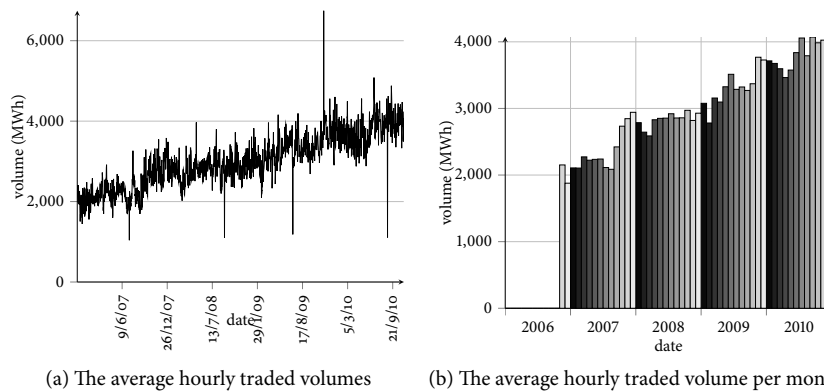


Figure 1.4: The traded volumes of the APX day ahead market for the period 22/11/2006 - 9/11/2010

Market power is the ability of single market participants (producers) to influence the market clearing price, by strategically bidding, instead of bidding their true marginal costs, which is optimal in a competitive market. The work of [48] shows that strategic bidding can lead to increasing market prices; this has important implications for the design and governance of electricity markets. An example of exercising market power are given by [126], which show that on the Dutch electricity market in 2006 during many hours one or multiple producers were indispensable to serve the demand, which made them capable of setting the price. In [93] scarce availability of generation capacities result in the exercise of market power in the sense that a generation company could influence the market price by withdrawing one of its generators from the market. It shows that investments in generation could decrease market power. [36] shows that the interconnection of two markets in Italy (North and South) mitigates the market power of one large generation company, whereas [104] concludes that the integration of different markets can cause price disruptions, showing that interconnection between markets does not always lead to improvements. The work of [80] on double-sided auctions includes active bidding of retailers on different markets, which shows a decrease in market power of generation firms and results in more stable equilibria on these markets. Other incentives to reduce market power are presented in [132] and [110]; the latter prevents large generators to use market power, where the first concentrates on social welfare in market mechanisms. To illustrate that electricity markets can function well, [55] shows that there is no evidence of exercised market power in the Scandinavian Nord Pool market in different periods of time. Also, [23] presents that long term markets mitigate market power on short term spot markets.

Related to the discussion of exercising market power is the choice for an auction mechanism. Different auction mechanisms for day ahead markets are studied. We mention two of them: Uniform Pricing (UP) and Pricing as Bid (PAB). In an UP mechanism all generation companies that win the auction, get paid a uniform price,

i.e. the market clearing price. In a PAB mechanism each auction winning generation company gets paid its own offered price. The discussion between the choice for UP and PAB concentrates on the fairness of the received price and strategic behaviour of producers [25, 47, 100, 132]. In UP some generation companies with low marginal costs receive a price that is well above this cost, such that eventually consumers pay unnecessarily high prices. On the other hand, the UP mechanism gives incentives for the producers to offer electricity at their true marginal costs, while PAB gives an incentive to bid strategically.

1.1.3 DEVELOPMENTS IN THE ELECTRICITY SUPPLY CHAIN: THE EMERGENCE OF THE SMART GRID

Managing the electricity supply chain does not solely consist of traditional producers and consumers acting on the electricity market and awaiting the realtime control of network operators and power generation companies. Increasingly, distributed generation, distributed storage and demand side load management is applied in the electricity supply chain. This development has strong influences on the way the different areas (production, consumption, storage, transmission and distribution, and management and control) of the traditional supply chain are managed and balanced, and leads to a growing need for decentralized intelligence in the electricity supply chain and, thus, to the emergence of the Smart Grid. Distributed generation, distributed storage and demand side load management display into a massive amount of dispersed controllable appliances, for which decision making is necessary. To enable such a dispersed decision making in an electricity system, that is very dependent on balance, asks for communication and management systems. The complete infrastructure, consisting of measuring, communication and intelligence, that enables the large-scale introduction of distributed energy entities is called a Smart Grid. The key motives for the change towards a Smart Grid are improved energy efficiency and sustainability of the electricity supply. It results in a bidirectional electricity infrastructure, since the traditional consumption side now also has possibilities to produce electricity.

In the following, we shortly sketch some effects in the different areas of the electricity supply chain.

Production

In the field of production, distributed generation is increasingly applied. This generation emerges in two general types: sustainable distributed generation and energy efficiency improving generation.

Examples of - less controllable - sustainable production are wind turbines (see e.g. [21, 22, 51, 85, 129]) and solar panels (e.g. [24, 52]). The generation capacity of different types of sustainable generation is limited by the geographical environment of the local area/country. Within these geographical restrictions a lot of research focuses on location planning of wind and solar generation (e.g. [129] studies the influence of atmospheric conditions on wind power, [21] searches for good geo-

graphical locations of wind turbines and [52] combines large scale solar generation in deserts with a supergrid in Europe). Depending on each countries situation, a certain mixture of sustainable generation is desirable, which leads to a specific shift towards renewable energy for each country. In general, this shift towards renewable energy brings along more fluctuating and less controllable generation. To allow a large share of sustainable generation, advanced control methodologies are therefore necessary to reduce the fluctuation. An example of such a control system is the integration of wind turbines and Compressed Air Energy Storage (CAES) [22], to reduce fluctuations in generation. Realtime excess or slack of energy is captured by controlling the air pressure in large caves, which allows storage of large amounts of energy.

An example of energy efficiency improving generation that is controllable is Combined Heat and Power (CHP). Such controllable generation is also the focus of this thesis. Although research is performed on different possibilities for small-scale CHP (25 - 200 kW) [16], we limit ourselves to CHP with output at the kW level on a domestic scale (microCHP). An initial summary of the potential for microCHP in the USA is given by [122]. The study of [50] concludes that a reduction of 6 to 10 Mton of CO_2 is possible in the year 2050 by applying microCHP in the built environment; [103] concludes that CO_2 savings between 9% and 16% for 1 kW microCHPs are possible, which offers a significant reduction compared to other possible domestic measures. A microCHP produces both heat and electricity for household usage at the kW level; the electricity can be delivered back to the electricity grid or consumed locally. The control of the microCHP is heat led, meaning that the heat demand of the building defines the possible production of heat and, simultaneously, the possible electricity output. Combined with a heat buffer, the production of heat and electricity can be decoupled and an operator has flexibility in the times that the microCHP is producing, which creates a certain degree of freedom in electricity production. This operational freedom gives us flexibility in control. Realtime operating strategies, showing the potential of control, are given in [37, 61, 69].

The output of a single distributed generator is in general much smaller than that of common power plants. Wind turbines generate in the order of MW, microCHPs in the order of kW. However, the total potential is large when applied on a large scale.

Consumption

At the consumption side of the electricity supply chain, developments in domestic consumer appliances lead to more flexibility in local control. For example, Heating, Ventilating and Air Conditioning (HVAC) systems offer large possibilities in managing electricity consumption [128]. Controllable washing machines, dryers, fridges and freezers add up to about 50% of the total electricity demand of a household [35]. Also heat pumps [75] are introduced to supply domestic heat demand, by transferring energy from the soil or the outside air.

This development means that the total load profile of a household gives room for adjustment by a control system, as opposed to the traditional uncontrollable consumption. Such control systems are referred to as demand side (load) management. Next to this controllability, there is the possibility to improve the energy efficiency of consumer appliances. In this context, consumer awareness is an important factor. The awareness of class labels during the purchase of energy efficient appliances is increasing, but, as in many other fields, it is still mostly money driven [92]. The paper of [99] analyzes the effect of policies on the consumer behaviour that can lead to both energy saving and an increase in energy efficiency. They show that self-monitoring can be a good option to increase awareness and thus aim for energy saving behaviour and that financial compensation for the relative high threshold for taking action towards energy saving behaviour has a better effect than taxing individuals for their energy usage.

Storage

Electricity storage is in principle the most helpful tool to control balance in the electricity supply chain. The temporary fluctuations in demand and supply can be managed much easier, when large buffers are available to put excess energy in and to withdraw energy from when there is additional need for energy. So far however, it is not used at a large scale. This is mostly due to its relatively high costs, in combination with efficiency losses and life time cycles. New storage techniques are emerging though. At a domestic scale, electricity storage can be combined with a power supply system as in [9]. The emergence of the electrical car brings along the opportunity to use the battery as a storage device, rather than only charging the battery, when the car is parked. Since on any time, 83% of the cars in California are parked, even during commuting hours [76], this gives the opportunity to form a Vehicle to Grid system, which could help the voltage/frequency control in the grid [71, 76]. At a larger scale, CAES can help control the fluctuation of wind parks, as well as pumped hydro-electric energy storage (the possibilities to exploit both systems in Colorado are described in [86]).

Transmission and Distribution

The increased flexibility in the generation of electricity and in the usage of controllable consumer appliances and storage, may have effects on the transmission and, in particular, the distribution grid. The bidirectional electricity flow gives both an increased attention towards load and congestion management and may ask for technical improvements in the infrastructure (e.g. a smart metering infrastructure has to be clearly defined and implemented).

On a nationwide scale, the interconnection between countries is developing. An example is the NordNed cable between Norway and The Netherlands [57]. [78] shows a smart MV/LV-station that improves power quality, reliability and substation load profile. It anticipates on the smart grid and bidirectional electricity flow. The work presented in [124] is oriented to maximize the amount of local

generation capacity while respecting the load limitations of the distribution network, whereas [59] demonstrate a software tool for alternative distribution network design strategies.

Management and control

As mentioned before, the introduction of distributed (sustainable) generation and the increased use of intelligent consumption and storage devices, demands for advanced energy monitoring and control. The introduction of smart metering is a first step towards intelligent control. Realtime load balancing and congestion management in distribution networks are mentioned before. A large system that is in use for years in the traditional electricity supply chain is SCADA (Supervisory Control And Data Acquisition), that, in combination with grid protection systems, secures the actual generation of electricity. In this system, human operated control rooms oversee and steer, in combination with the help of computer programs, the realtime generation. The mathematical basis of these computer programs is described in the Unit Commitment Problem. For the existing literature on Unit Commitment, we refer to Chapter 2.

The potential for Smart Grids is extensively studied. The study of [52] to create a supergrid in Europe and the northern part of Africa is already mentioned. An overview of distributed generation with a large share of renewable sources in Europe is given by [54, 123]; [121] gives an extensive analysis of the possibilities for distributed generation in Australia. For The Netherlands, [113] explains that a transition to smart grids offers many opportunities and high potential benefits for The Netherlands.

Strategic planning, regarding the location and type of generation and infrastructural possibilities, also plays a role in management systems. Different use cases of different countries, regarding strategic planning for advanced local energy planning, are studied in [72]. [97] offers modelling software for strategic decisions; a grid infrastructure can be made by selecting generators and other components (transformers, storage, etcetera) for which a global analysis is made.

Several ICT oriented methodologies are proposed to control and manage (a part of) the new Smart Grid [35, 46, 83, 84, 96], in addition to the already existing management systems that aim at dispatching generation (i.e. Unit Commitment), load balancing and congestion management. Some of these methodologies are especially focusing on specific objectives; [46] applies stochastic dynamic programming to facilitate a single generator with multiple storage possibilities, and [35] concentrates on micro energy grids for heat and electricity. The work of [84] uses a Multi Agent System (MAS) approach to manage power in an environment of hybrid power sources, based on an electrical background and thus especially focusing on electrical behaviour. From a policy point of view, [81] investigates investment policies of wind, plug-in electric vehicles, and heat storage compared to power generation investments, and studies the influence of the unreliability of wind generation. As an example of more generic energy control methodologies, we refer to [83] and [96]. The PowerMatcher of [83] proposes a MAS approach for supply and demand

matching (SDM). The TRIANA methodology of [96], of which this thesis forms a part, uses a hierarchical control structure in which, at several levels, energy supply chain problems are solved using a three step strategy: prediction, planning and realtime control.

1.2 FLEXIBLE AND CONTROLLABLE ENERGY INFRASTRUCTURE

The previous subsection shows that the request for sustainable generation and the emergence of distributed, more energy efficient, generation, storage and load side management leads to a change of the electricity supply chain towards a Smart Grid. In this context there is a substantial difference between controllable appliances (microCHP/micro gas turbines/heat pumps) and noncontrollable generation (solar/wind). To compensate for fluctuating noncontrollable generation, a certain share of generation in the complete electricity supply should be controllable and also actively controlled. A large part of this thesis focuses, from a mathematical point of view, on a specific emerging technique that can be controlled to some extent: microCHP. MicroCHP control can manifest in several ways. For example, individual control of microCHP operation can be aimed at profit maximization or cost minimization for a household. In a developing Smart Grid, a (two-way) variable pricing scheme for the use of electricity may be implemented, that in general asks a high price for the consumption of electricity during peak hours and lowers the price during baseload hours. In this case the operation is steered towards high priced hours, such that the electricity that is delivered back to the grid brings in the most money, or the demand in high priced hours is supplied locally, such that the imported electricity and its associated costs are minimized. A microCHP can also be used to provide electricity in case of blackouts (islanded operation). The last two types of control however, are not considered in this work. We focus on combined optimization of the planned operation of a large amount of microCHPs in a large-scale Energy Cooperation: a so-called Virtual Power Plant.

1.2.1 VIRTUAL POWER PLANT

A Virtual Power Plant (VPP) combines many small electricity generating appliances into the concept of one large, virtual and controllable power plant. This VPP can be comparable to a normal power plant in production size. However, the comparison ends here. Due to the geographical distribution of generators, the physical electricity production from a VPP has a complete other dimension than the production from a large generator that is located at a single site. The wide-spread distribution of generators asks for a well-controlled generation method. Instead of controlling one large generator, which has a limited number of options (i.e. not generating, generating at full power, and several decidable generation levels in between), all generators in a VPP can be individually steered. These generators must be scheduled or planned to generate at different times of the day in such a way, that the combined electricity production of all generators matches a given generation profile that resembles the production of a normal power plant.

We consider a VPP that consists of microCHP appliances. Although the steering of such a VPP is more complex than the steering of a normal power plant, the increase in energy efficiency due to the usage of both heat and electricity (95% compared to the 35%-50% of conventional power plants) shows the added value of such a VPP. The planned dispatch of generation depends on the objective of the controlling entity of the VPP. We focus on operating on the day ahead electricity market; compared to a conventional power plant the flexibility of the VPP is not deemed large enough to offer balancing capacity. In Chapter 2 additional information on the choices for our VPP are given.

1.3 PROBLEM STATEMENT

Many challenges exist in the evolving energy infrastructure. In mathematics, these challenges are usually called *problems*. We conform to this notation and use the term problem in the remainder of this thesis for the challenges we try to tackle.

Research focus

Planning problems in the energy supply chain can be divided into long term and short term problems. The long term problems are strategic decision problems, varying from location planning of power plants [73] or windmill parks [21] to portfolio selection problems [90] or long term generation contracts [74]. These problems treat the strategic planning of the production capacity of a certain stakeholder. On a shorter notice of time, the available production capacity has to be operated in an optimal way. In this thesis, we consider short term planning problems in the energy supply chain. We consider planning problems for a Virtual Power Plant, and a generalized energy planning problem with a focus on domestic, distributed generation, storage and demand side management.

The Virtual Power Plant case focuses on household sized appliances; miniCHPs and small biomass/biogas installations are not the primary focus, but they could be modelled as well in the general energy planning problem. We introduce the microCHP planning problem as the main problem for our VPP. For these small-sized microCHP appliances, *scalability* is a most demanding task. It should be possible to eventually plan the operation of millions of microCHPs. Together with scalability, we demand *feasibility* of the planned operation in two aspects. First, each individual microCHP should be operated, such that the basic heat demand in households is supplied, without harming the comfort of the consumers. Secondly, the combined electricity generation of all microCHPs has to fulfill desired bounds on the total output, either resulting from network constraints or market desires. *Limited computational capacity* is a natural requirement for both scalability and feasibility.

For the Virtual Power Plant we consider discrete planning problems and briefly sketch the influence of demand uncertainty. Furthermore a connection is laid between the ability to find a certain production output for a Virtual Power Plant by solving a planning problem and the practical problem of actually acquiring this

production profile as the settled result of an electricity market. We present a way of acting on a day ahead electricity market and discuss the influence of two market clearing mechanisms: Uniform Pricing and Pricing as Bid. In the case of the Virtual Power Plant Pricing as Bid may give an incentive to actively bid on the market, since our VPP has no operational fuel costs attached (see the definition of a business case in Chapter 2).

The general energy planning problem gives an extension of the Unit Commitment, with special attention to distributed energy appliances. This problem includes the microCHP planning problem and other types of distributed generation, distributed storage and demand side management possibilities. Since this general energy planning problem deals with different elements within the electricity supply chain, the goals for these participating elements may differ. Therefore the general energy planning problem combines multiple (possibly decentralized) objectives.

1.4 OUTLINE OF THE THESIS

In this introductory chapter an overview of the background of the electricity supply chain is given. Based on this introduction, Chapter 2 elaborates on some research areas and developments, that deserve an extended background information. In Chapter 3 the microCHP planning problem is studied in detail, and heuristics are developed to solve this problem. Chapter 4 treats the positioning of the planning problem in the TRIANA methodology, and Chapter 5 discusses a way of acting on electricity markets. A general energy planning problem is presented in Chapter 6. Conclusions and recommendations for future work are depicted in Chapter 7.

CONTEXTUAL FRAMEWORK

ABSTRACT – This chapter gives extended background information on topics that are closely related to our work. First we treat the Unit Commitment Problem, which gives the general mathematical description of the dispatch of electricity generation by a set of power plants. We also discuss recent developments in this field, which show a shift towards integrating relatively new electricity markets and a focus on stochastic influences of demand uncertainty. Secondly we give some details on Virtual Power Plants that are based on microCHP appliances and discuss a business case for such a Virtual Power Plant. Thirdly, the TRIANA 3-step control methodology for decentralized energy management, developed at the University of Twente, is presented. Fourthly, we present an energy flow model, that serves as a reference point for energy balancing.

This chapter builds upon the introduction that is given in Chapter 1. We give additional background information that further specifies the field to which the contribution of this thesis applies. First we discuss the Unit Commitment Problem. In Chapter 6 we extend this basic problem formulation by adding distributed production, storage and demand side load management. Secondly we show related work on Virtual Power Plants that consist of microCHP appliances. A business case for such a Virtual Power Plant is given, which forms the basic background for the developed planning methods and market participation within this work. Next we give an overview of the 3-step control methodology for Smart Grids (TRIANA), that embeds the planning problems that are presented in this thesis in a complete (domestic) Smart Grid management system, consisting of prediction, planning and realtime control possibilities. An energy flow model, that underlies this TRIANA methodology, is also discussed, since it gives a better understanding of the realtime balancing aspects of energy management (and electricity management in particular).

2.1 UNIT COMMITMENT

The Unit Commitment Problem (UCP) gives a mathematical formulation of an optimization problem that is related to energy generation. For literature overviews of the UCP we refer to [102, 115]. In the UCP, deterministic or stochastic energy demand has to be supplied by a number of generators. The UCP determines the commitment of specific generators during certain time windows (i.e. a binary decision whether generators are used to supply (part of) the demand or not) and determines the generation level of the committed generators in these time windows. To our knowledge the term Unit Commitment was first treated in [77].

In this section we first describe the original Unit Commitment Problem, meant to be used by a single generation company that has several generators (power plants) available. Then we describe the developments in the field of Unit Commitment that coincide with the emergence of the Smart Grid.

2.1.1 TRADITIONAL UNIT COMMITMENT

Originally, the UCP is seen as a decision support tool for a generation company. Such a generation company often used to be also the distribution system operator (DSO) and the only electricity retailer in a certain area; i.e. the generation company used to be a monopolist. The main task of this generation company simply is to supply all demand in the area. The complete demand of the area is relatively inelastic; the consuming behaviour of the area does not depend much on the electricity price (at least not in the price range that the electricity retailer is allowed to ask). Since revenues are not subject to much uncertainty (in the monopolistic case), the objective for the generation company in this case is to minimize costs that are associated with generation. An important aspect of this task is to predict the demand. High quality predictions are useful for the generation company; the more accurate the prediction is, the less adjustments are needed for the production that is planned for this prediction, and the better the cost-benefit optimization of the generation company can be planned by solving the Unit Commitment Problem. The Unit Commitment Problem (UCP) minimizes total costs (or maximizes total revenue/profit) for a set of generators, that are subject to a set of constraints on the generation. Main features of the UCP are unit commitment (the decision to actively participate in the production process of a certain time interval) and economic dispatch of committed units [28] (the decision to produce at a specific generation level in a certain time interval), whereby a large amount of possible additional requirements have to be taken into account. Several of these additional requirements deal with time: power plants have startup costs and ramp rates for example. Startup costs aim at using the same committed units for subsequent time intervals (long run periods are in general good for the energy efficiency of power plants). Ramp rates indicate the maximum increase/decrease of the generation level of power plants, reflecting that a generator that is producing at full capacity cannot always be stopped within an hour. Similarly, the full capacity cannot be immediately reached, if the generator is currently not committed.

We demonstrate the classical UCP by giving an example. This example considers a generation company as depicted in Figure 2.1. Figure 2.1a shows that this

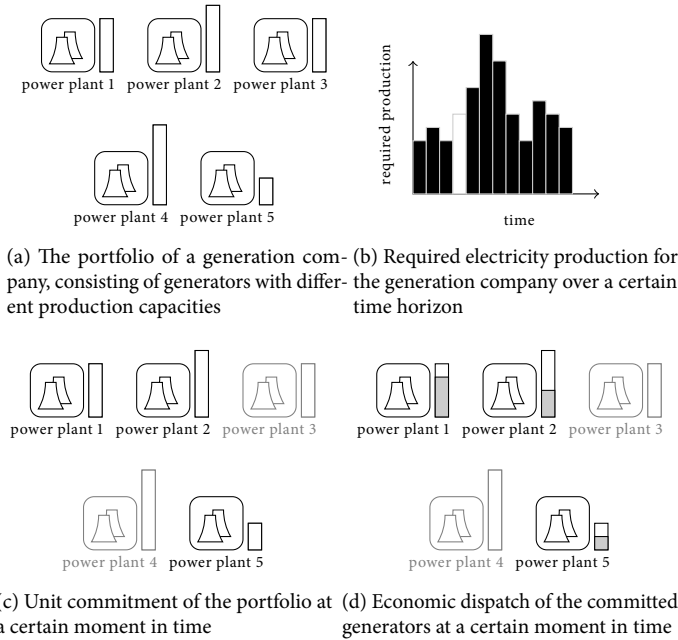


Figure 2.1: Classical Unit Commitment for a generation company

generation company consists of 5 different power plants with a different production capacity, indicated by the height of the rectangle next to each power plant. In Figure 2.1b the (predicted) demand of the area is depicted, varying over a certain time horizon. Of course, in the real world the actual demand is a continuous function. However, in the UCP and in the energy planning problems that are discussed in this thesis, the demand is aggregated over (hourly or even smaller) time intervals. In the context of the UCP this means that the planned production for the generation companies portfolio is a rough sketch for the actual production. However, note that the realtime adjustments that are needed, are relatively close to the aggregated demand and in general do not lead to severe problems for the power plants. The Unit Commitment Problem has to match the demand in each time interval by committing power plants (or generators or units) (see Figure 2.1c) and determining the generation level for committed units (see Figure 2.1d). The first task is called *unit commitment* and the second task *economic dispatch*. In the example these two subfigures show the unit commitment and the economic dispatch for the time interval that is coloured white in Figure 2.1b: we choose to commit power plant 1, 2 and 5, and produce the required 300 MWh as in Table 2.1.

Rather than looking at the numbers of the example, it is important to distinguish the important parameters and constraints in the Unit Commitment Problem that

unit commitment	demand	power plant				
		1	2	3	4	5
production/consumption (MWh)	-	1	1	0	0	1
	300	150	100	0	0	50

Table 2.1: Unit Commitment for a certain time interval

define the eventual commitment and dispatch. In the following, we list them, prioritizing on the most common ones in a descriptive way. A mathematical formulation of the UCP can be found in Chapter 6.

Cost minimization

One of the most common objectives is cost minimization. The operational costs of a power plant consist of fuel costs, maintenance costs and startup/shutdown costs. Fuel costs are often modelled by a quadratic function that depends on the generation level, since the usage of fuel increases more than linearly in the increase of the generation level. Maintenance costs are modelled by a linear function of the generation level, and are usually significantly lower than fuel costs. Startup costs and shutdown costs can be significant in the sense that they do have an influence on the unit commitment decisions in time.

Revenue/profit maximization

As we will see in a later subsection, profit or revenue maximization becomes a more attractive objective in a liberalized electricity market. In this case, operational costs are carefully considered against market clearing prices and cleared quantities; i.e, due to competition, the demand for a single generation company becomes more elastic in a liberalized market. This gives an incentive to shift the focus towards profit maximization.

Demand matching

The most clear constraint deals with the production requirements for the set of generators. In the UCP the sum of the production should always exceed the demand. Overproduction can be dumped, in case this is necessary.

Reserve matching

Next to matching the demand, it is an additional requirement to have a so-called spinning reserve. This spinning reserve represents the unused part of the capacity of already committed generators, and, as such, provides possible additional generation capacity that can be dispatched directly, in case of unexpected deviations in demand.

Minimum runtime/offtime constraints

Only committed units can (and should) produce electricity. Once a unit is committed it is desired to let it run for a period of time that is usually longer than a single time interval. The startup costs may be effective in deriving this property. In some cases, hard constraints are used to require that a unit stays committed for a given number of intervals.

Ramp rates

Generators are technically limited in the speed at which they can adjust their generation level. Ramp rates define the maximum increase and maximum decrease of the generation level in a single time interval.

Capacity limitations

Of course each generator has a given production capacity; it is obvious that its generation level cannot exceed this capacity.

Crew scheduling

Some (long-term) variants of the UCP take crew scheduling into account. This deals with assigning employees to power plants, with a focus on maintenance scheduling and including operational crew constraints.

Network limitations

The electricity transmission and distribution system can also be taken into account. Usually it is assumed that the grid capacity exceeds the available production capacity, but this may not be the case. For example, the layout design of the network in relation to the source of distributed demand can introduce capacity problems and even blackouts [34]. In this case network constraints should be added. Especially the interconnection between different areas (different markets) can be a bottleneck.

2.1.2 RECENT DEVELOPMENTS IN UNIT COMMITMENT

After the liberalization of the electricity market, generation companies and electricity retailers were strictly separated. This leads to changes in the Unit Commitment Problem. A summary of recent developments can be found in [66]. Whereas the generation company can be regarded as the only player in the UCP, now the emergence of other players in a competitive market leads to systematic changes that could be reflected in the UCP in different ways. Considering the market mechanism, the competitive auction system directly influences the required production of a single generation company and, thereby, its primary constraints. Also, ways of acting on an electricity market can be merged with the classic UCP. The influence of competitive auction mechanisms is also expressed by an increasing demand uncertainty, which leads to the development of stochastic problem formulations, in

which several stochastic scenarios are considered. Finally the decentralization of the electricity supply chain leads to new types and sizes of generators, and therefore to an increasing size of the problem instances.

Unit Commitment and electricity markets

The implementation of electricity markets leads to more competition and thus to changes in the way energy planning problems in general are treated. This influence cannot only be noticed in the short term UCP, but also long term market effects are seen. The work in [73] considers power generation expansion planning, which in essence is the problem of locating new power plants for a generation company. The shift from having a monopoly to competition leads to a change from inelastic demand to elastic demand (also on the long term), since a generator has an increased risk of being out-competed in a couple of years. Therefore it is much more crucial to take location, primary energy source and future expectations into account, when planning new power plants.

On the short term, generation companies are pushed towards an active role in offering market bids, consisting of price and quantity pairs. The production is not matched to inelastic demand but to offered amounts on a short term (day ahead) market, as in [41, 42, 119, 133]. Varying fuel costs are also taken into account in [119]. The interconnection of different regional markets is also studied [101]. In this case export and import between four different areas are optimized in a UCP framework.

Stochastic Unit Commitment

There have always been inaccuracies in the prediction of demand (and thus the prediction of the required generation) in the UCP. These inaccuracies however, could be relatively easily 'repaired' in realtime, due to the relation between the amount of demand uncertainty and the available flexibility in the generation capacity.

Now all generating companies have to act on (long, medium or short term) electricity markets. This acting on electricity markets introduces price uncertainty. Besides that, in the change towards Smart Grids the use of distributed generation shows that the generation capacity of many generators has decreased. This results in a stronger impact of demand uncertainty, which cannot easily be repaired anymore by the committed generators, and leads to the introduction of stochastic Unit Commitment Problems.

The stochasticity of the demand (and of market prices) can be considered in two ways. In [66] probabilistic constraints are used to model demand uncertainty. In most related work [40, 41, 42, 60, 105, 116, 119] scenario trees are developed and the expected profit, revenue or costs are optimized. Scenario trees consist of possible variations on predicted outcomes in a certain time horizon. Each scenario has a certain probability of occurrence, such that the expected value of the problem can be calculated.

paper	type of generation	total capacity (MW)	# generators	# intervals	stochastic
[40]	hydro/thermal	13000	32	168	yes
[116]	hydro/thermal	-	6	12	yes
[42]	coal/hydro/other	3834	20	168	yes
[60]	hydro/thermal	13000	32	168	yes
[105]	hydro	-	8	48	yes
[119]	thermal	-	33	168	yes
[101]	thermal	12020	104	24	no
[38]	thermal/microCHP	20	5010	48	no

Table 2.2: Problem instances of related work

Decentralized Unit Commitment

The decentralization of the electricity supply chain gives room to study smaller-scale generation. However, most work in the UCP still focuses on relatively large-scale generation. Table 2.2 gives details on the problem instances that are studied in a selection of papers. The average generation capacity per generator remains in the order of hundreds of MW, which is still relatively large. The final row in the table shows our contribution, which focuses on distributed low-scale generation with large numbers of generators. This type of problem brings along a focus shift towards feasibility. By feasibility we mean the ability to find a solution (not necessarily the optimal solution) that satisfies all constraints. Feasibility is extensively discussed in Chapter 3.

2.2 A VIRTUAL POWER PLANT OF MICROCHP APPLIANCES

In this section we discuss a Virtual Power Plant, consisting of microCHP appliances (as indicated in the previous chapter). From an economic and policy point of view, there are some concerns regarding the large scale introduction of microCHP [62, 68]. The policy analysis of [62] describes possible conflicts between policy instruments to support microCHP and other energy efficiency measures (i.e building insulation). They conclude that simultaneous support for energy efficiency measures (e.g. insulation) and microCHP can be justified, but care must be taken to ensure that the heat-to-power ratio and capacity of the micro-CHP system are appropriate for the expected thermal demand of the target dwelling. The study of [68] concludes that individual households lack incentives to switch from conventional boiler systems to microCHP; however, from the viewpoint of a centrally organised entity, there is large potential to operate a Virtual Power Plant. We propose a business case in which such a central entity has control over the individual generators.

2.2.1 EXISTING APPROACHES

There have been different studies to the introduction of a Virtual Power Plant. The dissertation of [114] shows the concept and the controllability of a VPP from an electrical point of view. The minimal power output in this case is on a miniscale (tens/hundreds of kW). The economic possibilities of a VPP with microCHP systems are studied by [68, 112]. In the work of [112] the difference between the virtual

generation capacity in summer and winter periods leads to the conclusion that a VPP cannot replace a conventional power plant in the sense of supplying continuous baseload, but is a mere competitive entity during the daily dispatch of electricity. Short term economics of VPPs are studied in [82]; the study concludes that generators may differ in the form they bid within a VPP (either true to marginal costs or auction oriented/strategic bidding). The combination of both forms is crucial for a good operation of a VPP.

The term Virtual Power Plant is not only used in the literature for a cooperation of small-sized electricity generators, but also for an artificial financial option to increase the performance of the functioning of the electricity market. In this description of a Virtual Power Plant, a VPP is defined as a tradeable option, which gives the right to produce electricity at certain time periods [53]. In this case, a Virtual Power Plant is an auctioned right to generate electricity, that not necessarily needs to be exercised when the specified time period arrives. We want to stress that we do not propose our VPP as such an option, but we want to explicitly incorporate the duty to commit to the generation levels that are auctioned for our VPP.

In the real world, several examples of VPPs exist. In Germany, [87] is an example of a VPP in practice, that uses Volkswagen motors to generate 19 kW electric and 31 kW thermal power. The company Lichtblick pays rent for the used space of the installation, gives an environmental bonus for the production of electricity in the form of a compensation for the price that the household has to pay for the heat generation, and pays most of the installation costs. In return they have the right to operate the appliance, between the comfort limits set by the additional heat buffer. In The Netherlands, the concept of the Multi Agent System oriented PowerMatcher is tested in a field trial consisting of 9 microCHPs [67].

2.2.2 BUSINESS CASE

The VPP in our business case consists of microCHP appliances with a fixed output of 1 kW electric power. All generated electricity is auctioned on a day ahead market. This means that the total generated electricity of all microCHPs is sold and not only the total measured export (i.e. the generation minus the electricity immediately used in home). To differentiate between the exported electricity and the total generated electricity, measuring equipment needs to be installed at each microCHP appliance in each household.

We propose an ownership construction as in the Lichtblick case, meaning that the operational control lies with a centralized entity, in return for a compensation for installation and possessed space. The costs for heating remain a household responsibility, minus an annual compensation for the contribution to the environment via high efficient electricity generation. The fact that such a construction already is used in practice, shows that households are willing to accept this form of loss of control, as long as this has some financial advantages and as long as this does not lead to inconveniences in heat supply. A consequence of this setting is that the operational costs, related to market participation, from the viewpoint of the centralized owner are zero.

Note that our VPP is not intended to be used as balancing dispatch power or ancillary services (e.g. for congestion management), but only to act on a day ahead electricity market. The maximum capacity of the VPP results from the predicted heat demand of all households and is close to the minimum amount of generation that is needed for supplying this heat demand. The small difference results from the use of the heat buffer. In general this difference is too small to act as an ancillary service; if no service is needed, we still need to produce heat. Although most microCHP appliances have an additional burner that only produces heat, in general it is not desired to use that burner, since this results in loss of energy efficiency and torments the basic principle of the introduction of microCHP.

Focus shift in our work related to UCP

Recent contributions to the Unit Commitment Problem show a shift towards market inclusion and stochastic influences. We concentrate on large scale decentralization of electricity generation with a certain flexibility in the timing of the individual operation of generators, but with fixed generation levels when the binary decision to produce or not is made.

In the Unit Commitment Problem for large sized generators the transition to realtime control allows for relatively easy up- and downgrading of the generation levels of the committed power plants. For this reason the optimization objective can focus on the economic dispatch and can take stochastic variations on the demand into account. In our problem scenario trees are not easily implemented, since the operation of the VPP depends on the individual heat demand of households. This would give an exponential scenario space in the number of appliances, where we have already feasibility problems when solving a single scenario as explained in Chapter 3. Considering this we focus on the deterministic variant of the problem, and repair demand uncertainty in a realtime step by applying the TRIANA approach.

2.3 A THREE STEP CONTROL METHODOLOGY FOR DECENTRALIZED ENERGY MANAGEMENT

To make a real world large-scale implementation of a Smart Grid possible, this implementation needs to be controllable and manageable. TRIANA is such a control methodology that focuses on decentralized energy systems and is developed at the University of Twente [29, 94]. This methodology consists of three steps (see Figure 2.2), which are taken in order to assure the ability to control different objectives for different stakeholders in the Smart Grid. These steps are:

- prediction;
- planning;
- realtime control.

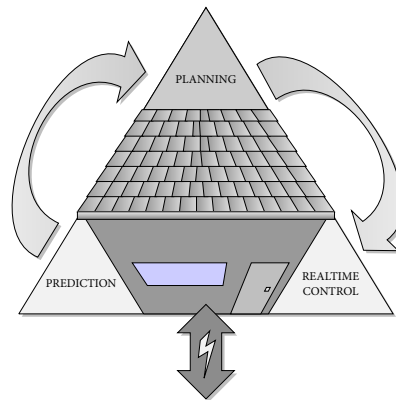


Figure 2.2: The three step approach

A simulator is built to evaluate the consequences of the shift towards distributed generation, distributed storage and demand side load management [30, 95]. The basic energy flow model that underlies this simulator is presented in Section 2.4. This model is organized in such a way, that balance of energy flows is the central requirement. Before presenting the energy model, in this section we give a short overview of the three steps of the control methodology. We start with the potential of the management system.

2.3.1 MANAGEMENT POSSIBILITIES

There are several optimization objectives within a domestic Smart Grid. On a local level, the energy flow within a house can be optimized towards lowering import peaks or working towards minimizing the transport of energy (using as much electricity locally, when it is locally produced). Also price driven objectives can be incorporated, meaning that demand side management is applied to schedule local controllable consumption goods towards variable electricity pricing schemes. For example, controllable washing machines could be scheduled at periods when the electricity price is low and electrical cars could be controlled to be charged during cheap time intervals.

On a global level, houses can cooperate in a Virtual Power Plant, as is depicted in this thesis. In this case the local (in-home) control is driven by global objectives. When working with global objectives, the local household has minimum comfort requirements that may not be violated. For example, the operation of a microCHP in combination with a heat buffer may not lead to a situation in which heat demand cannot be delivered. Optimization objectives on a global level may be to peak shave the total electrical demand that a neighbourhood or village draws from the distribution network. Peak shaving consists of minimizing the maximum load that occurs in a complete time horizon. In the case of a neighbourhood that is equipped with heat pumps, which draw a large amount of electrical power to heat

implement different scenarios using various generation, storage and demand side management techniques.

This setup of the methodology offers many possibilities for applying energy supply chain management. We mainly focus on the case of a Virtual Power Plant consisting of microCHP appliances, but we do also consider combined optimization of distributed generation, distributed storage and demand side management.

Prediction

To be able to act as a Virtual Power Plant on a day ahead electricity market or, more generally, to be able to take future parameters into account, predictions are necessary. Since the operation of different appliances takes place on a household level, we need predictions on this household level too. For our specific Virtual Power Plant we are interested in two types of prediction. For a microCHP appliance it is necessary to have information on the heat usage in a household; the work in [29] gives an overview of the prediction of heat demand in local households. Besides these heat demand predictions, an estimation of electricity prices is also necessary, to model the market behaviour.

If heat demand predictions are done locally, this leads to a scalable prediction system, where each household individually predicts heat demand for a complete day without the necessity to communicate with each other. This prediction results in a certain degree of scheduling freedom for a microCHP, when this microCHP is combined with a heat buffer. The scheduling freedom represents the ability to operate the microCHP (or in general any other generator, buffer or consuming appliance) with a certain flexibility, while still meeting the consumers comfort requirements (in this case respecting the heat demand at all time intervals, by maintaining the heat buffer within its operational heat levels). One possible way to perform the heat demand prediction of a household is to use a neural network (see [29]). Neural networks are generally used whenever a clear causal mapping between input parameters and behaviour is unknown. This neural network consists of a set of input parameters, which in the case of predicting heat demand may be: the heat demand data of one up to several days before the regarded day, predicted windspeed information for the regarded day and the day before, and outside temperature information for the regarded day and the day before. The challenge for a neural network now is to select the right input parameters and to find the weighed combination of these input parameters, such that the prediction becomes most accurate. An important result from [29] is that continuous relearning in a sliding window approach showed good results. This means that a short term history of data (only a couple of weeks) is used each day to update the combination of input parameters, to adjust the prediction to time varying behaviour. The heat demand of the previous day and the demand of exactly a week before showed to be also of importance in choosing the right parameters. A Simulated Annealing heuristic is used to search on the set of possible input parameters to find a good combination of input parameters. More details on the quality of the heat demand prediction are given in Chapter 4.

Market clearing prices on the day ahead market may be predicted by using a short term history of market clearing prices. [17] presents price forecasting using a combination of neural networks, evolutionary algorithms and mutual information techniques. The work in [44] shows price forecasting of a day ahead market; they conclude that time series techniques outperform neural networks and wavelet-transform techniques. Important aspects are that the mean and variance of the market clearing price are non-constant and have a relative high volatility. Next to predicting the market clearing price an indication (prediction) of the variation of this price is therefore also important, when we want to act on an electricity market. In a short term history seasonal influences on the development of the market clearing price are marginal. Based on a prediction of the mean and variance of the market clearing price, estimates of market bids can be calculated. More information regarding this subject is given in Chapter 5.

Planning

Predictions are necessary to derive the operational possibilities of distributed energy management. Based on this operational flexibility a planning can be made, which is the subject of this thesis. The necessity of having a planning in the case of a Virtual Power Plant is evident. Without a planning, the controlling entity of the VPP lacks information on the bids that need to be made on an electricity market; the controller does not know whether it is possible to offer a certain amount of electricity in a certain time interval.

In other energy related optimization problems, the planning step can also be a helpful tool to cooperate with a realtime control scheme. Realtime control without any knowledge about the future can lead to disastrous failures in meeting certain requirements or objectives. For example, realtime control of charging a group of electrical cars, solely based on electricity price signals, can lead to interruptions in the power supply, when all cars start charging at the same moment in time, which probably exceeds the available capacity of the distribution network. In our TRIANA methodology, we choose to implement knowledge about the future via a combination of predicting and planning the operational freedom of appliances in the domestic Smart Grid.

Realtime control

As a final step in the management methodology, we apply realtime control. Due to uncertainty in predicted parameters, it is in general not always possible to exactly follow the planned operation. To overcome this problem, realtime control tries to optimize the energy management problem in an online fashion [94].

In this realtime control step actual decisions are taken for all involved elements in the problem that is under consideration. For the VPP case, this means that the real decision to generate electricity is taken in this step, while focusing on matching the planned operation, and simultaneously respecting the heat demand requirements of each household. In such realtime decision making, the balance in the energy

flow model should always be respected. The realtime control is based on generic cost functions, that are related to the decision freedom for each of the controllable entities in the energy flow model. These cost functions and their influence on realtime control in the VPP case are discussed in Chapter 4.

The realtime control mechanism also has the possibility to take some future time intervals into account, while making a decision for the current time interval. In this Model Predictive Control (Rolling Horizon) way, the realtime control step can anticipate on possible future obstacles.

2.4 ENERGY FLOW MODEL

The TRIANA methodology is an energy management approach that offers a framework to test and simulate different control mechanisms for distributed generation, distributed storage and demand side load management, but it is for example also capable of testing the economic dispatch of large-scale generation in power plants. The main advantage of the underlying model of the simulator, that enables the user to define such a wide range of energy related scenarios, is a generic way of preserving energy balance in a given time interval. This so-called energy flow model, extensively discussed in [94], serves as the basis for the different flows of energy in the planning and realtime control step of the TRIANA method. It also provides a lot of insight in the architecture of the Smart Grid.

In this section we therefore pay attention to modelling the energy infrastructure as a flow network. First we present the basic elements of the model, the corresponding balancing constraints and the decision freedom in the model. Then we show the resulting energy flow model for the example of the generation company in Section 2.1 and we conclude with a general model of an extended example of a Smart Grid.

Note that the energy flow model depicts the situation in an energy infrastructure for a certain (short) time interval. We use the term energy infrastructure, since we distinguish between different types of energy (e.g. gas, heat, electricity). A simulation scenario consists of a series of subsequent energy flow models with dependencies between elements of the current time interval and the elements of the next time interval. These elements are the usual distributed energy management elements, e.g. elements for production, consumption, storage, measuring and communication. In each flow model balance has to be found, by using the decision freedom for the different elements. Elements are classified by different types. We divide the elements E into consuming, exchanging, buffering, converting and source elements: $E = E_{cons} \cup E_{ex} \cup E_{buf} \cup E_{conv} \cup E_{source}$. These elements, together with a special set of pools P , form the nodes in the energy flow graph $G = (V, A)$, i.e. $V = E \cup P$ ($E \cap P = \emptyset$). The set of directed edges A of the graph consists of two types of arcs: $A = A_{EP} \cup A_{PE}$. Hereby, each directed edge (arc) $a_{ep} \in A_{EP}$ denotes an energy flow from a node $e \in E$ to a node $p \in P$ and an arc $a_{pe} \in A_{PE}$ denotes a flow between a node $p \in P$ and a node $e \in E$. Note that in general G is a sparse graph. The directed edges that occur in the model for the different elements are described below. Since we speak about energy flows in Wh, flows are always nonnegative. As a

consequence of the chosen structure of the edge set A , different elements $e_1, e_2 \in E$ are never directly connected to each other nor are different pools $p_1, p_2 \in P$; energy is always transferred from elements to pools of a specific energy type and vice versa. Pools are a means for transportation and keeping track of energy flows of a specific energy type (e.g. heat, electricity or gas); they offer an interconnection of several elements for a certain type of energy. A nice property of the simulator is that the energy type of an element is always type checked with the pool it wants to be connected too; in the formulation we omit these type checks.

A consuming element consumes energy of certain energy types, which means that it allows for multiple flows from different pools to the element. The decision to consume is either fixed for a normal non-controllable consumption appliance, or has some freedom for intelligent controllable appliances. Consuming elements are regarded as sinks in the flow network.

Exchanging elements are mainly used to connect different operational levels in the infrastructure and consist of two bidirectional arcs, to model possible bidirectional flows between the two different 'worlds' that are separated by an exchanging element. Transformers are modelled by using exchanging elements, as well as gas and electricity connections of a household. Exchanging elements form good reference points for measuring and controlling net flows on strategic locations in the infrastructure. The flow model demands that the flow of the incoming arcs minus the flow of the leaving arcs of an exchanging element is zero.

Converting elements can have multiple incoming arcs and multiple leaving arcs. These elements represent different types of generators, that consume possibly different types of (primary) energy and convert the corresponding energy to other forms of energy. Loss is also considered as a form of energy, and is eventually consumed by loss consuming elements. In this way, efficiency calculations can be easily executed, and it allows the model to again demand that the sum of incoming energy flows minus the sum of leaving energy flows of a converting element is zero. In general the decision freedom of a converting element is determined by the different ways that the converter can be operated. Based on these (possibly) different operational modes, the relations between incoming and leaving energy flows are fixed by the model, where flexibility in the energy efficiency of different operational modes is included.

Buffering elements represent energy buffers, and are the only type that allows an internal state. This state keeps track of the buffer level, which determines the operational freedom of the buffer. A buffering element can have multiple incoming arcs and multiple leaving arcs. Again, balance is preserved by requiring that the sum of incoming arcs minus the sum of leaving arcs minus the increase in state (which may also be negative) is zero. The decision freedom of a buffering element can be found in the possible range in which the internal state can be altered.

The different elements are coupled to each other via pools. Each pool has a corresponding energy type (e.g. gas, heat, electricity) and may only have incoming arcs from and leaving arcs to elements that have a leaving/incoming energy flow of the same type. Within a pool balance is required.

Source elements only allow flows from this source to a pool and are the primary

form of energy that enter in the model. Since balance is preserved in the complete model, the total energy flow from the sources into the model should equal the total consumption plus the total increase/decrease in the buffering states.

In a simulation scenario, the decision freedom in a certain time interval is determined by the decisions of the precious intervals and possibly the internal state of the elements. For each interval, the space of possible decisions is searched to optimize for some objectives, while preserving the balancing requirements in the graph. The (possibly conflicting) objectives of realtime control and the use of generic cost functions for decision making are discussed in Chapter 4.

Example of a generation company

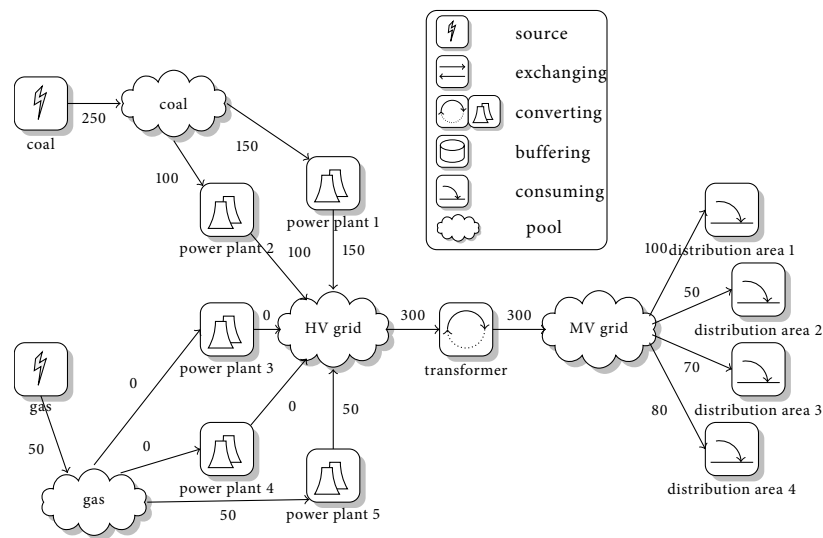


Figure 2.4: An energy flow model of the example of the generation company

In Figure 2.4 the infrastructure of Section 2.1 is given in a flow graph, by using different types of nodes to stress the differences between generation, transportation and consumption. Two different types of energy sources are available: power plant 1 and 2 run on coal, and power plant 3, 4 and 5 are gas-fired. They convert coal or gas into electricity flows. The power plants are connected to the high voltage grid, which is connected to the medium voltage grid, via a transformer. The demand of the example is split into different demands for 4 distribution areas, that are connected to this medium voltage grid. The flow graph shows that a balance occurs at every node in the graph, with the exclusion of source and consuming elements.

Example of a Smart Grid infrastructure

Figure 2.5 shows an example of a Smart Grid infrastructure. As in Figure 2.4, directed edges show the electricity flow in the network. Some edges are bidirectional, indicating that a flow in both ways is possible. In the graph such a bidirectional flow is represented by two opposite arcs with nonnegative flow. The large power plants are connected to the high voltage grid; smaller generation (e.g. windmill/solar panel parks, biogas installations, etcetera) is connected to the medium voltage grid. Note that this smaller generation is not directly coupled to the medium voltage grid, but via an additional electricity node, which is connected to an exchanging node, called 'new generation'. This 'exchanger' expresses the introduction of a new, lower level in the model of the smart grid and functions as a separator between a possible optimization problem on the higher level and the local commitment problem on the lower level (e.g. the smaller generation). The exchanger can be seen as a communication means between higher and lower order planning problems. This division into levels is further explained in Chapter 6.

Compared to the model of Figure 2.4, distribution areas like villages are now modelled in more detail. In the previous model it sufficed to consider the connection of a village to the medium voltage grid, since only aggregated demand is taken into account. In the extended model a village is connected to the lower voltage grid and an exchanger is used to specify a lower level. In this lower level, a next level is introduced for the houses to model their own generation/consumption characteristics. Within the model (e.g. within the houses) different types of energy (i.e. gas and heat) are combined. This is one of the strengths of the extended model. In the model presented in Figure 2.5 we include the modeling of a microCHP. It is convenient to use a heat buffer next to this microCHP to guarantee the heat supply in the house and to partially decouple heat consumption from the generation of heat (and electricity). In the model, gas import information is stored in the gas exchanger. The energy efficiency of generation can be modeled by adding energy losses. In the example, the loss flow of the microCHP has a fixed ratio to the heat and electricity generation; the loss flow of the heat buffer is determined by the state of the buffer. In a similar way, the efficiency of each type of generation can be modeled. However, for simplicity this is left out of Figure 2.5.

2.5 CONCLUSION

This chapter gives additional background on the focus of this thesis. We give an overview of the Unit Commitment Problem as a basic reference point of the type of electricity/energy planning problems that we study. Next, the concept of Virtual Power Plants is further explained and a business case is given, which shows the type of Virtual Power Plant that we consider. The planning of the generation output of this Virtual Power Plant is part of a 3-step control methodology for decentralized energy management, called TRIANA. An important aspect of TRIANA is an energy flow model, which focuses on the balance requirement of electricity networks.

THE MICROCHP PLANNING PROBLEM

ABSTRACT – This chapter treats the microCHP planning problem, which models the planned operation of a Virtual Power Plant consisting of microCHP appliances that are installed with additional heat buffers. Since it is a new type of planning problems in the electricity supply chain, it is extensively modelled and studied for its complexity. The microCHP planning problem is proven to be \mathcal{NP} -complete in the strong sense. Based on this complexity result and on initial computational results for an Integer Linear Programming formulation as well as a dynamic programming formulation, heuristics are developed to solve the problem. A first heuristic is a local search method that is based on the dynamic programming formulation for an individual house. This heuristic restricts the search for the optimal solution to general moves in the space domain (i.e. the set of different microCHP appliances). A second heuristic stems from approximate dynamic programming and concentrates more on time dependencies. A third heuristic uses a column generation technique, where the planned operation of individual microCHPs is represented by a column. This heuristic gives the most promising results for implementation in a real world setting.

In this chapter we focus on planning the operation of a Virtual Power Plant that completely consists of microCHPs, by defining the microCHP planning problem. Within the mathematical formulation we already take into account the connection between the microCHP planning problem and the research questions that are answered in the following chapters. This connection is mainly expressed by the bounds on the total electricity generation of the Virtual Power Plant.

The focus of this chapter is on modelling and solving the planning problem for a large number of microCHPs. The microCHP planning problem is treated as an

Parts of this chapter have been published in [MB:5], [MB:8], [MB:6], [MB:9], [MB:7], [MB:4], [MB:21] and [MB:20].

example of a difficult planning problem in the field of distributed energy generation, due to its strong dependency in both time and space. Similar algorithms can be derived for other types of generators and heating systems, e.g. gas turbines, heat pumps, etcetera. It serves as a standalone planning problem for a Virtual Power Plant and as an important starting point to treat combined planning problems in the changing energy supply chain, as explained in Chapter 6. In Section 3.1 the general problem of planning a group of microCHPs is introduced. Before we show different methods to solve the microCHP planning problem, we first draw some attention to the complexity of the problem in Section 3.2. We formulate two types of optimization problems for the microCHP-based Virtual Power Plant. In Section 3.3 an Integer Linear Programming formulation of the microCHP planning problem is given and in Section 3.4 a dynamic programming formulation, which may be used to solve small instances to optimality. These first results, in combination with the theoretical complexity of the microCHP planning problem, show the urge to develop efficient methods, i.e. methods that find solutions in reasonable time and which are close enough to the (possibly unknown) optimal solution. Such methods are presented in Sections 3.5-3.7. Finally a conclusion is drawn in Section 3.8.

3.1 PROBLEM FORMULATION

In this section we describe the requirements for a group of microCHPs to operate correctly. Next to these requirements, several optimization objectives are indicated that could be of interest for an operator or planner of this group of microCHPs. Together these requirements and optimization objectives form input to the mathematical planning problem for a group of microCHPs. The term *planning* reflects to the series of decisions to let (a/multiple) microCHP(s) run at sequential time periods or not. The formal definition of these planning problems is postponed to Section 3.2, where a general notion of complexity is explained and the complexity of the microCHP planning problem in particular is treated.

3.1.1 MICROCHP AS AN ELECTRICITY PRODUCER

Combined Heat and Power appliances on a domestic scale (microCHP appliances) consume natural gas and produce both heat and electricity at a certain heat to electricity rate. The electrical output is in the order of kiloWatts (kW), which means that it is suitable for use on a household scale.

MicroCHP is considered as one of the possibilities to implement the decentralization of energy production (see Chapter 1 and 2). It has a relatively high energy efficiency compared to that of large(r) power plants, which shows the main advantage of this type of distributed generation. The important benefit in the energy efficiency originates from the more efficient use of the heat, since produced heat in a power plant cannot be transported/used as efficiently (if it is not lost already in the production process) as on domestic scale. However, this means that the principle focus of Combined Heat and Power production on a domestic scale should be on the efficient storage/consumption of heat in order not to lose this advantage. Therefore,

microCHP mainly can be seen as a replacement for current boiler systems, and secondly as a domestic electricity generator.

There are several possible technical realizations of a microCHP, such as Stirling engines [43], rankine cycle generators [109], reciprocating engines [117] and fuel cells [5], where Stirling engines are nearest to full market exposure.

3.1.2 REQUIREMENTS

The requirements for the operation of a group of microCHPs can be divided into three sets: appliance specific characteristics, operational (time dependent) requirements for each microCHP and cooperational requirements on groups of microCHPs. For a list of used variables and parameters we refer to the list of symbols.

Appliance specific characteristics

The microCHP generation characteristics are given by a set of parameters that describe the behaviour of the microCHP, once its way of operation has been decided. To be used in a domestic setting the order of magnitude of the production of heat and electricity by the microCHP should be such that the operation within the house is allowed (according to local grid policy) and such that the appliance is able to fulfill the heat demand. This means that local heat demand can be supplied completely by the microCHP (in combination with a heat buffer) and that the electricity production does not exceed the maximum output that may be delivered back to the electricity grid. This supply is namely bounded by regulations set by the national government. This combination of heat supply requirements and electricity supply limitations results in a limited freedom for technology development. By this we mean that the ratio between the heat and electricity generation of the different microCHP technologies is more or less decided by environmental factors. Naturally this ratio is also influenced by the technological possibilities itself. From the viewpoint of the planner, we can assume that the electricity to heat ratio is fixed and known for a certain generation technology and can be used as given input for the planning problem.

Figure 3.1 shows the electricity output profile for an example run of a microCHP based on a Stirling engine. It can be seen, that there is no one-to-one relation between the microCHP being switched on and the power output. In general, a run can be roughly divided into three phases:

- a startup phase, in which, after some grid tests, the engine is started and the power output slowly increases to its maximum output value;
- a constant phase, in which the power output balances around the maximum output value;
- a shutdown phase, in which the engine is slowed down.

Roughly the same division into phases yields for the heat output.




Figure 3.1: Electricity output of a microCHP run

These appliance specific characteristics are modelled by the following parametric behaviour. We define a maximum power output that corresponds to the average electricity output in the constant phase. Furthermore, startup and shutdown behaviour is described via corresponding output functions. Finally, an electricity to heat ratio is defined to give a direct relationship between the production of electricity and heat.

Operational requirements

The highest energy efficiency is reached in the constant phase. For this reason, and to prevent wearing of the system, longer runs are preferred over shorter ones. This leads to the requirement of having a minimum time that the microCHP has to run, once switched on. For similar reasons the microCHP has to stay off for a minimum amount of time, once switched off. Naturally, these minimum runtimes and minimum offtimes are larger than or equal to the startup and shutdown periods that are required for an efficient use of the microCHP, since we want to run at maximum power for at least some time during each run (and of course for as long as possible).

If the heat consumption would be directly supplied by the microCHP, the decisions to run the microCHP are completely determined by the heat demand. As a result often short runs of the microCHP would occur. This is the reason why microCHPs are in general combined with a heat buffer. This additional heat buffer allows to decouple production from consumption up to a certain degree and, therefore, to make a planning possible.

Based on the above considerations, the planning for a house with a microCHP and a heat buffer is heat demand driven, where the requirement is to respect certain lower and upper limits of the heat buffer in order to be able to supply the domestic heat demand at all times in a feasible planning. Note that there is a strong time dependency between operational decisions; decisions in certain time periods have

a large impact on possible decisions in future time periods. E.g. switching on a microCHP now leads to a certain minimum amount of heat generation and, therefore, increases the heat level in the heat buffer. This may have as a consequence that in certain future time periods the microCHP cannot run, since it cannot get rid of the produced heat without spoiling heat to the near environment, which is an option that we do not allow.

Cooperational requirements

Once houses are collaborating in a larger grid, the aggregated power output of the different houses adds a global electricity driven element to the planning problem. The group of houses can act as a so-called Virtual Power Plant (VPP) by producing a certain electricity output together. This output may be partially consumed by the houses themselves, but part of it may also be delivered to the electricity network. The aggregated electricity production is not always free to be chosen; there may be several constraints on this aggregated power output. This global electricity driven requirement can be specified by a desired lower bound and a desired upper bound for the aggregated electricity output. These bounds can be determined by (governmental) regulations, capacity limitations of the underlying grid or desired operational achievements such as causing stability and reliability in the grid. Also these bounds can origin from actions that were taken on an electricity market.

The electricity retailer of the households may act on a short term electricity market in advance (e.g. for 24 hours ahead) or on a realtime market. As the prices of electricity on these markets vary over time, it may be beneficial to steer the fleet to produce more electricity in expensive periods. The retailer may consider to bid the expected overall production profile of the group/fleet on the market and operate the fleet according to the cleared outcome of this bid. This resulting profile somehow will depend on the prices of the market, but for the planner the most important question is whether he is able to reach this profile with the fleet or not, since a deviation of the realized planning the next day leads to (huge) costs on the balancing market. This requirement can be specified as operating the group of microCHPs in such a way, that the aggregated electricity output lies between desired bounds.

The cooperational requirements represent the second direction of dependency in the microCHP planning problem; next to time dependency the problem deals with dependency in space. The interaction between the different types of requirements is depicted in Figure 3.2. In the left part of the figure, the solution space X_1^1 represents the space that is formed by respecting the appliance specific characteristics for house/microCHP 1. The time dependent operational requirements are given by solution space X_2^1 , and the intersection $X_{1,2}^1 := X_1^1 \cap X_2^1$ shows the feasible solutions regarding appliance characteristics and time dependent behaviour. In the middle part of the figure, the spaces $X_{1,2}^i$ are combined for houses $i = 1, 2, 3$, leading to the solution space $Y_{1,2} := X_{1,2}^1 \times X_{1,2}^2 \times X_{1,2}^3$. The right part of the figure shows this space $Y_{1,2}$ and the subspace $\tilde{Y}_{1,2} \subseteq Y_{1,2}$, which includes the cooperational requirements.




Figure 3.2: Solution space for the microCHP planning problem

3.1.3 OPTIMIZATION OBJECTIVES

Next to the requirements also the goals/objectives for the optimization problem of planning a group of microCHPs need to be specified. In general there are two kinds of objectives for a Virtual Power Plant: maximizing the profit on an electricity market or minimizing the deviation from the given bounds on the aggregated electricity output. We do not consider other objectives as e.g. in [70], where microCHPs are optimized for their individual profit.

Maximizing the profit on an electricity market

Given (a prediction of) the prices on an electricity market the planner searches for the optimal operation of all microCHPs, such that all requirements are met and the aggregated electricity output is maximized for the given prices. The base for this objective is the solution space $\tilde{Y}_{1,2}$.

Minimizing the deviation from the given bounds on aggregated electricity output

As a second type of optimization objective we do not consider the direct optimization on an electricity market, but the feasibility of the problem is inspected. The nature of the combination of the two-dimensional dependencies in time and space namely makes it sometimes really difficult in practice to even find a solution that respects all requirements. In such cases we may allow a planner to soften some of the cooperational requirements on the aggregated electricity output, meaning that the base for this problem now gets the solution space $Y_{1,2}$. We minimize the violation of these cooperational requirements by minimizing the deviation from these softened cooperational bounds as objective. Although this objective does not optimize for an electricity market directly, the electricity market can still be indirectly taken into account via the (softened) cooperational bounds.

3.2 COMPLEXITY

Since the invention of the computer in the last century a lot of progress has been shown in solving computationally intensive problems. Both in hardware and in software many advances have resulted into an increasing computational performance.

Regarding the developments in hardware, Moore's law, stating that the number of transistors that can be placed on an integrated circuit doubles roughly every two years, has been followed until now quite accurately. This law has comparable effects for the developments in processing speed and memory capacity for example, which leads to an exponential growth in the capability to compute. Nevertheless, it is of importance that methods/algorithms developed for given problems are efficient in the way that the number of steps to be executed gets minimized. The focus in the following subsection is on this algorithmic side of software development and thus, on the complexity of problems.

3.2.1 COMPLEXITY CLASSES

Complexity classes are introduced to make a classification possible that distinguishes problems that are in general very difficult to solve from problems that are easier to solve. Difficulty in this sense can be loosely described by the relation between the amount of calculations that is needed to find a solution and the input size of the problem instance. It is worthwhile to note the difference between this notion of complexity classification and the difficulty of solving specific problem instances. For some problem instances namely, instance specific properties can be used to derive some relations that make an efficient solution method possible. However, complexity is determined by the weakest possible problem instance; if there is some instance that does not satisfy the specific properties, this efficient solution method cannot be applied to the problem in general.

Optimization problems and decision problems

So far, we only mentioned the term *difficulty* as a loose description of complexity. To give a more precise definition, we first describe the difference between a (combinatorial) optimization problem and a decision problem. Then we discuss the difference between the two complexity classes \mathcal{P} and \mathcal{NP} .

An *optimization problem* is given by a set of *feasible solutions* X that satisfies problem specific constraints and an *objective function* f on this set X . The optimization problem asks for a feasible solution $x \in X$ that returns the optimal value of the objective function f , i.e. an *optimal solution* to the underlying problem. A *decision problem* does not search for an optimal solution to a problem. Instead it poses a question that needs to be answered with a simple 'yes' or 'no'. An optimization problem can be easily transformed into a decision problem by introducing a certain bound K and asking for feasible solutions x that also respect the additional constraint ' $f(x) \leq K$ ' or ' $f(x) \geq K$ ', where the inequality depends on the optimization direction (\leq for a minimization problem and \geq for a maximization problem). In this way the decision variant of an optimization problem asks whether a solution exists that is equal to or better than a bound K : is the problem feasible under the additional constraint?

\mathcal{P} vs \mathcal{NP}

The complexity classes \mathcal{P} and \mathcal{NP} refer to the complexity of decision problems rather than optimization problems. The class \mathcal{P} consists of all decision problems that can be *solved* in polynomial time. This means that a deterministic algorithm exists that can solve all problem instances in polynomial time in the input size of the instance. The class \mathcal{NP} consists of all decision problems that can be solved in polynomial time by a non-deterministic algorithm. The statement that it ‘can be solved’ may be a bit misleading in this context of non-determinism. Namely, non-determinism means that, for an instance that can be answered with ‘yes’, a guessed solution can be *verified* for its correctness by a polynomial time algorithm. The difficulty of guessing a (correct) solution is not taken into account.

For all decision problems in the class \mathcal{P} the guessing and verification are combined in the polynomial time algorithm, showing that $\mathcal{P} \subseteq \mathcal{NP}$. One of the most important remaining open problems (rewarded with a million dollar prize, see [8]) is whether $\mathcal{P} = \mathcal{NP}$ or $\mathcal{P} \neq \mathcal{NP}$, i.e. can all solutions that can be verified in polynomial time also be found in polynomial time or not?

An important factor in this open problem is the notion of \mathcal{NP} -complete problems. A problem is \mathcal{NP} -complete if all other problems in \mathcal{NP} can be *reduced* to this problem, where reduction means a transformation from the original problem into the other problem in polynomial time. This states that this \mathcal{NP} -complete problem is at least as hard as all other problems in \mathcal{NP} ; if a polynomial time algorithm can be found for an \mathcal{NP} -complete problem, then $\mathcal{P} = \mathcal{NP}$. The other way around, if $\mathcal{P} \neq \mathcal{NP}$, then no \mathcal{NP} -complete problem can be solved in polynomial time.

The first decision problem that was proven to be \mathcal{NP} -complete was the SATISFIABILITY problem [45]:

SATISFIABILITY

INSTANCE: Given is a set of boolean variables B , and a boolean expression b on these variables using \vee , \wedge , \neg and/or parentheses.

QUESTION: Is there a truth assignment for the variables in B such that the boolean expression b is truth (i.e. satisfied)?

For the proof of Cook we refer to [45], where the boolean expression b is considered in disjunctive normal form, or to [56], where b is considered in conjunctive normal form. Based on this proof a long list of \mathcal{NP} -complete problems has been formed, of which a classical overview has been given by Garey and Johnson [56]. To prove that a decision problem is \mathcal{NP} -complete, one has to perform the following actions. First, the decision problem needs to be in \mathcal{NP} . Then a known \mathcal{NP} -complete problem needs to be reduced to this decision problem, which means that a polynomial transformation is found from the \mathcal{NP} -complete problem to the decision problem under consideration. Any \mathcal{NP} -complete problem can be used as a starting point for proving \mathcal{NP} -completeness. However, usually one of the basic \mathcal{NP} -complete problems is chosen.

Guidelines for solving new problems

The above complexity classification for decision problems is transferred to optimization problems by calling an optimization problem \mathcal{NP} -hard, if its corresponding decision problem is \mathcal{NP} -complete. The complexity classification can be used as a guidance on how to treat a given optimization problem. It is not likely to find an efficient exact algorithm for an \mathcal{NP} -hard problem. However, the size or the properties of relevant practical instances may be such that an exact algorithm may be applicable. If exact algorithms are not helpful for these practical instances since the size of these instances gets too large, another approach is to use heuristics to find solutions that are close to the optimum. The focus in developing heuristics is twosided: they should provide quality solutions in reasonable time. Bounds for the computation time are often provided by the time that is available for solving practical instances. Since the optimal solution is often unknown (otherwise we would not need heuristics) it is difficult to measure the quality of a solution. However, for some well defined problems it can be proven that a specific heuristic never leads to a solution that deviates more than a fixed factor from the optimal solution. This heuristic is called a ρ -approximation, since the objective value $f(x)$ of the constructed solution x is kept within a factor ρ of the optimal value OPT ($OPT \leq f(x) \leq \rho OPT$ for a minimization problem and $\rho OPT \leq f(x) \leq OPT$ for a maximization problem).

An example: the Traveling Salesman Problem

To clarify the above concepts a bit more, we consider the well known Traveling Salesman Problem. The Traveling Salesman Problem (TSP) deals with a salesman who has to visit n cities, including his hometown as a starting and finishing point. The distance between two cities i and j is given by $d_{i,j}$. The objective of the TSP is to minimize the total distance of a tour that visits all cities. The decision variant of the Traveling Salesman Problem is defined by (see also [56]):

TRAVELING SALESMAN PROBLEM

INSTANCE: Given is a set C of n cities, distances $d_{i,j} \in \mathbb{Z}^+$ for all arcs (i, j) between cities $i, j \in C$, and a bound $B \in \mathbb{Z}^+$.

QUESTION: Is there a tour of all cities in C with a total distance no more than B ; i.e. does an ordering $(\pi(1), \dots, \pi(n))$ exist such that

$$\sum_{i=1}^{n-1} d_{\pi(i), \pi(i+1)} + d_{\pi(n), \pi(1)} \leq B?$$

This decision problem is shown to be \mathcal{NP} -complete [56]. In the following we present some specific methods to show that such a hard problem can be approached from different angles and that practical results can still be achieved for such a hard problem. First we show an exact algorithm that has the lowest known time complexity bound. Then we give another exact solution method by describing the TSP by an Integer Linear Programming (ILP) formulation. Furthermore we

show a heuristic method, and the combination of this heuristic with other solution techniques into a computer program that is fully dedicated to solving TSPs.

- **Exact algorithm (Held-Karp algorithm/Bellman algorithm [32, 63])**

The number of possible tours for the TSP equals $(n - 1)!$, since the starting city can be chosen arbitrarily, which leaves $(n - 1)!$ choices for the remaining $n - 1$ cities. If we consider the symmetric TSP, this number equals $\frac{(n-1)!}{2}$. One of the existing exact algorithms that solves the TSP has been proposed by [63] and [32]. This algorithm is currently still known to have the lowest time complexity of $\mathcal{O}(n^2 2^n)$ [130]. The idea of this on dynamic programming based method is to avoid calculating all possible tours. Instead, only relevant subpaths are taken into consideration in the following way. Without loss of generality city 1 is chosen as the starting point for the dynamic programming method. States (S, j) are given by a subset of cities $S \subseteq C \setminus \{1\}$ and a city $j \in S$ that represents the last city visited in the shortest path from city 1 to j through all cities in S . The value $v(S, j)$ belonging to state (S, j) denotes the length of this shortest path. The algorithm calculates the value $v(S, j)$ by looking at the values $v(S \setminus \{j\}, i)$ for subpaths ending in $i \in S \setminus \{j\}$. Initially, $v(\{i\}, i) = d_{1,i}$ for all $i \in C \setminus \{1\}$. Then in several phases in which the size of each subset incrementally expands, the recursive equation $v(S, j) = \min_{i \in S \setminus \{j\}} v(S \setminus \{j\}, i) + d_{i,j}$ is used to calculate the shortest path for the corresponding subsets. Finally, the shortest tour $v(C)$ is the shortest path from 1 to any other city i , that visits all cities in $C \setminus \{1\}$ and returns to city 1: $v(C) = \min_{i \in C \setminus \{1\}} v(C \setminus \{1\}, i) + d_{i,1}$. This algorithm is summarized in Algorithm 1.

Algorithm 1 Exact algorithm for the Traveling Salesman Problem

```

 $v(\{i\}, i) = d_{1,i} \forall i \in C \setminus \{1\}$ 
 $s = 2$ 
while  $s < |C|$  do
  for all  $S \subseteq C \setminus \{1\}, j \in S, |S| = s$  do
     $v(S, j) = \min_{i \in S \setminus \{j\}} v(S \setminus \{j\}, i) + d_{i,j}$ 
  end for
   $s = s + 1$ 
end while
 $v(C) = \min_{i \in C \setminus \{1\}} v(C \setminus \{1\}, i) + d_{i,1}$ 

```

- **ILP formulation**

An alternative way to achieve an exact solution method, is to model the given problem as an Integer Linear Programming (ILP) formulation. Using binary decision variables $x_{i,j}$ indicating whether arc (i, j) is part of the tour ($x_{i,j} = 1$) or not ($x_{i,j} = 0$), the following formulation (3.1)-(3.6) models the TSP as an

ILP:

$$\min \sum_{i=1}^n \sum_{j=1}^n d_{i,j} x_{i,j} \quad (3.1)$$

$$s.t. \sum_{i=1}^n x_{i,j} = 1 \quad \forall j \in \{1, \dots, n\} \quad (3.2)$$

$$\sum_{j=1}^n x_{i,j} = 1 \quad \forall i \in \{1, \dots, n\} \quad (3.3)$$

$$y_i - y_j + n x_{i,j} \leq n - 1 \quad \forall i \in \{1, \dots, n\}, j \in \{2, \dots, n\} \quad (3.4)$$

$$x_{i,j} \in \{0, 1\} \quad \forall i \in \{1, \dots, n\}, j \in \{1, \dots, n\} \quad (3.5)$$

$$y_i \in \mathbb{Z}^+ \quad \forall i \in \{1, \dots, n\} \quad (3.6)$$

In Equation (3.1) the objective function is to minimize the sum of arc lengths of the chosen arcs $\sum_{i=1}^n \sum_{j=1}^n d_{i,j} x_{i,j}$. Equations (3.2) and (3.3) demand that each city has one incoming arc and one leaving arc, which corresponds to the requirement to visit each city exactly once. These equations (3.2) and (3.3) are necessary restrictions for having a tour, but they are not sufficient restrictions. These restrictions namely also allow for disjoint nonempty subtours, which are impossible to follow in practice by a salesman. Equation (3.4) prevents the existence of disjoint nonempty subtours, modelled as in [91]. The idea of this equation is to create an ordering for the n cities, where city 1 is the initial city, and force the salesman to visit the cities in this order. This leads to $n - 1$ moves forward in the ordering, which leaves one move from the final city to the initial city to complete the tour. This final move to city 1 plays a crucial role in the proof of the existence of exactly one subtour. Equation (3.4) namely defines the following relationship:

$$\forall j \neq 1 : x_{i,j} = 1 \Rightarrow y_i < y_j \quad (3.7)$$

Now assume that a subtour $T = \{i_1, i_2, \dots, i_k, i_1\}$ exists where city 1 is not part of the subtour. By (3.7) this gives $y_{i_1} < y_{i_2} < \dots < y_{i_k} < y_{i_1}$, which is a contradiction. So a subtour can only exist when starting and ending in city 1. Thus, no feasible solution with two or more subtours can exist, since at least one subtour would not contain city 1.

- **Lin-Kernighan heuristic**

The Lin-Kernighan heuristic [89] provides a method that iteratively tries to improve a given tour. To improve an existing tour so-called k -opt moves are used. In general a k -opt move consists of replacing k arcs from a feasible tour by k new arcs in such a way that connectivity of the complete graph is preserved. A typical k -opt heuristic for the TSP searches for shorter tours using a specific k -opt move. For $k = 2$, Figure 3.3 shows a feasible and an infeasible 2-opt move, where the dashed arcs are replaced by the dotted arcs. The 2-opt move in Figure 3.3a preserves the connectivity of the complete




Figure 3.3: A feasible and an infeasible 2-opt move

graph and thus results in a feasible tour, whereas the move in Figure 3.3b is not a feasible move, as it results in two disconnected subgraphs. This means that for $k = 2$ the move is completely determined once the two arcs that are to be removed have been chosen. In Figure 3.4 we present the four feasible 3-opt moves for $k = 3$. A k -opt TSP heuristic uses the set of feasible k -opt




Figure 3.4: Feasible 3-opt moves

moves for each combination of k arcs in a local search strategy.

The basis for the Lin-Kernighan heuristic is to use not just one specific value for k , but to allow different k -opt moves in one neighbourhood for a local search strategy. This is done by applying a specific way to construct k -opt moves of variable length. Only feasible moves are allowed, since the heuristic does not want to ‘repair’ broken tours. The way to construct these variable k -opt moves is by sequentially breaking an arc and adding a new arc. Initially one arc (v_1, v_2) is removed and a new arc (v_2, v_3) (that does not exist already) is added. In the i th step city $v_{2(i+1)}$ is chosen, the arc $(v_{2i+1}, v_{2(i+1)})$ is removed and a new arc $(v_{2(i+1)}, v_{2(i+1)+1})$ to a next city $v_{2(i+1)+1}$ is added. The crucial step in the sequential construction is that the next arc that will be broken is the unique existing arc $(v_{2i+1}, v_{2(i+1)})$ incident to v_{2i+1} that *allows the tour to stay connected if the arc $(v_{2(i+1)}, v_1)$ would be added*. This means that each arc that is broken should allow the possibility to complete a connected tour with a single addition of an arc. As the next arc that is actually added, any arc $(v_{2(i+1)}, v_{2(i+1)+1})$ can be chosen (where $v_{2(i+1)+1}$ has not been considered during the construction before), including the option to complete the tour via $(v_{2(i+1)}, v_1)$. Figure 3.5 gives a summary of the se-

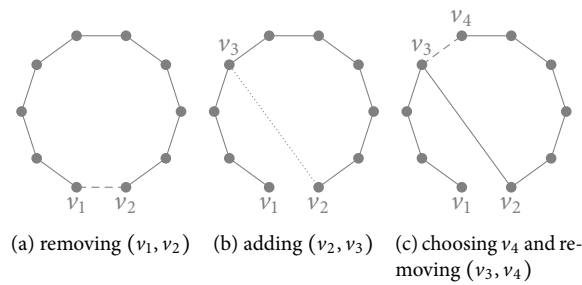


Figure 3.5: Sequential construction of k -opt moves

quential construction, where the choice for v_4 is specified. Note that the other neighbour of v_3 cannot be chosen, since a connected tour cannot be completed with one arc. The sequential construction continues until the tour returns to v_1 , or when the last removal and addition do not improve (i.e. decrease) the tour length. In this case the last added arc is replaced by the arc to v_1 . Note that not all feasible k -opt moves can be constructed in this way (e.g. Figure 3.4b cannot be constructed, since all reductions from this 3-opt move lead to infeasible 2-opt moves that cannot be created sequentially). To partially compensate for these lacking moves, the first choice (for v_4) may be non-sequential, in which case of course the eventual move cannot be a 2-opt move and some rearrangements are necessary to keep an eventual tour connected. The heuristic uses specific options to search for new arcs; we refer to [89] for more details of the algorithm and to [64, 65] for implementation details.

In Table 3.1 the geographical distances between cities, based on the geographical distance calculation defined by [14], are given for a small example to demonstrate the behaviour of the Lin-Kernighan heuristic. Figure 3.6a shows the location of the capital cities of the 12 provinces of The Netherlands. An initial tour given in Figure 3.6b is improved by applying a 2-opt move (Figures 3.6c and 3.6d) and a 3-opt move (Figures 3.6e and 3.6f). As before, the dashed arcs are replaced by the dotted arcs. The final tour that is found also represents the optimal tour for this instance. This tour is also printed in bold in Table 3.1.

- **Concorde**

Concorde [4] is a computer program that is created to solve TSP instances. It includes the Lin-Kernighan heuristic to find feasible solutions, but foremost it consists of a branch-and-cut method that solves an ILP formulation of the TSP, where it uses elaborate cutting techniques to improve on the lower bound. In general, to solve an ILP formulation, the principle of a branch-and-cut method can be applied. In an iterative way Linear Programming relaxations

	Groningen	Leeuwarden	Assen	Zwolle	Lelystad	Arnhem	Utrecht	Haarlem	Den Haag	Middelburg	Den Bosch	Maastricht
GRO	-	53	26	85	109	145	161	160	199	277	192	271
LEE	53	-	57	79	81	136	133	120	160	240	172	262
ASS	26	57	-	61	91	120	141	146	183	260	168	245
ZWO	85	79	61	-	42	61	83	100	131	204	108	188
LEL	109	81	91	42	-	66	53	59	92	169	92	185
ARN	145	136	120	61	66	-	58	99	112	168	54	128
UTR	161	133	141	83	53	58	-	47	55	122	47	144
HAA	160	120	146	100	59	99	47	-	41	121	91	186
DHA	199	160	183	131	92	112	55	41	-	81	82	168
MID	277	240	260	204	169	168	122	121	81	-	119	162
DBO	192	172	168	108	92	54	47	91	82	119	-	97
MAA	271	262	245	188	185	128	144	186	168	162	97	-

Table 3.1: Geographical distances between the capital cities of the 12 provinces of The Netherlands

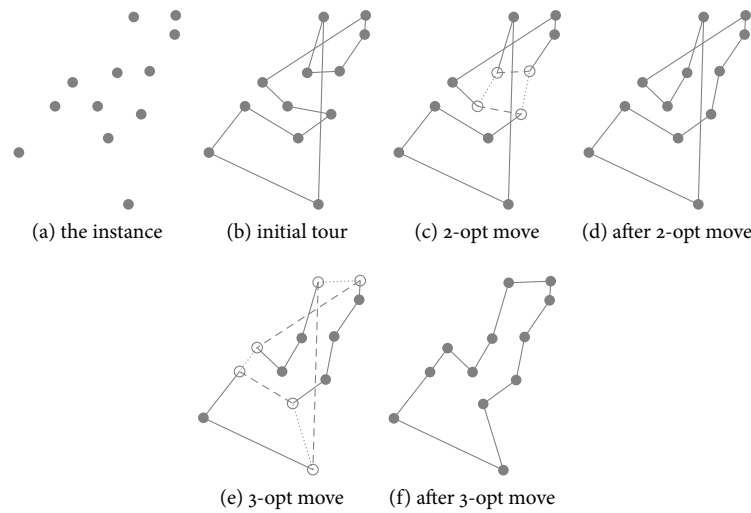


Figure 3.6: Example: the capital cities of the 12 provinces of The Netherlands

of the ILP are solved (e.g. by use of the simplex algorithm). If the solution to this LP relaxation is not a completely integer solution, so-called cuts can be added, which are additional inequalities derived from extra information from the LP-relaxation. The addition of these cuts is combined with a normal branch-and-bound strategy, which consists of adding constraints that break fractional solutions in two separate branches, followed by a search through the created tree of LP problems until an integer solution is found that is globally optimal.

In the mentioned ILP formulation (3.1)-(3.6) we eliminate subtours by ex-

plicitly using Equation (3.4). By use of a solver (CPLEX 12.2) cutting planes are automatically selected and the problem is solved. Opposite to this ILP formulation, cutting planes are specifically designed in the Concorde for solving TSPs, based on the work of [49]. The basic LP formulation of the Concorde only consists of (a variant of) Equations (3.2), (3.3) and (3.5), and cutting planes are added to find a feasible tour. The Concorde consists of various types of cuts and a way of selecting between them in a branch-and-cut framework. Below the most used cut is explained. For a more detailed explanation of this cut and a description of the other cuts we refer to [19].

The basic idea behind the most used cut is to eliminate subtours. The basic LP relaxation of (3.1)-(3.6) without subtour elimination is:

$$\min \sum_e d_e x_e \quad (3.8)$$

$$s.t. \sum_e (x_e | i \in e) = 2 \quad \forall i \in \{1, \dots, n\} \quad (3.9)$$

$$0 \leq x_e \leq 1 \quad \forall e, \quad (3.10)$$

where x_e defines the selection of an undirected edge $e \in E$ (E is the set of edges in the complete graph on the city set C) and $i \in e$ means that $i \in C$ is incident with e . Equation (3.9) requests that each city is incident with two edges (which surely has to be valid in a feasible solution). The binary constraint on the choice for selecting an edge is relaxed in Equation (3.10). Concorde uses now different heuristics to find cutting planes that remove subtours. To explain this we describe an important property of a subtour. We define the set $S \subset C$ as a strict subset of C . Any strict subset S must have two or more connections to the cities that are not in S :

$$\sum_e (x_e | e \cup S \neq \emptyset, e \cup C \setminus S \neq \emptyset) \geq 2 \forall S \subset C, S \neq \emptyset, \quad (3.11)$$

where $e \cup X$ means that some city in the set X is incident with e . This restriction (3.11) is called the subtour inequality. Several heuristics have been developed that find subsets $S \subset C$ that do not fulfill the subtour inequality (i.e. $\sum_e (x_e | e \cup S \neq \emptyset, e \cup C \setminus S \neq \emptyset) < 2$). Corresponding cutting planes (3.11) are then added.

In Figure 3.7 the four methods are compared to each other for their computation time. This comparison is done on a desktop computer (3.00 GHz and 2.00 GB RAM). We implemented the Held-Karp algorithm in C++ in combination with an SQL database to overcome large memory problems. The ILP formulation is implemented in AIMMS modelling software [1] using CPLEX 12.2. We use the implementation of the Lin-Kernighan heuristic by [6] and the Concorde TSP solver from [4].

We compare several instances from the publicly available TSP library TSPLIB [14]. The size of these instances varies between 14 and 3795 cities. In addition to




Figure 3.7: Comparison of runtimes for TSP instances

this set, the instance of the capital cities of the 12 provinces of The Netherlands is used (see Table 3.1).

The Held-Karp algorithm has the lowest known time complexity. The number of states that has to be evaluated is completely determined by the size of the problem (although the actual number of calculating steps may vary per evaluated state), which results in very predictable computation times. The ILP formulation shows to be a faster exact algorithm in practice than the Held-Karp algorithm, although no guarantee can be given that this is always the case. The especially designed TSP solver Concorde improves this practical computation time by a large amount. This shows that in practice often quite large instances can be solved to optimality, although no guarantee can be given that this solution is computed in reasonable time. The Lin-Kernighan heuristic results in comparable results to the Concorde (which is no surprise), with a side remark that the optimal tour is not found for the largest problem (consisting of 3795 cities).

Outline for solving the microCHP planning problem

The above example indicates that a mathematical problem can be solved in different ways, varying from exact algorithms to heuristics. We treat the planning problem for a group of microCHPs in a similar way. First we show the complexity of the microCHP planning problem. Next we develop solution techniques for this problem. We explore the possibilities for solving this problem by looking at exact formulations and heuristics.

3.2.2 3-PARTITION

In the previous subsection we presented an overview of the complexity classes \mathcal{P} and \mathcal{NP} and we gave an example of an \mathcal{NP} -complete problem, including computational results for different methodologies that can be applied to such a problem. In this subsection another classical \mathcal{NP} -complete problem is introduced, which we use to prove that the planning problem for a group of microCHPs is \mathcal{NP} -complete itself. This problem is called 3-PARTITION and has the following form, as described by [56]:

3-PARTITION

INSTANCE: Given is a set A of $3m$ elements, a bound $B \in \mathbb{Z}^+$, and a size $s(a) \in \mathbb{Z}^+$ for each $a \in A$ such that $\frac{B}{4} < s(a) < \frac{B}{2}$ and $\sum_{a \in A} s(a) = mB$.

QUESTION: Can A be partitioned into m disjoint sets A_1, A_2, \dots, A_m such that, for $1 \leq i \leq m$, $\sum_{a \in A_i} s(a) = B$?

The decision problem consists of the question whether m bins of size B can be exactly filled with the given $3m$ elements. These elements have an integer size that is larger than $\frac{B}{4}$ and smaller than $\frac{B}{2}$; elements have to be completely assigned to exactly one bin. When four or more elements are assigned to a certain bin, this can never be part of a feasible solution to the 3-PARTITION problem, since the sum of the sizes in this particular bin is strictly larger than B in this case. When two or less elements are assigned to a certain bin, this can never be part of a feasible solution, since the sum of the sizes in the bin is now strictly smaller than B . This leads to the observation that all bins must contain exactly 3 elements to allow the possibility of having a feasible solution to the 3-PARTITION problem. The name of the problem originates from this observation: all elements need to be partitioned in disjoint sets of 3 elements, such that these sets all have equal sums of the element sizes.

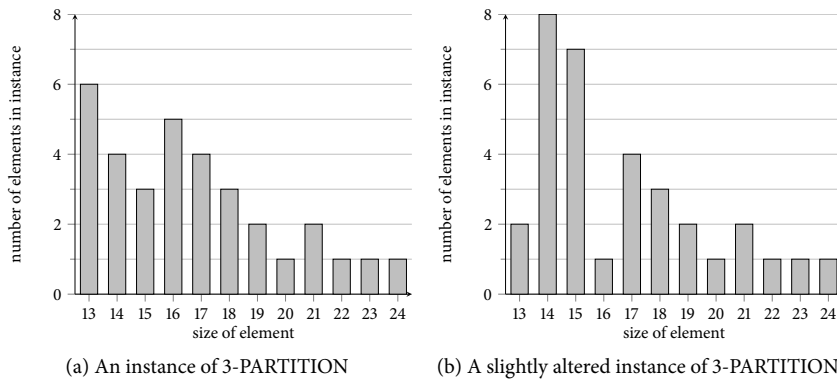


Figure 3.8: Two instances of 3-PARTITION

As an example of 3-PARTITION, we formulate an instance, which consists of 33 elements and a bound $B = 50$ for the 11 bins that have to be filled. The size $s(a)$ of each element a can be picked from the following set of allowed element sizes: $s(a) \in \{13, 14, \dots, 24\}$. The numbers of elements for each size in this instance are shown in Figure 3.8a. The sum of all element sizes equals 550, which at least does not exclude the existence of a 3-PARTITION. The question remains whether a feasible partitioning can be found.




Figure 3.9: One of 16 feasible partitions in the given 3-PARTITION example

Figure 3.9 shows a solution to this particular instance, where the distribution of the elements over the bins is depicted. It turns out that in total 16 possible solutions exist for this instance. If we now alter the instance slightly by removing 4 elements of size 13 and 4 of size 16, and adding 4 elements of size 14 and 4 of size 15, the sum of all element sizes does not change and neither does the number of elements as can be seen in Figure 3.8b. However, for this slightly altered instance no feasible solution exists. This small example shows the essence of the difficulty of 3-PARTITION.

3.2.3 COMPLEXITY OF THE MICROCHP PLANNING PROBLEM

Until now, the microCHP planning problem has been only described in words. In Section 3.3 we give a more detailed description and a mathematical modelling of this planning problem. Here we give a simplified version of one of the mentioned versions of the decision problem leaving out details on how the inputs are precisely generated. We show that already this simple version is \mathcal{NP} -complete in the strong sense.

The microCHP planning problem considers N microCHPs (houses). Each of these microCHPs has a finite set of (feasible) local production patterns. More

formally, for each house $n = 1, \dots, N$ a set of production patterns C_n is given. Each pattern $p \in C_n$ is a $\{0,1\}$ vector of dimension N_T , specifying the use of the microCHP in the different time intervals, whilst fulfilling all local (household) constraints of the planning problem (i.e. p is a feasible solution for the standalone household problem of house n). In this way, the constraints of the local houses are already incorporated in the sets C_1, \dots, C_N , and the only constraint that is left for the global planning problem is to respect the global predefined electricity production bounds $P^{upper} = (P_1^{upper}, \dots, P_{N_T}^{upper})$ and $P^{lower} = (P_1^{lower}, \dots, P_{N_T}^{lower})$. To formalize these constraints, let $pe(p)$ be the vector of generated electricity, corresponding to the production pattern p (note, that $pe(p)$ is independent of the actual house for which p is used as pattern!). To respect the production bounds P^{upper} and P^{lower} , for each house $n = 1, \dots, N$ a production pattern $p_n \in C_n$ has to be chosen such that $P_j^{lower} \leq \sum_{n=1}^N pe(p_n)_j \leq P_j^{upper}$ for each $j \in \{1, \dots, N_T\}$. Summarizing, we get the following decision problem:

The microCHP planning problem

INSTANCE: Given is a collection of sets C_1, C_2, \dots, C_N of N_T -dimensional binary production patterns, an electricity generation function pe and target electricity production bounds $P^{upper} = (P_1^{upper}, \dots, P_{N_T}^{upper})$ and $P^{lower} = (P_1^{lower}, \dots, P_{N_T}^{lower})$.

QUESTION: Is there a selection of production patterns $p_n \in C_n$ for each $n = 1, \dots, N$, such that $P_j^{lower} \leq \sum_{n=1}^N pe(p_n)_j \leq P_j^{upper}$ for each $j \in \{1, \dots, N_T\}$?

Figure 3.10 gives an example of the output of the microCHP planning problem. For a 24 hour time horizon it depicts the planned on/off operation in time (horizontally) for 10 different microCHP appliances (vertically). Looking at the




Figure 3.10: An example of the output of the microCHP planning problem

combined generation we see that in this example at any moment in time not more than 5 microCHP appliances are switched on simultaneously.

In the following we prove the complexity of this microCHP planning problem.

Theorem 1 *The microCHP planning problem is \mathcal{NP} -complete in the strong sense*

Proof The problem whether a feasible match exists between the production bounds (P^{lower}, P^{upper}) and the sum of possible electricity production patterns of all houses is proven to be \mathcal{NP} -complete in the strong sense by reducing 3-PARTITION to the microCHP planning problem.

First, it is clear that the microCHP planning problem belongs to \mathcal{NP} , since feasibility can be verified within polynomial time, once production patterns are chosen for each microCHP. The task that is left to do is to reduce 3-PARTITION to the microCHP planning problem. To do this we construct a specific instance of the microCHP planning problem and show that this instance corresponds to a general instance of 3-PARTITION, and that this transformation is done in pseudopolynomial time. Note that it is sufficient to use a pseudopolynomial reduction to prove \mathcal{NP} -completeness in the strong sense.

The specific instance of the microCHP planning problem that corresponds to a general instance of 3-PARTITION is as follows. First, the time horizon consists of $2mB$ time intervals. Next, for each element $a \in A$ of the 3-PARTITION problem, a cluster C_a is created with $m(B - s(a) + 1)$ production patterns. So we have $N = 3m$ houses. Each of the $m(B - s(a) + 1)$ patterns in cluster C_a has a sequence of $s(a)$ consecutive 1's at time intervals (see Figure 3.11). The dark gray areas correspond to sequences of 1's and light gray areas to sequences of 0's. Note, that the patterns are chosen such that only production in the periods $[(2i+1)B, 2(i+1)B]$, $i = 0, \dots, m-1$ is possible for the created houses. If MR is chosen as the smallest element of the 3-PARTITION instance and if the heat demand is such that at the end of the day the microCHP had to run for $s(a)$ time intervals in house a , all production patterns p are feasible for the microCHP model (note that MO is not important, since each pattern contains only one run). The production function is defined by $pe(p)_j = E_{\max} p_j$ (meaning that startup and shutdown periods are ignored), and the target production plan by:

$$P_j = P_j^{upper} = P_j^{lower} = \begin{cases} E_{\max} & (2i+1)B < j \leq 2(i+1)B \\ & \text{for some } i \in \{0, \dots, m-1\} \\ 0 & \text{otherwise.} \end{cases} \quad (3.12)$$

This choice implies that for each house a now exactly one planning pattern from C_a must be chosen. Due to the definitions of P_j and pe , these patterns must be chosen such that two patterns never overlap and in all intervals within the m periods $[(2i+1)B, 2(i+1)B]$, $i = 0, \dots, m-1$ of length B , exactly one pattern has to be active. This comes down to assigning to each interval $[(2i+1)B, 2(i+1)B]$ non overlapping patterns of total length exactly B . Since furthermore for each house exactly one pattern is used in this process, a feasible solution of the microCHP planning problem instance exists if and only if 3-PARTITION has a solution. Thus, the constructed instance of the microCHP planning problem corresponds to a general instance of 3-PARTITION.

The used reduction is clearly pseudo-polynomial in the size of the 3-PARTITION instance, but, as mentioned, this is sufficient to prove the result of the theorem. ■

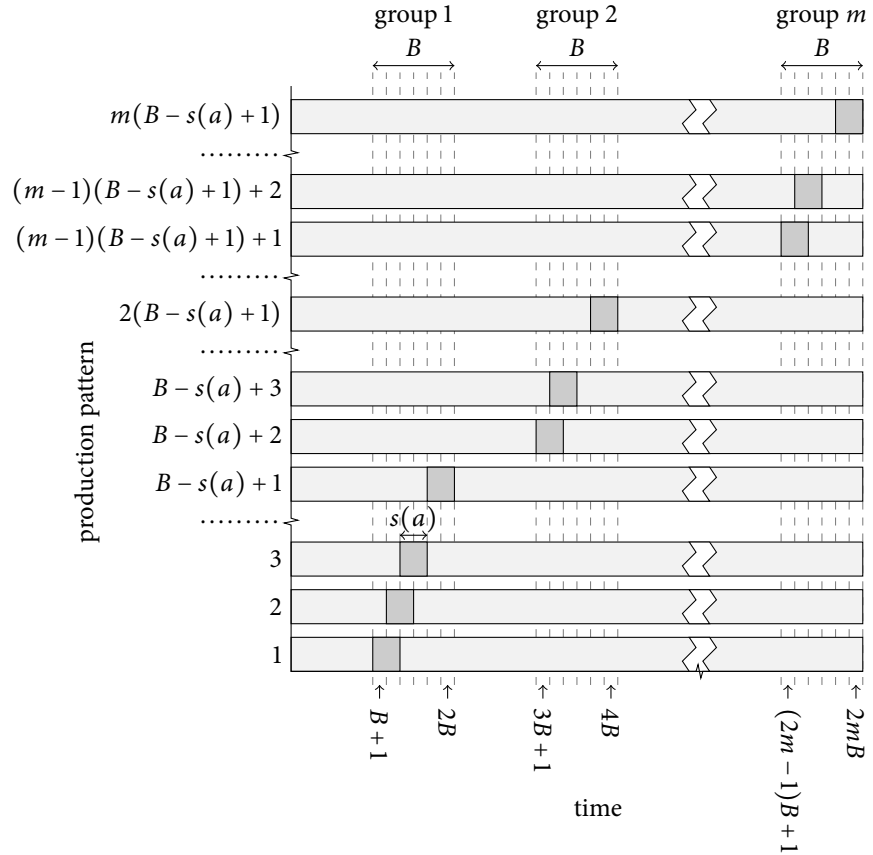


Figure 3.11: The cluster C_a , consisting of $m(B - s(a) + 1)$ production patterns for the house corresponding to the element a of length $s(a)$.

The construction in the proof is limited to only one run per day for each house and the minimum runtime depends on the smallest element a , which does not represent a very realistic instance. In real world instances, a microCHP has multiple runs on a single day, due to a large heat demand and a relatively small heat buffer, that does not allow to produce the complete heat demand in a single long run. To indicate that also real world instances include the properties, which make the microCHP planning problem hard, we construct a more realistic but also more complicated instance that broadens the limitations that are used in the proof. For this example we use each element a of 3-PARTITION in $B - s(a) + 1$ houses; each of them containing $m + 1$ production patterns, and in total we use $\sum_{i=1}^{|A|} B - s(a_i) + 1$ houses as in Figure 3.12. Each house n has a basic pattern p_n^b , representing the runs

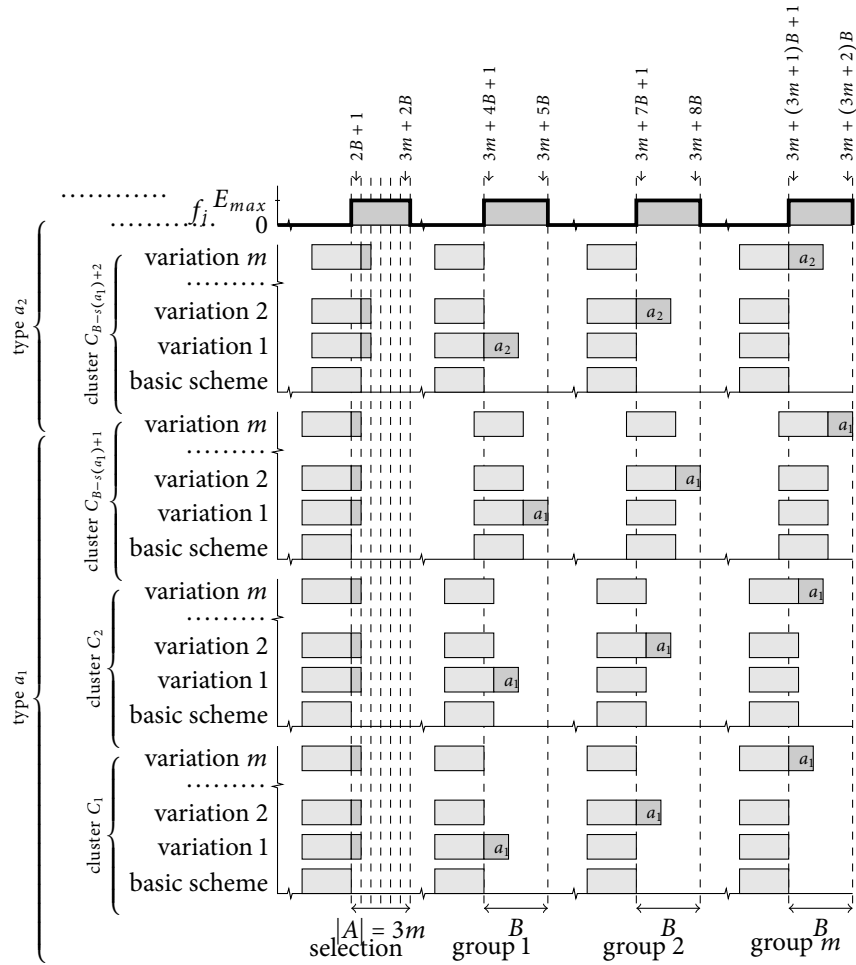


Figure 3.12: Production patterns in a more realistic example

of a normal day within a time horizon of $3m + (3m + 2)B$ time intervals. Next to the basic pattern, each house has m variations on this basic pattern, in which this basic pattern is copied and some adjacent production is done, as in Figure 3.12. We assume that heat demand and buffer level constraints are fulfilled, and that there is enough space left in the heat buffer to run for the additional $s(a) + 1$ time intervals for the given house. The periods $[0, B]$ and $[3m + (3i - 1)B, 3m + 3iB]$, $i = 1, \dots, m$ are left idle in all patterns. Production is allowed in the periods $[B, 3m + 2B]$ and $[3m + 3iB, 3m + (3i + 2)B]$, $i = 1, \dots, m$, where a run of length MR is positioned precisely in front of the runs of length $s(a)$ and the run of length 1. Obviously, these

runs fulfill minimum runtime and offtime constraints if we choose $MR = MO \leq B$. The first run of the patterns in each cluster C_a has a special form. For each cluster we want to select a variation pattern that has additional generation compared to the basic pattern. To derive this, we designed a so-called selection section of length $|A| = 3m$ (see Figure 3.12). In the selection section exactly one 1 is added at the same time interval, for each cluster of microCHPs corresponding to the same $a \in A$. The target production plan is defined in a similar way as Equation (3.12):

$P_j = P_j^{upper} = P_j^{lower} = \sum_{n=1}^N p_n^b + f_j$, where $pe(p)_j = E_{\max} p_j$ (startup and shutdown periods are neglected again) and

$$f_j = \begin{cases} E_{\max} & 2B < j \leq 3m + 2B \text{ or} \\ & 3m + (3i + 1)B < j \leq 3m + (3i + 2)B \text{ for some } i \in \{1, \dots, m\} \\ 0 & \text{otherwise.} \end{cases} \quad (3.13)$$

Equation (3.13) is given in the top of Figure 3.12. Due to the definition of P_j and the design of the selection section exactly one varied pattern belonging to a must be chosen from the $m(B - s(a) + 1)$ variations based on the element a . Thus, only one of the corresponding $B - s(a) + 1$ houses does not select its basic pattern. Therefore all elements a are chosen exactly once, and they must fill the m periods of length B in the same way as in the given proof. This example shows that we can construct also a more realistically structured instance that has a direct correspondence to 3-PARTITION.

3.2.4 OPTIMIZATION PROBLEMS RELATED TO THE MICROCHP PLANNING PROBLEM

As mentioned before we consider two types of optimization problems that are related to the decision problem shown to be \mathcal{NP} -complete in the previous section.

In the first type of optimization problem we want to maximize the profit that is made on an electricity market with (predicted) prices $\pi = (\pi_1, \dots, \pi_{N_T})$.

Maximizing the profit on an electricity market

INSTANCE: Given is a collection of sets C_1, C_2, \dots, C_N of N_T -dimensional binary production patterns satisfying the operational requirements of the corresponding households, an electricity generation function pe , target electricity production bounds $P^{upper} = (P_1^{upper}, \dots, P_{N_T}^{upper})$ and $P^{lower} = (P_1^{lower}, \dots, P_{N_T}^{lower})$ and an electricity price π .

Maximizing the profit on an electricity market (continued)

OBJECTIVE: Maximize the profit that can be made on the electricity market while satisfying domestic (operational) and fleet (cooperational) constraints:

$$\max \sum_{j=1}^{N_T} (\pi_j \sum_{n=1}^N pe(p_n)_j)$$

where $P_j^{lower} \leq \sum_{n=1}^N pe(p_n)_j \leq P_j^{upper}$ for each $j \in \{1, \dots, N_T\}$
and $p_n \in C_n$ for each $n \in \{1, \dots, N\}$.

In the second type of optimization problem we introduce slack and excess variables sl and ex that measure the deviation from the bounds on the target electricity production. The sum of slack and excess over the full planning horizon is minimized, while respecting the adjusted cooperational requirements.

Minimizing the deviation from the given bounds on aggregated electricity output

INSTANCE: Given is a collection of sets C_1, C_2, \dots, C_N of N_T -dimensional binary production patterns $p = (x_1, \dots, x_{N_T})$ satisfying operational requirements, a corresponding electricity generation function pe , target electricity production bounds $P^{upper} = (P_1^{upper}, \dots, P_{N_T}^{upper})$ and $P^{lower} = (P_1^{lower}, \dots, P_{N_T}^{lower})$.

OBJECTIVE: Minimize the deviation from the target electricity bounds P^{upper} and P^{lower} :

$$\min \sum_{j=1}^{N_T} (sl_j + ex_j)$$

where $\sum_{n=1}^N pe(p_n)_j - ex_j \leq P_j^{upper}$ for each $j \in \{1, \dots, N_T\}$ and

$P_j^{lower} \leq \sum_{n=1}^N pe(p_n)_j + sl_j$ for each $j \in \{1, \dots, N_T\}$ and
 $p_n \in C_n$ for each $n \in \{1, \dots, N\}$.

3.3 AN INTEGER LINEAR PROGRAMMING FORMULATION

In this section we model the two versions of the microCHP planning problem by an Integer Linear Programming (ILP) formulation. This ILP formulation is used to explain the different requirements of the underlying problem in more detail. After modelling the problem as an ILP, we discuss some small benchmark instances and the solutions to these instances and draw conclusions on the applicability of ILP in practical situations.

In practice a decision maker is completely free to instantaneously switch on or switch off a microCHP at any moment in time. However in our model, we discretize the time and allow a decision maker only to switch on or off the microCHP for complete time intervals. The discretization of the time horizon on the one hand leads to a simpler model, but on the other hand, the short term electricity market also works with time intervals, hence a discretization of time also matches the context the problem is used in. More precisely, we divide the planning horizon $[0, T]$ of the microCHP planning problem into N_T time intervals $[t_k, t_{k+1}]$ of equal length $\frac{T}{N_T}$. The decision to have a microCHP on or off is made for a complete interval $[t_k, t_{k+1}]$. As a consequence of this, we introduce decision variables x_j^i for the intervals j and microCHPs i :

$$x_j^i = \begin{cases} 1 & \text{if the } i\text{th microCHP is on during interval } j \\ 0 & \text{if the } i\text{th microCHP is off during interval } j, \end{cases} \quad (3.14)$$

where interval j is the interval $[t_{j-1}, t_j]$, $j = 1, \dots, N_T$. A solution to the operational planning problem of a single house i is a vector $x^i = (x_1^i, \dots, x_{N_T}^i) \in X_{1,2}^i$, where $X_{1,2}^i \subseteq \{0, 1\}^{N_T}$ is the N_T -dimensional space of possible binary decision variables respecting appliance specific and operational constraints. In case the objective is profit maximization, a solution to the microCHP planning problem is a combination of domestic solutions $x = (x^1, \dots, x^N) \in \tilde{Y}_{1,2}$ and in case the objective is to minimize the deviation from the target electricity production bounds, it is a vector $x = (x^1, \dots, x^N) \in Y_{1,2}$.

In the following we transform this general description of a solution to constraints formulated by linear inequalities using additional (integer) variables. To start we request that the variables x_j^i are binary decision variables:

$$x_j^i \in \{0, 1\} \quad \forall i \in I, \forall j \in J. \quad (3.15)$$

We use the notation I to represent the set of houses $I = \{1, \dots, N\}$ and J for the set of intervals $J = \{1, \dots, N_T\}$. Whenever an equation is not applied to all intervals in the planning horizon or to intervals that are situated outside the planning horizon, this is explicitly denoted. We furthermore define binary parameters \bar{x}_j^i that represent the given behaviour of the microCHP in the short term history before the start of the planning period (i.e. $j = 0, -1, -2, \dots$). This information is used to guarantee a correct transition between a current (realization of a) planning and the first couple of intervals of the planning horizon. Next we discuss the three types of requirements for the planning of the microCHPs.

Appliance specific constraints

A microCHP appliance has specific startup and shutdown behaviour and a heat to electricity ratio (as explained in Section 3.1.2), that define the heat and electricity output of a run. We have to model this behaviour by linear constraints. For this,

let the parameter G_{\max}^i characterize the heat generation for a time interval if the microCHP of house i is running at full power, and let a value α^i specify the ratio between electricity and heat generation. Furthermore, each microCHP has two vectors: $\hat{G}^i = (\hat{G}_1^i, \dots, \hat{G}_{N_{up}^i}^i)$, giving the loss of the heat generation during the startup intervals and $\check{G}^i = (\check{G}_1^i, \dots, \check{G}_{N_{down}^i}^i)$, giving the heat generation that still occurs during the shutdown intervals, where N_{up}^i and N_{down}^i give the number of intervals that it takes to startup and shutdown respectively. The heat generation g_j^i in time interval $j \in J$ for house $i \in I$ is now given by:

$$g_j^i = G_{\max}^i x_j^i - \sum_{k=0}^{N_{up}^i-1} \hat{G}_{k+1}^i start_{j-k}^i + \sum_{k=0}^{N_{down}^i-1} \check{G}_{k+1}^i stop_{j-k}^i \quad \forall i \in I, \forall j \in J, \quad (3.16)$$

where $start_j^i$ and $stop_j^i$ are additional binary start and stop variables, indicating if in an interval the decision is made to start the microCHP or to turn it off. The generation of electricity e_j^i follows from g_j^i by:

$$e_j^i = \alpha^i g_j^i \quad \forall i \in I, \forall j \in J. \quad (3.17)$$

The binary variables $start_j^i$ and $stop_j^i$ are not additional decision variables, but variables depending on the decision variables x_j^i . To ensure that the variables $start_j^i$ and $stop_j^i$ are consistent with the x -variables, constraints (3.18)-(3.25) are added. If necessary ($j < 1$), the run history \bar{x} is used in these equations by defining $x_j^i = \bar{x}_j^i$.

$$start_j^i \geq x_j^i - x_{j-1}^i \quad \forall i \in I, j = 2 - MR^i, \dots, N_T \quad (3.18)$$

$$start_j^i \leq x_j^i \quad \forall i \in I, j = 2 - MR^i, \dots, N_T \quad (3.19)$$

$$start_j^i \leq 1 - x_{j-1}^i \quad \forall i \in I, j = 2 - MR^i, \dots, N_T \quad (3.20)$$

$$stop_j^i \geq x_{j-1}^i - x_j^i \quad \forall i \in I, j = 2 - MO^i, \dots, N_T \quad (3.21)$$

$$stop_j^i \leq x_{j-1}^i \quad \forall i \in I, j = 2 - MO^i, \dots, N_T \quad (3.22)$$

$$stop_j^i \leq 1 - x_j^i \quad \forall i \in I, j = 2 - MO^i, \dots, N_T \quad (3.23)$$

$$start_j^i \in \{0, 1\} \quad \forall i \in I, j = 2 - MR^i, \dots, N_T \quad (3.24)$$

$$stop_j^i \in \{0, 1\} \quad \forall i \in I, j = 2 - MO^i, \dots, N_T \quad (3.25)$$

Note that the parameters MR^i and MO^i which are used in Equations (3.18)-(3.25), are not defined yet. For now it suffices to know that $N_{up}^i \leq MR^i$ and $N_{down}^i \leq MO^i$, which implies that the necessary $start$ and $stop$ variables for (3.16) are at least specified. The parameters MR^i and MO^i are explained below as part of the operational constraints. To characterize them, in some cases we need additional information on the short term history of the $start$ and $stop$ variables, resulting in the use of MR^i and MO^i instead of N_{up}^i and N_{down}^i .

Table 3.2 shows how constraints (3.18)-(3.23) force the variables $start_j^i$ and $stop_j^i$ to take their correct values, depending on the x_j^i variables. The four possible

x_{j-1}^i	x_j^i	eq. (3.18)	eq. (3.19)	eq. (3.20)	$start_j^i$	eq. (3.21)	eq. (3.22)	eq. (3.23)	$stop_j^i$
0	0	≥ 0	≤ 0	≤ 1	0	≥ 0	≤ 0	≤ 1	0
0	1	≥ 1	≤ 1	≤ 1	1	≥ -1	≤ 0	≤ 0	0
1	0	≥ -1	≤ 0	≤ 0	0	≥ 1	≤ 1	≤ 1	1
1	1	≥ 0	≤ 1	≤ 0	0	≥ 0	≤ 1	≤ 0	0

Table 3.2: The construction of *start* and *stop* variables from consecutive *x* variables

combinations of x_j^i and x_{j-1}^i result in the given right hand sides of the three start and three stop constraints. These right hand sides determine the correct values for $start_j^i$ and $stop_j^i$, when we also respect the binary requirements of Equations (3.24) and (3.25).

Operational constraints

Contrary to other electricity generators (especially compared to the operation of a power plant) the electrical output of a microCHP is completely determined by the decisions to switch the appliance on or off; an operating range does not exist. Given a feasible sequence of binary decision variables x , the appliance specific constraints describe a direct and unique output for the microCHP. To force x to be a feasible sequence we have to respect the minimum runtime and minimum offtime requirements, as well as the correct functioning of the heat buffer. The minimum runtime constraint demands that the microCHP has to run for at least MR^i consecutive intervals, once a choice is made to switch it on. The minimum offtime constraint demands that the microCHP has to stay off for at least MO^i consecutive intervals, once a choice is made to switch it off. As we have mentioned before in Section 3.1.2, it is completely natural to demand that $N_{up}^i \leq MR^i$ and $N_{down}^i \leq MO^i$.

The minimum runtime constraint can be modelled by (3.26), which forces the decision variable x_j^i to be 1 if one start occurs in the previous $MR^i - 1$ intervals, since x_j^i is only allowed to take the values 0 and 1. Likewise, equation (3.27) forces the decision variable x_j^i to be 0 if one stop occurs in the previous $MO^i - 1$ intervals. Again, if needed the given *start* and *stop* variables from the past (following from the given \bar{x} values) are used.

$$x_j^i \geq \sum_{k=j-MR^i+1}^{j-1} start_k^i \quad \forall i \in I, \forall j \in J \quad (3.26)$$

$$x_j^i \leq 1 - \sum_{k=j-MO^i+1}^{j-1} stop_k^i \quad \forall i \in I, \forall j \in J \quad (3.27)$$

Note, that after a start of the microCHP, it takes at least MR^i intervals before a stop may occur. Since furthermore between two consecutive starts one stop occurs, we never can have more than one start in MR^i consecutive intervals. Similar reasoning learns that we never can have more than one stop in MO^i consecutive intervals.

To specify the constraints resulting from the heat demand, we introduce variables hl_j^i specifying the heat level in the buffer of house i at the beginning of interval j . For the first interval, this level is given by the initial heat level BL^i (equation (3.28)). The heat demand of house i is characterized by a heat demand vector $H^i = (H_1^i, \dots, H_{N_T}^i)$. Next to the parameter BL^i to describe the initial heat level in the buffer, a value BC^i to describe the buffer capacity and a value K^i to describe the heat loss parameters for the buffer are used. This heat loss is assumed to be constant for all intervals, since we assume that the temperature range in which the heat buffer is operated is not too large. The change of the heat level in interval j is given by the amount of generated heat (g_j^i) minus the heat demand (H_j^i) and the loss parameter (K^i) (see equation (3.29)). Finally, the capacity of the heat buffer has to be respected (equation (3.30)).

$$hl_1^i = BL^i \quad \forall i \in I \quad (3.28)$$

$$hl_j^i = hl_{j-1}^i + g_{j-1}^i - H_{j-1}^i - K^i \quad \forall i \in I, \forall j \in J \setminus \{1\} \cup \{N_T + 1\} \quad (3.29)$$

$$0 \leq hl_j^i \leq BC^i \quad \forall i \in I, j \in J \cup \{N_T + 1\} \quad (3.30)$$

Cooperational constraints

The equations (3.26)-(3.30) give the constraints for a feasible domestic decision sequence. The total electricity output of the group of microCHPs is specified by lower and upper bound vectors $P^{lower} = (P_1^{lower}, \dots, P_{N_T}^{lower})$ and $P^{upper} = (P_1^{upper}, \dots, P_{N_T}^{upper})$ for the production pattern of the fleet. The constraints on the global production pattern can be formulated as follows:

$$\sum_{i=1}^N e_j^i \leq P_j^{upper} \quad \forall j \in J \quad (3.31)$$

$$\sum_{i=1}^N e_j^i \geq P_j^{lower} \quad \forall j \in J. \quad (3.32)$$

In the above form, constraints (3.31) and (3.32) are hard constraints and demand that the total production aggregates to an amount that lies between the lower and upper bounds. These constraints are used when the optimization objective is to maximize profit on an electricity market as in the profit maximization problem defined in Section 3.2.4.

When we relax this problem to the deviation minimization problem of finding a total production that is the closest to the given bounds, we need slightly modified constraints. For these constraints we introduce slack and excess variables sl_j and

ex_j :

$$\sum_{i=1}^N e_j^i - ex_j \leq P_j^{upper} \quad \forall j \in J \quad (3.33)$$

$$\sum_{i=1}^N e_j^i + sl_j \geq P_j^{lower} \quad \forall j \in J \quad (3.34)$$

$$ex_j \geq 0 \quad \forall j \in J \quad (3.35)$$

$$sl_j \geq 0 \quad \forall j \in J. \quad (3.36)$$

The excess and slack variables account for the deviation from the range $[P_j^{lower}, P_j^{upper}]$ instead of the deviation from the points P_j^{lower} and P_j^{upper} . Equations (3.35) and (3.36) are necessary to prevent that values within this range are pulled towards the boundaries.

Objectives and optimization problems

In the previous all constraints for the two planning problems have been specified. Now we deal with the objective functions. For the profit maximization problem we have given the electricity prices on an electricity market, specified by a price vector $\pi = (\pi_1, \dots, \pi_{N_T})$. The objective function is to maximize the profit on this electricity market:

$$z_{\max} = \max \sum_{j=1}^{N_T} \sum_{i=1}^N \pi_j e_j^i. \quad (3.37)$$

For the deviation minimization problem the objective is given by the minimization of the total slack and excess:

$$z_{\min} = \min \sum_{j=1}^{N_T} sl_j + ex_j. \quad (3.38)$$

This objective demands the slack and excess variables to take their minimal values such that (3.33) and (3.34) are respected.

The profit maximization problem (*Maximizing the profit on an electricity market*) is now defined by objective (3.37) and constraints (3.15)-(3.32). The deviation minimization problem (*Minimizing the deviation from the given bounds on aggregated electricity output*) is given by objective (3.38) and constraints (3.15)-(3.30) and (3.33)-(3.36). These optimization problems are studied in more detail in the following sections.

The size of the problem is determined by the planning horizon, specified by the number of intervals N_T , and the number of houses forming the fleet, denoted by N . The ILP problem has $N \times N_T$ binary decision variables x_j^i and $O(N \times N_T)$ constraints and depending variables, and the existence of constraints in both time and space clearly shows the two-dimensionality of the problem.

The input of the microCHP planning problem consists of numbers specifying the dimensions of the problem and data specifying characteristic behavior within the problem. We use different sets of benchmark instances to test varying solution methods for the problems that we have described in the previous subsection. At this point we give an overview of these benchmark sets and we indicate the main differences between them. First we give a global comment on the dimensions of the problem and on the type of data that is used. Then we describe the different benchmark sets.

Dimensions

As mentioned before, the two dimensions of the problems are time and space. Although they are both of importance in the structure of the problem, the nature of these dimensions in practice may ask for a slight focus shift towards space (i.e. the number of microCHPs in the problem).

The microCHP planning problem concentrates on planning for a time horizon of one day, i.e. 24 hours. Since short term electricity markets work with bidding blocks of one hour in The Netherlands [2] the interval length of the planning problem should comply to this hourly basis. According to [131] an interval length of 5 minutes “seems a reasonable compromise to give good accuracy with reasonable data volume”, for the evaluation of electrical on-site generation. This interval length of 5 minutes is used to allow for a large variation that is usually present in the electrical load profiles of houses. For the planning problem however, the electrical production of the microCHP is more stable in its output, due to the requirement to run for at least a minimum time MR , which is typically set to half an hour. This indicates that the planning problem itself does not need to deal with variable load and accompanying fluctuating electricity import/export. If measurement technology is available to account for all locally generated electricity, as mentioned in the business case in Section 2.2.2, it is possible to auction all locally generated electricity on the market, instead of auctioning the measured import/export of houses. The heat demand that needs to be fulfilled is predicted in hourly intervals. In this setting of hourly heat demand and half hourly generation requirements the need for an interval length of 5 minutes may be relaxed and half an hour seems a more appropriate interval length. Since the heat demand is predicted in hourly intervals, we also study instances with an hourly interval length. To allow for some flexibility in the local assignment of production, we also use instances with an interval length of 15 minutes. This gives the planner more opportunities to set the starting point of a microCHP run and more possibilities to apply longer runs. Based on the above we use three different interval lengths in the planning problem: 15 minutes, 30 minutes and 60 minutes. These interval lengths correspond to $N_T = 96$, $N_T = 48$ and $N_T = 24$ intervals.

The number of microCHPs in the problem is subject to more variation. To verify the functional correctness of the different solution methods and compare them to each other, instances with a small number of microCHPs are used (i.e. $N \leq 10$).

However, for practical use a solution method needs to be scalable. Therefore we also use instances where we have $N = 25$, $N = 50$, $N = 75$ and $N = 100$ to analyze scalability aspects of the different methods, and sizes $N = 1000$ to $N = 5000$ to further evaluate promising methods. Instances where $N \leq 10$ are referred to as small instances; instances where $10 < N \leq 100$ as medium instances; and instances where $N > 100$ as large instances.

Data

An instance creation tool has been designed that works independently of the already specified choices for N and N_T in the previous paragraph and, thus, can be used to generate a wide range of instances for the microCHP planning problem. The specific characteristics of instances of the planning problem are described by several parameters, which all have been introduced in the ILP formulation. These parameters can be divided into appliance specific parameters and problem defining parameters. For the appliance specific parameters we usually use values corresponding to the following setting. The microCHP behaviour is modelled according to the use of a Stirling engine developed by Whispergen [15, 43]. However, other microCHPs can be modelled as well. The minimum runtime and the minimum offtime are both set to half an hour. Startup and shutdown periods are 12 minutes and 6 minutes respectively; the electrical output is assumed to increase/decrease linearly in these periods to/from the maximum generation of 1 kW of electrical energy and 8 kW of heat. The values for all parameters are chosen such that they are consistent with these periods (e.g. $N_{up}^i = \lceil \frac{12}{il} \rceil$, where il is the interval length in minutes). The vectors representing the loss of heat generation and the additional heat generation are calculated based on the losses/gains resulting from the 12/6 minutes startup/shutdown periods. We model a heat buffer by specifying a certain range $[BLL^i, BUL^i]$ of the heat capacity HC^i of this buffer; the heat level should stay between the lower heat level BLL^i and the upper heat level BUL^i . This interval may be smaller than the actual capacity of the heat buffer: $BC^i = BUL^i - BLL^i \leq HC^i$. By demanding that the planning stays to this tightened range we leave some flexibility to accommodate for minor fluctuations in realtime. As standard heat buffer we reserve 10 kWh, which corresponds to a heat buffer of around 150 l [79].

The heat demand for the houses is usually given by an hourly prediction of heat usage [29]. In general, we assume that the heat demand consists of central heating and hot tap water demand. This heat profile of a house during winter has two peaks¹, typically one in the morning and one in the evening. To offer benchmarking instances that can be used by other planning methods, we generate reproducible hourly heat demand data. The creation of this heat demand data is explained in detail in Appendix A. The idea of this data creation tool is that we define two periods (one between 7 a.m. and 11 a.m. and one between 6 p.m. and 10 p.m.), during which two peak demands occur. In a winter day the average daily heat demand is assumed

¹Derived from gas usage patterns in The Netherlands

production pattern variant	small instances									
	1	2	3	4	5	6	7	8	9	10
lower bound (%)	0	0	0	0	0	0	0	10	20	tight
upper bound (%)	100	90	80	70	60	50	40	100	100	tight
production pattern variant	medium instances									
	11	12	13	14						
lower bound (%)	0	10	20	25						
upper bound (%)	75	50	40	35						

Table 3.3: Electricity production bounds, based on percentages of possible electricity production

to be 54 kWh, which is typical for a cold day in The Netherlands. Therefore, we aim to create heat demand data that has an average daily demand of 54 kWh.

Another important characteristic of an instance is its definition of the desired production bounds P_j^{lower} and P_j^{upper} . We use two ways of defining these total production bounds.

For the profit maximization problem, we derive P_j^{lower} and P_j^{upper} using constant percentages of the total maximally possible electricity output of the group of houses. The used percentages are given in Table 3.3. The last variant (variant 10) for the small instances gives the tightest combination of lower and upper bounds: the highest lower bound for which a feasible solution is found to the profit maximization problem (variant 1, 8 or 9) is combined with the lowest upper bound for which a feasible solution is found (variant 1-7). For the medium instances we use bounds, specified by the percentages in Table 3.3. The large instances are described in more detail in Chapter 6.

For the profit maximization problem we use flat bounds, which correspond to the objective of minimizing production peaks and thus to the requirements of stability and reliability. For the deviation minimization problem we use fluctuating electricity demand patterns, to see whether the different methods are able to follow such fluctuating bounds on the total production. These fluctuating electricity bounds have two properties. First they are subject to some kind of variation. Secondly, the upper and lower production bounds are relatively close to each other. This indicates that we concentrate more on the ability to follow a predetermined total electricity pattern and less on the flexibility of total generation in certain time intervals. Production bounds that fulfill both properties are created by making use of curves that result from a sine function plus some constant for both the upper and the lower production bounds. The variability of the desired pattern is determined by the sine curve. The relative closeness of these bounds results from the fact that we use the same sine function for both types of bounds and the constants are chosen such that the total maximally and minimally possible production coincides with the bounds on the total desired production. An explanation in more detail is given in Section 3.7.4.

k	l	1	2	3	4	5	6	7	8	9	10	μ
1		1.147	1.092	-	-	-	-	-	-	-	1.092	1.110
2		1.236	1.208	1.016	1.016	1.016	1.016	-	-	-	1.016	1.075
3		1.197	1.197	1.106	1.106	1.002	-	-	-	-	1.002	1.102
4		1.183	1.164	1.128	1.114	1.021	1.021	-	1.009	1.009	0.949	1.066
5		1.164	1.149	1.120	1.060	1.060	-	-	1.118	-	1.023	1.099
6		1.163	1.150	1.130	1.092	1.048	1.027	-	1.139	-	1.021	1.096
7		1.156	1.145	1.137	1.109	1.069	0.972	0.925	1.150	-	0.924	1.065
8		1.156	1.145	1.130	1.114	1.080	1.032	0.919	1.152	1.069	0.902	1.070
9		1.153	1.143	1.121	1.098	1.072	0.993 ^[2]	-	1.150	1.113	0.976 ^[5]	1.091
10		1.176	1.162	1.143	1.122 ^[1]	1.095	1.037 ^[3]	0.948 ^[4]	1.173	1.028	0.945	1.083
μ		1.173	1.156	1.115	1.092	1.051	1.014	0.931	1.127	1.055	0.985	
	[1]	terminated by solver, upper bound 1.144										
	[2]	terminated by solver, upper bound 1.058										
	[3]	terminated by solver, upper bound 1.075										
	[4]	terminated by solver, upper bound 0.989										
	[5]	terminated by solver, upper bound 0.999										

Table 3.4: Objective value for instances $I(k, l)$

Benchmark instances

For the small and medium instances a set of 200 heat demand profiles is generated using Algorithm 5. Using this heat demand set, a subset of these demand profiles is selected to participate in the different problem instances. This selection is simply based on the order in which the houses appear in the creation process of the heat demand. The notation $I(k, l)$ is used to represent an instance with $N = k$ microCHPs and production pattern variant l . For the small instances that aim at profit maximization for the VPP $k, l \in \{1, \dots, 10\}$ and for the medium instances $k \in \{25, 50, 75, 100\}$ and $l \in \{11, \dots, 14\}$. For example, $I(3, 1)$ is a small instance, consisting of 3 microCHPs and total production bounds of 0% and 100% (meaning that a completely independent planning for the 3 microCHPs can be made).

For the large instances an initial set of 5000 heat demand profiles is created. The accompanying choice for production bounds is given in Chapter 6.

3.3.3 ILP RESULTS

For the ILP formulation we use the small instances to give an indication of the practical computational time that is needed to find optimal solutions. The ILP formulation is modelled in AIMMS modelling software using the commercial CPLEX solver (version 12.2).

The normalized objective values for the instances $I(k, l)$ (i.e. the objective values divided by the number of houses), calculated by the ILP approach, are given in Table 3.4. If an instance does not have a feasible solution this is denoted by a dash (-).

Some of the instances with a large number of houses and tight production pattern constraints were terminated by the ILP solver, due to slow convergence towards the best found solution. For these instances, where the ILP solver did not find the optimal solution, the upper bound on the objective, given by the solver,

k	l	1	2	3	4	5	6
1		0.08	0.06	–	–	–	–
2		0.22	0.31	0.70	0.48	0.51	0.95
3		1.33	1.41	1.75	2.50	1.17	–
4		1.34	1.81	2.45	2.05	8.27	7.36
5		5.00	6.06	9.33	45.28	57.41	–
6		6.78	4.88	20.06	38.64	221.17	254.84
7		7.89	16.77	25.11	47.64	839.84	6373.31
8		27.36	43.89	60.72	109.48	396.69	3302.30
9		129.39	200.31	332.52	2382.53	2858.05	7143.67*
10		461.98	1174.94	873.14	6879.08*	17285.17	6704.20*
μ		64.14	145.04	147.31	1056.41	2407.59	3398.09
σ		137.81	348.21	275.37	2185.10	5330.13	3087.73

k	l	7	8	9	10	μ	σ
1		–	–	–	0.08	0.07	0.01
2		–	–	–	1.00	0.60	0.28
3		–	–	–	1.22	1.56	0.46
4		–	5.98	2.78	9.86	4.66	3.05
5		–	28.19	–	467.30	88.37	155.84
6		–	25.52	–	2326.13	362.25	748.13
7		3052.55	9.75	–	1745.06	1346.44	2037.84
8		9918.91	39.03	97.66	17999.19	3199.52	5753.94
9		–	130.49	1270.13	16265.91*	3412.56	5016.03
10		8648.84*	79.89	1765.80	5757.88	4963.09	5080.93
μ		7206.77	45.55	784.09	4457.36		
σ		2982.89	41.38	755.25	6570.37		

Table 3.5: Computational time (in seconds) for instances $I(k, l)$

is also presented below the table. This gives an indication that the ‘large’ small instances are close to the largest ones that can be solved to optimality by the ILP approach. Note that the upper bound of $I(10, 4)$ can be lowered by looking at the objective value of $I(10, 3)$. The average objective values show that the tighter the fleet constraints are, the less money can be earned.

In Table 3.5 the computational times are given, where we only show the times corresponding to the feasible instances. A star (*) denotes an instance that is terminated by the solver premature, without determining the optimality of the solution. The computational times grow extremely fast if the number of houses grows and/or the production pattern bounds get more tight. Also note the large variance in these times under a fixed number of houses or a fixed production pattern variant.

3.3.4 CONCLUSION

The Integer Linear Programming formulation that is presented in this section gives a clear overview of the dependencies in time and the dependencies in space (microCHPs). The discretization of the problem is modelled such, that the granularity of the time horizon can be chosen; the continuous variant of the problem is approached if the interval length goes to 0.

However, solving a fine-grained problem instance (to optimality) is out of question. Results are presented for this formulation for small instances, with a

limited number of microCHPs ($N \leq 10$) and a relatively large interval length of one hour. These results are used as a comparison benchmark for heuristics.

3.4 DYNAMIC PROGRAMMING

In the previous section we have given an ILP formulation of the microCHP planning problem, which gives an intuitive model of the underlying problem and gives us basic insights into the difficulties of the problem. To improve on the computational speed of finding a solution we develop different heuristics. An interesting approach for these heuristics is the use of dynamic programming.

Dynamic programming is one of the techniques that is applied to large optimization problems that are structured in such a way, that they can be divided into subproblems that are easier to solve. In Section 3.2.1 we presented the Held-Karp algorithm; this is a good example of a dynamic programming method, since it possesses the main ingredients of dynamic programming. In general, a dynamic programming method consists of so-called phases, which are ordered sets of states. For each state in a phase several decisions may be possible, all of which indicate a transition from this state to a state in the successive phase. A cost is associated with each decision; this cost only depends on the current state and the corresponding state transition and is independent of previous states or decisions. This means that all information that is necessary to derive this cost is available in the description of the state. Furthermore, each state has a certain value. This value represents the optimal sum of costs leading to the state, either calculated starting from the first phase or from the final phase. The first way of value calculation is called forward dynamic programming and the latter is called backward dynamic programming. The nice property of dynamic programming is that the value of each state has to be calculated only once. In an iterative way all phases are visited in the order they appear (forwards or backwards) and all states are updated using the values of the states in the neighbouring phase and the costs associated with the decisions via a recursive function.

In the Held-Karp algorithm the phases are determined by the size of the subsets of cities, and the state is given by a subset of cities and the city that is currently visited as the last of these cities. Figure 3.13 shows the structure of the state space belonging to an instance of 4 cities numbered 1 to 4. One can easily see that a state transition only occurs between states of neighbouring phases. The Held-Karp algorithm updates all phases subsequently in the way we explained in Section 3.2.1. An important fact to remember is that the state (S, j) allows any order of visiting the cities in S as long as we end up in j , which reduces the amount of states enormously and also shows the independence between a decision that has to be made for this state and the historic decisions leading to this state. This shows that it is extremely important to define relevant states in which as much information is compressed as possible.

In the following we use these requirements for a dynamic programming method to develop a basic dynamic programming algorithm that solves the microCHP plan-




Figure 3.13: The structure of dynamic programming by example of the Held-Karp algorithm

ning problem exactly. We formulate a state description that compresses historic decision paths, hereby reducing the state space enormously. Although we do not expect that this basic dynamic programming method is applicable to real life problems, it forms the heart of a local search method that is explained in Section 3.5.

3.4.1 BASIC DYNAMIC PROGRAMMING

Before we explain our choice for the description of a state, we first observe the following. Since the problem consists of N microCHPs and N_T time intervals we have 2^{NN_T} combinations of possible binary decisions. A straightforward choice for

a state description is to denote the state by a matrix $A = \begin{pmatrix} a_{11} & \dots & a_{1N_T} \\ \vdots & \ddots & \vdots \\ a_{i1} & \dots & a_{iN_T} \end{pmatrix}$ or

a matrix $B = \begin{pmatrix} b_{11} & \dots & b_{1j} \\ \vdots & \ddots & \vdots \\ b_{N1} & \dots & b_{Nj} \end{pmatrix}$ consisting of all made decisions $a_{ij} = b_{ij} = x_j^i$.

Matrix A belongs to a phase that is based on the number of microCHPs (for a given microCHP a complete decision path is given) and matrix B to a phase that is based on time (for a given interval the decisions for all microCHP are given). A natural choice for recursion is to incrementally add a decision path for a microCHP to A to create states in the next phase or to add all decisions for the next time interval to B . This leads to $\mathcal{O}(2^{N \times N_T})$ possible states.

The first choice that we have to make in order to specify a state is to determine the basis for the different phases: do we follow the idea of matrix A or B ? Since we deal

with a two-dimensional problem we essentially have the choice between two phase indicators: time and space. If the choice would fall onto the latter one, this would make the description of a state complicated. In this case a state transition should describe the decisions for one microCHP for the complete time horizon. This is not easily represented in another way than by using a vector of length N_T . A state could be described by the total production of microCHPs $1, \dots, i$ for the N_T intervals, which does not improve on the order of magnitude of the states. Furthermore, the choice for using space as a phase descriptor does not naturally correspond to the feasibility checks that have to be performed in order to see whether a state transition does not violate any appliance, operational or cooperational requirement. The appliance and operational constraints namely depend strongly on the short term behaviour in time, whereas the cooperational requirements cannot be guaranteed until all microCHPs are planned. Because of these reasons we focus on time as our phase descriptor. In the following a description is given of the state representation, where we first focus on a single microCHP and then combine them into a dynamic programming method that deals with a group of microCHPs.

Dynamic programming for a single microCHP

A disadvantage of the straightforward state representation by matrix B is that it takes the complete decision history explicitly into account. Since rows in B are not mutually exchangeable due to different underlying heat demand (or possibly different generator types) we cannot reduce the state B by compressing rows. However, we can compress columns. To see this we focus on a single row of matrix B , i.e. the operation of a single microCHP appliance from the start of the time horizon up to interval j . Figure 3.14 shows two possible representations of decision paths until interval j : this corresponds to states $(0, 0, 1, 1, 1, 1, 1, 0, 0, 0, 1, 1, 1, 1, 1, 0, 0, 1, 1, 1, 1, 0, 0, 1, 1, 1)$ and $(1, 1, 1, 1, 0, 0, 0, 0, 1, 1, 1, 1, 1, 1, 0, 0, 0, 1, 1, 1)$ in B . Of course a decision for interval j




Figure 3.14: Two possible representations of decision paths until interval j

is influenced by the current operation of the microCHP, i.e. the information that the appliance is currently running for 3 intervals. However, the explicit description of the complete run history before the current run is unnecessary. The total production of this history namely only depends on the amount of intervals that the appliance

is planned to run and the startup and shutdown behaviour. As long as production runs fulfill the appliance specific and operational requirements, historic production can be described by the number of completed runs and the total amount of intervals in which the microCHP is on. This historic production of two decision paths is equal when the number of completed runs is equal and when the total amount of intervals in which the microCHP is on is equal. This is the case in Figure 3.14 and the state space can be reduced by merging $(0, 0, 1, 1, 1, 1, 1, 0, 0, 0, 1, 1, 1, 1, 0, 0, 1, 1, 1, 1)$ and $(1, 1, 1, 1, 0, 0, 0, 1, 1, 1, 1, 1, 0, 0, 0, 1, 1, 1)$ into one state.

Following the above idea, we describe a state σ_j^i of a single microCHP i by a 3-tuple (A_j^i, B_j^i, C_j^i) , which represents the situation at the begin of interval j . More precisely, we have:

- A_j^i , expressing the number of consecutive intervals that the on/off state of the microCHP is unchanged looking back from the start of the current interval j (positive values indicate that the microCHP is running and negative values indicate that the microCHP is off);
- B_j^i , expressing the total number of intervals the microCHP has been running from the beginning of the planning period until the start of the current interval j ;
- C_j^i , expressing the number of runs of the microCHP which have already been completed.

The number of possible states per phase for a given house i is bounded by N_T^3 . In the DP we get $N_T + 1$ phases corresponding to the start of the intervals $j = 1, \dots, N_T + 1$, where the final phase corresponds to the state at the *end* of the planning horizon (after interval N_T). The two possible representations of decision paths from Figure 3.14 are represented by state $(3, 13, 2)$. An example of the possible decisions is given in Figure 3.15, where it appears to be infeasible to switch off the appliance.




Figure 3.15: State changes from $(3, 13, 2)$ with corresponding costs

We apply backwards dynamic programming to the formed state space. For each state σ_j^i in phase j a value function $F_j^i(\sigma_j^i)$ is introduced, which expresses the maximal profit which can be achieved in the intervals j, \dots, N_T if the microCHP is in state σ_j^i at the begin of interval j . The calculation of $F_j^i(\sigma_j^i)$ depends on the

possible actions in state σ_j^i and the values of the value function for some states in phase $j + 1$. The possible actions are to either have the on/off state unchanged or to change it. If we leave the state unchanged (no start or stop) we get as new state in interval $j + 1$:

$$\hat{\sigma}_j^i := \begin{cases} (A_j^i + 1, B_j^i + 1, C_j^i) & \text{if } A_j^i > 0 \\ (A_j^i - 1, B_j^i, C_j^i) & \text{if } A_j^i < 0. \end{cases}$$

If we change the on/off state, we have:

$$\check{\sigma}_j^i := \begin{cases} (-1, B_j^i, C_j^i + 1) & \text{if } A_j^i > 0 \\ (1, B_j^i + 1, C_j^i) & \text{if } A_j^i < 0. \end{cases}$$

This leads to the following recursive expression for $F_j^i(\sigma_j^i)$:

$$F_j^i(\sigma_j^i) := \max\{c_j^i(\sigma_j^i, \hat{\sigma}_j^i) + F_{j+1}^i(\hat{\sigma}_j^i), c_j^i(\sigma_j^i, \check{\sigma}_j^i) + F_{j+1}^i(\check{\sigma}_j^i)\},$$

where $c_j^i(\sigma, \sigma')$ denotes the cost associated with the choice corresponding to the transition from σ to σ' . The calculation of these costs is similar to the calculation of the values e_j^i used in Section 3.3 plus some feasibility checks on the state transitions and can be done in constant time. If a decision is infeasible we set the cost $c_j^i(\sigma, \sigma') = -\infty$. When we define $F_{N_T+1}^i(\sigma_{N_T+1}^i) = 0$ for all possible states $\sigma_{N_T+1}^i$ in phase $N_T + 1$ we can recursively calculate $F_1^i(\sigma_1^i)$ and deduce a corresponding optimal decision vector x^i . Since there are $\mathcal{O}(N_T^3)$ state tuples and there are N_T time intervals to evaluate, the dynamic programming approach of the single house model has runtime $\mathcal{O}(N_T^4)$.

Dynamic programming of a group of microCHPs

As stated before, the combination of the dynamic programming formulations for different microCHPs cannot be merged, since local information remains to be extractable. A state in the dynamic programming formulation for the group of houses has to be specified by a vector of states for the individual houses; $\sigma_j := (\sigma_j^1, \dots, \sigma_j^N)$. From each state σ_j we have 2^N possible actions that can be taken (existing of N binary choices to leave the state unchanged or not in each house). Note that a state transition is only feasible if, next to the individual feasibility checks on the house states, the state vector (of the combined houses) also fulfills the cooperational constraints of the given interval.

To formalize the dynamic programming for the group of houses, we denote by $D_j(\sigma_j)$ the maximal cooperational profit that can be achieved in the intervals j, \dots, N_T if, at the begin of interval j , the state of house i is given by σ_j^i , for $i = 1, \dots, N$. Due to the cooperational constraints a state transition from σ_j to σ_j' may not be allowed even if all individual state transitions $(\sigma_j^i, \sigma_j'^i)$ are allowed for the individual houses. Therefore we cannot simplify the dynamic programming

by calculating the individual house dynamic programming formulations independently and merging the results. So we need to calculate the complete dynamic programming, which has an exponential runtime of $\mathcal{O}(N_T^{3N+1})$, since the state space explodes by the possible combinations of houses in each phase of the dynamic programming ($\mathcal{O}(N_T^{3N})$).

3.4.2 RESULTS

The dynamic programming methods for the single house and for the group of houses are implemented in C++. For the group of houses we make use of an SQL database to store the produced values for all states, due to the exponential increase in the number of states.

Results for small instances

Table 3.6 shows the results for the small instances regarding profit maximization. The instances again are represented by the number of microCHPs and the variant describing the total production bounds. We give the (optimal) objective values for profit maximization z_{\max} (divided by the number of houses N), the computational time in seconds and the memory usage of the database (in MB). The table clearly shows the exponential growth in the state space, which is visible in the memory usage and the increase in computational time.

The results of this basic dynamic programming approach are used to validate the results of the ILP model. The same objective values result from both methods in all cases that were not terminated by the ILP solver, which indicates that the planning problem is correctly implemented. In addition, the basic dynamic programming formulation gives the optimal solutions to the prematurely terminated instances of the ILP formulation. In two cases both solutions are equal; for the three remaining instances, the optimal solution lies between the upper bound and the current solution of the ILP solver.

3.4.3 CONCLUSION

The basic dynamic programming formulation gives a structured description of the state space of the microCHP planning problem. A state consists of a vector of individual microCHP states, which in their turn are 3-tuples representing the historic decision path until a given time interval.

Using this representation, small instances can be solved to optimality. The results of the ILP formulation are validated and prematurely interrupted solutions are improved. However, the computational times and memory usage indicate that solving realistically sized instances by the DP approach is intractable in practice.

instance		solution			instance		solution		
N	variant	$\frac{z_{max}}{N}$	time (s)	mem. (MB)	N	variant	$\frac{z_{max}}{N}$	time (s)	mem. (MB)
1	1	1.147	0.015	-	6	10	1.021	256.013	18.43
1	2	1.092	0.015	-	7	1	1.156	1167.600	59.70
1	10	1.092	0.016	-	7	2	1.145	1173.185	59.70
2	1	1.236	2.280	0.03	7	3	1.137	1187.634	59.70
2	2	1.208	2.745	0.03	7	4	1.109	1140.431	59.70
2	3	1.016	3.045	0.03	7	5	1.069	1089.978	59.70
2	4	1.016	2.994	0.03	7	6	0.972	932.687	59.66
2	5	1.016	3.147	0.03	7	7	0.925	867.137	59.47
2	6	1.016	2.777	0.03	7	8	1.150	1117.568	59.70
2	10	1.016	2.695	0.03	7	10	0.924	837.076	59.47
3	1	1.197	3.875	0.12	8	1	1.156	5937.227	285.87
3	2	1.197	3.474	0.12	8	2	1.145	5849.868	285.87
3	3	1.106	3.813	0.12	8	3	1.130	5711.468	285.87
3	4	1.106	3.434	0.12	8	4	1.114	5762.172	285.87
3	5	1.002	3.569	0.12	8	5	1.080	5676.447	285.87
3	10	1.002	3.240	0.12	8	6	1.032	5344.593	285.87
4	1	1.183	13.150	0.79	8	7	0.919	4596.758	285.42
4	2	1.164	13.895	0.79	8	8	1.152	5679.480	285.87
4	3	1.128	13.027	0.79	8	9	1.069	5337.957	285.64
4	4	1.114	13.020	0.79	8	10	0.902	3864.537	285.19
4	5	1.021	12.173	0.79	9	1	1.153	36806.901	1762.07
4	6	1.021	12.045	0.79	9	2	1.143	36385.480	1762.07
4	8	1.009	11.000	0.75	9	3	1.121	37538.051	1762.07
4	9	1.009	10.780	0.75	9	4	1.098	35902.998	1762.07
4	10	0.949	9.753	0.75	9	5	1.072	34578.857	1762.07
5	1	1.164	35.890	2.39	9	6	0.999	32823.057	1761.94
5	2	1.149	35.779	2.39	9	8	1.150	34938.175	1762.07
5	3	1.120	35.784	2.39	9	9	1.113	31966.586	1761.40
5	4	1.060	34.148	2.39	9	10	0.976	28798.542	1761.27
5	5	1.060	34.720	2.39	10	1	1.176	372792.485	15458.30
5	8	1.118	32.564	2.38	10	2	1.162	373217.551	15458.30
5	10	1.023	31.638	2.38	10	3	1.143	378439.755	15458.30
6	1	1.163	293.977	18.44	10	4	1.122	378338.190	15458.30
6	2	1.150	305.999	18.44	10	5	1.095	361314.246	15458.28
6	3	1.130	295.154	18.44	10	6	1.044	345850.900	15458.19
6	4	1.092	288.180	18.44	10	7	0.956	311694.388	15452.69
6	5	1.048	275.603	18.44	10	8	1.173	365599.919	15458.30
6	6	1.027	267.682	18.43	10	9	1.028	301642.878	15443.78
6	8	1.139	284.076	18.44	10	10	0.945	236609.514	15438.16

Table 3.6: Results for the basic dynamic programming method

3.5 DYNAMIC PROGRAMMING BASED LOCAL SEARCH

In the previous section the basic dynamic programming approach is introduced. This method gives a fast technique for solving single microCHPs, but the computational effort ‘explodes’ when the number of microCHPs increases. In this section we develop a local search based heuristic which uses the single microCHP DP as a subroutine.

In case we optimize for the electricity market (i.e. maximize the profit), the dynamic programming method for a single microCHP can be seen as a function f^i on the price vector π :

$$f^i(\pi) \rightarrow x^i. \tag{3.39}$$

The function in Equation (3.39) gives an optimal local planning for the single house and can be calculated in runtime $\mathcal{O}(N_T^4)$, given the electricity market price vector

π (and of course the data of house i). However, we also may apply this function to any other vector that differs, except from its length, from π . In this way the function might not return the optimal decision path for the house in relation to the market prices, but it still returns a locally feasible path, satisfying appliance specific and operational constraints. We might want to use such artificial prices to explore different operational decision paths for individual microCHPs.

This observation forms the basis of a local search method. In the following we explain the separation of the two dimensions we are dealing with and propose a method for which the running time can be controlled to some extent.

3.5.1 SEPARATION OF DIMENSIONS

Similar to the function f^i for microCHP i , the dynamic programming method for the group of houses can be seen as a function d on the price vector π :

$$d(\pi) \rightarrow (\tilde{x}^1, \dots, \tilde{x}^N).$$

This function $d(\pi)$ maximizes a certain objective function and outputs N vectors consisting of the planning in the N corresponding houses. We call these vectors \tilde{x}^i the optimal decision paths. Whereas $f^i(\pi)$ finds a solution in polynomial time, $d(\pi)$ needs exponential time to be evaluated. Since this is not feasible in practice, a heuristic is developed to find a solution to the microCHP planning problem that is both feasible and, hopefully, close to the optimum solution, and can be found in reasonable time.

Due to the cooperational restrictions on the total electricity production we cannot (in general) solve the function d by individually solving N functions f^1, \dots, f^N in parallel for the price vector π . In general $d(\pi) \neq (f^1(\pi), \dots, f^N(\pi))$; however, we presume that there exist some vectors v^1, \dots, v^N such that $d(\pi) = (f^1(v^1), \dots, f^N(v^N))$. In this case each $f^i(v^i)$ forces each corresponding microCHP to plan the production that is found in the optimal solution according to $d(\pi)$ (i.e. the decision paths \tilde{x}^i).

In the following we discuss the assumption that we can find a vector v^i with $f^i(v^i) = \tilde{x}^i$ for each microCHP i , i.e. we want to show that any possible individual production plan can be reached by cleverly designing the vector v^i . For a dynamic programming formulation in which the choice for the cost that is associated with each state transition is completely free, it is obvious that any decision path can be constructed, for instance by letting all state transitions which are not in the desired solution have a cost of $-\infty$ (in a maximization problem) and 1 otherwise.

In our case however the cost determination is prescribed by the instance. The cost is determined by a multiplication of the artificial price vector v^i and the electricity production that belongs to a state transition (and is $-\infty$ when the state transition is infeasible). This means that we cannot define state transition costs individually, but we must focus on all possible state transitions for a certain phase simultaneously. This indicates that there might be a chance that desired decision paths are preveiled by other paths and might never be found, if we only have the option to steer via v^i . We show that this is not the case. Note that the cost of $-\infty$ of infeasible state

transitions (heat demand violation, runtime/offtime violations) can be neglected in the following reasoning, since they can never outperform the decision path \tilde{x}^i in the solution of $d(\pi)$. Let \tilde{e}^i be a vector of electricity production corresponding to the optimal solution \tilde{x}^i of $d(\pi)$. Since \tilde{e}^i results directly from the decision path \tilde{x}^i , we want to force this path to be taken by focusing on the electrical output \tilde{e}^i only. This vector has a unique structure of zeroes, possibly interrupted by positive and/or negative values. If the electricity output in \tilde{e}^i is always nonnegative, then we could define v^{i*} simply to be:

$$\tilde{v}_j^i = \begin{cases} 1 & \text{if } \tilde{e}_j^i > 0 \\ -M & \text{if } \tilde{e}_j^i = 0, \end{cases} \quad (3.40)$$

where $M > 0$ is chosen large enough to supersede the contribution of the positive electricity output multiplied by 1. For the vector defined in (3.40), any other decision path $x^i \neq \tilde{x}^i$ would result in a loss in the objective value, since a contribution of $1 \times \tilde{e}_j^i > 0$ is lost (in case $e_j^i = 0$ where $\tilde{e}_j^i > 0$) or a contribution of $-M \times e_j^i < 0$ is earned (in case $e_j^i > 0$ where $\tilde{e}_j^i = 0$).

If negative electricity output is also allowed, we define:

$$\tilde{v}_j^i = \begin{cases} 1 & \text{if } \tilde{e}_j^i > 0 \\ -M & \text{if } \tilde{e}_j^i \leq 0. \end{cases} \quad (3.41)$$

Again these 'prices' focus on determining the correct start and stop moments of the microCHP control, which determines the correct decision path. Late and early starts and stops again have negative contributions to the objective value, once M is securely chosen.

This shows that we can control the output of the dynamic programming formulation of a single microCHP completely by the price vector, in the sense that we can derive all feasible decision paths by using artificial price vectors. This idea of controlling the output of the dynamic programming method for a single microCHP leads to the local search heuristic of this section.

3.5.2 IDEA OF THE HEURISTIC

The idea of the heuristic method is the following. If we discard the cooperational constraints in first instance, we can calculate the group planning by separating it into N single house dynamic programming methods. This separation of dimensions reduces the runtime to $\mathcal{O}(N \times N_T^4)$. Now we reintroduce the cooperational constraints as a feasibility check on the output of this calculation. This combination of calculating separate dynamic programs and performing a feasibility check results in a new structure: a certain total electricity production as output and a yes/no answer whether this production is allowed by the bounds on the combined electricity production. The basis of the heuristic method now is to use this structure of separately calculated dynamic programs and a feasibility check on the sum of these individual dynamic programs, by iteratively searching on the sets of artificial

vectors for each microCHP in an effective way until a combination of artificial vectors is found, where the feasibility check leads to a positive answer. The search for artificial vectors starts from the price vector π . In this way we may expect that the resulting solution is somehow close to the optimum.

Note that if we would have chosen to take space as our candidate for the phases of the DP, then the separation of dimensions would have been a problem. The task of combining different outcomes of single dynamic programs would introduce heavy feasibility problems in time, which are harder to ignore than a possible infeasibility in the cooperational constraints. This gives rise to the notion of weak constraints in the sense of the cooperational constraints; we may allow some small violations from the desired total electricity output. In fact, one of the optimization problems is especially aimed at minimizing the deviation from this weak constraint.

3.5.3 DYNAMIC PROGRAMMING BASED LOCAL SEARCH METHOD

The idea behind the local search method presents the structure of the heuristic: we use separate dynamic programming formulations for individual microCHPs and simply combine the output of these dynamic programming formulations to perform a feasibility check on the cooperational constraints. The searching part of the heuristic consists of the search for input vectors v^i that result in solutions 'close' to the optimal solution of the microCHP planning problem. Therefore it is of importance to define local moves in the search method and to define stopping criteria, since we also want to limit the computational effort of the heuristic. We propose the following local moves and stopping criteria.

Local moves

In Section 3.4 we proposed the dynamic programming formulation for the group of microCHPs, where all possible combinations of production vectors in individual houses are coded by the state space. In this heuristic we need a way to search through these possible combinations, since the dependence between different house productions is lost when calculating separate house DPs. Since we do not want to change the state definition in the house DP to compensate for this loss of (cooperational) information (this would lead to the original DP for the group of microCHPs or similar state expansions), the only way of applying a search can go via the input of the individual DPs. Since $f^i(v^i)$ depends on the (artificial) price vector v^i we change the price vector of the house DPs in our search. Of course the value of the objective function for the output of the group of microCHPs is still calculated with the original price vector π .

Starting with a price vector $v^i = \pi$ for each house i , we iteratively adjust the price vectors based on the result of the DPs for the individual houses using their current price vectors. We try to remain as close as possible to the original price vector, in the hope to stay close to the optimal value for the objective function. In each iteration the price v_j^i of interval j for each microCHP i is locally adjusted if:

- P_j^{upper} is violated and the microCHP of house i is decided to be on in the current solution;
- P_j^{lower} is violated and the microCHP of house i is decided to be off in the current solution.

In the first case we want to make time interval j less attractive for production by generator i . This can be reached by reducing the price v_j^i . In the second case we want to make time interval j more attractive for production. To achieve this we increase the price v_j^i . To test this approach, we have chosen for the following simple updating scheme:

- in the first case v_j^i is multiplied with a factor a , where $0 < a < 1$;
- in the second case v_j^i is multiplied with a factor $2 - a$.

All other prices remain unchanged.

Stop criteria

The method stops when a feasible solution is found or when a maximum number of iterations $MaxIt$ is reached. If the maximum iteration count $MaxIt$ is reached and we did not find a feasible solution, the solution with the smallest error value err is given as a best approximation to the fleet constraints. This error err is the absolute sum of the mismatch to the upper and/or lower bounds P_j^{upper} and P_j^{lower} :

$$err := \sum_{j=1}^{N_T} \left(\max \left(\sum_{i=1}^N e_j^i - P_j^{upper}, 0 \right) + \max \left(P_j^{lower} - \sum_{i=1}^N e_j^i, 0 \right) \right).$$

In Algorithm 2 a summary of the algorithm is given. Note that the basic structure of this heuristic may also be applied to other Dynamic Programming formulations which allow a decomposition of the state, leading to a simplified version, consisting of a set of individual DPs.

3.5.4 RESULTS

Below we present the results of the local search method for both the small instances and the medium instances. In the local search method for the small instances we set the parameters $a = 0.9$ and $MaxIt = 100$. In the medium instances we again choose $MaxIt = 100$. As multiplication factor a we now use the values 0.9, 0.7, 0.5, 0.3 and 0.1.

Results for small instances

The quality of the local search method is verified by comparing its objective values and computational times to the ones of the ILP approach given in Tables 3.4 and 3.5. The local search method is only applied to the feasible instances as found by solving

Algorithm 2 Local search on the microCHP planning problem

Input: price vector π , lower and upper bounds P^{lower} and P^{upper} , $v^i := \pi$ for all houses i

repeat

solve $f^i(v^i)$ for all i resulting in solution $x = (x^1, \dots, x^N)$;

calculate total production $(\sum_{i=1}^N e_1^i, \dots, \sum_{i=1}^N e_{N_T}^i)$ of solution x ;

for all j **do**

if $\sum_{i=1}^N e_j^i > P_j^{upper}$ **then**

for all i with $x_j^i = 1$ **do**

$v_j^i \leftarrow av_j^i$

end for

end if

if $\sum_{i=1}^N e_j^i < P_j^{lower}$ **then**

for all i with $x_j^i = 0$ **do**

$v_j^i \leftarrow (2 - a)v_j^i$

end for

end if

end for

until solution x is feasible or *MaxIt* is reached

the ILP. Table 3.7 gives the details of the solutions for the small instances that are found by applying the local search method, where the objective value divided by the number of microCHPs, the computational time, the number of iterations and the error value are presented. As a first verification, the local search method produces optimal results for all instances $I(k, 1)$ as should be the case, since independent DPs can be used in case of no network restrictions. In 15 of the 78 instances no feasible solution is found; the corresponding deviation from the bounds is denoted in the table. It is noteworthy that in one case a feasible solution is found in the 100th iteration ($I(8, 7)$).

When we categorize the results by the number of microCHPs and by the production pattern variants, we obtain average results, as in Table 3.8. On the left hand side averages are taken over all (feasible) production pattern variants and on the right hand side averages over all (feasible) numbers of houses. We define the quality of the objective value to be the quotient of the local search objective value and the optimal objective value $Q := \frac{z_{max,ls}}{z_{max,opt}}$. This quality is presented in Table 3.8, as well as the computational time and the percentage of infeasible solutions. The average quality \bar{Q} of all instances is 0.95. No trend can be identified between the number of houses and the quality of the local search method. The production pattern variant has an effect on the quality. An explanation for this behavior is that

instance		solution				instance		solution			
N	variant	$\frac{z_{max}}{N}$	time (s)	iter	err	N	variant	$\frac{z_{max}}{N}$	time (s)	iter	err
1	1	1.147	0.015	0	0	6	10	0.942	0.141	71	0
1	2	1.092	0.015	3	0	7	1	1.156	0.015	0	0
1	10	1.092	0.016	3	0	7	2	1.137	0.016	4	0
2	1	1.236	0.015	0	0	7	3	1.137	0.016	4	0
2	2	1.208	0.015	3	0	7	4	1.079	0.016	6	0
2	3	0.937	0.078	100	200	7	5	1.068	0.031	12	0
2	4	0.937	0.078	100	400	7	6	0.904	0.078	39	0
2	5	0.937	0.078	100	600	7	7	0.893	0.219	100	900
2	6	1.016	0.016	10	0	7	8	1.093	0.266	53	0
2	10	1.016	0.016	10	0	7	10	0.839	0.203	100	150
3	1	1.197	0.015	0	0	8	1	1.156	0.016	0	0
3	2	1.197	0.015	0	0	8	2	1.138	0.016	4	0
3	3	1.097	0.016	12	0	8	3	1.087	0.016	6	0
3	4	1.097	0.016	12	0	8	4	1.087	0.016	6	0
3	5	0.863	0.031	22	0	8	5	1.073	0.031	12	0
3	10	0.863	0.031	22	0	8	6	0.875	0.063	29	0
4	1	1.183	0.015	0	0	8	7	0.838	0.218	100	0
4	2	1.068	0.015	12	0	8	8	1.151	0.016	4	0
4	3	1.050	0.016	12	0	8	9	0.986	0.266	100	1050
4	4	1.103	0.015	13	0	8	10	0.881	0.219	100	7600
4	5	0.939	0.078	31	0	9	1	1.153	0.016	0	0
4	6	0.794	0.047	32	0	9	2	1.137	0.031	4	0
4	8	0.931	0.125	100	1250	9	3	1.092	0.016	6	0
4	9	0.931	0.141	100	2850	9	4	1.092	0.031	6	0
4	10	0.822	0.141	100	1500	9	5	1.057	0.031	11	0
5	1	1.164	0.015	0	0	9	6	0.842	0.078	25	0
5	2	1.083	0.016	6	0	9	8	1.148	0.015	4	0
5	3	1.063	0.016	6	0	9	9	0.960	0.250	100	650
5	4	1.054	0.047	22	0	9	10	0.843	0.188	70	0
5	5	1.054	0.031	22	0	10	1	1.176	0.016	0	0
5	8	0.978	0.172	100	400	10	2	1.161	0.016	4	0
5	10	0.856	0.062	44	0	10	3	1.098	0.016	6	0
6	1	1.163	0.015	0	0	10	4	1.098	0.015	6	0
6	2	1.096	0.015	6	0	10	5	1.094	0.031	6	0
6	3	1.096	0.016	6	0	10	6	0.963	0.109	33	0
6	4	1.092	0.031	9	0	10	7	0.871	0.109	37	0
6	5	0.967	0.047	20	0	10	8	1.170	0.031	7	0
6	6	0.940	0.063	23	0	10	9	0.968	0.297	100	2100
6	8	0.925	0.203	100	500	10	10	0.849	0.297	100	6850

Table 3.7: Results for the dynamic programming based local search method

houses	\mathcal{Q}		time (s)			production pattern	\mathcal{Q}		time (s)		
	μ	σ	$\mu(ILP)$	$\mu(Is)$	infeasible %		μ	σ	$\mu(ILP)$	$\mu(Is)$	infeasible %
1	1.000	0.000	0.07	0.015	0.00	1	1.000	0.000	64.14	0.015	0.00
2	0.967	0.038	0.60	0.042	42.86	2	0.979	0.028	145.04	0.017	0.00
3	0.951	0.063	1.56	0.021	0.00	3	0.962	0.024	147.31	0.023	11.11
4	0.916	0.062	4.66	0.066	33.33	4	0.980	0.022	1056.41	0.029	11.11
5	0.942	0.059	88.37	0.051	14.29	5	0.955	0.047	2407.59	0.043	11.11
6	0.937	0.057	362.25	0.066	12.50	6	0.891	0.068	3398.09	0.065	0.00
7	0.969	0.032	1346.44	0.096	22.22	7	0.929	0.025	7206.77	0.182	33.33
8	0.958	0.047	3199.52	0.088	20.00	8	0.936	0.067	45.55	0.118	42.86
9	0.946	0.064	3412.56	0.073	11.11	9	0.912	0.030	784.09	0.239	100.00
10	0.961	0.038	4963.09	0.094	20.00	10	0.913	0.057	4457.36	0.131	40.00

Table 3.8: Average results for small instances

<i>houses</i>	z_{ls}/N	<i>time (s)</i>	<i>iterations</i>	<i>error</i>	<i>infeas. (%)</i>
25	1.007	1048	80.3	20588	75.0
50	1.026	1982	79.1	36063	75.0
75	1.040	2869	78.7	59734	75.0
100	1.031	3831	78.8	71050	75.0
<i>production pattern</i>	z_{ls}/N	<i>time (s)</i>	<i>iterations</i>	<i>error</i>	<i>infeas. (%)</i>
11	1.165	859	16.8	0	0.0
12	0.971	2951	100.0	9340	100.0
13	0.984	2962	100.0	65654	100.0
14	0.984	2958	100.0	112440	100.0
<i>intervals</i>	z_{ls}/N	<i>time (s)</i>	<i>iterations</i>	<i>error</i>	<i>infeas. (%)</i>
24	0.953	1	75.3	27163	75.0
48	1.023	243	80.5	39252	75.0
96	1.103	7053	81.9	74162	75.0

Table 3.9: Results for medium instances with $a = 0.9$

the local search method has more difficulty in finding a feasible solution under tighter network constraints, resulting in larger deviations from the original price vector. This original vector is used in the objective value, which results in worse results. This is also shown in the percentages of infeasible solutions (violating the electricity constraints) that are found by the local search method. The network variant has more influence on this percentage than the number of houses. If we look at the deviation from the electricity bounds (given by the error *err*), the solutions are relatively close to these bounds. Therefore we included these infeasible solutions in all calculations and comparisons.

Results for medium instances

For the medium instances, we are interested in the behavior of the local search method dependent on the following three instance parameters: the size of the group of houses, the production pattern variant, and the number of intervals in a planning for 24 hours. The criteria we use to evaluate the behavior are the objective value, the computational time, the number of iterations the local search method needs, the error and the percentage of infeasible solutions. The results in Table 3.9 and 3.10 are, for a given value of one of the parameters, the averages over all combinations which are derived from the two other parameters.

The results achieved with the value $a = 0.9$ (as applied to the small instances) are given in Table 3.9. The computational time per house and the error per house decrease slightly when the number of houses increases. For 100 houses the error corresponds to 0.7 kWh over/underproduction per house. For production pattern variant 11 the method always finds a feasible solution (in a few iterations), while for the variants 12, 13 and 14 no feasible solution is found (and the method stops after 100 iterations). However, note that these production constraints are more tight than in the small instances and, thus, there is quite a chance that no feasible solution may exist. Regarding the number of intervals the computational times grow fast. The

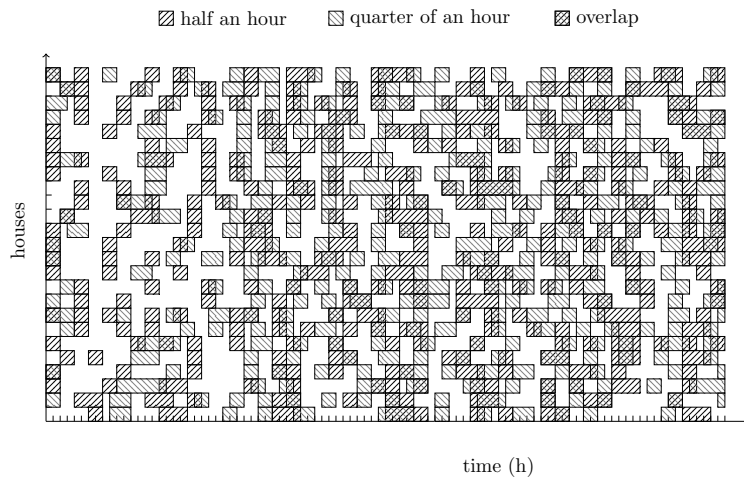
<i>houses</i>	z_{is}/N					<i>time (s)</i>				
	0.9	0.7	0.5	0.3	0.1	0.9	0.7	0.5	0.3	0.1
25	1.007	0.994	0.996	0.977	0.979	1048	1007	1011	997	1012
50	1.026	0.981	0.984	0.976	0.949	1982	1915	1912	1898	1925
75	1.040	1.001	0.975	0.966	0.970	2869	2719	2741	2746	2737
100	1.031	0.981	0.962	0.976	1.026	3831	3628	3569	3686	3647
	<i>iterations</i>					<i>error</i>				
	0.9	0.7	0.5	0.3	0.1	0.9	0.7	0.5	0.3	0.1
25	80.3	76.9	76.0	70.4	71.2	20588	21498	16103	17807	16377
50	79.1	77.0	72.8	71.8	72.3	36063	41865	33840	33758	28492
75	78.7	76.6	71.4	72.3	69.8	59734	60839	44525	47031	41800
100	78.8	76.6	72.1	72.4	70.8	71050	78146	68117	63392	56813
<i>production pattern</i>	z_{is}/N					<i>time (s)</i>				
	0.9	0.7	0.5	0.3	0.1	0.9	0.7	0.5	0.3	0.1
11	1.165	1.090	1.033	1.005	1.005	859	405	417	519	535
12	0.971	0.964	0.966	0.964	0.971	2951	2949	2900	2886	2867
13	0.984	0.942	0.951	0.966	0.961	2962	2956	2957	2958	2956
14	0.984	0.961	0.967	0.959	0.987	2958	2959	2957	2963	2963
	<i>iterations</i>					<i>error</i>				
	0.9	0.7	0.5	0.3	0.1	0.9	0.7	0.5	0.3	0.1
11	16.8	7.1	8.7	9.9	9.6	0	0	0	0	0
12	100.0	100.0	83.7	77.0	74.6	9340	13148	12204	13381	10079
13	100.0	100.0	100.0	100.0	100.0	65654	69146	58121	48221	38450
14	100.0	100.0	100.0	100.0	100.0	112440	120054	92259	100386	94951
<i>intervals</i>	z_{is}/N					<i>time (s)</i>				
	0.9	0.7	0.5	0.3	0.1	0.9	0.7	0.5	0.3	0.1
24	0.953	0.955	0.947	0.952	0.946	1	1	1	1	1
48	1.023	0.953	0.939	0.926	0.939	243	235	197	191	180
96	1.103	1.060	1.052	1.043	1.058	7053	6716	6726	6803	6810
	<i>iterations</i>					<i>error</i>				
	0.9	0.7	0.5	0.3	0.1	0.9	0.7	0.5	0.3	0.1
24	75.3	75.1	75.1	75.2	75.2	27163	16775	16588	17969	17041
48	80.5	77.0	65.8	61.3	58.8	39252	44048	26431	23150	25548
96	81.9	78.2	78.4	78.8	79.2	74162	90937	78919	80372	65021

Table 3.10: Results for medium instances and varying a

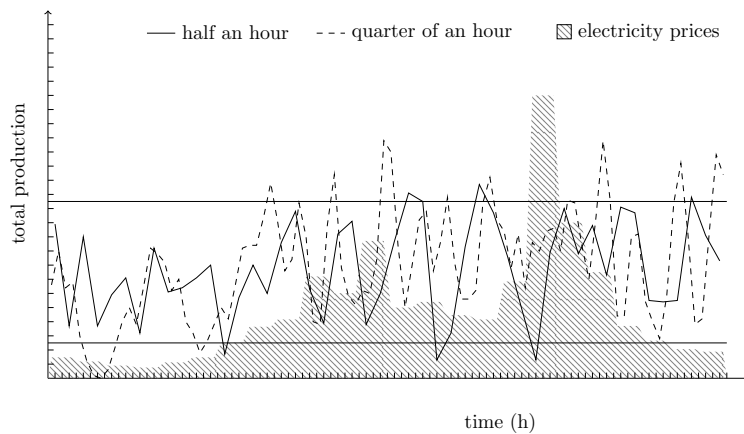
objective value increases as the number of intervals increases; however, the error increases accordingly, so the convergence is slower for a larger number of intervals.

Next, since optimal objective values are unknown for these instances, the solutions of different updating schemes of the price vector are compared to each other. In this comparison, the focus is in first instance on the ability to find a feasible solution and the objective value is only of secondary interest. The results for using the values 0.9, 0.7, 0.5, 0.3 and 0.1 for the parameter a are given in Table 3.10. The different updating schemes perform similar. If the focus is more on minimizing the error, the values 0.5, 0.3 and 0.1 are advantageous. For these values of a for some instances with production pattern variant 12 the local search method could find feasible solutions. If the focus is on the objective value, $a = 0.9$ gives better results against a slightly higher number of iterations and computational time.

Figure 3.16 shows a comparison of the detailed planning of a fleet of 25 houses and production pattern variant 12. A planning based on half an hour intervals is compared to a planning based on intervals of a quarter of an hour. 202.5 run hours are planned for the half an hour based planning and 210.75 run hours for the quarter of an hour planning. Figure 3.16a shows that only 74.5 of these run hours of



(a) Two planning results using $a = 0.9$ for production pattern 12 and 25 houses



(b) Fleet behaviour

Figure 3.16: The detailed planning of a case with a different number of intervals

the two plannings do overlap. In Figure 3.16b the total generation is plotted against the background of the original price vector. This example emphasizes that making a planning for 15 minutes intervals clearly leads to different results compared to a planning for 30 minutes intervals (both in total as for individual houses), although the minimum runtime and offtime stay fixed on 30 minutes.

In general we can state that an increase in the number of houses leads to a better fit for the fleet to the given production bounds (i.e. the amount of electricity per house outside the bounds decreases). Concerning the amount of iterations,

the largest improvement in objective value is reached within the first 25 iterations. Remaining iterations only lead to slightly better objective values. As a general comment, it is hard to flatten the total output profile over the whole day, when the aggregated heat demand profile deviates too much from the desired production bounds for a too long period.

3.5.5 CONCLUSION

In this section a local search method is developed to solve the microCHP planning problem. Small instances are tested to verify the quality of this heuristic method in comparison to the (optimal) solutions by solving the ILP or basic DP formulation; the local search method results in a 5% loss in objective value and a 99.0% gain in computation time compared to the ILP formulation and a 99.9999% gain in computation time compared to the basic DP formulation. Furthermore, the local search method is tested for the medium instances, to see whether it is applicable in practice. Considering the fact that, in practice, we can unfold one calculating entity per house, a planning for 100 houses, 96 intervals and using 100 iterations can be made within 2.3 minutes. In our experience we find that the maximum number of iterations $MaxIt$ can easily be reduced with a factor 4, since most best solutions are found within the first 25 iterations. Using this reduction a planning can be made in about half a minute. Regarding feasibility, a feasible solution for the small instances is not found in 19% of the cases, where the ILP formulation did find a solution. Depending on the value of a , for 67% to 75% of the medium instances no feasible solution is found (note that the production bounds for the medium instances are more tight). However, it may be that for a larger percentage of these instances even no feasible solution may exist at all.

3.6 APPROXIMATE DYNAMIC PROGRAMMING

In the local search method of Section 3.5 and the column generation-like technique that is proposed in Section 3.7 we use a separation of the two dimensions in the microCHP planning problem. The reason for this separation is that the size of the basic dynamic programming formulation is too large for practical instances to be solved to optimality in reasonable time. Approximate Dynamic Programming (ADP) offers another approach to treat this difficulty caused by the size of the DP. The idea of this heuristic is explained in Section 3.6.1. Section 3.6.2 shows an initial attempt to apply ADP to solve the microCHP planning problem, see also [127].

3.6.1 GENERAL IDEA

In general, a dynamic programming formulation models a process for which several decisions have to be made. A state represents a possible outcome after a choice for (a part of) the decision variables. In the way we apply DPs to the planning problems in this thesis, states are grouped in so-called phases, where a phase represents the progress of the determination of all decision variables, measured by the amount of

decision variables that are fixed by the corresponding states. State transitions only occur between states of subsequent phases; so, a state transition models the choice for a specific decision (in the microCHP planning problem this decision consists of the on/off decisions for all N microCHPs in a specific time interval). For each state σ_ϕ in phase ϕ the value $V(\sigma_\phi)$ gives (for a maximization problem) the maximum objective value that can be reached in the remaining phases by making a decision for each of the remaining ‘open’ decision variables, assuming that one starts in the situation described by the state σ_ϕ . This value is calculated using a backwards updating structure, in which the value of each state in phase ϕ is determined from the values $V(\sigma'_{\phi+1})$ for the states in phase $\phi + 1$ and the costs that are associated with the state transitions from the state σ_ϕ in phase ϕ to the states in phase $\phi + 1$, where infeasible state transitions are penalized with a cost of $-\infty$. The value of $V(\sigma_1)$ gives the optimal value for the considered problem, assuming that σ_1 is the initial state at the beginning of the planning process, and the decisions that result in this value can be found by backtracking the corresponding state transitions that result in this value. This sequence of decisions is also called the optimal decision path. For solving a problem to optimality by a DP, all state transitions need to be considered.

Although considerable effort is put into reducing the number of states in each phase by cleverly setting up the definition of a state (as we have seen in the Held-Karp algorithm and in the development of a DP formulation for the microCHP planning problem), the size of a DP may still be too large.

Approximate Dynamic Programming [107, 108] may be a helpful tool to solve such large DPs. It approximates the value of states by evaluating only a small part of the DP transition graph instead of accurately calculating the correct value by evaluating the complete DP transition graph. To reach a satisfying result, an ADP method needs to focus on two important aspects. An ADP method wants to search only a small, but relevant, part of the state space and it wants to have an effective way of using information that results from this search into determining an approximation of the value function for the states in the different phases.

First, the ADP method uses sample paths. A sample path is a chosen sequence of decisions, as depicted in Figure 3.17. This sample path is used in the process of updating the approximation of the value function V . Iteratively sample paths are created until the approximation of the value function is such, that a sample path is found that is close to the optimal decision path. Of course we want to have sample paths that are helpful in this updating process, i.e. we want to have sample paths that are close to the optimal decision path. Since this optimal decision path is unknown, the task of creating sample paths needs a good mixture of intensification and diversification. Intensification concentrates on using an approximation of the value function to create sample paths, in which sample paths are either completely determined by this approximation of the value function or chosen with a certain probability that is based on this approximation of the value function. Diversification is used to escape from staying in a local area in the state space, by determining completely arbitrary sample paths. Note that the creation of a sample path should be given by a very simple heuristic; we want to avoid using computationally intensive

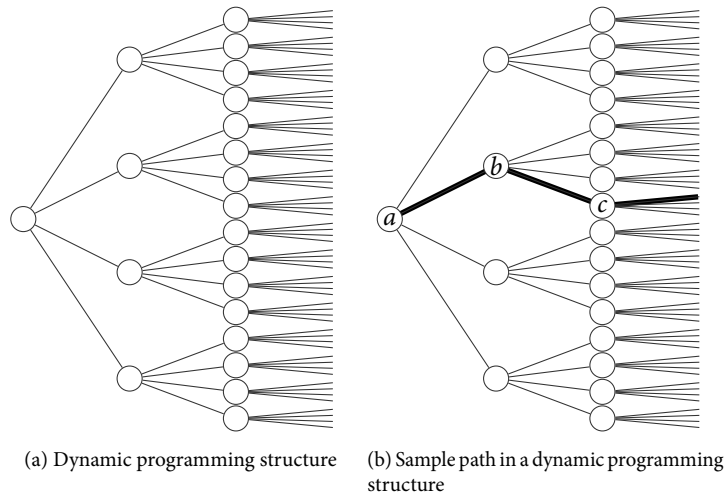


Figure 3.17: A (partial) transition graph of a DP formulation and a sample path through this structure

approximated value functions for all possible state transitions, which would resemble to solving the original DP, which we want to avoid. Therefore we stress that a good implementation of ADP uses relatively few iterations, thereby analyzing only a fraction of all possible sample paths, where these sample paths are representative for the wide range of options that a decision maker has.

The second important aspect of an ADP method is the approximation of the value function. This value function has influence in determining the desire to choose a certain state in the creation process of new sample paths. States with a large estimated value are more likely to be visited in subsequent sample paths than states with a smaller estimated value. This approximation needs to extract relevant information from the sample paths to update the estimated value of the desire to be in a certain state for all phases of the DP transition graph. This information is not only based on the state transition between subsequent states, but may also depend on the approximated values of states in successive phases and the result of the sample path in these phases.

The general idea of ADP is illustrated by an example. Although this example stems from a complete different research topic, it treats sample paths in a natural way and uses these sample paths to update value function approximations.

Example: analyzing the transformation of leukemic stem cells by gene overexpression [7]

Leukemia is caused and maintained by the presence of leukemic stem cells, which behave differently from the normal hematopoietic stem cells. Leukemic stem cells

are formed out of normal hematopoietic stem cells by a multistep transformation process, in which various mutations occur. These cells consist of hundreds or thousands of different gene subtypes, that have several different activities for each cell. Activities are e.g. differentiation, proliferation and cell death. Leukemic stem cells differ from normal cells in their activity. To be able to treat these cells we might be interested to unravel the different activities between leukemic stem cells and normal hematopoietic stem cells: which gene expression (which denotes the role of the gene in determining activities) results in activities that make a cell a leukemic cell?

By overexpressing (giving a gene more importance in determining cell activities) genes $g \in G = \{1, \dots, |G|\}$ one can stress the role of the corresponding genes and, thus, try to create a leukemic cell. Let the overexpression of a gene be represented by e_g . The 'resemblance' function $r(e_g)$ gives the extent to which the overexpression of this gene makes the cell resemble a leukemic stem cell. Suppose that there is a maximum total amount E which can be used to overexpress the set of genes G :

$$\sum_{g \in G} e_g \leq E. \quad (3.42)$$

In practice at most four different genes are overexpressed in a single test. The objective of our DP may then be to overexpress genes in such a way, that the resemblance $\sum_{g \in G} r(e_g)$ is maximized, while the total overexpression is smaller than or equal to E .

A dynamic programming formulation of this problem can be given by using the different genes as phases and using the different possibilities for overexpression for each gene as state transitions between two consecutive phases. A state s_f in phase f is defined as the amount of overexpression E_{s_f} which still can be used in the following phases ($E_{s_f} \leq E$) - it is not necessary to know how the other amount is used in the first phases. The costs that are associated with state transitions are the resemblance values $r(e_f)$ for state transitions from $s_{f-1} = E_{s_{f-1}}$ to $s_f = E_{s_f}$ between phase $f - 1$ and phase f (where $e_f = E_{s_{f-1}} - E_{s_f}$), for $E_{s_f} \geq 0$; otherwise the cost is $-\infty$. The value $V(s_f)$ of a state s_f maximizes the remaining resemblance of genes $f, \dots, |G|$; in a recursive way $V(s_1)$ can be calculated, which gives the optimum solution of overexpressed genes.

The ADP approach of solving this problem concentrates on sample paths and an approximation of the value of states. A sample path is simply defined as a vector of overexpression for all genes, which has an accompanying resemblance value. Based on this resemblance value for a specific gene expression $r(e_g)$ and the total resemblance of the remaining genes $g + 1, \dots, |G|$, it can be verified whether it is worth to increase or decrease the overexpression of gene g and, thus, whether it is increasingly/decreasingly worth to visit the corresponding state; i.e. the contribution of the corresponding state transition to the objective value can be compared to the contribution of the remaining decisions in the sample path, leading to a value that is above or below the average contribution. In this way a value function approximation can be updated and the creation of gene expression sample paths can be steered

towards a certain direction that increases the total resemblance to a leukemic stem cell.

3.6.2 APPROXIMATE DYNAMIC PROGRAMMING BASED HEURISTIC TO SOLVE THE MICROCHP PLANNING PROBLEM

In the microCHP planning problem we have $N_T + 1$ phases, which are related to time intervals, and $\mathcal{O}(j^3)$ states in phase $j \in \{1, \dots, N_T + 1\}$. A state σ_j describes the taken decisions for all N houses up to the interval j : it is a vector of length N of state tuples for individual houses. Translated to sample paths, a single edge (state transition) (σ_j, σ_{j+1}) in the DP transition graph consists of a decision for all N microCHPs in time interval j . Between successive phases there are 2^N possible state transitions for each state in the considered time interval. The value function $D_j(\sigma_j)$ gives the maximal cooperational profit that can be achieved in the intervals j, \dots, N_T if, at the begin of interval j , the state is given by σ_j .

The creation of sample paths

A decision in creating a sample path consists of a combination of binary choices for N microCHPs. Since there are 2^N possible choices we want to reduce the amount of options that we consider in creating a sample path. In a first implementation [127] we only consider an ordered set of microCHPs. This order is based on the necessity to have the corresponding microCHP running: the earlier it has to produce, the higher in the order. For this ordered set we allow the following type of state transitions: given an integer k , choose the first k microCHPs to be running and the remaining microCHPs to be off. This limits the state transitions to $N + 1$ choices per phase. In a forward calculation we choose the state transition that leads to the state with the largest approximated value function. The value $V(\sigma_j) := D_j(\sigma_j)$ is approximated by the function $\tilde{V}(\sigma_j)$.

Value function approximation

Value function approximation is an iterative process using iterations t in which the approximated value function $\tilde{V}_t(\sigma_j)$ is updated, based on the previous value and information that is extracted from a sample path. New information can be taken into account in different ways. In our implementation, the extent to which new information influences the value function approximation is determined by a factor α : the value function approximation $\tilde{V}_t(\sigma_j)$ in iteration t is updated as follows:

$$\tilde{V}_t(\sigma_j) := (1 - \alpha)\tilde{V}_{t-1}(\sigma_j) + \alpha v(\omega), \quad (3.43)$$

where $v(\omega)$ is a value which is somehow extracted from the sample path ω . A possible choice for this function is given below.

The approximation of the value function in a certain state and phase has to keep in mind that the value represents the maximum cooperational profit that can be

achieved in the remaining intervals. Therefore we base our function ($v(\omega)$) on the following properties:

- the electricity profit that can be made by applying a certain state transition,
- an estimation of the profit that can be made in the remaining time intervals (including an approximation of the total remaining production capacity),
- the approximation of the value function of states in the subsequent phase,
- penalty costs for the validation of heat constraints, operational and cooperational violations.

The first three properties concentrate on taking the best possible state transition given that we are in a certain state. The fourth property focuses on the desire of being in a certain state. Altogether these properties aim to increase the approximated value of a certain state, if this state and a corresponding state transition do not violate any constraint of the microCHP planning problem and have a relatively high contribution to the objective value. For a detailed description of the heuristic we refer to [127].

Results

Table 3.11 shows preliminary results of the Approximate Dynamic Programming based heuristic applied to the small instances. The table shows that for 91% of the small instances feasible solutions are found. This is an improvement when compared to the dynamic programming based local search method. Note that for the presented results, for each instance the factor α is chosen such that the resulting objective value is optimized and the deviation from the production bounds is minimized. The computational times that are presented, consist of the time that it takes to solve the instance using the found factor α , whereby the process of finding this value of α is not taken into account.

The results for the medium instances are summarized in Table 3.12. Compared to the results for the dynamic programming based local search method, we see that for an increasing number of instances a feasible solution is found, thereby decreasing the average deviation from the production bounds. If we differentiate the error to the production pattern variant, we see a trend of increasing errors when the production bounds get more tight. Regarding the results, differentiated to the number of time intervals, we see that all instances with intervals of half an hour have a feasible solution. However, when the interval length changes to 15 minutes, the results show large deviations and a large percentage of infeasible solutions. This is unsatisfactory, since we want at least to get close to the results for the instances with half an hourly time intervals.

3.6.3 CONCLUSION

In this section we have sketched a way to apply Approximate Dynamic Programming to the microCHP planning problem. It proposes a method that uses the original

instance			solution			instance			solution					
N	variant	$\frac{z_{max}}{N}$	time (s)	err	N	variant	$\frac{z_{max}}{N}$	time (s)	err	N	variant	$\frac{z_{max}}{N}$	time (s)	err
1	1	1.126	0.343	0	6	10	1.004	2.041	0					
1	2	1.071	0.345	0	7	1	1.132	3.074	0					
1	10	1.071	0.345	0	7	2	1.132	2.872	0					
2	1	1.225	0.683	0	7	3	1.132	2.762	0					
2	2	1.197	0.620	0	7	4	1.101	3.033	0					
2	3	0.996	0.681	0	7	5	1.059	2.775	0					
2	4	0.996	0.637	0	7	6	0.948	2.240	0					
2	5	0.996	0.672	0	7	7	0.914	2.582	0					
2	6	0.993	0.622	0	7	8	1.131	3.099	0					
2	10	0.993	0.622	0	7	10	0.891	3.108	0					
3	1	1.189	0.978	0	8	1	1.134	3.700	0					
3	2	1.189	0.960	0	8	2	1.134	3.838	0					
3	3	1.102	0.957	0	8	3	1.109	3.800	0					
3	4	1.102	0.947	0	8	4	1.082	3.963	0					
3	5	0.905	0.970	0	8	5	1.070	3.713	0					
3	10	0.905	0.970	0	8	6	1.024	3.713	0					
4	1	1.173	1.383	0	8	7	0.887	3.536	0					
4	2	1.096	1.124	0	8	8	1.132	3.561	0					
4	3	1.072	1.122	0	8	9	1.078	3.804	1500					
4	4	1.107	1.356	0	8	10	0.894	3.835	1500					
4	5	0.985	1.350	0	9	1	1.133	4.427	0					
4	6	0.985	1.350	0	9	2	1.133	4.449	0					
4	8	1.042	1.397	700	9	3	1.111	4.522	0					
4	9	1.042	1.402	1500	9	4	1.088	4.666	0					
4	10	0.971	1.460	750	9	5	1.062	4.452	0					
5	1	1.159	1.902	0	9	6	0.974	4.670	0					
5	2	1.119	1.881	0	9	8	1.130	4.671	0					
5	3	1.114	1.854	0	9	9	1.109	4.602	1550					
5	4	1.052	1.746	0	9	10	0.948	4.698	0					
5	5	1.052	1.854	0	10	1	1.150	5.529	0					
5	8	1.117	1.892	0	10	2	1.150	5.606	0					
5	10	1.004	1.864	0	10	3	1.129	5.060	0					
6	1	1.159	1.945	0	10	4	1.108	5.412	0					
6	2	1.126	2.053	0	10	5	1.090	5.384	0					
6	3	1.126	2.031	0	10	6	1.024	5.047	0					
6	4	1.080	2.047	0	10	7	0.940	4.937	0					
6	5	1.006	2.061	0	10	8	1.151	5.189	0					
6	6	1.005	2.043	0	10	9	0.979	5.088	3750					
6	8	1.138	2.044	0	10	10	0.927	5.816	0					

Table 3.11: Results for the small instances for the Approximate Dynamic Programming method

structure of the basic Dynamic Programming formulation to perform a heuristic on. This heuristic uses sample paths, which are fixed decision plans for all microCHPs, and extracts information from the objective values that belong to these sample paths to approximate the real value of the optimal solution.

3.7 COLUMN GENERATION

The heuristics that are developed in the previous sections treat the twodimensional aspect of the microCHP planning problem simultaneously, by concentrating on both feasibility (satisfying cooperational constraints) and profit maximization. Since this treatment shows increasing difficulties in the context of feasibility for the medium instances (e.g. an increasing error for an increasing number of intervals in the local

<i>houses</i>	z_{ls}/N	<i>time (s)</i>	<i>error</i>	<i>infeas. (%)</i>
25	1.019	22.3	2342	41.7
50	1.024	68.7	4042	41.7
75	1.026	137.6	13281	41.7
100	1.026	242.2	15175	50.0
<i>production pattern</i>	z_{ls}/N	<i>time (s)</i>	<i>error</i>	<i>infeas. (%)</i>
11	1.191	121.8	0	0.0
12	1.063	117.4	1796	41.7
13	0.951	116.5	6742	66.7
14	0.891	115.1	26302	66.7
<i>intervals</i>	z_{ls}/N	<i>time (s)</i>	<i>error</i>	<i>infeas. (%)</i>
24	1.002	34.9	10353	75.0
48	1.052	83.1	0	0.0
96	1.018	235.0	15777	56.3

Table 3.12: Results for the medium instances for the Approximate Dynamic Programming method

search method), we shift our focus towards feasibility. The heuristic that we propose in this section is based on the framework of column generation [58]. The main advantage of column generation in general is that it offers a technique that can be separated into tractable parts. It aims at reducing the state space of the problem in a natural way, which can be best explained by an example.

Example: minimizing waste in a glass company [12]

Suppose we have a glass company that manufactures different types of windows $w \in W = \{1, \dots, N_W\}$ that differ in length and height. This company has several production lines that produce standardized glass plates, from which the different types of windows are to be cut. Each window type has a certain customer demand d_w , that needs to be fulfilled. Given this demand, the glass company wants to minimize the number of used glass plates that have to be produced (and thus to minimize the glass loss/waste), such that all demand can be cut from these plates.

To solve this problem we can define a cutting pattern $p \in P$ to represent a specific way to cut a glass plate. Such a pattern p consists of nonnegative integer numbers of windows for all types ($p = (s_{1,p}, \dots, s_{N_W,p})$), such that these windows can be cut from the glass plate. For each cutting pattern we have to choose a nonnegative value x_p , which specifies how many glass plates are produced using this cutting pattern. The constraints for the variables x_p are:

$$\sum_{p \in P} s_{w,p} x_p \geq d_w \quad \forall w \in W \quad (3.44)$$

$$x_p \in \mathbb{N} \quad \forall p \in P. \quad (3.45)$$

Of course, the total number of used glass plates has to be minimized:

$$\min \sum_{p \in P} x_p. \quad (3.46)$$

The optimization problem is then given by (3.46), (3.44) and (3.45).

The number of possible cutting patterns increases significantly when the number of different window types increases, which could result in increasing computational times to solve practical problem instances. To overcome this, the column generation technique aims at using only a limited set of cutting patterns $P_{lim} \subset P$, instead of the complete large set of feasible cutting patterns P , and increasing this set P_{lim} by adding patterns that could improve the current solution. The optimization problem that uses such a limited set of patterns has the following form:

$$\begin{aligned}
 \min \quad & \sum_{p \in P_{lim}} x_p \\
 \sum_{p \in P_{lim}} s_{w,p} x_p & \geq d_w & \forall w \in W \\
 x_p & \in \mathbb{N} & \forall p \in P_{lim}.
 \end{aligned} \tag{3.47}$$

In a solution to the optimization problem, the constraints (3.44) bound the minimum number of necessary plates. When we increase the right hand side of (3.44) by one, for some constraints the objective value increases by a certain amount. This amount is called the shadow price λ_w of this constraint. The value λ_w results from the LP-relaxation of (3.47). New cutting patterns are now created by looking for the combination of window types y_w that fit in a glass plate and maximize their combined weighed influence:

$$\begin{aligned}
 \max \quad & \sum_{w \in W} \lambda_w y_w \\
 s.t. \quad & y_w \in \mathbb{N} & \forall w \in W \\
 & (y_1, \dots, y_{N_W}) \in P.
 \end{aligned} \tag{3.48}$$

If $\sum_{w \in W} \lambda_w y_w > 1$ a cutting pattern $p_{new} = (s_{1,p_{new}}, \dots, s_{N_W,p_{new}}) = (y_1, \dots, y_{N_W})$ can be added to the set P_{lim} , that could improve the current solution of the limited optimization problem (3.47). As long as we find new cutting patterns that improve the current solution, we can continue solving the main problem (3.47) and sub problem (3.48) iteratively.

3.7.1 GENERAL IDEA

The column generation technique divides a problem into two problems: a main problem that selects patterns from pattern sets for each microCHP to obtain a certain objective, and a sub problem that generates new patterns for different microCHPs that are attractive for the main problem in the sense that they offer some value to the current pattern sets.

The microCHP planning problem offers a natural framework to apply column generation. As we have shown in previous sections, the separation of dimensions is a crucial step in the search for a heuristic that attains good results in practice. Recall that the dependencies in time are represented by hard constraints and the dependencies in space by weak constraints. Therefore we consider a pattern to be a

sequence in time of electricity generation that corresponds to a feasible sequence of binary on/off decisions for a microCHP. In addition to that, it is a promising approach to consider a pattern in this way in the context of scalability requirements. Namely, since the time horizon is fixed for each pattern, they can be easily combined into patterns at a higher hierarchical level. This structure is discussed in more detail in Chapter 6.

The main disadvantage of the heuristics that are presented in the previous sections, is that they combine the search for a solution that maximizes the profit of the group of microCHPs with the search for a solution that minimizes the deviation from the weak cooperational constraints. The column generation method gives a more direct way of focusing on either one of these conflicting optimization objectives. Especially, we use the column generation technique to concentrate on minimizing the deviation from the weak cooperational bounds, since this problem has not been solved that good as we aim at.

The basic idea for the column generation approach is depicted in Figure 3.18. Figure 3.18a shows four pattern sets S_i for microCHP $i = 1, \dots, 4$, which each consist of six patterns. For each of these sets S_i one pattern is selected and the combined electricity generation is plotted in the middle of the figure. It shows some deviation from the cooperational bound set by $P = P^{upper} = P^{lower} = 2$. Based on this deviation, new patterns are generated for each set as can be seen in Figure 3.18b. Figure 3.18c shows that with these extended pattern sets a solution is found such that the cooperational bound P can always be followed.

3.7.2 PROBLEM FORMULATION

The problem of minimizing the deviation from the cooperational bounds (desired production pattern) is given by:

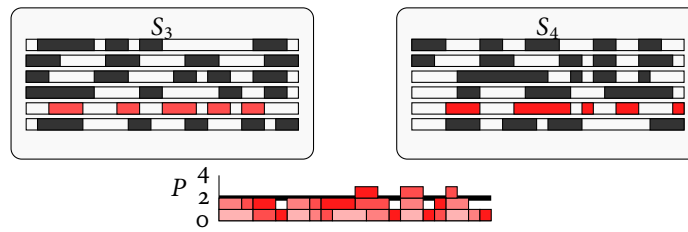
$$z_{\min} = \min \sum_{j=1}^{N_T} (sl_j + ex_j), \quad (3.49)$$

and by the constraints (3.15)-(3.30) and (3.33)-(3.36). Constraints (3.15)-(3.30) derive feasible generation patterns for individual microCHPs and constraints (3.33)-(3.36) represent the desired production bounds.


When we transform this ILP formulation into ILP formulations that can be used by the column generation approach, we need to introduce the notion of patterns, which are currently not included in the ILP formulation. Next we show the main problem of the column generation approach, followed by the sub problem of generating new patterns. Finally we give an overview of the complete algorithm.

Patterns


The indicator set of patterns $\mathbb{P} = \{1, \dots, N_{\mathbb{P}}\}$ represents all possible production patterns in a horizon of 24 hours of the type of microCHP that is considered, regardless of heat demand requirements or total desired electricity production, but



(a) Pattern sets for four microCHPs and the selection of patterns to minimize the deviation from P



(b) The extension of the pattern sets by finding new promising patterns



(c) Extended pattern sets for four microCHPs and the selection of patterns to minimize the deviation from P

Figure 3.18: The idea of the column generation technique applied to the microCHP planning problem

including appliance restrictions. For each pattern $p \in \mathbb{P}$ a corresponding electricity generation vector $pe_p = (pe_{p,1}, \dots, pe_{p,N_T})$ can be deduced. Note that the set of patterns can be extremely large.

Since each microCHP has to respect individual requirements (due to local heat demand and heat buffer requirements), the feasibility of patterns may differ between microCHPs: if a pattern respects the constraints (3.15)-(3.30) for one microCHP, it does not necessarily respect the constraints of another microCHP/house. Therefore, we cannot use a single pattern set from which one pattern has to be chosen for each microCHP, but we need to define pattern sets for each individual microCHP. This set of feasible patterns for microCHP i is denoted by $F_i \subset \mathbb{P}$ which takes the local constraints into account of the building where microCHP i is installed.

Main problem

First we construct an ILP formulation, which includes the notion of patterns such that a formulation that is equivalent to the original ILP given by (3.49), (3.15)-(3.30) and (3.33)-(3.36) is derived. Then we subtly adapt this formulation into one that acts as the main problem in our column generation heuristic.

In general, the offered bounds on the market (i.e. the desired production pattern of the total fleet of microCHPs) are represented by upper and lower bound vectors $P^{upper} = (P_1^{upper}, \dots, P_{N_T}^{upper})$ and $P^{lower} = (P_1^{lower}, \dots, P_{N_T}^{lower})$. A production profile for a microCHP is defined as a vector $pe_p = (pe_{p,1}, \dots, pe_{p,N_T})$. The problem is to pick exactly one pattern p for each microCHP, such that the sum of all production patterns falls between the lower and upper bound of the desired production pattern in all time intervals. For this selection decision, we introduce a binary decision variable $y_{i,p}$ indicating whether production profile pe_p is chosen for generator i (in this case $y_{i,p} = 1$) or not ($y_{i,p} = 0$). Of course we may only select locally feasible patterns (from the sets F_i). This results in the following Integer Linear Program (ILP) formulation:

$$\min \sum_{j=1}^{N_T} (sl_j + ex_j) \quad (3.50)$$

$$\sum_{i=1}^N \sum_{p \in F_i} pe_{p,j} y_{i,p} + sl_j \geq P_j^{lower} \quad \forall j \in J \quad (3.51)$$

$$\sum_{i=1}^N \sum_{p \in F_i} pe_{p,j} y_{i,p} - ex_j \leq P_j^{upper} \quad \forall j \in J \quad (3.52)$$

$$\sum_{p \in F_i} y_{i,p} = 1 \quad \forall i \in I \quad (3.53)$$

$$sl_j, ex_j \geq 0 \quad \forall j \in J \quad (3.54)$$

$$y_{i,p} \in \{0, 1\}. \quad (3.55)$$

In Equations (3.51) and (3.52) slack and excess variables sl_j and ex_j are introduced to calculate the deviation from the desired (and predefined) production pattern

(P^{upper}, P^{lower}). The sum of slack and excess variables is minimized in Equation (3.50). Finally, Equation (3.53) requires that exactly one pattern is chosen for each generator.

A feasible planning is achieved when the sum of slack and excess variables equals 0. If no feasible planning can be found, the objective value is a measure of the deviation from the desired production pattern.

The problem formulated by equations (3.50)-(3.55) takes into account only locally feasible production patterns from the sets F_i . These sets, however, are still very large and it is not an option to generate these sets explicitly. For this reason a column generation technique is developed.

The column generation technique starts with a relatively small set of feasible patterns $S_i \subset F_i$ for each microCHP i . By looking at only a small set of patterns the above ILP problem can be solved relatively fast. However, this comes with a possible loss of patterns that are needed for a high quality solution. The group might perform better, when some specific feasible production patterns from F_i would be in the feasible pattern sets S_i of the individual microCHPs. Unfortunately, we do not know on beforehand which patterns are in the final solution. Therefore it is the idea of the column generation technique to improve the current solution step by step, by searching for the patterns which hopefully improve the current solution by a high value, and by adding these patterns to the (small) feasible pattern set S_i of the corresponding microCHP. We have chosen to expand the pattern set S_i by at most one pattern per iteration as the heuristic evolves.

The column generation technique uses a main problem and sub problems (as indicated in Algorithm 3). The main problem is similar to equations (3.50)-(3.55), with the only difference that the set F_i is replaced by S_i :

$$\min \sum_{j=1}^{N_T} (sl_j + ex_j) \quad (3.56)$$

$$\sum_{i=1}^N \sum_{p \in S_i} pe_{p,j} y_{i,p} + sl_j \geq P_j^{lower} \quad \forall j \in J \quad (3.57)$$

$$\sum_{i=1}^N \sum_{p \in S_i} pe_{p,j} y_{i,p} - ex_j \leq P_j^{upper} \quad \forall j \in J \quad (3.58)$$

$$\sum_{p \in S_i} y_{i,p} = 1 \quad \forall i \in I \quad (3.59)$$

$$sl_j, ex_j \geq 0 \quad \forall j \in J \quad (3.60)$$

$$y_{i,p} \in \{0, 1\}. \quad (3.61)$$

Sub problem

The second phase of the column generation technique consists of creating new patterns that can be added to the current pattern sets S_i for each microCHP in the main problem. These new patterns should contribute to the existing sets in the sense that they should give possibilities to decrease the objective value in the first

phase (i.e. the sum of slack and excess). A new pattern pe_g is only added 1) if it is a locally feasible pattern ($g \in F_i$) and 2) if it may improve the existing solution. We follow the intuition of the column generation approach to use shadow prices that result from the LP-relaxation of the main problem to determine candidate patterns.

Let λ_j represent the shadow prices for equations (3.57) and (3.58), obtained from the dual of (3.56)-(3.61). Following the idea of duality, a new pattern g is a good candidate to improve the existing solution if:

$$\sum_{j=1}^{N_T} \lambda_j (pe_{g,j} - pe_{c,j}) > 0, \quad (3.62)$$

where c represents the chosen pattern in the current solution of the main problem (i.e. $y_{i,c} = 1$). In practice, λ_j appears to be either -1 , 0 or 1 . If equation (3.57) is strictly respected, $\lambda_j = 1$ and the new pattern is encouraged to generate more electricity in this time interval than in the selected pattern. On the opposite, if equation (3.58) is strictly respected, $\lambda_j = -1$ and the new pattern is encouraged to generate less electricity in this time interval than in the selected pattern. In this way the newly generated pattern is optimized towards the output of the main problem. However, this does not necessarily mean that this pattern can be automatically selected in the new solution of the main problem, since newly added patterns of other microCHPs (by solving these sub problems) could lead to different choices for this specific microCHP. So, the main problem has to be solved completely at each visit.

The second requirement (g is locally feasible) is formalized by the ILP formulation of the sub problem (3.63)-(3.79), which follows from the ILP formulation in Section 3.3. Noteworthy to mention is the objective of maximizing the added value of the electricity generation pattern in (3.63) and the notational change in equation (3.66). We also point out that for each microCHP i this sub problem has to be solved individually.

$$\max \sum_{j=1}^{N_T} \lambda_j (pe_{g,j} - pe_{c,j}) \quad (3.63)$$

$$x_j^i \in \{0, 1\} \quad \forall j \in J \quad (3.64)$$

$$g_j^i = G_{\max}^i x_j^i - \sum_{k=0}^{N_{up}^i - 1} \check{G}_{k+1}^i start_{j-k}^i + \sum_{k=0}^{N_{down}^i - 1} \check{G}_{k+1}^i stop_{j-k}^i \quad \forall j \in J \quad (3.65)$$

$$pe_{g,j} = \alpha^i g_j^i \quad \forall j \in J \quad (3.66)$$

$$start_j^i \geq x_j^i - x_{j-1}^i \quad j = 2 - MR^i, \dots, N_T \quad (3.67)$$

$$start_j^i \leq x_j^i \quad j = 2 - MR^i, \dots, N_T \quad (3.68)$$

$$start_j^i \leq 1 - x_{j-1}^i \quad j = 2 - MR^i, \dots, N_T \quad (3.69)$$

$$stop_j^i \geq x_{j-1}^i - x_j^i \quad j = 2 - MO^i, \dots, N_T \quad (3.70)$$

$$stop_j^i \leq x_{j-1}^i \quad j = 2 - MO^i, \dots, N_T \quad (3.71)$$

$$stop_j^i \leq 1 - x_j^i \quad j = 2 - MO^i, \dots, N_T \quad (3.72)$$

$$start_j^i \in \{0, 1\} \quad j = 2 - MR^i, \dots, N_T \quad (3.73)$$

$$stop_j^i \in \{0, 1\} \quad j = 2 - MO^i, \dots, N_T \quad (3.74)$$

$$x_j^i \geq \sum_{k=j-MR^i+1}^{j-1} start_k^i \quad \forall j \in J \quad (3.75)$$

$$x_j^i \leq 1 - \sum_{k=j-MO^i+1}^{j-1} stop_k^i \quad \forall j \in J \quad (3.76)$$

$$hl_1^i = BL^i \quad (3.77)$$

$$hl_j^i = hl_{j-1}^i + g_{j-1}^i - H_{j-1}^i - K^i \quad \forall j \in J \setminus \{1\} \cup \{N_T + 1\} \quad (3.78)$$

$$0 \leq hl_j^i \leq BC^i \quad j \in J \cup \{N_T + 1\} \quad (3.79)$$

If constraint (3.62) is satisfied, the pattern g is added to the set S_i .

To speed up the computational time that is needed to find the optimal solution of the sub problem for microCHP i , we may change the objective (3.63) into an objective that aims at the binary commitment of the microCHP instead of on the actual electricity generation:

$$\max \sum_{j=1}^{N_T} \lambda_j (x_j^i - x_{i,c,j}), \quad (3.80)$$

where $x_{i,c,j}$ is the binary commitment of the chosen pattern c . The idea is that, by focusing on unit commitment rather than on production, the side effects of production (startup/shutdown) diminish, which might have a positive influence on the outcome of the planning process.

Algorithm

To summarize, the solution method is given in Algorithm 3. Initially, the pattern sets of microCHPs i consist of a single pattern in each set; this pattern is optimized to maximize its local profit. In each iteration, the main problem is solved first, after which for each microCHP the sub problem is solved and new feasible and improving patterns are added. Based on experience we set the maximum running time of finding a solution for the main problem on 60 seconds and for the sub problem on 10 seconds for the electricity generation based objective and on 1 second for the binary decision based objective. Note that for both sub problems the resulting electricity patterns are always used in the main problem (instead of the binary decision variables). The stopping criteria are twofold. For the routine to continue, we demand that at least one sub problem leads to an improvement. In addition to that we require that the main problem always finds an improvement for the global objective value. If one of these requirements is not satisfied, the heuristic stops.

Algorithm 3 Column generation

```

init  $S_i$  for all  $i$ 
solve main problem
for all  $i$  do
  solve sub problem
end for
while stopping criteria not met do
  for all  $i$ :  $\sum_{j=1}^{N_T} \lambda_j (pe_{i,g,j} - pe_{i,c,j}) > 0$  do
     $S_i \leftarrow S_i \cup g$ 
  end for
  solve main problem
  for all  $i$  do
    solve sub problem
  end for
end while

```

3.7.3 RESULTS

In this section we show the results for the small and medium instances. We distinguish two variants of the column generation technique. The first variant uses the sub problem that aims at electricity generation (3.63), whereas the second variant uses objective (3.80). The column generation approach is modelled in AIMMS using CPLEX 12.2.

Results for the small instances

Table 3.13 shows the results for the small instances, where the local objective is oriented at electricity generation, whereas Table 3.14 gives the results for when the objective is based on the on/off decisions. Problem instances $I(8,1)$, $I(9,1)$ and $I(10,1)$ for this second variant show slightly worse solutions in comparison to the optimal values. Although the desired production pattern is completely free in these instances, the local objective that focuses on the on/off decisions rather than on the electricity output gives the reason for this difference.

Both tables show some instances for which the deviation from the desired generation bounds is not 0. However, in comparison with the results for the local search method the amount of deviation remains relatively small, especially for the 'larger' small instances of 8, 9 and 10 microCHPs. This trend is continued for the medium instances.

Results for the medium instances

Table 3.15 and 3.16 present the average results of the medium instances, categorized by the number of houses, the production bounds variant and the number of time in-

instance		solution				instance		solution			
N	variant	z_{\max}^N	time (s)	iterations	mismatch	N	variant	z_{\max}^N	time (s)	iterations	mismatch
1	1	1.147	0.00	1	0	6	10	0.882	0.52	2	0
1	2	0.709	0.00	2	0	7	1	1.156	0.26	1	0
1	10	0.709	0.00	2	0	7	2	0.937	0.53	2	0
2	1	1.236	0.00	1	0	7	3	0.818	0.52	2	0
2	2	0.915	0.00	2	0	7	4	0.818	0.53	2	0
2	3	0.884	0.00	2	0	7	5	0.918	0.52	2	0
2	4	0.884	0.00	2	0	7	6	0.923	1.06	4	0
2	5	0.884	0.00	2	0	7	7	0.885	1.06	4	100
2	6	0.884	0.00	2	0	7	8	1.069	0.51	2	0
2	10	0.884	0.00	2	0	7	10	0.871	1.04	4	0
3	1	1.197	0.00	1	0	8	1	1.156	0.26	1	0
3	2	1.197	0.00	1	0	8	2	0.818	0.55	2	0
3	3	0.959	0.26	2	0	8	3	0.818	0.55	2	0
3	4	0.959	0.26	2	0	8	4	0.818	0.51	2	0
3	5	0.959	0.25	3	100	8	5	0.948	0.53	2	0
3	10	0.959	0.25	3	100	8	6	0.944	0.52	2	0
4	1	1.183	0.00	1	0	8	7	0.875	1.56	5	500
4	2	0.929	0.25	2	0	8	8	0.971	0.53	2	0
4	3	1.004	0.27	2	0	8	9	0.960	0.80	3	0
4	4	1.004	0.26	2	0	8	10	0.888	1.84	5	1100
4	5	0.910	0.27	2	0	9	1	1.153	0.27	1	0
4	6	0.956	0.27	2	0	9	2	0.811	0.78	2	0
4	8	0.920	0.79	4	700	9	3	0.811	0.80	2	0
4	9	0.920	0.78	4	1500	9	4	0.811	0.80	2	0
4	10	0.946	0.52	3	2150	9	5	0.927	0.78	2	0
5	1	1.164	0.25	1	0	9	6	0.910	0.53	2	0
5	2	0.991	0.53	2	0	9	8	1.097	0.53	2	0
5	3	0.931	0.53	2	0	9	9	1.021	1.33	3	0
5	4	0.904	0.52	2	0	9	10	0.892	1.08	3	0
5	5	0.904	0.51	2	0	10	1	1.176	0.52	1	0
5	8	1.036	0.81	3	400	10	2	0.912	0.78	2	0
5	10	0.943	0.53	3	350	10	3	0.937	0.80	2	0
6	1	1.163	0.26	1	0	10	4	0.839	0.78	2	0
6	2	1.015	0.52	2	0	10	5	0.973	0.80	2	0
6	3	1.015	0.53	2	0	10	6	0.939	0.77	2	0
6	4	1.003	0.53	2	0	10	7	0.903	1.58	4	0
6	5	0.979	0.53	2	0	10	8	1.037	0.78	2	0
6	6	0.918	0.53	2	0	10	9	1.020	2.92	5	100
6	8	1.052	0.78	3	0	10	10	0.893	3.73	7	50

Table 3.13: Results for the column generation method applied to the small instances (local objective is electricity generation)

tervals in the planning horizon. Both tables show a large decrease in the error value, which may result from the more direct search towards minimizing the deviation from the desired aggregated electricity bounds. The error increases linearly in the number of houses, which is no surprise. Compared to the error development in the local search method and the approximate dynamic programming based method, we now see behaviour that can be explained when we differentiate the error to the production pattern variant or to the number of intervals. If we differentiate the error to the production pattern variant, this shows the natural trend that the tighter the bounds are the larger the error is. However, this error is much smaller than for the local search method. When we increase the number of intervals (and thus, increase the flexibility of the planning problem), we now see an improvement

instance		solution				instance		solution			
N	variant	z_{\max}^N	time (s)	iterations	mismatch	N	variant	z_{\max}^N	time (s)	iterations	mismatch
1	1	1.147	0.00	1	0	6	10	0.911	0.80	4	450
1	2	1.050	0.00	2	0	7	1	1.156	0.25	1	0
1	10	1.050	0.00	2	0	7	2	1.080	0.52	2	0
2	1	1.236	0.00	1	0	7	3	0.964	0.53	2	0
2	2	0.890	0.00	2	0	7	4	0.841	0.53	2	0
2	3	0.995	0.00	2	0	7	5	1.015	0.52	2	0
2	4	0.995	0.00	2	0	7	6	0.915	0.80	3	0
2	5	0.995	0.00	2	0	7	7	0.920	1.56	6	100
2	6	0.995	0.00	2	0	7	8	0.929	0.26	2	0
2	10	0.995	0.00	2	0	7	10	0.927	1.04	4	250
3	1	1.197	0.00	1	0	8	1	1.154	0.25	1	0
3	2	1.197	0.00	1	0	8	2	1.041	0.53	2	0
3	3	0.899	0.00	2	0	8	3	0.938	0.53	2	0
3	4	0.899	0.00	2	0	8	4	0.831	0.51	2	0
3	5	1.017	0.25	3	150	8	5	0.957	0.53	2	0
3	10	1.017	0.25	3	150	8	6	0.894	0.53	2	0
4	1	1.183	0.00	1	0	8	7	0.876	1.33	5	0
4	2	0.892	0.00	2	0	8	8	0.945	0.51	2	0
4	3	1.090	0.25	2	0	8	9	0.977	1.06	4	0
4	4	1.000	0.26	2	0	8	10	0.878	3.98	7	550
4	5	0.883	0.26	3	0	9	1	1.151	0.26	1	0
4	6	0.984	0.26	2	0	9	2	1.018	0.53	2	0
4	8	0.825	0.53	4	700	9	3	1.050	0.52	2	0
4	9	0.914	0.53	4	1500	9	4	0.915	0.51	2	0
4	10	0.946	0.52	4	1450	9	5	0.931	0.53	2	0
5	1	1.164	0.27	1	0	9	6	0.909	0.53	2	0
5	2	0.863	0.26	2	0	9	8	1.003	0.51	2	0
5	3	0.902	0.26	2	0	9	9	1.046	0.78	3	0
5	4	0.941	0.25	2	0	9	10	0.905	1.33	4	0
5	5	0.941	0.25	2	0	10	1	1.174	0.26	1	0
5	8	1.056	0.52	3	400	10	2	1.010	0.53	2	0
5	10	0.942	0.78	4	0	10	3	1.010	0.53	2	0
6	1	1.163	0.26	1	0	10	4	0.832	0.53	2	0
6	2	0.999	0.53	2	0	10	5	0.900	0.80	2	0
6	3	0.938	0.52	2	0	10	6	0.930	0.53	2	0
6	4	1.041	0.26	2	0	10	7	0.911	1.34	4	0
6	5	0.917	0.52	2	0	10	8	1.107	0.52	2	0
6	6	0.868	0.26	2	0	10	9	0.886	6.88	6	50
6	8	0.941	0.53	3	0	10	10	0.876	4.54	7	50

Table 3.14: Results for the column generation method applied to the small instances (local objective is binary commitment)

(a general decrease) of the error, which we did not see before in the local search method. The error decreases for the switch from 24 to 48 intervals and only shows a minor increase for the switch from 48 to 96 intervals, which is an improvement when compared to the results for the approximate dynamic programming based method. This minor increase can origin from the time limit of 60 seconds on the main problem, which results in prematurely abortion of solving the main problem. Note that the sub problems are often solved faster than the limits of 10 respectively 1 seconds. However, in the 9 extra seconds for the electricity based objective on average significant improvements are made for 48 intervals and for 96 intervals, which justifies the choice for this increased computational time limitation.

When we compare the two variants of the column generation approach, the

<i>houses</i>	z_{ls}/N	<i>time (s)</i>	<i>iterations</i>	<i>error</i>	<i>infeas. (%)</i>
25	0.896	316.0	3.8	978	41.7
50	0.873	622.5	3.9	1998	41.7
75	0.901	939.1	3.9	3025	41.7
100	0.878	1223.7	3.9	3925	41.7
<i>production pattern</i>	z_{ls}/N	<i>time (s)</i>	<i>iterations</i>	<i>error</i>	<i>infeas. (%)</i>
11	0.840	362.1	1.8	0	0.0
12	0.951	612.5	2.8	1283	33.3
13	0.900	809.6	4.0	3367	33.3
14	0.856	1317.1	7.0	5277	100.0
<i>intervals</i>	z_{ls}/N	<i>time (s)</i>	<i>iterations</i>	<i>error</i>	<i>infeas. (%)</i>
24	0.956	22.8	3.4	7141	75.0
48	0.848	891.7	4.0	116	25.0
96	0.856	1411.4	4.3	189	25.0

Table 3.15: Results for medium instances (local objective is electricity generation)

<i>houses</i>	z_{ls}/N	<i>time (s)</i>	<i>iterations</i>	<i>error</i>	<i>infeas. (%)</i>
25	0.898	104.6	4.2	950	33.3
50	0.884	128.0	3.8	1919	33.3
75	0.901	191.8	3.8	3059	41.7
100	0.907	237.4	3.8	3917	41.7
<i>production pattern</i>	z_{ls}/N	<i>time (s)</i>	<i>iterations</i>	<i>error</i>	<i>infeas. (%)</i>
11	0.846	56.5	1.8	0	0.0
12	0.986	77.0	2.6	1283	33.3
13	0.899	116.3	3.8	3367	33.3
14	0.860	412.1	7.4	5194	83.3
<i>intervals</i>	z_{ls}/N	<i>time (s)</i>	<i>iterations</i>	<i>error</i>	<i>infeas. (%)</i>
24	0.969	38.5	3.7	7141	75.0
48	0.859	181.2	3.9	19	12.5
96	0.864	276.6	4.1	224	25.0

Table 3.16: Results for medium instances (local objective is binary commitment)

variant with the focus on the binary decision variables in the sub problem shows a clear advantage in computational effort. However, note that this advantage diminishes as the heuristic is parallelized on calculating entities close to each microCHP appliance. Since the performance of both heuristics is similar for the objective value, the number of iterations and the error value, and since the binary commitment variant even shows a smaller percentage of infeasible solutions, we prefer this variant over the variant with a focus on the actual electricity output. An explanation for the (slightly) worse performance of this electricity led variant could be that too much attention is drawn to the apparently negligible startup and shutdown behaviour of the microCHP.

The developed heuristics in this chapter show that feasibility of the problem of maximizing profit, while satisfying global electricity bounds, is not easily reached. We want to illustrate the benefits of the column generation technique with respect to this feasibility aspect. Therefore we derive lower bounds on the guaranteed minimal (absolute) deviation (which we call mismatch) between possible and desired generation for a special type of problem instances. Then we show that the column generation technique finds solutions that are equal or at least very close to these lower bounds. In this example we focus on the problem of minimizing the mismatch instead of maximizing the profit, since the main objective of this section is to show the feasibility aspects of the problem.

Special type of problem instances

The special type of instances we consider in this section is the set of problem instances for which we have no startup and shutdown behaviour: the electricity generation has a one to one correspondence to the binary on/off decisions. We choose for this set of problem instances to clarify the principle effect of the column generation technique, which is to minimize the deviation from the global electricity bounds. In this setting the computational results show that the solution that is found is close to or even equal in many instances to a derived lower bound on the mismatch for these instances.

In principle it is also possible to derive lower bounds for other types of instances. However, the startup and shutdown behaviour has an undesirable effect. Namely, this side effect influences the proposed calculation of the lower bounds in such a way that we cannot identify the origin of the gap between these lower bounds and the found mismatch: is this gap mostly due to a weak estimate of the lower bound or due to the inability of the column generation method to find a good mismatch? For this reason we neglect startup and shutdown effects in this section. In this way we can concentrate completely on the binary commitment of microCHP appliances.

Simplifications for the special type of problem instances

The formulation of the sub problem (3.63)-(3.79) can be simplified for the special type of instances. For the main problem it suffices to know that the patterns have been checked for feasibility before; these feasible patterns are fixed input data for the main problem. The feasibility check is simplified by using two parameter sets, specifying in each interval j the minimum number of intervals the microCHP generator i should have run ($MinOn_{i,j}$) and the maximum number of intervals the generator could have run ($MaxOn_{i,j}$) up to and including the current interval j . These parameters $MinOn_{i,j}$ and $MaxOn_{i,j}$ are derived from the same heat demand profiles as we used in the medium instances for the profit maximization objective. The calculations that we perform to derive these parameters exclude startup and shutdown behaviour, as this is not included in these special instances,

but include the use of a heat buffer (heat demand and periodical heat loss). Technical runtime/offtime constraints of the microCHP are automatically fulfilled by using time intervals of half an hour. If the starting patterns in S_i are chosen feasible, all possible solutions in the eventual set are feasible, since the newly generated patterns are always feasible.

For the considered special case the sub problem of the column generation is now given by the following ILP formulation for microCHP i using half an hour intervals:

$$\max \sum_{j=1}^{N_T} \lambda_j (pe_{g,j} - pe_{c,j}) \quad (3.81)$$

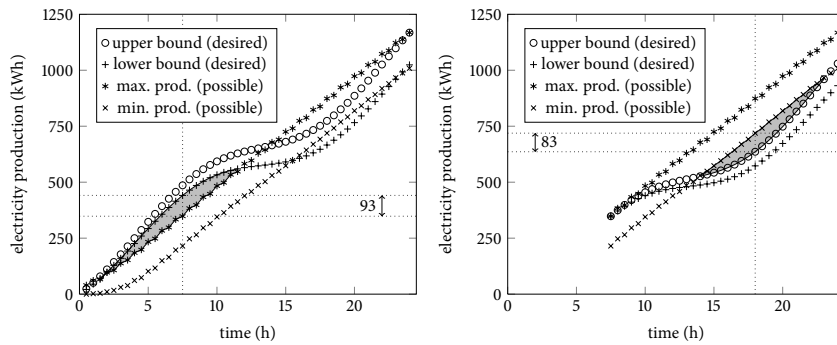
$$\sum_{k=1}^j pe_{g,k} \leq \frac{1}{2} MaxOn_{i,j} \quad \forall j \in J \quad (3.82)$$

$$\sum_{k=1}^j pe_{g,k} \geq \frac{1}{2} MinOn_{i,j} \quad \forall j \in J \quad (3.83)$$

$$2pe_{g,j} \in \{0,1\} \quad \forall j \in J, \quad (3.84)$$

where from all locally feasible patterns the one is chosen that maximizes the added value to the main problem. The factors 2 and $\frac{1}{2}$ are used, since we use time intervals of half an hour and $MaxOn_i$ and $MinOn_i$ are defined in time intervals. If constraint (3.62) is satisfied, the pattern $pe_{i,g}$ is added to the set S_i .

Lower bound calculation



(a) The total desired production and the total possible production result in a first phase lower bound calculation and the resulting lower bound improvement

Figure 3.19: The calculation of the lower bound of the group planning problem

The lower and upper bounds P^{lower} and P^{upper} (representing the desired production pattern) and the possible production bounds $MaxProd_i$ and $MinProd_i$

derived from $MaxOn_i$ and $MinOn_i$ ($MaxProd_i := \frac{1}{2}MaxOn_i$ and $MinProd_i := \frac{1}{2}MinOn_i$) form the basic input parameters of a problem instance. To derive a theoretical lower bound z_{LB} for the objective, we only look at these parameters. Since we have a minimization problem and the sum of slack and excess variables cannot be negative, the lower bound z_{LB} is at least 0.

The calculation of the lower bound works in phases. In each phase a minimal guaranteed mismatch (slack or excess) z_{LB}^{extra} is found and added to the current lower bound.

In the first phase, the additional value of the lower bound z_{LB}^{extra} equals:

$$z_{LB}^{extra} = \max_j \begin{cases} \sum_{k=1}^j P_k^{lower} - \sum_i MaxProd_{i,j} \\ \sum_i MinProd_{i,j} - \sum_{k=1}^j P_k^{upper} \\ 0. \end{cases} \quad (3.85)$$

This value equals the maximum deviation of the aggregated possible production from the aggregated desired production pattern over all intervals. An example of the results for this phase is shown in Figure 3.19a, where the aggregated minimal mismatch per time interval is given by the gray area. A maximum difference is found between the maximal possible production and the minimal desired production at 7.5 hours, with a value of 93. So, in this example, the theoretical lower bound has now improved from 0 to 93.

The first value of j for which a positive z_{LB}^{extra} is found is the starting point r for the calculation of the next phase. This starting point is important in two ways. First, the mismatch in previous intervals cannot be undone, since we only look at intervals $j > r$. Secondly, the starting point r offers a natural reset point; we can take our losses up to this interval (i.e. the mismatch in the previous intervals) and start with a renewed mismatch calculation. This reset point requires that the sum of desired maximum (minimum) production upto and including interval r can be replaced by the maximum (minimum) possible production upto and including interval r . Resetting to other values is either not allowed (in this case these total productions are larger (smaller) than the maximum (minimum) possible production and, therefore, not possible at r) or would increase the value of z_{LB}^{extra} . Considering this second option, these values are not fully incorporated in the current lower bound. More precisely, if these values are realized in a planning, the achieved mismatch upto and including r would increase by the difference to the maximum (minimum) possible production. So, $\sum_{k=1}^j P_k^{lower}$ can be replaced by $\sum_i MinProd_{i,r} + \sum_{k=r+1}^j P_k^{lower}$ and $\sum_{k=1}^j P_k^{upper}$ by $\sum_i MaxProd_{i,r} + \sum_{k=r+1}^j P_k^{upper}$. Again we look for mismatch in the

future (time intervals $j > r$):

$$z_{LB}^{extra} = \max_{j>r} \begin{cases} (\sum_i MinProd_{i,r} + \sum_{k=r+1}^j P_k^{lower}) - \sum_i MaxProd_{i,j} \\ \sum_i MinProd_{i,j} - (\sum_i MaxProd_{i,r} + \sum_{k=r+1}^j P_k^{upper}) \\ 0. \end{cases} \quad (3.86)$$

In the example, the second phase calculation is shown in Figure 3.19b, where an additional theoretical lower bound z_{LB}^{extra} of 83 is found. The theoretical lower bound is now: $z_{LB} = 93+83 = 176$. This process can now be iterated until no further positive values occur in the calculation of (3.86). Note that at each reset point the ‘direction’ of mismatch changes: based on the definition of r , we know that an additional mismatch in the same direction cannot occur, since the desired bounds are reset by the maximum value of the previous iteration. However, an additional mismatch in the other direction may always occur, although this additional mismatch per iteration is bounded by the smallest additional mismatch in the previous iteration.

Scenario

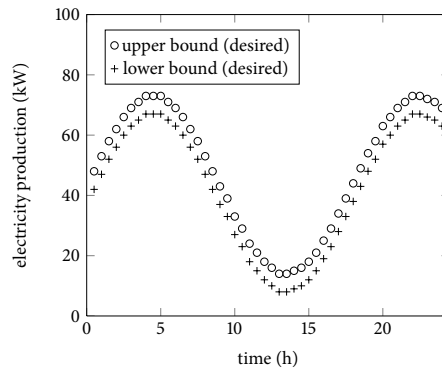


Figure 3.20: An example of a desired production pattern; a sine with amplitude 30 and period 18

We set up a scenario to answer the following kind of problem: the defined instances should provide a framework to test the quality of the column generation technique.

To support this question, we focus on variation in the offered/desired patterns and keep the possible production the same for the different instances. The variation is created by using sine functions, where we vary both in amplitude and in period.

The instances consist of a group of 100 microCHPs and 48 time intervals in a horizon of one day. The group size is too small to be able to act on the electricity

market at the moment, since the microCHP generates at the 1 kW level. However, this size gives a good indication of the possibilities of the planning method. Decisions are made on an half an hour basis. This is more fine grained than required, since the day ahead electricity market works on an hourly basis. However, using this setting the planning problem gets more realistic (and offers more possibilities for variation). The production patterns can be simply converted to hourly blocks when we want to transform the planning to the electricity market.

The maximum and minimum possible numbers of runtime intervals $MaxOn_i$ and $MinOn_i$ differ per microCHP and are derived directly from the heat demand in the medium instances. As mentioned before, they remain the same in all instances; variation is applied to the desired production profile. The aggregated values of the possible production are shown in Figure 3.19a.

The initial patterns in the sets S_i are derived from $MaxOn_i$ and $MinOn_i$. The microCHP sub problem starts with two patterns, one resulting from the earliest possible time intervals that the microCHP can be switched on, and one resulting from the latest possible time intervals that the microCHP has to be switched on. In the first case the microCHP stays on as long as the buffer is not at its upper limit and in the second case the microCHP stays off as long as the buffer is not at its lower limit.

Upper and lower bounds of the desired production are defined, based on a sine function and a constant. The sine function is given (and equal) for both the upper and the lower bound. The constant for the upper bound is maximized such that the total aggregated desired production stays within the limit given by the total maximally possible production of all microCHPs. Likewise, the constant for the lower bound is minimized, such that the total desired production is larger than or equal to the total minimally possible production. In other words, the upper bound P^{upper} is derived from the highest integer value of μ^{upper} for which a given amplitude amp (in kW) and period per (in hours) result in a total desired production that is still feasible, when only looking at the total possible production:

$$\max \mu^{upper} \quad (3.87)$$

$$\sum_j P_j^{upper} \leq \frac{1}{2} \sum_i MaxOn_{i,N_T} \quad (3.88)$$

$$P_j^{upper} = \frac{1}{2} rnd(amp \times \sin(f(per) \times j)) + \mu^{upper} \forall j, \quad (3.89)$$

where $f(per)$ is the frequency corresponding to the given period per and $rnd()$ is a rounding function that converts to the nearest integer. Likewise, the lower bound P^{lower} results from the lowest sine curve fitting in the possible minimum production:

$$\min \mu^{lower} \quad (3.90)$$

$$\sum_j P_j^{lower} \geq \frac{1}{2} \sum_i MinOn_{i,N_T} \quad (3.91)$$

$$P_j^{lower} = \frac{1}{2} rnd(amp \times \sin(f(per) \times j)) + \mu^{lower} \forall j. \quad (3.92)$$

The final time interval in Figure 3.19a shows that the lower and upper bound of the example fit within the possible total production domain. Figure 3.20 gives the resulting individual values (in kW) of this example with amplitude 30 and period 18.

Using the sketched approach, an instance can be defined as a pair $I'(amp, per)$ and a solution as a tuple $(I'(amp, per), z_{LB}, z_{found})$. For the instances, we choose $amp \in \{0, 1, \dots, 40\}$ and $per \in \{2, 3, \dots, 24\}$.

Results

Figure 3.21 shows the calculated lower bounds for the instances in a surface plot. The found solutions of the column generation technique are plotted on top of that surface plot. The results show that a tight match to the lower bounds is found

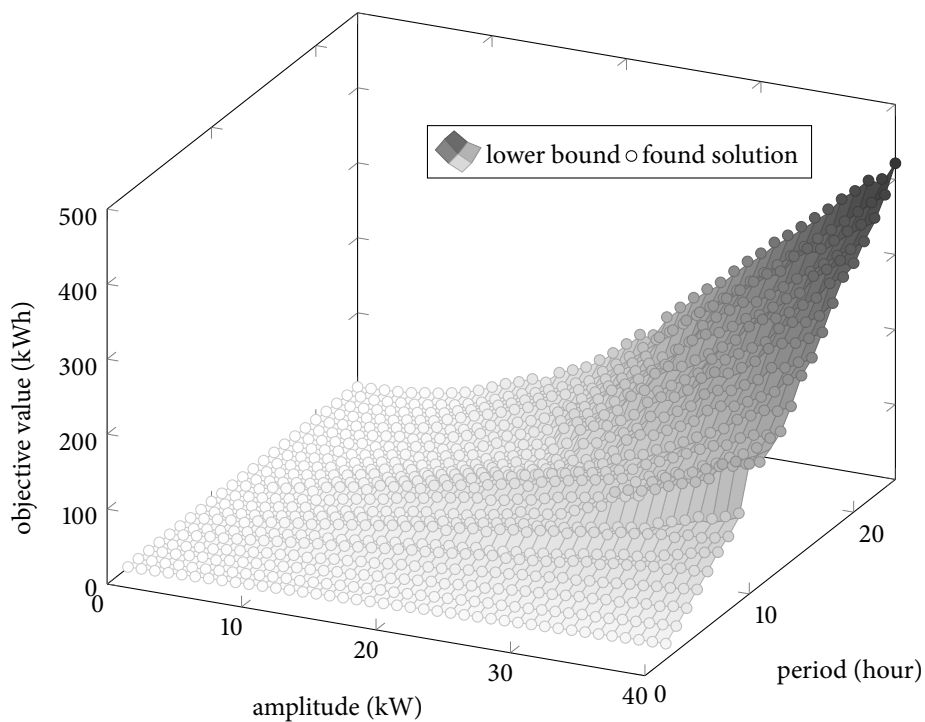


Figure 3.21: Calculated lower bounds and solutions derived from the column generation technique, for sines with varying amplitude and period

for all instances, which shows the strength of the column generation heuristic. Besides that, the lower bound value depends on the combination of both period and amplitude. A slow repetitive periodic behaviour of the desired aggregated

electricity production and a large amplitude of the desired production function lead to large lower bounds. On the other hand, a large amplitude combined with a short sine period (i.e. a fast repetitive behaviour of the sine) results in a small value for the lower bound, which is validated by the results. This indicates for the Virtual Power Plant case that we may ask relatively fluctuating production over the time horizon, as long as the running average is close to the average possible production. Positive/negative spikes in certain time intervals should be compensated for by negative/positive spikes in time intervals that are close to the time interval under consideration.




Figure 3.22: Computation times related to the number of iterations for the column generation technique

Figure 3.22 gives the computational times related to the number of iterations (i.e. the number of newly generated patterns) for the column generation technique. This figure shows a linear relation in the number of intervals, showing that the computational effort for the sub problem does not increase when the number of iterations grows. In the solution method we use a small modification: we use an LP-relaxation of the main problem during the iterations and solve the main problem normally as a final stage after termination of the iterative process. The influence of this final stage is visible in the computational times: the time limit of 60 seconds is (sometimes) reached in this final stage; and if so, it occurs only in this stage and not during earlier iterations.

Remark on the results

Based on the results in the previous section one might think that the calculated lower bound is always reached in the optimum. However, this is not the case. To show this a simple counterexample is constructed in Figure 3.23.

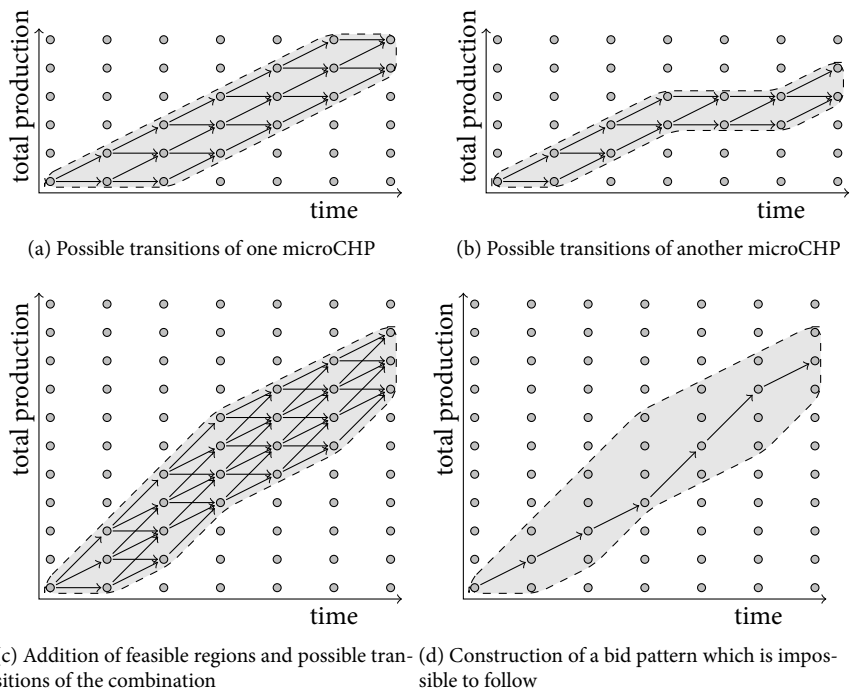


Figure 3.23: A counterexample for the natural fleet bounds

Figure 3.23a and 3.23b show the possible decision paths within the natural bounds (the gray area) for two households equipped with a microCHP; Figure 3.23c shows the combined decisions, including capacity constraints, for which the given decision path in Figure 3.23d is impossible to follow. This decision path stays within the gray area, indicating that the lower bound on the deviation from the possible production bounds is 0. Although the two generators may run simultaneously in the fourth or in the fifth time interval, it is impossible to have the two generators running simultaneously in both time intervals subsequently, due to the limited possibilities for the second household in the fourth and fifth time intervals. This counterexample shows that the lower bound is not always reached.

3.75 CONCLUSION

In this section a column generation technique is developed for the microCHP planning problem. This heuristic offers a special focus on minimizing the total deviation from the desired aggregated production bounds (mismatch) for a group of microCHPs. This method outperforms the local search method when we look at this deviation for the small and medium instances. Furthermore, we investigate a special type of problem instances, and show that the found mismatch is close or

equal to a calculated lower bound.

3.8 CONCLUSION

This chapter introduces the microCHP planning problem, which consists of the problem to plan the operation of domestic combined heat and power generators in a cooperational setting of a Virtual Power Plant. Locally these microCHP generators need to satisfy heat demand from the household, while globally the aggregated electricity output of the microCHPs needs to fulfill a desired production pattern. The operation of the microCHP itself is restricted to binary decisions to switch the appliance on or off; the related electricity output is then completely determined. In the microCHP planning problem, the profit of the Virtual Power Plant on an electricity market is maximized and/or the total deviation from the desired aggregated electricity output is minimized.

A mathematical description of the problem is given and it is shown that the problem is \mathcal{NP} -complete in the strong sense. Exact formulations by modelling the problem as an Integer Linear Programming or a dynamic programming model show that practical instances are indeed difficult to solve in limited computational time. Therefore, three heuristics are proposed. A local search method, based on the dynamic programming formulation, shows a large improvement in computational time; the deviation from the desired bounds however asks for improvements. An approximate dynamic programming approach shows interesting first results, but needs further evaluation on larger problem instances. A column generation technique offers a nice framework to minimize the deviation from the desired aggregated electricity output. For simplified instances it is shown, based on a lower bound calculation, that this method can solve this deviation (close) to optimality.

For the different approaches mainly only a basic variant is developed to explore the different concepts. Further research towards a real world implementation are necessary. Both the local search method and the column generation method are appropriate in the context of scalability. The division in global aggregation/optimization problems and local optimization problems offers a framework that is scalable.

EVALUATION OF THE MICROCHP PLANNING THROUGH REALTIME CONTROL

ABSTRACT – This chapter presents a short evaluation of the impact of demand uncertainty on the microCHP planning problem. It also covers the other two steps in the TRIANA methodology, being the prediction step and the realtime control step. In the context of the microCHP planning problem, the quality of local predictions and the ability to cope with realtime fluctuations in demand are sketched. Finally, possibilities of reserving heat capacity in heat buffers are depicted.

The translation of the planned production of a Virtual Power Plant a day ahead to the realtime control of the production process on the actual day has to deal with different obstacles. The main cause of these obstacles is the uncertainty that comes along with the predicted input of the planning process. This uncertainty can be found in the realtime behaviour of predicted parameters, such as the demand and supply of individual appliances or households, but also in predicted parameters as the prices of the electricity market. Whereas the latter type of predicted parameters may have financial consequences, the first type of uncertainty can initiate a snowball effect and can eventually lead to large deviations from the planning (e.g. causing difficulties in the distribution/transmission grid) and can have economical/electrical consequences (blackouts).

Price uncertainty occurs for example at day ahead markets, which are analyzed in more detail in Chapter 5. This uncertainty differs from demand uncertainty (of heat, in our case) in the sense that price uncertainty reveals itself on beforehand when the day ahead prices are settled during the clearing of the market. This allows

the operator of a Virtual Power Plant to consider the influences of this uncertainty (i.e. the outcome of the bidding process of the electricity market) and anticipate to this by a renewed execution of the planning method. Uncertainty of heat demand is revealed in an online setting, meaning that only during a certain time interval the exact heat demand of this time interval is known.

In general there are two possibilities to cope with this demand uncertainty in the transition from a planning to a practical realization. As a first option, to relieve the realtime control, stochastic influences could be incorporated in the planning step already. This can be done e.g. by using probabilistic constraints or by designing (demand) scenario trees that take demand uncertainty into consideration. Scenario trees are most common in stochastic unit commitment approaches [40, 41, 42, 60, 105, 116, 119]. On the other hand, we could also deal with demand uncertainty in the realtime control step. In this case, the planning serves as a guideline, which the realtime control has to follow as close as possible.

We choose for this second option by using a combination of prediction/planning and realtime control that accounts for demand uncertainty. This choice is accompanied by the nature of the demand; we study large amounts of appliances with individual demands, each with its own uncertainty, instead of centralized demand. In this case scenario trees are not helpful. Although this choice shifts the responsibility for coping with demand uncertainty to the realtime control step, the planning step can aid in the sense that a heat buffer reservation can be made for capturing (part of) the demand uncertainty.

The focus of this thesis is on planning problems. However, in this chapter we give a short overview of the other two steps in the TRIANA methodology, for sake of completeness. Hereby the focus is on results related to the microCHP use case. The quality of the prediction step is crucial for the extent to which realtime control needs to be able to cope with demand uncertainty. Therefore, we focus on the quality of the prediction in Section 4.2 and on the ability to cope with realtime demand uncertainty in Section 4.3. The reservation of heat capacity in Section 4.4 shows the possibilities that a discrete planning can offer to deal with realtime demand uncertainty.

4.1 REALTIME CONTROL BASED ON PLANNING AND PREDICTION

The TRIANA 3-step control methodology for Smart Grids introduced in Chapter 2 consists of three major steps in which a distributed energy infrastructure is optimized and controlled. As a first step a prediction is needed for the demand of different types of energy consumption/production up to a very small scale (i.e. at a household scale). This prediction serves as basic input for the planning step, which is the second step in the control methodology. In this step the possibilities for production, storage and consumption are optimized, for example towards the objectives that are presented in Chapter 3. The prediction and planning steps are executed in advance; in general one day before the actual demand/supply takes place, a prediction and a planning is made. The third step of the TRIANA methodol-

ogy is to manage the actual demand/supply online: decisions are required for each appliance at a given time interval, for that same time interval. When the prediction is perfect, then the appliances can be realtime controlled by simply following the planning outcome. However, when the prediction is not perfect, the planned operation cannot always be followed. Therefore a realtime control method is needed, that reacts to this realtime deviation from the predicted values. Of course, this realtime control wants to stick to the planning as close as possible. In Section 4.2 we analyze the implications of the quality of the heat prediction (demand uncertainty) to the amount of flexibility that we want to have in the microCHP planning problem (the rest capacity of the heat buffer that is not available for the planning problem). Section 4.3 summarizes results obtained by the realtime control step for the microCHP use case. Additionally, an evaluation of necessary heat capacity reservations to compensate for demand uncertainty is given in Section 4.4.

4.2 PREDICTION

Prediction of local (household) electricity demand is done in different ways (e.g. [31, 120]). For a prediction of the heat demand in households, which is most interesting for the microCHP case, we give a short overview of the work of [29]. This prediction is done by a neural network approach. Important input parameters are the heat demand data of one upto several days before the regarded day, predicted windspeed information for the regarded day and the day before, and outside temperature information for the regarded day and the day before. Continuous relearning in a sliding window approach shows good results.

The quality of the heat demand prediction can be measured by calculating the Mean Absolute Percentage Error (MAPE) and the Mean Percentage Error (MPE). The MAPE is defined as follows:

$$MAPE = \frac{1}{24} \sum_{j=1}^{24} \frac{|H_j^{pred} - H_j^{actual}|}{F_j} \quad (4.1)$$

$$F_j = \begin{cases} H_j^{actual} & \text{if } H_j^{actual} \neq 0 \\ \frac{1}{24} \sum_{k=1}^{24} H_k^{actual} & \text{otherwise.} \end{cases} \quad (4.2)$$

The MPE is defined as follows:

$$MPE = \frac{1}{24} \sum_{j=1}^{24} \frac{H_j^{pred} - H_j^{actual}}{F_j} \quad (4.3)$$

$$F_j = \begin{cases} H_j^{actual} & \text{if } H_j^{actual} \neq 0 \\ \frac{1}{24} \sum_{k=1}^{24} H_k^{actual} & \text{otherwise.} \end{cases} \quad (4.4)$$

The quality of the prediction now is characterized by the total error E_{total} , which is defined as:

$$E_{total} = MAPE + |MPE|. \quad (4.5)$$

house	Sunday		Monday		Tuesday		Wednesday	
	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE
1	0.43	-0.17	0.66	-0.09	0.61	0.06	0.64	-0.03
2	0.85	0.30	0.69	0.00	0.97	0.51	0.46	-0.12
3	0.20	0.01	0.16	-0.01	0.16	-0.08	0.23	-0.06
4	0.39	-0.14	0.31	-0.03	0.50	-0.02	0.50	-0.08
house	Thursday		Friday		Saturday			
	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE
1	0.71	-0.01	0.53	-0.16	0.47	-0.10		
2	0.68	0.03	0.47	-0.19	0.86	0.35		
3	0.19	-0.06	0.21	-0.06	0.22	0.01		
4	0.40	-0.22	0.50	-0.04	0.45	-0.05		

Table 4.1: The results for MAPE and MPE using Simulated Annealing for different weekdays for 4 houses

In the training process of the neural network, the mean squared errors are minimized. By using a validation set, the evaluation of the training is measured by calculating the sum of MAPE and absolute MPE. Different selections of input parameters for the neural network are searched in a Simulated Annealing framework. The best results using this framework for some real world data are presented in Table 4.1. The prediction has an average MAPE of 0.48, which is the average of the MAPE values corresponding to 4 houses and 7 different types of weekdays of the table. Likewise, the average MPE is -0.02 . If we consider individual hours, we mispredict the average hourly heat demand by almost 50%, but if we look at the overall prediction for a complete day, we are almost correct. This prediction is not ideal, but of a quality which may be sufficient for the planning process, since often the error is that a peak in demand is predicted in a time interval next to the real peak; i.e. mainly variations within small time differences occur.

As we can see from the average MPE (which is around 0) as a measure for the selected input data for prediction, this selection has a tendency to underestimate the prediction, since the actual heat demand values are used as a denominator in the MPE calculation.

4.3 REALTIME CONTROL

The basis for the realtime control consists of the energy model that has been presented in Chapter 2. This model gives a flow problem formulation of different types of energy for a single time interval, in which balance plays a crucial role. Balance within this energy flow model guarantees a match between supply and demand. Also it resembles strong similarities with the way the grid infrastructure is organized, which makes it easy to incorporate these network constraints. However, balance can often be realized in many different ways, since there exist many elements for which a decision has to be made.

To select the best option that balances the model, a decision problem is solved for the given time interval. This problem has an objective that differs from the earlier presented objectives of the microCHP planning problem, by minimizing artificial total costs that are derived from certain cost functions for each element in the flow

model. This change in objective has two reasons. First, the balancing problem has to be solved in realtime, which poses a stronger time limit on the realtime control than on the planning process. Therefore we want to focus on balancing constraints only, and relax other possible constraints by using cost functions. Secondly, the energy flow model is not exclusively aimed at incorporating the planned operation of a fleet of microCHPs, but also at the possible inclusion of different types of generation, storage and consumption. To include and combine these - possibly conflicting - objectives, the balancing problem uses generalized cost functions for the elements for which a decision has to be made. These cost functions consist of artificial costs against consumed, produced or stored amounts of energy. Of course the cost functions can depend on the outcome of the planning, which means that the cost function can vary over time. The objective of the balancing problem is to minimize these artificial costs. It is of importance in the determination of these cost functions that infeasible state changes for the different elements are penalized in the cost function by large artificial costs, such that infeasibility is prevented (unless it is impossible to find a balance without including such a high cost).

The optimization problem of minimizing the artificial costs, while balancing the energy flow model can be summarized as follows:

Minimizing the costs of balancing the energy flow model

INSTANCE: Given is an energy flow model, consisting of a graph $G = (V, A)$, where $V = E \cup P$ ($E \cap P = \emptyset$) and cost functions $f_e(x_e)$ associated with decisions x_e for elements $e \in E$, whereby x_e results in flows a_{ep} to pools $p \in P$ and in flows a_{pe} from pools $p \in P$ that are in balance for element e .

OBJECTIVE: Minimize the total cost functions $\sum_{e \in E} f_e(x_e)$, such that balance is preserved for all pools $p \in P$: $\sum_{e \in E} (a_{ep} - a_{pe}) = 0$.

Model predictive control

In different time intervals the cost functions for the same appliance can vary. This gives the possibility to take planning decisions into account and to follow these planning decisions as good as possible. However, this process focuses only on the current interval. It may be worthwhile to anticipate on future time intervals as well, since this may prevent the realtime controller to take relatively good decisions for the current interval, that lead to severe problems in later time intervals. That means that it may be a good idea for the realtime control method to deviate from the planning at a current time interval, although it is currently possible to follow it, in order to be more close to the planning in later time intervals, compared to when the current planning would have been followed.

This setting of looking ahead in time is called model predictive control (MPC). It simply consists of minimizing the total costs of sequential balancing problems,

and using only the results for the decision variables in the first considered time interval to perform realtime operation.

Several use cases in [94] show that the realtime control step is able to follow a planning upto a large extent. Furthermore, it is shown that the addition of MPC can result in an improved ability to cope with realtime fluctuations.

4.4 EVALUATION OF HEAT CAPACITY RESERVATION

Although the realtime control method shows that we can follow a plan for a VPP of microCHPs quite closely, even if we do not reserve free space in the heat buffer, we are interested in the amount of heat buffer space that we would have to reserve, to be fully able to compensate for demand uncertainty by perfectly following the planning. This means that we want to achieve a feasible operation in realtime, while sticking exactly to the planned operation. We perform this way of evaluating heat buffer reservation in the presence of demand uncertainty on two of the tightest medium instances, being problem instance $I(100, 14)$ for 24 and for 48 intervals.

We use the same heat demand profiles and heat buffers of these instances as we defined them in the previous chapter. For the heat buffers, the planning uses a capacity of 10 kWh. The demand profiles now represent the predicted heat demand H_{pred}^i , for the different houses $i \in I$. We introduce demand uncertainty to these predicted heat demand profiles. This is done by applying a normally distributed deviation u_j^i to the different hours j of the predicted heat demand of house i with mean μ and standard deviation σ . These i.i.d. variables $u_j^i \in \mathcal{N}(\mu, \sigma)$ are added to the predicted heat demand to create real heat demand H_{real}^i artificially: $H_{real,j}^i = H_{pred,j}^i + u_j^i$. The parameters μ and σ are chosen such that the average MAPE and MPE of Table 4.1 are approximated. For different possible choices of $\sigma \in \{0, 200, 400, 600, 800, 1000\}$ Wh the corresponding choices for μ are found, such that values for MAPE and MPE are calculated that are the closest to the ones in Table 4.1. Note that these additional uncertainties are skewed in the sense that $\mu > 0$ if $\sigma > 0$, due to the underestimation of the prediction. To see the influence of an unskewed prediction that does not underestimate, we also apply normal distributions with $\mu = 0$ and the found values for σ in the MAPE/MPE approximation.

Figures 4.1 and 4.2 show the maximum excess, maximum slack and the total maximal necessary reserve capacity of all 100 heat buffers. The maximum excess ME (in kWh) is the largest excess (i.e. the overproduction that does not fit in the heat buffer) that occurs in all 100 houses. The maximum slack MS (in kWh) is the largest amount of heat demand that cannot be supplied, since there is too few production of heat, over all 100 houses. The maximum reserve capacity R (in kWh) is the sum of the maximum slack and the maximum excess (note that this can be larger than the maximum of the sum of slack and excess for all houses). If this reserve capacity R is applied to the heat buffers of the houses, such that the total capacity of the heat buffers equals $R + 10$ kWh, we can plan the operation of the heat buffer in the range $[MS, MS + 10]$, thereby not violating any form of real heat

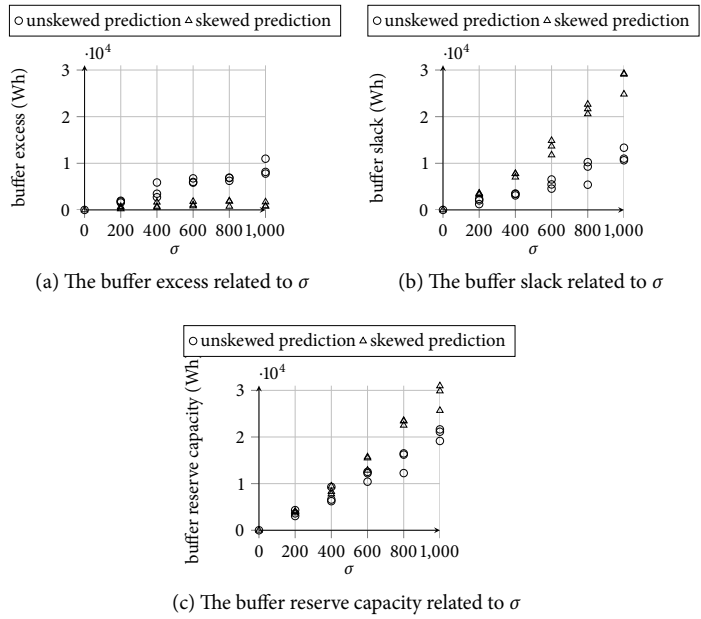


Figure 4.2: The necessary buffer reserve capacity for different values of MAPE and MPE for a planning using 48 intervals (hourly prediction!)

4.5 CONCLUSION

This chapter positions the microCHP planning problem in the TRIANA control methodology. It further explains the interaction between prediction, planning and realtime control. This interaction is necessary when we want to cope with demand uncertainty. TRIANA uses a model predictive control approach to cope with this uncertainty of the predicted demand, which performs well. Furthermore, we sketch the measures that a planner can take in reserving parts of the heat buffer capacity, such that the impact of demand uncertainty on the realtime control step is diminished.

AUCTION STRATEGIES FOR THE DAY AHEAD ELECTRICITY MARKET

ABSTRACT – This chapter discusses bidding strategies for a Virtual Power Plant that wants to operate on an electricity market. We distinct between two auction mechanisms: uniform pricing and pricing as bid. For both mechanisms bidding vectors are proposed that the VPP can offer to the market, such that the resulting quantity of the outcome of the auction is very close to the planned operation of the VPP, and such that the expected profit is maximized. For uniform pricing we propose a simple optimal strategy. In the case of pricing as bid we prove a lower bound on the expected profit that depends on the probability density function of the market clearing price.

A solution to the microCHP planning problem (treated in Chapter 3) consists of a planning of the operation of individual microCHP appliances. The aggregated electricity output of the planned operation of a group of microCHPs is of importance in the concept of a Virtual Power Plant. Such a Virtual Power Plant eventually wants to act on a (virtual) electricity market. In this chapter we treat the day ahead electricity market, since this market suits the Virtual Power Plant well, due to the short term notice on which heat demand predictions are made and due to the relative strict capacity requirements, which makes the Virtual Power Plant less suitable to act on a balancing market. We assume that a solution to the microCHP planning problem is available on beforehand; i.e. a distribution of total electricity generation over the time horizon of 24 hours is known, at the time the operator of the Virtual Power Plant starts acting on the day ahead market. Of course for this solution predictions of the prices on this electricity market may have been used as input and, thus, the predictions may have influenced the planning. The job of the operator of the VPP is to sell the planned production to the electricity market; this

is done by offering supply bids. This job is made difficult by the uncertainty of the actual market clearing prices for the different hours. This uncertainty should be accounted for in the supply bids in such a way that the probability of being allowed to supply is large. Namely, the planned electricity quantities need to be generated to a large extent anyhow, due to the local heat demand constraints of the individual households. If these quantities are not settled on the day ahead market, they will be accounted for and penalized on the balancing market.

In this chapter we concentrate on bidding strategies for different auction mechanisms for the day ahead balancing market. These bidding strategies take into account that we want a large probability of being allowed to supply the planned amount in each hour, and such that we maximize the price that we receive for these quantities. The auction mechanisms that we study are uniform pricing and pricing as bid.

Section 5.1 describes the general background of auction mechanisms on a day ahead electricity market. The specific requirements for the VPP to act on this market are discussed in Section 5.2. Next, bidding strategies for the auction mechanisms uniform pricing and pricing as bid are studied, where the focus is on quantity and price of the outcome of the auction; the quantity should be close to a desired amount and may have only a small variation, and the expected revenue (profit) is optimized for normally distributed market clearing prices. Section 5.3 shows the mechanism of uniform pricing and Section 5.4 the mechanism of pricing as bid. Finally conclusions are drawn on how to act on the electricity market in Section 5.5.

5.1 AUCTION MECHANISMS ON THE DAY AHEAD ELECTRICITY MARKET

Electricity trading is subject to similar market principles as other economic activities. In an electricity market demand and supply of electricity meet; based on demand curves, which show the price that the consumption side is willing to pay related to the quantity of traded electricity, and supply curves, which show the price that the generation side wants to receive related to the quantity of traded electricity, an electricity market price is settled. This market clearing price is found at the intersection of the aggregated demand curve and the aggregated supply curve. Figure 5.1 shows an example of supply and demand curves. In Figure 5.1a two supply curves for two different generators are plotted. The aggregated supply curve of these two generators is depicted in Figure 5.1c. However, in the practice of an electricity market a supply or demand curve consists of a limited number of price/quantity tuples (p, q) , that define stepwise supply or demand functions. Figure 5.1b and 5.1d show the same supply and demand curves of Figure 5.1a and 5.1c, but now they are stepwise approximated.

In electricity markets, the price elasticity of electricity demand is really low [88]. This results in steep demand curves, for which an example is shown in Figure 5.1c and 5.1d.

In the day ahead electricity market supply and demand curves are offered for 24 hours, resulting in 24 independent auctions. For each hour, each supplier and each

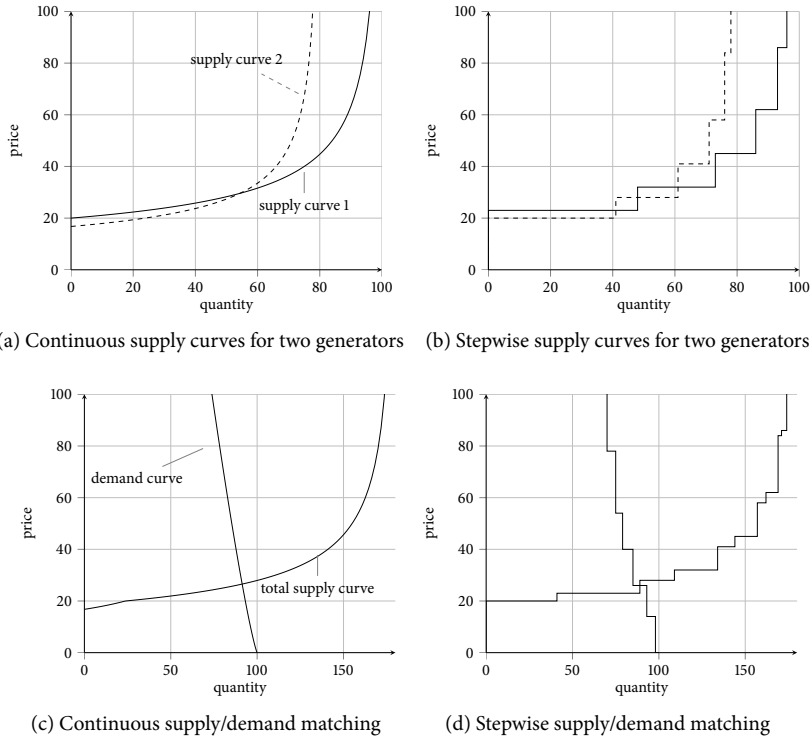


Figure 5.1: An example of supply/demand curves

retailer/consumer offers a vector of price and quantity tuples $\{(p_1, q_1), (p_2, q_2), \dots, (p_T, q_T)\}$; this is called a stepwise bid curve for a supplier/retailer. These prices and quantities have the properties that $p_1 < p_2 < \dots < p_T$ and $q_1 < q_2 < \dots < q_T$. These vectors are aggregated for both the supply and the demand side, and the market clearing price is found where both stepwise functions meet. For all suppliers which have at least one bid (p_k, q_k) in their bid curve for which p_k is below the market clearing price p , the bid with the price closest to p but lower than or equal to p is accepted. The generation of electricity is never partially dispatched, meaning that the cleared quantity is either 0 or equal to a certain quantity q_t that belongs to a bid in the set of bids $\{(p_1, q_1), (p_2, q_2), \dots, (p_T, q_T)\}$. In a practical situation the operator of any kind of power supply has to construct its bids in such a way that the operation of its assets is optimized. Usually this means that optimal bids are constructed based on the relationship between expected revenue and costs (see e.g. [18, 98, 132]). In the case of our Virtual Power Plant however, we do not consider costs, since operational fuel costs are the responsibility of the households in our business case. This lack of costs has its implications on the construction of bidding strategies. These strategies need to be applied to the 24 individual auctions for the

day. However, these auctions are simultaneously settled. This implies that we cannot base our strategy for the auction of hour s on the outcome of the auctions for the hours preceding s .

5.2 A VIRTUAL POWER PLANT ACTING ON A DAY AHEAD ELECTRICITY MARKET

A Virtual Power Plant can operate on an electricity market with its available capacity. Due to the daily nature of the underlying generation techniques in the VPP we concentrate on short term markets. The day ahead market suits the VPP well, since it expects a supplier to simultaneously place bid vectors for each hour for a complete day, which coincides with the aggregated generation for a complete day that has been planned. A balancing market (intraday market/spot market/strip market) is less suitable, since the online setting of this market offers too much risk for a supplier which has an almost fixed amount of supply at each moment in time.

In the following we shortly sketch how the planning problem can be used in a framework to find a desired total planned output that we want to auction. This framework is based on the insight that we derive from the lower bound calculation in Chapter 3. This insight can help us in the practical situation of the Virtual Power Plant, in which we would like to act on an electricity market. In this case we want to know that we can guarantee that a desired aggregated pattern can be reached by the individual generators. In an exploratory phase, a sketch of the aggregated output can be found, using the lower bound calculation as a guideline. The actual planning of the individual microCHPs can be postponed until a rough sketch is found that satisfies the (profit maximization) objective of the owner of the Virtual Power Plant, and that has a promising lower bound. Using this framework, we can find a detailed planning for all individual houses, that results in a total output that is desirable to auction on the market, and we can do this by saving a lot of computational effort.

Returning to the actions on the electricity market, we model price uncertainty as follows. In the related planning problems, the operation of the VPP on the day ahead market is indicated by using a (direct or indirect) objective value that maximizes the profit on this electricity market. However, the prices to which the planning problem optimizes are predictions of the electricity prices of the upcoming day. These predictions are subject to uncertainty. This uncertainty can be expressed by a probability density function $f(p)$ for the market clearing price. Yet these variable prices have some interesting properties that can be used to develop a way of placing bid vectors on the market.

The calculated planning is executed based on the expectation of the price, and is thus based on $f(p)$. This planning gives quantities for each hour that have to be sold. As this selling is the primary goal of the VPP that acts on the day ahead market, we have to guarantee almost for sure that the VPP gets the possibility to sell its generated electricity (we call this winning the auction). Based on experience we ask in our setup that 99% of the auctions should be won. In addition to that, we cannot deviate too much from the planned quantities. This means that the quantities in the

submitted bid vectors all have to be real close to the desired amounts. The received price is of secondary importance, but is still optimized to be as large as possible. To spread risks, in general the differences in prices belonging to different bids in the bid vector should be relatively large and depend on the density function $f(p)$.

In the situation of our Virtual Power Plant we deal with very specific limitations on the amount of electricity that can be offered. These limitations are obvious when a planning has already been made. In this case it is important to develop a strategy that respects the outcome of the planning process. To stay close to this outcome, the quantities q that can be offered to the market are limited. Therefore one goal of the bid construction focuses on guarantees on the amount of electricity that is cleared, i.e. the auctioned quantity is close to the amount of electricity that we want to sell. If the placement of bids is executed securely, we also want to maximize the price for which the previously mentioned quantity is sold. The second goal of the bidding process namely is to maximize the expected revenue for the Virtual Power Plant.

The output of the bidding process gives a definitive division of the available capacity over the time horizon, which might request for a renewed planning. If a new planning is not possible or infeasible, we want this assignment to correspond to the planned assignment, such that only small additional bids need to be offered to a balancing market. Large necessary adjustments are namely undesirable, since we may assume that the prices on the balancing market are not beneficial for a market player which has to offer almost fixed amounts of electricity to this realtime market. Ultimately, large adjustments might even not be tradeable on a balancing market and lead to a situation where the VPP cannot be operated properly. Therefore we want to prevent these large deviations occurring by designing auction strategies.

In the following we concentrate on electricity markets that are similar to the short term day ahead market as applied in The Netherlands [2]. Hourly market bids (p_t^s, q_t^s) for hour s on this type of day ahead electricity market are cleared simultaneously (independently) for a complete day (i.e. for 24 hours individual prices are settled for which electricity is traded). This implies that a bidding strategy for hour s cannot be adapted based on the outcome of the market clearing of hour $s - 1$. The descriptor $t \in \{1, \dots, T\}$ is used to distinguish between T different bids for the same hour (and for the same supplier).

5.2.1 THE BID VECTOR

Different bids for the same hour are required to vary in quantity, since the quantity is the cumulative quantity for a single supplier; without loss of generality we require quantities q_t^s to be strictly increasing. A minimum difference in quantity is set to 0.1 MWh for the APX Power NL day ahead auction [2] for hour s :

$$q_{t+1}^s \geq q_t^s + 0.1. \quad (5.1)$$

Next to this necessary constraint, we additionally require that the prices of different bids in the same hour are different, since we want each bid to be meaningful. Namely, the occurrence of two bids for the same price and different quantities makes the

bid with the lowest quantity redundant. Moreover we require that prices p_i^s are also strictly increasing:

$$p_{t+1}^s > p_t^s. \tag{5.2}$$

The combination of increasing prices accompanied by increasing quantities can be simply explained by the natural desire of a supplier to offer at least the same amount when prices increase. Furthermore, for the prices on the day ahead market an interval $[-p_{\max}, p_{\max}]$ is given.

An example of a set of bids for one hour is depicted in Figure 5.2. It shows




Figure 5.2: A price/supply curve for one hour on the day ahead market

a price/supply curve for one supplier. In this case six bids form a step function. Similar step functions are given for other (supplying and demanding) actors on the day ahead market. Based on these functions the auction is cleared. In its simplest form, aggregate supply and demand curves are formed and the intersection of both curves gives a market clearing price p , as indicated in the previous section. On the day ahead market all products are settled against this price p ; the largest bids that have a price below or equal to p have won the auction. If the market clears at price $p \in [p_t^s, p_{t+1}^s)$ then the supplier has to deliver q_t^s . In the example of Figure 5.2 this is shown for a market clearing price of $p = 37$; the corresponding quantity (0.4) can be easily read from the step function. Note that the market clearing price is ‘only’ the settlement price, from which all quantities for the different market participants can be deduced. This price is not necessarily equal to the price that each participant receives for its settled quantity. In a *uniform pricing* mechanism the settlement price equals the participants price, whereas in a *pricing as bid* mechanism the participants price could be lower than the settlement price.

5.2.2 PRICE TAKING

We assume that the operator of the Virtual Power Plant is a price taker in the sense that its influence on the (oligopolistic) market is negligible and that the constructed bids do not affect the market clearing price. This assumption is reasonable for small

sized VPPs; e.g. a cooperation consisting of 100000 microCHPs reaches a maximal market share during winter of about 1%. For such a price taker the market clearing price is considered as given. The supplier has no influence on this price, so the density function $f(p)$ remains unchanged by the actions of the supplier.

5.2.3 QUANTITY OUTCOME OF THE AUCTION

The operator of the VPP wants to settle a quantity that is close to its desired quantity. In the following we introduce a quantity interval $[Q, Q_{\max}]$ that resembles this closeness. For sake of simplicity, we may refer to Q as the desired quantity, although this value actually might be a little bit below the desired quantity. To this end we define two requirements:

- any positive outcome of the auction (winning the auction) should have a quantity that is larger than Q and close to Q ;
- the probability of having a positive outcome should be larger than a given value β .

Let the interval $[Q, Q_{\max}]$ define the domain from which the bid quantities may be chosen. To obtain the closeness requirement of quantities that result from the auction, we demand that Q_{\max} is maximally 10% larger than Q : $\frac{Q_{\max}}{Q} \leq 1.1$. We use a limited amount of T_{\max} bids, which implies that $Q_{\max} - Q \geq 0.1(T_{\max} - 1)$. We choose $Q_{\max} = Q + 0.1(T_{\max} - 1)$ such that we have the smallest possible domain. Later on we will use the following inequality for the minimum quantity of Q :

$$\begin{aligned} \frac{Q_{\max}}{Q} &\leq 1.1 \\ \Rightarrow \frac{0.1(T_{\max} - 1)}{Q} &= \frac{Q_{\max} - Q}{Q} = \frac{Q_{\max}}{Q} - 1 \leq 0.1 \\ \Rightarrow Q &\geq \frac{0.1(T_{\max} - 1)}{0.1} = T_{\max} - 1. \end{aligned} \quad (5.3)$$

In case we lose the auction we are of course not close to the desired quantity Q . To prevent the occurrence of this event we propose a probability β , which value represents the chance of winning the auction. The probability of winning the auction should be larger than this value. This is defined by the following equation:

$$\int_{p_1^i}^{p_{\max}} f(p) dp \geq \beta. \quad (5.4)$$

This means that we have an additional restriction for the price of the first (and lowest) bid.

5.2.4 MARKET CLEARING PRICE DISTRIBUTION

We focus on taking part in the day ahead electricity market of The Netherlands. For this APX day ahead market, we collected data from November 22, 2006 until

November 9, 2010. The average price is 48.87 €/MWh for the complete time horizon, with a minimum daily average of 14.83 €/MWh and a maximum of 277.41 €/MWh. In general no real trend in the development of the electricity prices can be found, other than that prices stabilize after a temporary peak in 2008. For short term time periods the prices remain relatively stable and are highly correlated to the prices of previous days. This led to the assumption that the prices on a short term history might follow a normal distribution with mean μ_s and standard deviation σ_s based on the prices of the previous days. Upto a history of 35 days, this assumption has been tested on the collected data, whereby our assumption was validated.

To show this behaviour in a graphical way, Figure 5.3 shows the acceptance rate of single bids whose hourly price is based on the hourly price of the previous day. This acceptance rate is determined by comparing the market clearing price for each hour s with the market clearing price of the previous day for the same hour, and using the latter price, multiplied with a percentage, as a bid for the current day. The figure shows the percentage of accepted bids to the percentage of the previous price, for all 24 hours. This figure resembles a cumulative normal distribution.




Figure 5.3: The acceptance rate of single bids whose hourly price is based on the hourly price of the previous day

5.3 BIDDING STRATEGIES FOR UNIFORM PRICING

We are interested in the expected profit that the supplier can make and in guarantees on the resulting quantity that has to be supplied. These guarantees on the resulting quantity are mentioned in the previous subsection, which means that we can now focus on the price forming. In most markets the price that each market participant receives equals the market clearing price. An often used mechanism for which this

holds is called uniform pricing. In uniform pricing all participants receive the same price for their settled quantities; this price is the market clearing price.

In the literature on uniform pricing mechanisms the construction of a set of bids is related to the usual costs associated with production. Constructed bids equal the marginal cost function in the situation of perfect competition. [98] shows that the bid construction still follows the marginal cost function when the problem is restricted to a limited numbers of bids. We also deal with a situation in which the number of different bids is limited, due to our closeness requirement and possibly due to rules that are determined by the organizer of the auction. In our situation however, we do not consider cost functions, but we further restrict the form of the bids by focusing on bounded quantities $q_t^s \in [Q, Q + 0.1(T_{\max} - 1)]$.

The lack of cost functions has its implications on the bid construction under the assumption of uniform pricing. The bid construction problem has the following form:

$$\max \sum_{t=1}^{T-1} \int_{p_t^s}^{p_{t+1}^s} p q_t^s f(p) dp + \int_{p_T^s}^{p_{\max}} p q_T^s f(p) dp \quad (5.5)$$

$$s.t. \int_{p_1^s}^{p_{\max}} f(p) dp \geq \beta \quad (5.6)$$

$$T \in \{1, \dots, T_{\max}\} \quad (5.7)$$

$$q_{t+1}^s \geq q_t^s + 0.1 \quad (5.8)$$

$$p_{t+1}^s > p_t^s \quad (5.9)$$

$$Q \leq q_t^s \leq Q + 0.1(T_{\max} - 1) \quad (5.10)$$

$$-p_{\max} \leq p_t^s \leq p_{\max}. \quad (5.11)$$

Equation (5.5) expresses that we want to maximize the expected profit by integrating the function $p q_t^s$ over the intervals $[p_t^s, p_{t+1}^s]$ for all bids (p_t^s, q_t^s) , $t = 1, \dots, T - 1$ and the function $p q_T^s$ over the interval $[p_T^s, p_{\max}]$ for the last bid (p_T^s, q_T^s) . The price winning constraint is given by (5.6) and we restrict to using at most T_{\max} bids (5.7). Equations (5.8) and (5.9) force that bids are strictly increasing and (5.10) and (5.11) that the limitations on quantity and price are followed.

The optimal auction strategy for uniform pricing is based on the fact that, for the integral $\int_{p_t^s}^{p_{t+1}^s} p q_t^s f(p) dp$, the price p that the participant receives is integrated over the corresponding interval, whereas the quantity q_t^s is fixed. This leads to the observation that, for $0 \leq a < b < c$, $k < l$ and $f(p)$ a positive function, we have:

$$\begin{aligned} \int_a^b p k f(p) dp + \int_b^c p l f(p) dp &= k \int_a^b p f(p) dp + l \int_b^c p f(p) dp \\ < l \int_a^b p f(p) dp + l \int_b^c p f(p) dp &= \int_a^c p l f(p) dp. \end{aligned} \quad (5.12)$$

Similarly, for $a < b < c \leq 0$, $k < l$ and $f(p)$ a positive function we observe:

$$\begin{aligned} \int_a^b pkf(p)dp + \int_b^c plf(p)dp &= k \int_a^b pf(p)dp + l \int_b^c pf(p)dp \\ &< k \int_a^b pf(p)dp + k \int_b^c pf(p)dp = \int_a^c pkf(p)dp. \end{aligned} \quad (5.13)$$

Equation (5.12) shows that it is not beneficial to have more than one bid with a positive price; (5.13) shows that it is not beneficial to have more than one bid with a negative price.

The optimal set of bids consists of one bid for positive prices, where the maximum amount Q_{\max} is offered, and possibly a second bid in case of negative prices, where the minimum amount Q is offered. The existence of one or two bids depends on the value of $p_1^s(\beta)$, where $p_1^s(\beta)$ results from an equality for the auction winning constraint (5.6): $\int_{p_1^s(\beta)}^{p_{\max}} f(p)dp = \beta$. If $p_1^s(\beta) \geq 0$, the optimal bid is $(p^*, q^*) = (0, Q_{\max})$. By applying this construction all positive contributions to the profit are maximally accounted for. If $p_1^s(\beta) < 0$, negative contributions (the influence of negative prices) should be minimized. In this case the optimal set of bids consists of two bids $(p_1^*, q_1^*) = (p_1^s(\beta), Q)$ and $(p_2^*, q_2^*) = (0, Q_{\max})$.

5.4 BIDDING STRATEGIES FOR PRICING AS BID

In the setting of uniform pricing there is no incentive for the operator of a VPP to submit a bid with a positive price due to (5.12). This situation changes when the auction mechanism would be pricing as bid. In this mechanism the VPP receives the price that it has bidden for the quantity that is settled. This changes (5.5) into:

$$\max \sum_{t=1}^{T-1} \int_{p_t^s}^{p_{t+1}^s} p_t^s q_t^s f(p) dp + \int_{p_T^s}^{p_{\max}} p_T^s q_T^s f(p) dp \quad (5.14)$$

Figure 5.4 presents the difference between the two auction mechanisms (uniform pricing and pricing as bid) in a graphical way. In this simple example we concentrate on a market clearing price with mean 50 and standard deviation 10. Figures 5.4a, 5.4b, 5.4c and 5.4d plot the corresponding functions $pf(p)$ and $p_t f(p)$ that are integrated on the domain $[0, 100]$, corresponding to the equations (5.5) and (5.14) respectively. Figures 5.4a, 5.4b, 5.4c and 5.4d show the effect of strategically bidding in the pricing as bid case. For different number of bids, that are uniformly distributed over the price domain (in case of 2 bids, we choose for $p_1 = 0$ and $p_2 = 50$, in case of 4 bids, we choose for $p_1 = 0$, $p_2 = 25$, $p_3 = 50$ and $p_4 = 75$, etcetera), we see the 'loss' in profit (the gray areas) diminish as the number of bids grows. In the following we derive lower bounds for the expected revenue for bid sets with different sizes T_{\max} . Namely, we want to use as few different bids as possible, since they determine the variability in the quantity outcome. We see that in such cases, the price setting does not follow a uniform distribution, as we used in this example. First we state some properties of the market clearing price distribution.

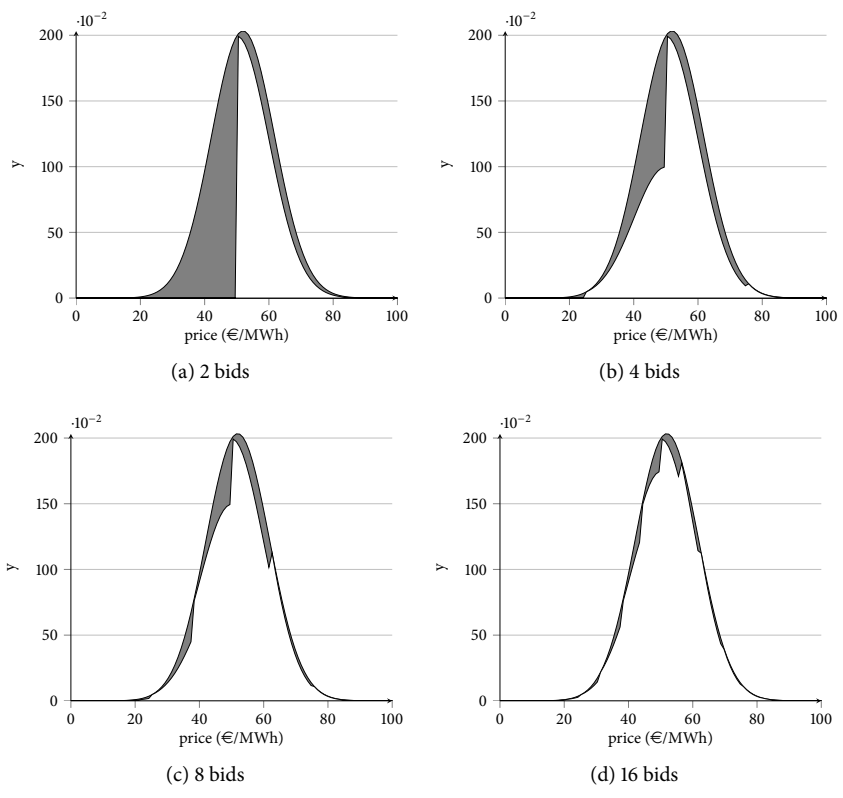


Figure 5.4: Graphical representations of the difference between uniform pricing and pricing by bid

5.4.1 NATURAL BEHAVIOUR OF THE MARKET CLEARING PRICE DISTRIBUTION

When we observe the short term history of the market clearing prices of individual hours, the clearing prices can be approximated by a normal distribution with mean μ_s and standard deviation σ_s . Negative prices are allowed on the day ahead market, but in practice they hardly occur (in our data it never occurred that prices became negative). However, the bid construction that we propose also deals with negative prices.

In the process of determining lower bounds for the expected revenue (profit) we have to evaluate the integral of the probability density function of a normal distribution $\int \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(p-\mu)^2}{2\sigma^2}} dp$, which cannot be evaluated by the use of elementary functions. For the cumulative distribution function $\Phi(x)$ of the standard normal

distribution $\mathcal{N}(0, 1)$ we use the approximation of [27]:

$$\Phi(x) = \begin{cases} 0.5 + 0.5\sqrt{1 - \frac{7e^{-\frac{x^2}{2}} + 16e^{-x^2(2-\sqrt{2})} + (7+\frac{\pi}{4}x^2)e^{-x^2}}{30}} & x \geq 0 \\ 0.5 - 0.5\sqrt{1 - \frac{7e^{-\frac{x^2}{2}} + 16e^{-x^2(2-\sqrt{2})} + (7+\frac{\pi}{4}x^2)e^{-x^2}}{30}} & x < 0. \end{cases} \quad (5.15)$$

We want to have a probability of winning the auction of $\beta = 0.99$. Note that for the approximation given in (5.15) we get $\Phi(-2.33) \leq 0.01$, which indicates that for a price $p_1^s \leq \mu_s - 2.33\sigma_s$ the auction winning equation is satisfied. From this description of p_1^s we deduce that for:

$$\frac{\mu_s}{\sigma_s} \geq 2.33, \quad (5.16)$$

a positive lowest bid price is possible, if we ask for a probability of winning of 99%. The relationship between μ_s and σ_s plays an important rule in the determination of the lower bound in the following section. The value 2.33 in (5.16) is used as an exemplary value in the proof and the discussion. Note that other values are also evaluated and negative prices are also discussed.

5.4.2 LOWER BOUNDS ON OPTIMIZING FOR PRICING AS BID

The goal of this subsection is to give an indication how good we can bid on the electricity market for the auction mechanism pricing as bid. The objective of acting on the day ahead market is to maximize the expected revenue. To achieve lower bounds on the expected revenue we use a special construction for the prices in the different bids that together form the offer. All prices can be written in the form $p_t^s := \mu_s + a_t\sigma_s$. By using this construction, the probability that a certain bid is accepted is independent from the choice for μ_s and σ_s and remains thus constant, since $\Phi\left(\frac{p_t^s - \mu_s}{\sigma_s}\right) = \Phi\left(\frac{\mu_s + a_t\sigma_s - \mu_s}{\sigma_s}\right) = \Phi(a_t)$. When we restrict the specific choices for the a_t values, we can consider corresponding lower bounds.

The lower bound depends on a minimum value for the quotient $\frac{\mu_s}{\sigma_s}$ and on the maximum amount of bids T_{\max} that is allowed in the bid construction. The resulting lower bound can be interpreted as the fraction between the expected revenue of the constructed bid and the revenue when the maximum amount Q_{\max} is sold for the average price μ_s . This revenue $Q_{\max}\mu_s$ is an approximation of the optimum expected profit for uniform pricing.

In the following we consider a very specific choice for a bidding strategy and prove a resulting lower bound for this case. We choose $\frac{\mu_s}{\sigma_s} \geq 2.33$, $T_{\max} = 4$, $a_1 = -2.33$, $a_2 = -1.20$, $a_3 = -0.39$ and $a_4 = 0.45$. For this we can prove the following result. The proof indicates how also for other cases a lower bound can be calculated.

Theorem 2 For a mean to standard deviation ratio $\frac{\mu_s}{\sigma_s} \geq 2.33$ and a maximum number of bids in an offer $T_{\max} = 4$ the quadruple offer $\{(\mu_s - 2.33\sigma_s, Q_{\max} - 0.3), (\mu_s - 1.20\sigma_s, Q_{\max} - 0.2), (\mu_s - 0.39\sigma_s, Q_{\max} - 0.1), (\mu_s + 0.45\sigma_s, Q_{\max})\}$ is at least 0.740 times the optimal expected profit of uniform pricing $Q_{\max}\mu_s$.

Proof We use two relations in this proof. The first one follows from $\frac{\mu_s}{\sigma_s} \geq 2.33$:

$$-\frac{\mu_s}{2.33} \leq -\sigma_s.$$

The second one results from (5.3) using $T_{\max} = 4$:

$$\begin{aligned} Q &\geq T_{\max} - 1 = 3 \\ \Rightarrow Q_{\max} &= Q + 0.1(T_{\max} - 1) = Q + 0.3 \geq 3.3 \\ \Rightarrow -\frac{Q_{\max}}{3.3} &\leq -1. \end{aligned}$$

The expected revenue of uniform pricing is $Q_{\max}\mu_s$. The lower bound on the expected revenue of pricing as bid is a direct result from applying the above two relations to the bids $(\mu_s - 2.33\sigma_s, Q_{\max} - 0.3)$, $(\mu_s - 1.20\sigma_s, Q_{\max} - 0.2)$, $(\mu_s - 0.39\sigma_s, Q_{\max} - 0.1)$ and $(\mu_s + 0.45\sigma_s, Q_{\max})$:

$$\begin{aligned} &\int_{\mu_s - 2.33\sigma_s}^{\mu_s - 1.20\sigma_s} (\mu_s - 2.33\sigma_s)(Q_{\max} - 0.3)f(p)dp + \\ &\int_{\mu_s - 1.20\sigma_s}^{\mu_s - 0.39\sigma_s} (\mu_s - 1.20\sigma_s)(Q_{\max} - 0.2)f(p)dp + \\ &\int_{\mu_s - 0.39\sigma_s}^{\mu_s + 0.45\sigma_s} (\mu_s - 0.39\sigma_s)(Q_{\max} - 0.1)f(p)dp + \\ &\int_{\mu_s + 0.45\sigma_s}^{p_{\max}} (\mu_s + 0.45\sigma_s)Q_{\max}f(p)dp \\ &= (\mu_s - 2.33\sigma_s)(Q_{\max} - 0.3)(\Phi(-1.20) - \Phi(-2.33)) + \\ &\quad (\mu_s - 1.20\sigma_s)(Q_{\max} - 0.2)(\Phi(-0.39) - \Phi(-1.20)) + \\ &\quad (\mu_s - 0.39\sigma_s)(Q_{\max} - 0.1)(\Phi(0.45) - \Phi(-0.39)) + \\ &\quad (\mu_s + 0.45\sigma_s)Q_{\max}(1 - \Phi(0.45)) \\ &= Q_{\max}\mu_s(1 - \Phi(-2.33)) + Q_{\max}\sigma_s(-2.33(\Phi(-1.20) - \Phi(-2.33)) - \\ &\quad 1.20(\Phi(-0.39) - \Phi(-1.20)) - 0.39(\Phi(0.45) - \Phi(-0.39)) + \\ &\quad 0.45(1 - \Phi(0.45))) + \mu_s(-0.3(\Phi(-1.20) - \Phi(-2.33)) - \\ &\quad 0.2(\Phi(-0.39) - \Phi(-1.20)) - 0.1(\Phi(0.45) - \Phi(-0.39))) + \\ &\quad \sigma_s(-0.3 \cdot -2.33(\Phi(-1.20) - \Phi(-2.33)) - 0.2 \cdot -1.20(\Phi(-0.39) - \\ &\quad \Phi(-1.20)) - 0.1 \cdot -0.39(\Phi(0.45) - \Phi(-0.39))) \\ &= 0.990Q_{\max}\mu_s - 0.505Q_{\max}\sigma_s - 0.111\mu_s + 0.142\sigma_s \\ &\geq 0.990Q_{\max}\mu_s - \frac{0.505}{2.33}Q_{\max}\mu_s - 0.111\mu_s + 0.142\sigma_s \\ &\geq 0.773Q_{\max}\mu_s - 0.111\mu_s \geq 0.773Q_{\max}\mu_s - \frac{0.111}{3.3}Q_{\max}\mu_s \\ &= 0.740Q_{\max}\mu_s. \end{aligned}$$

■

		$T_{\max} = 1$	$T_{\max} = 2$	$T_{\max} = 3$	$T_{\max} = 4$	$T_{\max} = 5$
$T = 1$	a_1	-2.33	-2.33	-2.33	-2.33	-2.33
	LB	0	0	0	0	0
$T = 2$	a_1	-	-2.33	-2.33	-2.33	-2.33
	a_2	-	-0.53	-0.49	-0.48	-0.47
		-	0.516	0.530	0.534	0.536
	LB	-	-	-	-	-
$T = 3$	a_1	-	-	-2.33	-2.33	-2.33
	a_2	-	-	-0.96	-0.94	-0.93
	a_3	-	-	0.10	0.13	0.14
		-	-	0.667	0.677	0.683
	LB	-	-	-	-	-
$T = 4$	a_1	-	-	-	-2.33	-2.33
	a_2	-	-	-	-1.20	-1.18
	a_3	-	-	-	-0.39	-0.36
	a_4	-	-	-	0.45	0.48
		-	-	-	0.740	0.748
	LB	-	-	-	-	-
$T = 5$	a_1	-	-	-	-	-2.33
	a_2	-	-	-	-	-1.36
	a_3	-	-	-	-	-0.68
	a_4	-	-	-	-	-0.05
	a_5	-	-	-	-	0.68
		-	-	-	-	0.783
	LB	-	-	-	-	-

Table 5.1: Lower bounds for different values of T_{\max} and different numbers of bids

The above proof uses specific values for a_t . These values are not randomly chosen. Instead, the values for a_t are found by an extensive search over the parameters a_t , such that $a_1 < a_2 < a_3 < a_4$ and $a_t \in \{-2.43, -2.42, \dots, 2.32, 2.33\}$, and the lower bound coefficient is maximized. This set is chosen, since we did not expect that the coefficients would be such that negative prices would occur in the bidding strategy. However, we allowed a small possibility of negative prices, but, as we will see, these negative prices did not occur. Note that we also allow values for which (5.6) gives a strict inequality.

In a similar way as above we have derived also lower bounds for a few other choices of T_{\max} and T . The corresponding best results are given in Table 5.1. The table shows the results assuming $\frac{\mu_s}{\sigma_s} \geq 2.33$.



The lowest value for a_1 equals -2.33 in all cases. This is logical, since negative prices are unnecessary for this special case. We can observe from this table that the lower bound belonging to a fixed number of T bids increases when T_{\max} increases. This is due to the increased flexibility for the quantities q_t^s in the interval $[Q, Q_{\max}]$. Although this is an interesting result, in each case it is worth to use the full quantity domain, i.e. using $T = T_{\max}$ different bids. In general, the lower bound increases with increasing T_{\max} . Already with 5 bids we are close to 80% of the expected value. However, the accompanying quantity domain increases too. Therefore a good trade-off between lower bound and quantity domain is needed when we construct an actual bid set.

5.4.3 COMPUTATIONAL RESULTS


The results of Table 5.1 are valid for all fractions $\frac{\mu_s}{\sigma_s} \geq \gamma$ with $\gamma = 2.33$. In this section we evaluate the behaviour of the lower bound when γ varies. A small value of γ allows for a relative high standard deviation and a large value of γ allows for relative

small standard deviations. We expect that we find better bidding strategies when γ increases. For this evaluation we use $T = T_{\max}$ for the five different values of T_{\max} , such that the quantity domain is completely used.


Figure 5.5 shows the bid construction (the assignment of values for a_t) and the lower bound, for $0.01 \leq \gamma \leq 50$ with steps of 0.01. The bid coefficients are denoted on the left y -axis and the lower bound is plotted on the right y -axis.

(a) $T = 5$ (b) $T = 4$ Figure 5.5: The behaviour of a_t for different values of γ

For $T_{\max} = 5$, Figure 5.5a shows that for $\gamma \geq 5.13$ it is beneficial to use a price



(c) $T = 3$



(d) $T = 2$

Figure 5.5: The behaviour of a_t for different values of γ (continued)

$p_1^s < \mu_s - 2.33\sigma_s$. The point $\gamma = 5.13$ is called the switching point, since it means that from this point on the price winning equation is no longer of influence for the price setting in the bids. All bid prices coefficients a_t are non-increasing functions on γ . However, this does not necessarily mean that the accompanying prices are non-increasing, since the mean and standard deviation of the price can have different values. For $T_{\max} = 4$, a similar plot is given in Figure 5.5b. The switching point for a_1^s now is on $\gamma = 5.99$. This point is 7.60 for $T_{\max} = 3$, 11.48 for $T_{\max} = 2$ and 40.22




Figure 5.5: The behaviour of a_t for different values of γ (continued)


for $T_{\max} = 1$. In all cases, the corresponding bid prices are never negative for values of γ that are larger than or equal to the switching points. Negative prices may only occur due to the price winning equation, which forces the behaviour of a_t for all T bids before the switching point. Note that this behaviour is also visible for $a_t^s, t > 1$.

The lower bounds for the different values of γ and T_{\max} are combined in Figure 5.6. Figure 5.6a shows a surface plot of the lower bound. A contour plot of this figure is given in Figure 5.6b. The contour lines are plotted with steps of 0.05. Especially for small values of γ it is beneficial to choose large values for T_{\max} .


In Table 5.2 the above bid construction is applied to the data of the APX day ahead market. For the market clearing price prediction we assume a normal distribution, where μ_s and σ_s are based on the values of a number of days in the short term history.

In the top of the table the average, maximum and minimum fractions $\frac{\mu_s}{\sigma_s}$ are given for varying history lengths of 7, 14, 21, 28 and 35 days. This average is decreasing with increasing history length, showing that the variation is larger for larger time periods. The average and minimum $\frac{\mu_s}{\sigma_s}$ fall in the steep part of Figure 5.6, which shows that it is extremely important to choose a bid construction that follows from a value of γ that is close to $\frac{\mu_s}{\sigma_s}$.

Based on the found fraction $\frac{\mu_s}{\sigma_s}$ for the market clearing price prediction, bidding strategies are developed for each hour. In the bid strategies the highest value of γ is chosen, such that $\gamma \leq \frac{\mu_s}{\sigma_s}$. For the different history lengths and varying T_{\max} , the results of the auction mechanism pricing as bid are given in the table. The average price gives the average received price. This price is compared to the average market clearing price, which results in a certain percentage of the market price



(a) The lower bound depending on γ and T_{\max}



(b) A contour plot of the lower bound

Figure 5.6: The lower bound for different values of γ and T_{\max}

		history (# days)				
		7	14	21	28	35
	average $\frac{\mu_S}{\sigma_S}$	7.01	6.10	5.68	5.42	5.22
	max $\frac{\mu_S}{\sigma_S}$	124.64	63.24	38.64	32.26	24.39
	min $\frac{\mu_S}{\sigma_S}$	0.47	0.41	0.38	0.39	0.39
T = 5	average price	43.02	42.81	42.65	42.60	42.48
	% of market price	88.03	87.58	87.27	87.16	86.91
	average Q excess	0.27	0.27	0.26	0.26	0.26
T = 4	average price	41.90	41.62	41.48	41.37	41.18
	% of market price	85.72	85.16	84.87	84.65	84.26
	average Q excess	0.20	0.20	0.20	0.20	0.20
T = 3	average price	40.10	39.52	39.22	39.00	38.73
	% of market price	82.05	80.86	80.25	79.79	79.24
	average Q excess	0.14	0.14	0.14	0.14	0.14
T = 2	average price	36.03	35.53	35.17	34.92	34.57
	% of market price	73.72	72.69	71.96	71.45	70.73
	average Q excess	0.07	0.07	0.07	0.07	0.07
T = 1	average price	22.33	20.81	19.64	18.78	18.03
	% of market price	45.68	42.59	40.18	38.43	36.89
	average Q excess	0.00	0.00	0.00	0.00	0.00

Table 5.2: The different bid strategies applied to the data of the APX market

that is reached. The average excess denotes the average amount by which the quantity exceeds the value of Q . The average excess increases almost linearly with the number of bids T . This value can be used to fit the domain $[Q, Q_{\max}]$ to the desired production of the VPP in a practical situation. The average price increases sublinearly with T . In Figure 5.7 the percentage of the average price compared to the market clearing price is depicted for an extended set of history lengths. This

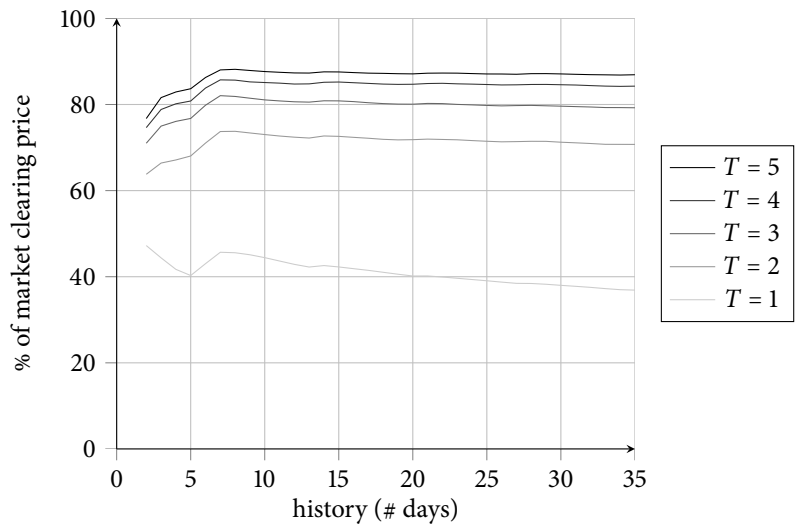


Figure 5.7: Evaluation of constructed bids for different history lengths




figure shows that a history of 7 or 8 days gives the best trade off between prediction accuracy and bid construction. Note that the average price is not completely equal to the average revenue, since the quantity for which the market is cleared is not given. However, we may assume that the behaviour is comparable, since the variation in quantity is limited.

5.5 CONCLUSION

This chapter shows methods to construct bids for two auction mechanisms on the day ahead electricity market. These methods are aimed to be used by a Virtual Power Plant as we describe it in Section 2.2.2. In comparison with existing approaches, our bid construction has the special form of having limited flexibility in the variation of the quantity-to-offer combined with the requirement of a very high minimum probability of winning the auction; bids are constructed in the absence of a cost function for the VPP.

For the auction mechanism uniform pricing, the bid construction is given by a unique bid for positive market prices and (possibly) an additional bid for negative prices, in case the probability of winning the auction cannot be satisfied with the first bid.

For the auction mechanism pricing as bid, the bid construction is given by successive bids (p_t, q_t) , for which the quantity q_t increases with the minimum required difference of 0.1 MWh and the price p_t is based on the predicted values for the market clearing price μ (mean price), σ (standard deviation of the price) and a coefficient a_t , such that $p_t = \mu + a_t \sigma$. The values of the different coefficients a_t are optimized for a given range of the fraction $\frac{\mu}{\sigma}$. Application of this form of bid construction to real world data shows that already 88% of the market clearing price can be reached as average settlement price, when at most 5 different bids are used.

THE GENERAL ENERGY PLANNING PROBLEM

ABSTRACT – This chapter treats the general energy planning problem as an extension of the Unit Commitment Problem. We add distributed generation, distributed storage and demand side management possibilities to this problem, thereby shifting the focus of this optimization problem towards the decentralization within the Smart Grid. The general energy planning problem differs from the UCP in size and in objective. We treat significantly more appliances and use a combination of objectives to include different types of generators and appliances. The general energy planning problem is solved using a hierarchical structure, in which the different elements are solved by using sub problems in levels. The general framework consists of creating patterns for single entities/appliances, combining patterns for such appliances on higher levels into so-called aggregated patterns, and using these aggregated patterns to solve a global planning problem. Two different case studies show the applicability of the method.

In Chapter 3 the microCHP planning problem has been introduced and treated. This problem gives a good example of the type of planning problems that arise in the field of distributed energy generation. It shows that for the combination of hard and weak constraints feasibility plays an important role in large scale, small sized generation: the planning of the operation of individual appliances cannot be neglected by aggregating groups of generators and only making a planning on this group level. A planning is necessary on the individual appliance level.

In the situation before the emergence of distributed energy generation, generators - even small sized ones (where small sized still means significantly larger than the kW level) - could be regarded as standalone entities in the portfolio of an energy

Parts of this chapter have been published in [MB:10].

supplier. Energy management then consisted of the problem of finding the optimal combination of the assignment of the available entities in the portfolio; i.e. solving the traditional Unit Commitment Problem (UCP). The large scale introduction of distributed energy generation, storage and load management asks for an extension to this Unit Commitment Problem. This extended problem is formulated as the general energy planning problem in this chapter. Due to the large differences in production capacity and the enormous amount of appliances (remind the practical intractability of the microCHP planning problem for instances with only a small amount of appliances) it seems unreasonable to attempt to solve this general energy planning problem to optimality when we treat all appliances simultaneously. Therefore we propose a leveled planning method, that plans the operation of generators, storage possibilities and consuming appliances in a hierarchical structure based on their location/size.

In general the technological developments in distributed generation, storage and demand side load management introduce more and more controllable entities that can be operated in different ways for given circumstances, which makes them suitable for use in a planning process. For instance, a microCHP in combination with a heat buffer is a controllable appliance, whereas a TV, although being controlled by the user, is an example of a non-controllable appliance in the context of the planning problem. The combination of microCHP and heat buffer allows for various operating patterns to supply a given heat demand for the time horizon of one day. A TV has exactly one completely determined electricity consumption pattern for a given user behaviour for the time horizon. This leaves no options for a planner, unless the user behaviour can be adapted, which is a situation that we do not desire. Therefore, in the general energy planning problem we focus on controllable appliances. The controllable appliances have a certain degree of freedom, which determines the flexibility with which these appliances can be used in the planning process. However, most of the considered appliances are less flexible than the generators in the UCP, which emphasizes the feasibility aspect of this extended problem: having only limited flexibility, global bounds on the total production need to be satisfied.

In this chapter we treat the combinatorial challenge of merging different distributed technologies in the energy supply chain with the already available elements in the existing infrastructure. We refer to this optimization problem as the general energy planning problem. In Section 6.1 we discuss the different application domains of the changing energy supply chain that each play an important role in the general energy planning problem. Then the problem is formulated in Section 6.2. A solution method that makes use of the available hierarchical structure in the energy supply chain is presented in Section 6.3. Section 6.4 shows a detailed study of exemplary case studies. Finally, conclusions are drawn in Section 6.5.

6.1 APPLICATION DOMAIN

In Chapter 1 various elements of the energy sector have been described that illustrate the partial decentralization of the energy supply chain. In the modelling process of the general energy planning problem we focus on the controllable decentralized elements of distributed production, distributed storage and demand side load management. We integrate these three types of elements in the existing framework for the conventional elements, which is controlled by an energy distribution management system that is based on the Unit Commitment Problem. So, the general energy planning problem is merely an extension of the UCP. Other elements described in Chapter 1 (e.g. photovoltaics (PV), windmills or the distribution and transmission grid itself) are not part of the central focus of this chapter. Solar panels and windmills for example are non-controllable in the context of the planning process, while the distribution and transmission grid is considered as a given infrastructure for the general energy planning problem. Although they are not part of the main design goal of the method which solves this problem, each of these elements may influence the objectives of the problem. In this way feedback may be given about design aspects of the electricity infrastructure, answering the question whether or not the capacity of the distribution and transmission grid suffices, or feedback about the allowable penetration rate of non-controllable electricity generation for a given infrastructure: how much solar panels or windmills can we allow, while still guaranteeing a reliable electricity supply?

The general energy planning problem combines different types of energy: heat, gas and electricity are examples of energy types that we have already seen in the microCHP planning problem. However, the driving factor of the objective is on electricity and its associated costs or revenues.

In the following subsections we describe application domains to which the basic Unit Commitment Problem may be extended. We show the flexibility that exists in each of the three types of decentralized elements. This flexibility has similarities with the flexibility in the microCHP planning problem; namely, the electrical outcome of the operation of local electricity consuming or producing appliances underlies a primal use of the corresponding appliances (e.g. the heat led operation of microCHPs). A microCHP is heat demand driven as is a heat pump, and a fridge or a freezer focuses primarily on controlling the temperature of the appliance. This shows that many decentralized elements have a twodimensional aspect. As a consequence they may have similar feasibility problems when combined in large groups of equivalent appliances as in the microCHP planning problem.

6.1.1 DISTRIBUTED GENERATION

Possibilities for distributed energy generation on a household scale (i.e. microgeneration) are abundant nowadays. We distinct between two types of appliances for microgeneration. First, microgenerators exist that are mainly installed to supply the heat demand of the household. There are different types of this kind of generation, of which we treat microCHPs and heat pumps. Other types of heat demand driven mi-

crogeneration (e.g. solar boilers) are not considered, since they are non-controllable for a planner. The second type of appliances consists of microgenerators that have the primary goal to produce electricity. On a household level, these generators (e.g. PV panels, small windmills) completely depend on renewable energy sources and are thus non-controllable in the planning process. Therefore, we concentrate on microCHPs and heat pumps, from which the microCHP has already been treated extensively in Chapter 3.

A heat pump [75] extracts heat from the immediate surrounding of a building. The heat is extracted from outside air or from a certain depth within the soil and transported through the air or through water. Electricity is used to provide the mechanical work that is needed to enforce the available heat of a certain temperature at the input of the appliance to achieve a higher temperature at the output. Part of the electricity that is needed to perform the heat transfer results in an additional heat generation that is used to increase the Coefficient of Performance (*COP*) of the heat pump. This *COP* is defined as the fraction between heat output and electricity input. The heat pump can also be operated in a reverse mode, meaning that it can be used to cool a building instead of heating it. In this case, the heat pump operates similar as a fridge/freezer.

When we model a heat pump, different aspects are of importance. The heat input for the heat pump is assumed to be unbounded (the soil or the outside air are represented by an infinite buffer). Note that in some countries on a long term (one or more years) the heat exchange with the surrounding environment is forced to be 0; i.e. energy neutrality is required, which forces the heat pump to use as much heat to provide heat demand in winter as it returns by cooling in the summer. For the short term operation for a single day we assume that this restriction has no influence on the possible operation of the heat pump. We model the electricity consumption of a heat pump by the variable e_j^i and the corresponding heat generation by the variable g_j^i for heat pump i and time interval j . Note that positive values for e_j^i correspond to electricity consumption, in contrast with the used variables in the microCHP case. A heat pump can operate at multiple modulation modes, which correspond to different levels of electricity consumption that result in differences in the heat output. These modulation modes are chosen from 0 kW (the heat pump is off) to 2 kW with steps of 0.4 kW. Using $COP = 4$ [11] this leads to a maximum heat generation of 8 kW, which corresponds to the amount a microCHP produces when it runs at maximum production. Converted to time intervals, E_{\max}^i represents the maximum possible electricity consumption (in kWh) in a time interval. Furthermore, let m_j^i be an integer variable which expresses the chosen modulation level of heat pump i in time interval j :

$$m_j^i \in \{0, 1, \dots, 5\}. \quad (6.1)$$

The electricity consumption then can be expressed by:

$$e_j^i = \frac{m_j^i}{5} E_{\max}^i. \quad (6.2)$$

The generated heat depends linearly on this consumed electricity:

$$g_j^i = COP \times e_j^i \quad \forall i, j. \quad (6.3)$$

The choice for $COP = 4$ coincides with usual Coefficients of Performance for heat pumps [11]. We require the heat pump to keep its chosen modulation level m_j^i constant for the duration of half an hour to prevent alternating behaviour on the short term. For the heat pump we assume negligible startup and shutdown times, which corresponds to the realtime behaviour of the heat pump. Furthermore, we assume that the heat output of the heat pump is connected to a heat buffer in a similar way as the microCHP is. In this way, the heat buffer offers a certain degree of freedom to the operation of the heat pump that is equivalent to the flexibility of a microCHP. The natural restriction to stay within the bounds of the heat buffer are given by the following equations:

$$hl_1^i = BL^i \quad \forall i \quad (6.4)$$

$$hl_j^i = hl_{j-1}^i + g_{j-1}^i - H_{j-1}^i - K^i \quad \forall i, j = 2, \dots, N_T + 1 \quad (6.5)$$

$$0 \leq hl_j^i \leq BC^i \quad \forall i, j = 1, \dots, N_T + 1, \quad (6.6)$$

where the heat buffer is modelled similar as in the microCHP case, using a begin level BL^i , a capacity BC^i , a heat loss K^i , the heat demand H_j^i and the variable hl_j^i that models the heat level at the start of time interval j .

Compared to the operation of microCHPs, we can model similar flexibility by using the same heat buffer sizes. However, due to the different modulation possibilities we obtain another degree of freedom in the operation. Of course the combined operation of different heat pumps is subject to cooperational constraints. These requirements form a desired aggregated electricity demand pattern, which is treated as part of the solution method for the general energy planning problem.

6.1.2 DISTRIBUTED STORAGE

Regarding energy storage related to local households, heat and electricity buffers can be distinguished. Heat buffers have already been treated in combination with the use of distributed generation techniques. As electricity buffers we consider batteries of a size equivalent to car batteries that become available with the introduction of electrical cars [33].

From a user perspective an electrical car battery is intended to be charged, such that the battery is full when the car is used for driving. Let the capacity of the battery of car i be denoted by CC^i . A typical value of the car battery is around 50 kWh [33]. The time period between the (planned) arrival and the planned departure can be used to schedule the charging process for the car. This time period can be partitioned in an uninterrupted set of increasing time intervals $\{t_a, \dots, t_d\}$ which is a subset of the complete set of time intervals $\{1, \dots, N_T\}$ of the planning horizon $[0, T]$. We use a binary parameter A_j^i to indicate the availability of the electric car of household i in time interval j ; if $A_j^i = 1$, the car is available for charging and if

$A_j^i = 0$ the car is unavailable. Let MC^i represent the maximum amount of electricity in kWh that the car of household i can be charged in one time interval. This value MC^i can result from the technical specifications of the car battery, but it is often the case that this technical maximum is too large for direct application within a house (the house would have to be equipped with a dedicated electric circuit to be able to reach this maximum). If this is the case, MC^i may be further limited by the technical constraints of the house. The decision variable c_j^i models the charging of the electrical car. To ensure that the variables c_j^i are consistent with the availability of the car, we use the following constraints:

$$0 \leq c_j^i \leq MC^i A_j^i \quad \forall i \in I, \forall j \in J. \quad (6.7)$$

This way, charging is prevented in case that the car is unavailable (c_j^i is forced to be 0); all other charging possibilities are still open. The battery level bl_j^i at the end of interval j depends on an initial battery level BBL^i at the arrival of the car and the charging decisions c_j^i . Formally, this is expressed by:

$$bl_{t_{a-1}}^i = BBL^i \quad (6.8)$$

$$bl_j^i = bl_{j-1}^i + c_j^i \quad \forall j \in \{t_a, \dots, t_d\}. \quad (6.9)$$

We assume here that the goal is that the battery has to be fully charged at the departure time:

$$bl_{t_d}^i = CC^i. \quad (6.10)$$

However, we also may define less strict requirements on the battery level at departure, which increases the flexibility of the planning.

Until now we have focused only on charging the car battery. However, as long as the car is available at the house and if the time of departure allows for it, we also may use the battery as an electricity supplier. In this (Vehicle to Grid [71, 76]) case constraint (6.7) changes into:

$$-\tilde{MC}^i A_j^i \leq c_j^i \leq MC^i A_j^i, \quad (6.11)$$

where the maximum amount that can be taken out of the battery in one time interval is given by \tilde{MC}^i . Note that discharging the car is denoted by negative values for c_j^i . Furthermore, the capacity limits of the battery cannot be exceeded:

$$0 \leq bl_j^i \leq CC^i \quad \forall j \in \{t_a, \dots, t_d\}. \quad (6.12)$$

Next to electrical cars, we also study batteries that are installed in houses. The model we use for these batteries originates from the model for the electrical cars ((6.7)-(6.12)), by setting $t_a = 1$ and $t_d = N_T$. In this case we can omit equation (6.7) and the availability parameter A_j^i . Besides this, we request the battery level at the end of the day to differ only slightly from the initial level at the start of the day, since we do

not want to use these batteries to compensate for large indiscrepancies. Therefore, (6.10) changes into:

$$0.8BBL^i \leq bl_{N_T}^i \leq 1.2BBL^i \quad \forall i \in I, \quad (6.13)$$

meaning that we want the total amount of energy in each battery i at the end of the planning horizon to be almost equal to the total amount of energy in the battery at the start of the planning horizon.

6.1.3 LOAD MANAGEMENT

The examples upto now have shown that there is a lot of interaction between distributed generation, distributed storage and local consumption. Therefore, it is not easy to draw a strict borderline between distributed generation, storage and demand side load management. In this context, note that load management is not only restricted to the appliances that we model below. Earlier described appliances, such as electrical cars and heat pumps, can be placed under the umbrella of load management too. However, the differences between pure consumption and consumption with additional restrictions (the generation of heat or the possible supply of electricity) make it worth to discuss them separately, as we did above.

As an example of controllable consuming appliances, we consider the operation of a freezer. The model of a freezer we present in the following is included in the case study of Section 6.4. Usually, a freezer has a very repetitive structure of cooling for a certain period, followed by a period where the freezer is switched off. This repetitive process is a result of the requirement, that the temperature of the freezer has to stay between a lower temperature T_{\min} and an upper temperature T_{\max} during operation. In our model, we choose $T_{\min} = -23^\circ\text{C}$ and $T_{\max} = -18^\circ\text{C}$. For modelling the freezer, furthermore a parameter T_{init}^i , representing the initial temperature of freezer i at the start of the planning horizon, is needed. The operation of the freezer can be expressed by binary decision variables d_j^i representing the decision to cool ($d_j^i = 1$) or not to cool ($d_j^i = 0$):

$$d_j^i \in \{0, 1\}. \quad (6.14)$$

To describe the cooling behaviour of the freezer, we specify the parameters for basic time intervals of 6 minutes ($\frac{1}{10}$ th of an hour). During such an interval, we assume that the temperature of the freezer increases with ΔT_{off} and decreases with ΔT_{on} when $d_j^i = 1$. For an interval of 6 minutes we choose $\Delta T_{off} = 0.1^\circ\text{C}$ and $\Delta T_{on} = 0.6^\circ\text{C}$, which corresponds to a cooling capacity of $0.1 \frac{^\circ\text{C}}{\text{minute}}$ (and an effective temperature drop of 0.5°C per basic time interval when the freezer is on). The operation of a freezer has to respect the temperature limits, which can be expressed by:

$$T_{\min} \leq T_{init}^i + j\Delta T_{off} - \Delta T_{on} \sum_{k=1}^j d_k^i \leq T_{\max} \quad \forall i \in I \forall j \in J. \quad (6.15)$$

We assume that the electrical consumption f_j^i of a freezer depends directly on d_j^i :

$$f_j^i = FC^i d_j^i, \quad (6.16)$$

where FC^i is the electricity consumption during one basic time interval when the freezer is on. We set $FC^i = 15$ Wh in 6 minutes intervals, which corresponds to a freezer with a power consumption of 150 W. We do not consider the influences of user interaction on the temperature level of the freezer.

If we integrate freezers in a use case, the flexibility of a freezer is bounded, similar to the operation of microCHPs. However, the regular temperature increasing behaviour and the corresponding regularity in the electricity consumption give a planner more possibilities to influence the decisions in later time intervals by shifting the operation (e.g. flattening the demand of a group of freezers is a promising objective).

6.2 THE GENERAL ENERGY PLANNING PROBLEM

In the previous section the constraints on the individual operation of some appliances are given. These appliances may be combined with the elements of the standard Unit Commitment Problem to form the framework of the general energy planning problem. This section sketches this framework of the general energy planning problem. Starting from the UCP we derive additional constraints to formulate the general energy planning problem.

6.2.1 THE UNIT COMMITMENT PROBLEM

The basis of the classic UCP is a set of generators. Each of these generators can produce electricity at different production levels against certain costs. The primal objective of the set of generators is to supply given electricity demands d_j , that are specified for time intervals j . Additionally, at each time interval a certain spinning reserve capacity r_j has to be available, which consists of (parts of) the currently unused capacity of the already committed (running) generators. The classic UCP focuses on operational costs or on the revenue/profit of the system of generators. For sake of simplicity, in the following we concentrate on the operational costs. Operational costs are depending both on the binary commitment variables u_j^i (specifying whether generator i is committed or not in time interval j) and on the production level x_j^i (specifying the electricity production of generator i in time interval j). In general the operational costs can be described by a function $f(u, x)$, where the variables u and x are indexed by time intervals and generators. Note that startup costs are incorporated in this notation. The decision problem for this set of generators is described as follows:

The Unit Commitment Problem

INSTANCE: Given is a set of N generators with capacities $x^{i,\max}, i = 1, \dots, N$, an electricity demand vector $d = (d_1, \dots, d_{N_T})$ and a spinning reserve vector $r = (r_1, \dots, r_{N_T})$. Furthermore, a bound K and a function $f_i(u^i, x^i)$ is given, which specifies for each vector pair (u^i, x^i) , whereby $x^i = (x_1^i, \dots, x_{N_T}^i)$ represents the production level and $u^i = (u_1^i, \dots, u_{N_T}^i)$ represents the binary unit commitment of generator i , the operational costs if generator i is operated in this way.

QUESTION: Is there a selection of unit commitment/operation level pairs (u^i, x^i) for all generators $i = 1, \dots, N$, such that $\sum_{n=1}^N x_j^n \geq d_j$,

$$\sum_{n=1}^N u_j^n x^{i,\max} - x_j^i \geq r_j \text{ and } \sum_{n=1}^N f_i(u^i, x^i) \leq K?$$

The description of unit commitment and production levels in this formal definition of the UCP is rather abstract. They become more clear when we sketch the optimization problem associated to the UCP and its operational costs, in which the objective of the UCP is to minimize $f(u, x) = \sum_{i=1}^N f_i(u^i, x^i)$. Also, some common constraints that are used in most descriptions of the UCP are given below. The total production has to satisfy the demand in each time interval. Moreover, additional spinning reserve capacity needs to be assigned to guarantee a certain amount of flexibility in the case of a higher-than-predicted demand or in the case of a failure of a committed generator. The possible production of a generator is restricted by lower and upper limits on the production level, as well as to ramp up and ramp down rates $s^{i,up}$ and $s^{i,down}$, which determine the speed with which generation can be adjusted. Another common constraint is that a generator has to stay committed for a certain number of consecutive time intervals, once it is chosen to generate (minimum runtime). Similarly, minimum offtimes are required once the decision is made to switch the generator off. These constraints are formulated in (6.17)-(6.25).

$$\min f(u, x) \quad (6.17)$$

$$s.t. \sum_i x_j^i \geq d_j \quad \forall j \quad (6.18)$$

$$\sum_i (u_j^i x^{i,\max} - x_j^i) \geq r_j \quad \forall j \quad (6.19)$$

$$u_j^i x^{i,\min} \leq x_j^i \leq u_j^i x^{i,\max} \quad \forall i, j \quad (6.20)$$

$$s^{i,down} \leq x_j^i - x_{j-1}^i \leq s^{i,up} \quad \forall i, j \quad (6.21)$$

$$u_j^i \geq u_{j-k}^i - u_{j-k-1}^i \quad \forall i, j, k = 1, \dots, t^{i,mr} - 1 \quad (6.22)$$

$$1 - u_j^i \geq u_{j-k-1}^i - u_{j-k}^i \quad \forall i, j, k = 1, \dots, t^{i,mo} - 1 \quad (6.23)$$

$$u_j^i \in \{0, 1\} \quad \forall i, j \quad (6.24)$$

$$x_j^i \in \mathbb{R}^+ \quad \forall i, j \quad (6.25)$$

Equation (6.18) requires that the total production satisfies the total electricity demand; equation (6.19) asks for a certain amount of spinning reserve r_j , i.e. the additional available generation capacity of already committed generators. The sum of the difference between the capacity $x^{i,\max}$ of committed generators and their current electricity generation needs to be larger than r_j in time interval j . The production boundaries of the generators $x^{i,\min}$ and $x^{i,\max}$ are modelled in equation (6.20). The ramp up and ramp down rates are taken into account in equation (6.21). Equations (6.22) and (6.23) state that the generator has to stay up and running (or stay switched off) once a corresponding decision to switch it on (or off) has been made within the last $t^{i,mr}$ ($t^{i,mo}$) time intervals. The decisions to commit a generator are binary decisions, where the production decisions are real numbers.

The operational cost function $f(u, x)$ includes two types of costs. First, it is desired to have long runs for committed power plants. Therefore, the start of a generator, which can be derived from the unit commitment variables u , is penalized with a certain penalty cost. By using these penalty costs, the UCP tries to avoid switching between the commitment of generators on a short term period. The second part of the cost function deals with the production levels of committed generators. Depending on the behaviour of the fuel costs related to the production level, quadratic cost functions are used to model the costs associated with these different production levels [102]. In our model these quadratic cost functions are approximated with piecewise linear cost functions, to incorporate this notion of quadratic costs in an ILP formulation.

Formulation (6.17)-(6.25) shows that the level of generation still has some flexibility, once the decision has been made to commit the corresponding generator. This flexibility can be used to prevent the additional use of currently uncommitted generators. However, as also the spinning reserve constraint has to be taken into account, the planning of the UCP cannot use this flexibility to its full extent.

6.2.2 THE GENERAL ENERGY PLANNING PROBLEM

In the general energy planning problem we include the developing energy infrastructure next to normal power plants. Especially, we focus on five distinct elements: microCHPs, heat pumps, electrical cars, batteries and freezers. In this section we sketch the influence of these elements on the UCP. Attention is given to the combined objective function of the general energy planning problem, as well as to the possibilities to steer the demand.

Next to the power plants and their usual operation that is given by the normal Unit Commitment Problem, we have decentralized appliances. We denote the set of these appliances by M . These appliances are somehow collected in a Virtual Unit (not to be confused with a Virtual Power Plant). On the one hand, this unit changes the requested demand profile (we use variables z^m for specifying the use of appliances $m \in M$ and a function h to describe how z influences the demand

in the different time intervals). On the other hand, it may be possible that some of the production within the Virtual Unit is offered to an electricity market. To cope with this option, we specify production by variables y^m , which depend on the unit commitment u^m and the actual use z^m of a subset of appliances that are able to generate electricity and to participate in acting on an electricity market. A function $g(p, y)$ that depends on the market clearing prices p and the production of (a part of) the Virtual Unit describes the profit that can be made. Finally, constraints specifying the correct use of the decentralized appliances have to be added.

The formulation of the general energy planning problem is given by equations (6.26)-(6.39), where the original UCP can be found in equations (6.27)-(6.32). Note that this is a mere modelling description of the behaviour of the different elements, and not a formulation of a specific form (like e.g. an ILP formulation).

$$\min f(u, x) - g(p, y) \quad (6.26)$$

$$s.t. \sum_i x_j^i + \sum_m y_j^m \geq d_j + \sum_m h_j(z^m) \quad \forall j \quad (6.27)$$

$$\sum_i (u_j^i x_j^{i,\max} - x_j^i) \geq r_j \quad \forall j \quad (6.28)$$

$$u_j^i x_j^{i,\min} \leq x_j^i \leq u_j^i x_j^{i,\max} \quad \forall i, j \quad (6.29)$$

$$s^{i,\text{down}} \leq x_j^i - x_{j-1}^i \leq s^{i,\text{up}} \quad \forall i, j \quad (6.30)$$

$$u_j^i \geq u_{j-k}^i - u_{j-k-1}^i \quad \forall i, j, k = 1, \dots, t^{i,mr} - 1 \quad (6.31)$$

$$1 - u_j^i \geq u_{j-k-1}^i - u_{j-k}^i \quad \forall i, j, k = 1, \dots, t^{i,mo} - 1 \quad (6.32)$$

$$u_j^m \geq u_{j-k}^m - u_{j-k-1}^m \quad \forall m, j, k = 1, \dots, t^{m,mr} - 1 \quad (6.33)$$

$$1 - u_j^m \geq u_{j-k-1}^m - u_{j-k}^m \quad \forall m, j, k = 1, \dots, t^{m,mo} - 1 \quad (6.34)$$

$$u^m \in \mathcal{H} \quad \forall m \quad (6.35)$$

$$y_j^m = l(u^m) \quad \forall m, j \quad (6.36)$$

$$z^m \in \mathcal{S} \quad \forall m \quad (6.37)$$

$$u_j^i, u_j^m \in \{0, 1\} \quad \forall i, m, j \quad (6.38)$$

$$x_j^i, y_j^m \in \mathbb{R}^+ \quad \forall i, m, j \quad (6.39)$$

We have a group of microCHPs that can operate as decentralized electricity producers, representing the part of the Virtual Unit that can operate on an electricity market. The combined generation of this group partially satisfies the electricity demand in the problem, but moreover this production can be offered to the day ahead market. To express this possibility we add a function $g(p, y)$ to the objective function, which represents the profit of the VPP of microCHPs based on the predicted electricity prices p and the electricity generation y .

In this general energy planning problem model an important change is the incorporation of demand side load management to adjust the distribution of the demand. This is formalized in equation (6.27) by the function $h_j(z)$, where z^m represents the tuple of controllable appliances in house m . We denote the space \mathcal{S}

in (6.37) to represent the feasible demand side management possibilities, which are constrained by equations (6.1)-(6.16) of the previous section.

Furthermore, the local generation y_j^m of the microCHP is taken into account in equation (6.27) too. The local generators have the same type of dependency constraints on runtime and offtime over time intervals (equations (6.33) and (6.34)) as the large generators (equations (6.31) and (6.32)). Next to these machine dependency constraints the generators also have user dependencies, resulting e.g. from the heat demand. Equation (6.35) uses the space \mathcal{H} to denote the feasible commitment options for the microCHPs. The generator output is completely determined by the commitment decisions, as in equation (6.36).

6.3 SOLUTION METHOD

The planning problem is already \mathcal{NP} -complete in the strong sense if only a group of microCHPs is considered. This complexity follows from the two-dimensional aspect of the problem (i.e. a strong dependency between generation in time intervals and a strong dependency between households due to the aggregated generation in the fleet). It is therefore practically intractable to solve the general energy planning problem (of which the microCHP planning problem is only a part of the problem) to optimality. In this section a heuristic method for the general energy planning problem is presented that uses the natural division into different production levels to separate the decisions that have to be made for the power plants, the decisions that have to be made for the local generators and the decisions to be made for demand side load management.

In Chapter 2 an energy model of the smart grid is given using a division into different levels. This division is based on the amount of energy the different generators produce and on the location. This division forms the base for a leveled approach to solve the general energy planning problem. In this section, this leveled approach for solving the planning problem is given, introducing patterns as building blocks for the method. First we elaborate on the hierarchical structure of the general energy planning problem; then we show the cooperation between different master and sub problems that are solved in order to find a solution to the general energy planning problem.

6.3.1 HIERARCHICAL STRUCTURE

Since it is hard to combine the commitment of large and small types of generation in one decision step, we divide the general energy planning problem in different smaller problems. An important aspect when dividing a problem in multiple parts, is to incorporate the given objective within the different sub problems in a proper way. We propose a hierarchical structure that naturally allows the planner to obtain information of sub problems on the smaller generation levels and to give feedback to more local problems on how to improve their local solution with respect to the global (original) problem. Figure 6.1 shows the proposed hierarchical division. In the top level we have the large power plants and aggregated generation that

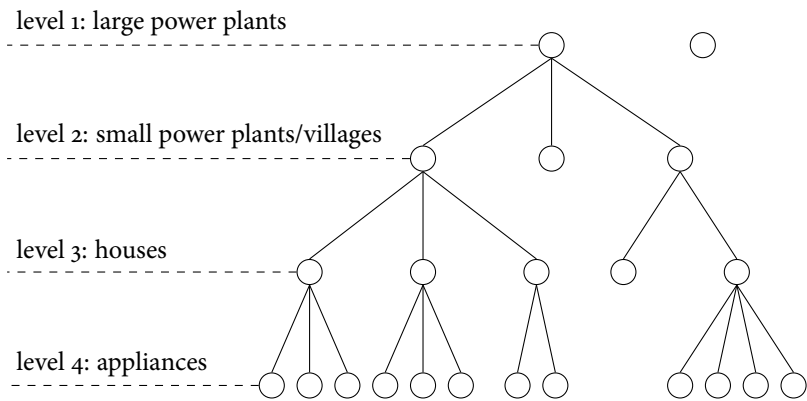


Figure 6.1: The hierarchical structure of the general energy planning problem

has equivalent capacity as a large power plant. The second level consists of small generators (e.g. biogas installations, small windmill parks) that produce significantly less than the large power plants, and the aggregated production/consumption of villages/cities. The third level is the house level, which operation is aggregated on the higher village level by using the exchanging elements of the energy model. On this house level single appliances are planned. In case only one controllable appliance is available at the house, this appliance is considered as a complete house, since it is the only controllable variable. In case of multiple controllable appliances a fourth level is introduced, which is the lowest level in the hierarchy.

6.3.2 SUB LEVELS AND SUB PROBLEMS

Each node in the hierarchy is considered as an entity in the solution method for the general energy planning problem. The original problem corresponding to the example in Figure 6.1 is depicted in Figure 6.2, where black nodes correspond to the elements for which a planning is needed and white nodes correspond to aggregation of information. The elements for which a planning is needed are the leaves in the




Figure 6.2: The general energy planning problem

hierarchical structure. The optimal solution to the general energy planning problem consists of specific production/consumption patterns for each black element. Intermediate nodes in the graph are unused in the original problem formulation, but are used in the heuristic method as communicating and aggregating entities.

The heuristic method is based on the notion of patterns for all considered elements. Hereby, a pattern consists of a vector of the electricity balance for each time interval, where a positive value corresponds to production and a negative value to consumption. To achieve problems which are better tractable we divide the problem in a master problem and various sub problems. A key property of these divided problems is that in each problem only a part of the original problem is optimized by creating (new) patterns for this part. The heuristic uses at all times only a small subset of the possible patterns which may exist for an element. Only this subset of patterns is considered for the elements in the planning process, meaning that we do a restricted search in the space of feasible patterns.

The master problem acts on the highest level of the considered problem instance. Figure 6.3a shows the elements that are used in this master problem. The black nodes correspond to elements for which a pattern has to be found, based on the original objective function of the general energy planning problem. Grey nodes serve as input for the master problem, i.e. the corresponding elements have to produce a set of patterns, which reflect possible overall patterns of the subproblems they are responsible for. This set is not changed during the solving process for the master problem in a given iteration and is thus a fixed input set during one iteration of the method.

Based on the achieved solution to the master problem, information can be derived that asks the lower (grey) elements to adjust their set of patterns. This information exchange is shown by the light grey nodes in Figures 6.3b-6.3f. In each sub problem patterns are created for the black nodes, based on this information from above, and possibly based on (limited) pattern sets from the grey nodes below. Note that in this setting in each master or sub problem decisions have to be made for only a limited amount of elements which are comparable in size. The objective for sub problems is to optimize the pattern that has to be created based on the information from above. Although the original objective function is invisible in the local sub problems, the local objectives are ultimately based on the original objective function.

6.3.3 PHASES AND ITERATIONS

The previous subsection shows the possible interaction between different sub problems and the main problem. In this subsection we sketch how these problems are solved sequentially (or in parallel) in different iterations. The general energy planning problem is solved in several phases, which describe subroutines in the general energy planning problem, using several iterations, which represent the amount of times that a certain subroutine is repeated.




Figure 6.3: The division into master and sub problems

Initial phase

In the initial phase subroutine, the master problem makes use of a rough estimation of the possibilities for local entities, by aggregating information from these local entities. This information is used to derive objective bounds for the local sub problems. Simultaneously, initial pattern sets are created for sub problems.

Method in progress

The subroutine ‘method in progress’ tries to improve on the matching problem, which wants to create for each sub problem a total pattern that equals the rough

estimation of the initial phase. When the solution method progresses, pattern sets of local sub problems are extended in order to improve the match to the local bounds. These extended pattern sets are used to solve the sub problems to achieve a new (and better) solution. This solution represents a new pattern one step higher in the hierarchical structure, i.e. it leads to new (combined) patterns on this higher level. Eventually, the rough planning in the highest level (the root) is approximated using the latest information.

This subroutine is an iterative process, in which the local sub problems can be repeatedly solved, based on new information from above. If a new solution in a sub problem has been determined that fulfills (partly) the requested changes, it is sent to a higher level, where the same process is repeated. We choose to continue this iterative process at each level, until no improvements on the requested changes occur in this level.

Final solution

If the solutions to the approximation of the local entities within the master problem show the desired behaviour, or if the iterative process in the previous subroutine is finished, the solution method stops. The master problem is solved using this latest information, where the rough planning is replaced by the best found approximation for the local entities. Depending on the quality of the final result, the root (top level node) may decide to repeat the complete planning process, starting with the initial phase. In this case information from the final result serves as additional input for the initial planning phase.

6.4 RESULTS

The general energy planning problem is tested for an instance that consists of 5000 houses. This number of houses corresponds to a small town or a large village, or a cluster of small villages. This amount of houses suffices for a thorough analysis of the behaviour of the lower levels of the hierarchical structure of the problem (i.e. level 2 and below). Since the planning heuristic is set up in a hierarchical way, the step towards including the first level when solving a problem instance with millions of houses at the lower levels is possible in theory, by adding the first level and solving the corresponding pattern matching problems. In practice we did not perform such a test, since the problem is currently being solved on a single computer, due to the unavailability of a network version of the modelling software AIMMS. However, note that the number of (local) generators is significantly larger than the problem instances that usually occur in the field of Unit Commitment (see Chapter 2).

We consider two case studies. In the first case study we focus on controlling a Virtual Power Plant consisting of 5000 microCHPs in combination with 10 small power plants. The second case study includes not only microCHPs, but also heat pumps, controllable freezers, electrical cars and batteries.

6.4.1 CASE STUDY I

To study the influence of generation on multiple levels in the electricity grid, we set up a case study with two or three levels. We start with two levels, to see the interaction between generators of different production capacity in a direct way. In the end we use an intermediate level to aggregate information from the lowest level and communicate this information to above. In this illustrative example we use 10 generators on the highest level, with a total production capacity of 15 MW. This capacity is divided over 5 generators with a capacity of 1 MW and a minimum production level of 0.5 MW, and 5 generators with a capacity of 2 MW and a minimum production level of 1 MW. The (absolute) ramp up and ramp down rates are equal to the minimum production for each power plant. Between the maximum and minimum production values the operator of the generator has flexibility to choose its power output, once the unit is committed. The minimum runtime and offtime are set to half an hour.

On the lowest level, we have 5000 houses containing a generator, leading to a total capacity of 5 MW. These generators are microCHPs with a production output of 1 kW. We neglect startup and shutdown times, meaning that the power output is a direct result from the decision to run the microCHP at a certain moment in time. As a consequence, there is no flexibility in the production level of committed low level units. Flexibility can only be found in the moments in time that the units are committed. However, these moments are constrained by the heat demand in the houses: the used heat demand profiles result in a maximum production of the fleet over the planning horizon of 39.8 MWh and a minimum production of 35.1 MWh, which is of the same order of committing a power plant for a complete day. The heat demand is defined in a similar way as in Chapter 3, using parameters *MaxOn* and *MinOn* describing the flexibility of the operation of a single microCHP. The minimum runtime and offtime are again set to half an hour.

In the case study we define four use cases to study the influence of introducing a fleet of microCHPs in the UCP. For each use case we use time intervals of 30 minutes length; the commitment is planned for a complete day, which comes down to 48 time intervals. The total daily demand for the group of houses is 114.2 MWh, with a peak of 8 MW and a base load of 2.5 MW. In this case study we do not consider demand side load management. We require a spinning reserve of 2 MW at all time intervals.

The objective function combines profit maximization for the fleet and operational costs for the power plants. The first use case is based on real prices from the APX day ahead market ¹. In the second use case we multiply all prices with -1 , which creates artificial negative values, to investigate to what extent the fleet changes its decisions. The third use case uses artificial prices that are based on the daily electricity demand; the higher the demand, the higher the price. This use case is defined to investigate if the fleet can behave in such a way that peak demand can be decreased and the demand for the power plants can be flattened. The fourth use

¹<http://www.apxindex.com>

case is the opposite of the third case, in the sense that prices are again multiplied with -1 ; the higher the demand, the lower the price.




Figure 6.4: The operational cost functions of the power plants

In Figure 6.4 the cost functions of the power plants are given. They are modeled as piecewise linear cost functions, to approximate quadratic operational cost functions (see e.g. [102]). Below certain production levels (625 kWh for the large power plants and 312.5 kWh for the small power plants) the cost functions of all power plants of the same size are equal, and power plants are mutually exchangeable. The start of a power plant is furthermore penalized with a cost of 1000.

We use different optimization problems in a structure as explained in Section 6.3. These different optimization problems are modeled as Integer Linear Programming formulations in AIMMS modeling software using CPLEX 12.2 as solver.

On the highest level, the operation of the power plants is optimized and a rough planning of the microCHPs is made, based on aggregated information from the operational flexibility of all households. A so-called fleet production f is introduced, which represents the total production of the fleet of microCHPs. This fleet production respects total maximum and minimum generation constraints which are bound by the sum of $MaxOn_m$ and $MinOn_m$ for all 5000 microCHPs m . Also in each time interval at most 5000 microCHPs can possibly generate, which gives an additional bound of 2500 kWh per half an hour interval. Using this aggregated information of the group of microCHPs, the master problem finds an overall rough planning of how much microCHPs are running in the different time intervals, combined with the operation of the power plants. Hereby, no individual planning of the microCHPs is carried out; it is only ensured that the restrictions resulting from the aggregated heat demand parameters and the production capacity of this VPP are taken into account. Next to the fleet production f for the VPP, we introduce operational costs c_j^i for the power plants. The form of the general energy planning problem that combines the microCHP planning problem with the Unit

Commitment Problem is summarized in the following ILP formulation:

$$\min \sum_{i,j} c_j^i + \sum_{i,j} (1000 \times start_j^i) - \sum_j \pi_j f_j \quad (6.40)$$

$$\sum_i x_j^i + f_j \geq d_j \quad \forall j \in J \quad (6.41)$$

$$\sum_i (u_j^i x_j^{i,max} - x_j^i) \geq r_j \quad \forall j \in J \quad (6.42)$$

$$start_j^i \geq u_j^i - u_{j-1}^i \quad \forall i \in I, \forall j \in J \quad (6.43)$$

$$c_j^i \geq A_r^i x_j^i + B_r^i \quad \forall i \in I, \forall j \in J, \forall r \in R \quad (6.44)$$

$$x_j^i \leq x_j^{i,max} u_j^i \quad \forall i \in I, \forall j \in J \quad (6.45)$$

$$x_j^i \geq x_j^{i,min} u_j^i \quad \forall i \in I, \forall j \in J \quad (6.46)$$

$$x_j^i - x_{j-1}^i \leq s^{i,up} \quad \forall i \in I, \forall j \in J \quad (6.47)$$

$$x_j^i - x_{j-1}^i \geq s^{i,up} \quad \forall i \in I, \forall j \in J \quad (6.48)$$

$$u_j^i \geq u_{j-k}^i - u_{j-k-1}^i \quad \forall i \in I, \forall j \in J, k = 1, \dots, t^{i,mr} - 1 \quad (6.49)$$

$$1 - u_j^i \geq u_{j-k-1}^i - u_{j-k}^i \quad \forall i \in I, j \in J, k = 1, \dots, t^{i,mo} - 1 \quad (6.50)$$

$$2 \sum_{k=1}^j f_k \leq \sum_m MaxOn_{m,j} \quad \forall j \in J \quad (6.51)$$

$$2 \sum_{k=1}^j f_k \geq \sum_m MinOn_{m,j} \quad \forall j \in J \quad (6.52)$$

$$f_j \leq 2500 \quad \forall j \in J. \quad (6.53)$$

The objective minimizes costs c_j^i and the total number of starts, and maximizes the profit $\pi_j f_j$ of the fleet. Constraint (6.43) determines the start of the operation of a power plant. In (6.44) the piecewise linear costs are calculated, using different linear inequalities indexed by r of the form $A_r^i x_j^i + B_r^i$. The use of (6.44) in combination with the objective of minimizing c_j^i is sufficient to model the approximation of quadratic operational costs. Equations (6.51)-(6.53) give the aggregated bounds on the fleet production. The factor 2 is used since we use time intervals of half an hour and $MaxOn_m$ and $MinOn_m$ are defined in time intervals.

The above ILP formulation gives, next to a planning of the power plants, a desired fleet production $\hat{f} = f$. This production \hat{f} needs to be approximated by solving the microCHP planning problem, where $P^{upper} = P^{lower} = \hat{f}$. We use such tight bounds to see to what extent the planning method is able to reach the rough planning exactly. For sake of completeness, the ILP formulations that give the pattern generation method for this particular desired aggregated production are given below. The sub problem that combines patterns for individual microCHPs

to generate aggregated patterns is given by:

$$\min \sum_{j=1}^{N_T} (sl_j + ex_j) \quad (6.54)$$

$$\sum_{m=1}^N \sum_{p \in S_m} pe_{p,j} y_{m,p} + sl_j \geq \hat{f}_j \quad \forall j \in J \quad (6.55)$$

$$\sum_{m=1}^N \sum_{p \in S_m} pe_{p,j} y_{m,p} - ex_j \leq \hat{f}_j \quad \forall j \in J \quad (6.56)$$

$$\sum_{p \in S_m} y_{m,p} = 1 \quad \forall m \in \{1, \dots, 5000\} \quad (6.57)$$

$$sl_j, ex_j \geq 0 \quad \forall j \in J \quad (6.58)$$

$$y_{m,p} \in \{0, 1\}, \quad (6.59)$$

where S_m represents the set of patterns that are generated for microCHP m . Initially each microCHP has two operational patterns. The first pattern is determined by postponing the generation to the latest possible time intervals, based on the values for $MinOn_m$ that determine the delayed operation that is feasible when the heat demand is taken into account. The second pattern follows the opposite idea: this pattern generates as early as possible, based on the values for $MaxOn_m$ (which satisfy the heat demand, by producing as early as possible). Naturally, minimum runtime and offtime requirements are satisfied, since we deal with intervals of half an hour. New patterns for individual microCHPs are created by solving the following sub problem:

$$\max \sum_{j=1}^{N_T} \lambda_j (pe_{g,j} - pe_{c,j}) \quad (6.60)$$

$$\sum_{k=1}^j pe_{g,k} \leq \frac{1}{2} MaxOn_{m,j} \quad \forall j \in J \quad (6.61)$$

$$\sum_{k=1}^j pe_{g,k} \geq \frac{1}{2} MinOn_{m,j} \quad \forall j \in J \quad (6.62)$$

$$2pe_{g,j} \in \{0, 1\} \quad \forall j \in J, \quad (6.63)$$

as in the column generation approach of Chapter 3. The factors 2 and $\frac{1}{2}$ again origin from the difference between the amount of energy, which is denoted in kWh, and the number of intervals per hour which is 2.

In case we use an additional level between the household level and the top level as we do at the end of this case study, the information exchange from the individual houses to the centralized operator of this general energy planning problem takes place in two steps. In case of this 2-step approach, the generated patterns of a limited number of houses are combined in a subgroup G_l , where patterns are combined to form aggregated patterns that are communicated to above. These aggregated

patterns need to fulfill a certain part of the total required production \hat{f} , denoted by \hat{f}^l . The formation of these aggregated patterns is given by the following sub problem, that combines individual patterns to create aggregated patterns for sub fleets:

$$\min \sum_{j=1}^{N_T} (sl_j^l + ex_j^l) \quad (6.64)$$

$$\sum_{m \in G_l} \sum_{p \in S_m} pe_{p,j} y_{m,p} + sl_j^l \geq \hat{f}_j^l \quad \forall j \in J \quad (6.65)$$

$$\sum_{m \in G_l} \sum_{p \in S_m} pe_{p,j} y_{m,p} - ex_j^l \leq \hat{f}_j^l \quad \forall j \in J \quad (6.66)$$

$$\sum_{p \in S_m} y_{m,p} = 1 \quad \forall m \in G_l \quad (6.67)$$

$$sl_j^l, ex_j^l \geq 0 \quad \forall j \in J \quad (6.68)$$

$$y_{m,p} \in \{0, 1\} \quad \forall m \in G_l. \quad (6.69)$$

The pattern set S_l of subgroup l consists of patterns w^l that are formed by aggregating patterns $w^l = (w_1^l, \dots, w_{N_T}^l) = (\sum_{m \in G_l} \sum_{p \in S_m} pe_{p,1} y_{m,p}, \dots, \sum_{m \in G_l} \sum_{p \in S_m} pe_{p,N_T} y_{m,p})$ that are found by solving (6.64)-(6.69). The total mismatch from the desired generation \hat{f} is found by selecting exactly one pattern from each pattern set S_l for all L subgroups, as follows:

$$\min \sum_{j=1}^{N_T} (sl_j + ex_j) \quad (6.70)$$

$$\sum_l \sum_{w^l \in S_l} w_j^l y_{w^l} + sl_j \geq \hat{f}_j \quad \forall j \in J \quad (6.71)$$

$$\sum_l \sum_{w^l \in S_l} w_j^l y_{w^l} - ex_j \leq \hat{f}_j \quad \forall j \in J \quad (6.72)$$

$$\sum_{w^l \in S_l} y_{w^l} = 1 \quad \forall l \quad (6.73)$$

$$sl_j, ex_j \geq 0 \quad \forall j \in J \quad (6.74)$$

$$y_{w^l} \in \{0, 1\}. \quad (6.75)$$

By combining the above ILP formulations in an iterative structure as sketched in Section 6.3 a detailed planning for the fleet of microCHPs is found. In a final step, the resulting f -values of this detailed fleet planning are added in the master problem (6.40)-(6.53) to calculate a unit commitment of the power plants based on the given planning for the fleet. In detail, constraints (6.51)-(6.53) are replaced by a constraint stating that the variables f are fixed to the values of the best found fleet planning bf : $f := bf$.

We refer to (6.40)-(6.53) as the rough planning. (6.60)-(6.63) is the problem that creates patterns. The ILP formulation (6.54)-(6.59) combines the created patterns in

one step (from individual patterns for microCHPs to global fleet patterns). Problems (6.64)-(6.69) and (6.70)-(6.75) represent the problem of creating patterns in a subgroup and of combining subgroup patterns in a higher level respectively. The resulting approach is summarized in Algorithm 4.

The algorithm continues until no improvement is found in the solution of the high level pattern combination problem. It iteratively adds patterns at a household level in case these patterns show the possibility to improve the global problem. The creation of patterns at a household level stops whenever the first pattern occurs that does not show a possibility to improve on the total mismatch or when the subgroup it belongs to is not active anymore. A complete subgroup remains active as long as improvements are found in its own pattern creation problem, meaning that the found new pattern should improve on the best mismatch to the desired bounds for the subgroup. When no improvement is found, the subgroup is no longer active. In case no subgroup is active anymore or if no improvement is found on the highest level by combining patterns on this level, the method stops. At the end, a final planning is calculated, which determines the unit commitment of the power plants, based on the found planning for the individual microCHPs.

Results and discussion

In this section we discuss the solutions of the four use cases, which are found by the planning method.




Figure 6.5: The solution of the UCP

Figure 6.5 shows the detailed solution of the UCP, where we do not use the 5000 houses and the complete demand is fulfilled against minimal operational costs. The commitment and corresponding generation patterns are given for the 10 power plants. This solution is used to validate the optimization model. At any time interval, the spinning reserve of 1000 kWh (2 MW for time intervals of 30 minutes length) is available within this solution. The ramp rates are taken into account; the

Algorithm 4 Case study 1

init $S_m := \{p_1^m, p_2^m\}$ for all microCHPs m : p_1^m is chosen such that $pe_{p_1^m, j}$ postpones the generation to the latest time intervals possible and p_2^m such that $pe_{p_2^m, j}$ generates as early as possible.

solve rough planning (6.40)-(6.53)

for all l **do**

$\hat{f}^l \leftarrow \frac{f}{L}$

end for

if $L > 1$ **then**

for all l **do**

solve create subgroup patterns (6.64)-(6.69)

$S_l \leftarrow S_l \cup w^l$

end for

solve combine subgroup patterns (6.70)-(6.75)

for all m **do**

solve create patterns on lowest level (6.60)-(6.63)

end for

else

solve combine patterns on the highest level (6.54)-(6.59)

for all m **do**

solve create patterns on lowest level (6.60)-(6.63)

end for

end if

while stopping criteria not met **do**

update best fleet planning bf

for all m : $\sum_{j=1}^{N_T} \lambda_j (pe_{g, j} - pe_{c, j}) > 0$ **do**

$S_m \leftarrow S_m \cup g$

end for

if $L > 1$ **then**

for all active l **do**

solve create subgroup patterns (6.64)-(6.69)

$S_l \leftarrow S_l \cup w^l$

end for

solve combine subgroup patterns (6.70)-(6.75)

for all active m **do**

solve create patterns on lowest level (6.60)-(6.63)

end for

else

solve combine patterns on the highest level (6.54)-(6.59)

for all active m **do**

solve create patterns on lowest level (6.60)-(6.63)

end for

end if

end while

update best fleet planning bf

solve final planning (6.40)-(6.53) where $f := bf$

minimum and maximum production constraints of the individual power plants are considered too. This can best be seen when a generator shuts down. The time interval before a generator shuts down, the production is reduced to the minimum production level, which happens to be equal to the (absolute) maximum ramp up and ramp down rates. In this case the generator may shut down in the next interval. We see that nine of the ten units are committed during the day; each power plant is started at most once.

	case 1		case 2	
	rough	detailed	rough	detailed
operational costs	158748	164846	163593	169325
# of starts	5	6	5	8
computational time (s)	32.28	1543.10	26.91	1798.19
mismatch (kWh)	2236.5		986	
mismatch/rough prod. (%)	5.6		2.8	
fleet production (kWh)	38926		35589	
	case 3		case 4	
	rough	detailed	rough	detailed
operational costs	154748	163230	154748	167730
# of starts	5	5	5	9
computational time (s)	61.72	1753.56	7.40	1451.00
mismatch (kWh)	3099.5		478.5	
mismatch/rough prod. (%)	7.8		1.2	
fleet production (kWh)	37927		38868.5	

Table 6.1: General energy planning problem results for the first case study

Table 6.1 shows the results for the general energy planning problem, where we incorporate the fleet of 5000 houses. The table shows the operational costs of the power plants, both for the initial planning with the rough fleet constraints (rough) and for the final result after applying the pattern generation process to the fleet and replanning the power plants using the elaborated fleet pattern (detailed). It also shows the number of starts for the power plants, again for the rough planning and for the final result. The computational time of the detailed result includes the computational time of the rough planning. Regarding the fleet planning, the final mismatch of the planning of the 5000 houses to the production plan f in the rough planning (i.e. absolute deviation from f) is given in kWh and in percentage of the total generation of this production pattern f . The resulting total generation of the 5000 houses is given in the last row of Table 6.1. Since we mostly use artificial prices for the electricity market, we do not show the profit maximization of the fleet in more detail.

The operational costs of the detailed planning are relatively close to the rough planning operational costs in all cases. This means that the commitment of the power plants is not altered too drastically after the elaborated planning of the fleet. Of course the final costs are higher, since the rough planning gives the optimal combination of power plant operation and fleet operation, whereby some constraints of the fleet are relaxed. The four use cases show that we are able to steer the fleet production by using different prices. The number of starts of the power plants increases in all cases, except for the third use case. In case 3 the power plants need only 5 starts in the final fleet planning. This is mostly due to the initial fleet

planning, which is aimed to reduce the peaks in the demand. In the realization of this planning the fleet has relatively much difficulties, since the mismatch from the rough planning is the highest of the four cases. Nevertheless this realization leaves enough possibilities for the power plants to find a planning that only needs 5 commitments. The big advantage is that the fleet does not interfere too much with the base load, which simplifies the continuity of the commitment in time intervals with low demand.

The computational time of the planning method stays below half an hour in all cases. This is acceptable for a practical application, especially since the column generation technique can be distributed over the smart grid in real life. The mismatch from the rough planning is below 8%. Figure 6.6 shows the development of the




Figure 6.6: The mismatch during the column generation for the four use cases

mismatch during the column generation method. The final solution is found after approximately 10 iterations. The existence of a mismatch can be partly explained by the relaxations in the rough planning, but also by the fact that we used a maximum runtime of 60 seconds for the pattern matching problem, which tries to select exactly one pattern for each house to minimize the mismatch from the rough planning. This maximum runtime results in a preliminary abortion of the solution method in all four cases. From this we may conclude that the fleet even may perform better, to the costs of higher computational time. Finally we see that the total fleet production approximates the upper production bound of 39.8 MWh in use cases 1, 3 and 4, whereas the total fleet production in case 2 is close to the lower production bound of 35.1 MWh. This indicates that the prices in case 2 are too negative, such that it is more cost effective to let the power plants produce more.

Figure 6.7 shows the unit commitment of the 10 power plants and the fleet production of the 5000 houses in the second use case. Figure 6.7a gives the rough fleet planning and Figure 6.7b gives the resulting, final fleet planning. We present the results of this use case in more detail to describe two phenomena that occur. Firstly, we see three short periods of commitment of the small power plants 1 and 3




Figure 6.7: The second use case in more detail (for a legend, see Figure 6.5)

in the final planning. These commitments are not necessary to fulfill the demand; the already committed power plants could have supplied this additional demand themselves, even against lower costs. However, in that case there would not be sufficient spinning reserve left in the committed power plants. For this reason, power plants 1 and 3 have to be committed during these periods. Secondly, we see many different generators committed in the final planning, in comparison with the rough planning. However, as we have stated before, below certain production levels (625 kWh for the large power plants and 312.5 kWh for the small power plants) the cost functions are equal, and power plants are mutually exchangeable. This could explain the larger number of committed generators.




Figure 6.8: Comparison of rough planning and final found solution for the planning of the local generators in the second use case

Finally, we present the rough fleet planning and the final fleet planning in one overview in Figure 6.8. This figure shows that the rough planning is matched to a large extent.

As a final part of this case study we extend the two leveled approach with an additional level in between the microCHP level and the highest level. We define group sizes of 5000, 1000, 500 and 100, which corresponds to 1, 5, 10 and 50 different groups. The desired production for each group \hat{f}^l is set to be equal to the total desired production for the fleet divided by the number of groups. Table 6.2 shows the results for these group sizes for the first case. The mismatch is the smallest

# fleets	total mismatch	iterations	time (s)
1	2237	11	1543.10
5	1582	31	10773.52
10	1896	35	16457.06
50	2561	20	11892.13

Table 6.2: Solving the microCHP planning problem using an intermediate level

when 5 groups are used, which corresponds to group sizes of 1000 microCHPs. The increase in computational effort for an increasing number of fleets is related to the number of times that an ILP problem is solved that needs to combine patterns and minimize mismatch. This type of ILP problem, that combines patterns to create new ones on a higher level, consumes the largest part of the computational time. Note that in practice, such problems can be solved in parallel, since each subgroup can be regarded as a stand alone entity.

6.4.2 CASE STUDY 2

The second case study builds upon the first one. It extends this case study by adding other types of distributed generation, distributed storage and load side management. In more detail, next to the power plants and the microCHP appliances of the previous section, we consider heat pumps, electrical vehicles, controllable freezers and storage capacities in the form of large batteries in households. This case study is intended to show the relation between microCHPs and heat pumps, to see whether a large share of heat pumps related to microCHPs has its impact on the profit maximization of the Virtual Power Plant. Next to this, we want to indicate whether the mismatch between the rough planning and the final planning as we have noticed in the previous case study can be compensated by using demand side load management and distributed storage. The electrical cars are added in one use case to see whether it is possible to allow a large share of electrical cars in a neighbourhood, without this leading to large peaks in electricity consumption. Also, the case study is used to show the potential of the proposed methods.

We use an equivalent framework as in the first case study to perform our evaluation, whereby we focus on 5000 houses, grouped in groups of 1000 houses, since the choice for this group size shows the smallest mismatch in the first case study. Opposite to having only local flexibility via the generation by a microCHP, we now have different options to shift with the generation and/or consumption of electricity.

We have chosen to use different forms of contribution to solve the general energy planning problem for the above options. These contributions differ in the desired operation for the group of each type of appliance. The operation of the set of microCHP appliances, the set of heat pumps and the set of electrical cars is required to follow a certain aggregated pattern, that is defined by solving a rough global planning problem using aggregated information on these appliances. The operation of the set of freezers and the set of batteries is not bound by globally defined cooperative patterns. Instead, their operation is completely free for each individual appliance, thereby reserving full potential to compensate for the mismatch that might be the result of the planning of the microCHPs, heat pumps and electrical cars.

The heat demand in this case study is equal to the heat demand in the first case study. This heat demand is either fulfilled by the production of a microCHP or by the production of a heat pump; we do not consider both appliances in the same house simultaneously. The total electricity demand of the group of 5000 houses is equal to the proposed electricity demand in the previous case study. This implies that the normal operation of the group of freezers is already part of this electricity demand. For this normal operation, different aggregated consumption patterns for the group of freezers are possible. From these patterns we have chosen the pattern, where the consumption of the total fleet of freezers is flattened as much as possible over the day.

To be able to investigate demand side management using the freezers, we need as starting point the normal operation of all individual freezers. The influence of demand side management is then determined by the change of the normal operation. For freezers the regularity in electricity consumption makes it relatively easy to derive a flat pattern as this normal operation. This aggregated pattern is formed in such a way that the freezer temperature at the end of the planning horizon is (almost) equal to the initial temperature. The creation of individual patterns that form this flat pattern is explained in the following. First we show the necessary changes for the switch from using 6 minutes intervals to 30 minutes intervals in the modelling description of the freezer planning. Then we give the details on the flat pattern that serves as the basic pattern for this case study.

As stated earlier, we use freezers with a power consumption of 150 W and allow binary decisions to cool or not to cool in 6 minutes intervals. To achieve a pattern in which initial and end temperatures are equal, we need on average to cool during $\frac{1}{6}$ th of the available time. Since the case study works with 30 minutes time intervals, we have to translate the proposed formulation (6.14)-(6.16) into a formulation that is valid for time intervals of half an hour and that accommodates operational flexibility in 6 minutes periods. This translation is done by introducing modulation in the operation of a freezer:

$$d_j^i \in \{0, 1, 2, 3, 4, 5\}, \quad (6.76)$$

where d_j^i represents the number of 6 minutes periods that freezer i is cooling in the 30 minutes time interval j . In addition to this change of (6.14) into (6.76), (6.15)

changes into:

$$\hat{T}_{\min} \leq T_{init}^i + j \times 5 \times \Delta T_{off} - \Delta T_{on} \times \sum_{k=1}^j d_k^i \leq \hat{T}_{\max} \quad \forall i \in I \forall j \in J, \quad (6.77)$$

and (6.16) remains unchanged. The bounds T_{\min} and T_{\max} are increased and decreased respectively, to accompany the shift from 6 minutes intervals to 30 minutes intervals:

$$\hat{T}_{\min} = T_{\min} + 4 \times \Delta T_{off} \quad (6.78)$$

$$\hat{T}_{\max} = T_{\max} - 4 \times \Delta T_{off}. \quad (6.79)$$

These new bounds are necessary to allow the cooling process to take place in any original 6 minutes interval without violating the original bounds T_{\min} and T_{\max} . The electricity consumption remains $f_j^i = FC^i d_j^i$, where we have $FC^i = 15$ Wh.

To create a flat pattern where initial and end temperatures remain unchanged, we should have an electricity consumption of $150 \times \frac{1}{6} \times \frac{1}{2} \times 5000 = 62500$ Wh in each time interval (150 W consumption in $\frac{1}{6}$ th of the time per half an hour interval times the number of freezers). Since 62500 cannot be divided by 15 (which is the multiplication factor of individual consumption), this flat pattern cannot be reached completely by the group of freezers. Therefore we change this average consumption slightly to the nearest multiple of 15, which is 62505. For the 5000 freezers we now propose the following operation patterns:

$$d_j^i = \begin{cases} 0 + \hat{d}_j^i & \text{if } (i + j) \bmod 6 = 0 \\ 1 + \hat{d}_j^i & \text{otherwise,} \end{cases} \quad (6.80)$$

where

$$\hat{d}_j^i = \begin{cases} 1 & \text{if } j = 1 + 3i \quad \forall i \in \{1, 3, \dots, 15\} \\ 1 & \text{if } j = -1 + 3i \quad \forall i \in \{2, 4, \dots, 16\} \\ 0 & \text{otherwise.} \end{cases} \quad (6.81)$$

This division of operation patterns follows the flat pattern exactly.

For the remaining available options (i.e. batteries, microCHPs, heat pumps and electrical cars) we do not define normal operational patterns. These options are not part of the normal electricity consumption pattern in this case study. They are either not desired to be used (batteries) or they have such a large power consumption/-generation parameters that we consider them as standalone entities (microCHPs, heat pumps and electrical cars).

Rough planning of the operation of microCHPs, heat pumps and electrical cars

In the first case study the operation of the group (fleet) of microCHPs is first roughly planned on an aggregated level, at which the generation level is of a similar order of

magnitude as that of the used power plants. In this case study we do the same for the combination of the fleet of microCHPs, the fleet of heat pumps and the fleet of electrical cars. Instead of determining generation possibilities f_{mchp} as is the case with microCHPs, in the case of heat pumps and electrical cars we have to determine electricity consumptions (f_{hp} and f_{ec}) for the fleet of heatpumps and electrical cars. The incentive for the fleet of microCHPs remains to make profit on an electricity market, while simultaneously the operational costs for the 10 power plants are minimized. For both the group of heat pumps and the group of electrical cars the incentive is to plan wisely to accommodate this simultaneous profit maximization and cost minimization. Therefore the objective function of the rough planning problem (6.40)-(6.53) remains unchanged, while we see a difference in constraint (6.41), where the electricity demand is increased (and/or decreased) by the use of heat pumps and electrical cars (note that discharging the battery of an electrical car is represented by negative values for $f_{j,ec}$):

$$\sum_i x_j^i + f_{j,mchp} \geq d_j + f_{j,hp} + f_{j,ec} \quad \forall j \in J. \quad (6.82)$$

Of course we also need to change the constraints that bound the total fleet production of the microCHPs and extend this formulation with the relevant requirements for the fleet consumption for the heat pumps and the electrical cars that are based on aggregated information. Let the set M_{mchp} denote the set of houses that are equipped with a microCHP, M_{hp} the set of houses with a heat pump and M_{ec} the set of electrical cars. The constraints that bound the fleet production for the microCHPs (constraints (6.51)-(6.53)) become:

$$2 \sum_{k=1}^j f_{k,mchp} \leq \sum_{m \in M_{mchp}} MaxOn_{m,j} \quad \forall j \in J \quad (6.83)$$

$$2 \sum_{k=1}^j f_{k,mchp} \geq \sum_{m \in M_{mchp}} MinOn_{m,j} \quad \forall j \in J \quad (6.84)$$

$$f_{j,mchp} \leq \frac{|M_{mchp}|}{2} \quad \forall j \in J. \quad (6.85)$$

For the fleet of heat pumps this is represented by the constraints:

$$\sum_{k=1}^j f_{k,hp} \leq \sum_{m \in M_{hp}} MaxOn_{m,j} \quad \forall j \in J \quad (6.86)$$

$$\sum_{k=1}^j f_{k,hp} \geq \sum_{m \in M_{hp}} MinOn_{m,j} \quad \forall j \in J \quad (6.87)$$

$$f_{j,hp} \leq |M_{hp}| \quad \forall j \in J. \quad (6.88)$$

We lose a factor of 2 compared to (6.83)-(6.85), since the *COP* of a heat pump is 4 which is half of the electricity to heat ratio $\alpha = 8$ for the microCHPs.

The constraints for the consumption of the electrical cars are based on the choice for the availability periods for the cars. We assume that the electrical cars arrive between 5 and 7 pm and leave between 5 and 7 am. Thus, they have to be charged overnight. For our one-day model this leads to a problem on how to define the initial levels of the battery. We request that overnight the cars have to be charged completely and that by day the cars consume an amount of electricity which is equal to 50 – 90% of the battery capacity, i.e. at the begin of the charging period the batteries have a load of 10 – 50%. Furthermore, we assume that each car has a repetitive behaviour (i.e. it shows the same arriving and departing behaviour for consecutive days). For the rough planning, we assume that half of the charging is done between arriving and the end of the day and the other half is done between the start of the day and departure. Using this information in a global setting, we pose the following constraints on the fleet consumption of the electrical cars, using f_{ec}^a to denote the consumption between arrival and the end of the planning horizon and f_{ec}^b to denote the consumption between the start of the planning horizon and departure. The maximum amount of electricity that a car can be charged per interval MC^m is set to 5 kWh.

$$f_{j,ec} = f_{j,ec}^a + f_{j,ec}^b \quad \forall j \in J \quad (6.89)$$

$$f_{j,ec}^a \leq \sum_{m \in M_{ec}} MC^m A_j^{m,a} \quad \forall j \in J \quad (6.90)$$

$$f_{j,ec}^a \geq - \sum_{m \in M_{ec}} MC^m A_j^{m,a} \quad \forall j \in J \quad (6.91)$$

$$\sum_{j \in J} f_{j,ec}^a = \sum_{m \in M_{ec}} \frac{CC^m - BBL^i}{2} \quad (6.92)$$

$$f_{j,ec}^b \leq \sum_{m \in M_{ec}} MC^m A_j^{m,b} \quad \forall j \in J \quad (6.93)$$

$$f_{j,ec}^b \geq - \sum_{m \in M_{ec}} MC^m A_j^{m,b} \quad \forall j \in J \quad (6.94)$$

$$\sum_{j \in J} f_{j,ec}^b = \sum_{m \in M_{ec}} \frac{CC^m - BBL^i}{2} \quad (6.95)$$

Compensating mismatch by using freezers and batteries

The rough planning of the operation of the group of microCHPs, heat pumps and electrical cars gives a planning that enables power plants, microCHPs, heat pumps and electrical cars to operate optimally in a combined setting. This sketch of the planning needs to be followed as good as possible by determining individual patterns for the different devices. This determination of individual patterns is done in a similar approach as shown in the previous case study, whereby the creation of new individual patterns is bounded by individual local constraints as shown in Section 6.1. Note that, although the process of matching the sum of individual patterns to the fleet patterns is executed sequentially in the calculations that are presented (since we have no network version of the solver), in practice this process

could be executed in parallel. After the planning of the microCHPs, the heat pumps and the electrical cars, the resulting mismatch from the rough planning is input for the planning of the group of freezers and the in-home batteries.

The normal freezer operation as proposed earlier in this section allows for a certain flexibility. Since the normal operation follows a flat pattern of exactly 62.505 kWh per interval, the maximal deviation from this pattern for each interval is bounded by the range $[-62.505, 312.495]$ (in kWh), since all freezers can maximally consume 375 kWh combined in one interval. Negative values indicate less consumption than in the basic flat pattern, positive values indicate more consumption. For the inhome batteries we define batteries with capacity $CC^m = 20$ kWh, maximum amount of (dis)charged electricity $MC^m = 2$ kWh per time interval and initial level $BBI^m = 10$ kWh. For these batteries the level at the end of the day may vary between 8 and 12 kWh, since we do not want to postpone difficulties in matching demand and supply to planning problems for consecutive days. Besides these constraints over the complete time horizon, the use of batteries offers far more flexibility in the amount of electricity that can be compensated per time interval: this compensation is bounded by the range $[-10000, 10000]$ in kWh. Positive values indicate consumption (charging the battery), negative values indicate supply (discharging the battery). Although the battery offers more flexibility, we prefer compensation by the group of freezers over the group of batteries, due to the limited amount of charge cycles that a battery can have. Therefore we penalize the use of batteries by a small cost.

Setup of scenarios in the case study

We study different setups for the 5000 houses in this case study. In these scenarios the APX prices are equal to the ones that are used in the first scenario of the previous case study. Furthermore, all houses are equipped with a freezer and a battery.

We consider a particularly cold day in winter. In this situation the houses have a significant heat demand, which is to be supplied by use of either a microCHP or a heat pump. The heat demand is equal to the heat demand in the previous case study. Table 6.3 shows the division of the number of microCHPs, heat pumps, electrical cars, freezers and batteries over the different scenarios.

scenario	1	2	3	4	5	6	7
# of power plants	10	10	10	10	10	10	10
# of microCHPs	5000	4000	3000	2000	1000	0	3000
# of heat pumps	0	1000	2000	3000	4000	5000	2000
# of electrical cars	0	0	0	0	0	0	1000
# of freezers	5000	5000	5000	5000	5000	5000	5000
# of batteries	5000	5000	5000	5000	5000	5000	5000

Table 6.3: Description of the scenarios in the second case study

Computational limits

The different optimization problems that are solved in the general energy planning problem are bounded in their maximal computational time. This leads to upper bounds on the computational time, both in case of sequential and in case of parallel computation. Here we give these upper bounds.

In the rough planning problem, as well as in the final planning problem, we limit the computational time to 60 seconds. As stated before we use fleets of 1000 houses in our hierarchical structure. Each pattern combination problem for these fleets, as well as the pattern combination problem that uses fleet patterns (on a higher level), also has a maximal computational time of 60 seconds. The creation of patterns (of individual microCHPs, heat pumps, electrical cars, freezers and batteries) at individual houses is limited by a maximal time of 1 second. In relation to these time limits, the practical communication time that is needed for exchanging information is neglected. The upper bound in case of sequential computation, related to the number of iterations k is then $60 + k \times (5000 \times 1 + 5 \times 60 + 60) + 5000 \times 1 + 60 = 5120 + 5360 \times k$, whereas the upper bound in case of parallel computation is $60 + k \times (1 + 60 + 60) + 1 + 60 = 121 + 121 \times k$.

Results and discussion

The different scenarios are implemented in AIMMS and solved by using CPLEX 12.2. Table 6.4 shows a summary of the important results of these scenarios. For the rough planning we show the operational costs (including startup costs) of the 10 power plants and the number of starts in the planning horizon. The total profit of the fleet of microCHPs is also depicted, as well as the computational time of this rough planning. After the rough planning is known, a detailed planning for individual appliances is found. For these detailed patterns the total mismatch from the rough planning is shown for the different types of appliances. The percentual mismatch compared to the total amount of roughly planned electricity production/consumption and the number of iterations that is used to find the smallest mismatch are also given. Next to this, the profit for the group of microCHPs after the detailed planning is shown. The computational times show the total elapsed time since the start of the planning process, including the time that is needed to find the rough planning. The last three rows in the table give the absolute amount of electricity that is planned for the freezers and inhome batteries to compensate for the total mismatch in each scenario. This value is found by using 1 planning iteration and the additional time that is needed for this leads to the total time depicted in the last row of the table.

The first scenario shows that the freezers and batteries cooperate in compensating the mismatch of the microCHPs. In one time interval the group of freezers and the group of batteries act in opposite directions (note that this only occurs in a few individual freezer patterns with minor total influence). Scenario 6 shows that the freezers and batteries can also compensate the mismatch of 5000 heat pumps completely, although this requests a larger compensation than in the first scenario and leads to more opposite behaviour of the freezers and the batteries. In scenarios

scenario		1	2	3	4	5	6	7	
rough	operational costs power plants	158748	198558	242586	287635	332666	377794	312586	
	# of starts power plants	5	6	6	7	8	9	6	
	profit maximization microCHPs (€)	6871.42	5940.25	4506.71	3011.28	1504.78	-	4506.71	
	profit (€/microCHP)	1.37	1.49	1.50	1.51	1.50	-	1.50	
	computational time (s)	61.49	60.92	10.13	14.54	22.03	7.16	4.93	
detailed	profit maximization microCHPs (€)	6398.00	5464.44	4222.11	2833.80	1408.13	-	4236.15	
	profit (€/microCHP)	1.28	1.37	1.41	1.42	1.41	-	1.41	
	microCHP mismatch (kWh)	1581.5	3066	1923	1287.5	806	-	1857	
	microCHP mismatch (%)	4.0	9.6	8.0	8.1	10.1	-	7.8	
	microCHP iterations	31	27	24	17	15	-	25	
	heat pump mismatch (kWh)	-	984	986	1025	2235	2744	2055	
	heat pump mismatch (%)	-	7.0	3.5	2.4	4.0	3.9	7.3	
	heat pump iterations	-	17	24	14	17	20	16	
	electrical car mismatch (kWh)	-	-	-	-	-	-	-	2921.167
	electrical car mismatch (%)	-	-	-	-	-	-	-	7.6
	electrical car iterations	-	-	-	-	-	-	-	16
computational time (s)	10773.52	6348.28	6695.47	4768.52	5456.48	7375.12	6595.46		
final	freezer compensation (kWh)	387.495	1500	1012	949.95	899.97	774.9	1999.995	
	battery compensation (kWh)	1194.065	2232	1426	1304.31	2215.85	2534.84	3734	
	computational time (s)	11101.59	6850.96	7155.94	5188.82	5808.47	8389.95	7055.46	

Table 6.4: General energy planning problem results for the second case study

2-5 and in scenario 7 the mismatch of the microCHP fleet planning and the heat pump fleet planning cancel each other out to some extent. This can be seen by comparing the sum of the mismatch for the microCHPs and the heat pumps to the sum of the compensation that is offered by the freezers and the batteries. This second sum is smaller than the first one, although a complete compensation is achieved: this indicates that part of the compensation is already offered by the microCHPs and heat pumps themselves. The remaining part, which is still the largest fraction of the overall total mismatch, is completely compensated by using freezers and batteries in all scenarios. This implies that the operation of the power plants as it was optimized in the rough planning does not need to be changed: the operational costs remain the same.




Figure 6.9: Operational costs related to additional electricity consumption

When the number of heat pumps increases the electricity consumption increases almost linearly, since the different households have heat demands that are

comparable. This increase in electricity consumption shows a linear increase in the operational costs of the power plants in Figure 6.9. In the case of the addition of electrical cars next to the usage of a set of heat pumps we see that the operational costs do not follow this linear trend. This is due to the flexibility that the electrical cars offer in the ability to increase the baseload that occurs in the night (see Figure 6.10). The ratio between peak demand and base demand is decreased, which leads to a demand that can be supplied more cost effective by the operation of the power plants.

Regarding profit maximization of the microCHPs we see that heat pumps can be planned in such a way that they offer a situation in which both the profit of the group of microCHPs and the operational costs of the power plants can be optimized. Moreover, the addition of heat pumps next to microCHPs leads to a small increase in the profit per microCHP, due to the flexibility of the heat pumps, which relieves the problem of minimizing operational costs for the power plants. The profit in scenario 7 is equal to that of scenario 3. Although this equality might indicate that the microCHPs focus on their own optimization in the combined problem of profit maximization and cost minimization, the underlying fleet pattern for the microCHPs differs on quite some time intervals: 26% of the two patterns has no overlap. This difference in patterns, in combination with the different operational planning of the heat pumps and the electrical cars, also helps optimizing the cost minimization of the power plants. The detailed planning leads to a decrease in profit, which is comparable for all scenarios (about 94% of the profit of the rough planning is reached).

The mismatch in kWh for the different fleets needs a short explanation. For the microCHP fleets this mismatch is at most 10.1% of the total roughly planned production. Note that the lower bound on the mismatch is 0 in all scenarios, which would lead to the question why we are not able to reach this lower bound in these scenarios. There are three explanations for this. First, we now aim at a single desired pattern f_{mchp} , instead of a band which is bound by upper and lower desired patterns. This leaves no room for compensation; once we find a mismatch in our planning, the effect of this mismatch prolonges in the next time interval(s). Second, the rough planning often plans at full fleet capacity or at zero production, which are the extreme points of the rough planning. Finally, the computational time of the problem of combining patterns at an intermediate level to form an aggregated pattern for each fleet of 1000 elements has a time limitation of 60 seconds. Due to this limitation, the problem of minimizing the mismatch from the desired fleet planning is interrupted, also in case no improvement of the best mismatch is found. In all scenarios this situation occurred. It is therefore a good advise to use the mismatch between the actual planning and the rough planning to give feedback towards making a new rough planning in an iterative setting, where the found mismatch can be used to tighten the global bounds on electricity production. However, for now, the results of the final part of the planning process, where the freezers and batteries are planned, show that these appliances are able to ‘repair’ the mismatch completely. Since a new rough planning would possibly lead to worse operational schedules for the 10 power plants, we stick with the results after a single iteration in

finding a rough planning.

The computational times show that the scenarios are solved using the sequential implementation in a couple of hours (1.4 to 3.1 hours). The computational bottleneck is in solving the pattern combination problem, which minimizes the deviation from the desired fleet planning, and not in the creation of patterns for the appliances. The computational times for this problem are limited at 60 seconds. By using an additional level in the general energy planning problem we could introduce a pattern combination problem, that has fleet patterns as input. By allowing 1000 different fleets we have pattern combination problems of a similar size, for which we propose a time limitation of again 60 seconds. In this setting, eventually a set of 1000000 houses could be planned, since a large part of the computations would occur in parallel.


To get a better feeling for the resulting planning, we show the outcome of the planning process in a graphical way in Figure 6.10. This figure shows the rough planning and the detailed planning for the scenario that includes 10 small power plants, 3000 microCHPs, 2000 heat pumps, 1000 electrical cars, 5000 freezers and 5000 batteries. The bottom half of each figure gives the consumption; the upper half gives the production of electricity. Figure 6.10a shows that the electrical cars fill the gap between peak- and baseload in the rough planning. Also the operation of the fleet of heat pumps and the fleet of microCHPs is planned in such a way that production and consumption cancel each other out for a large part. This leads to a situation in which the operation of the power plants can be scheduled by applying almost constant operation levels. The detailed planning in Figure 6.10b shows that the operation of the power plants is unchanged. The influence of shifting the consumption of the freezers and the availability of batteries both on the consumption and the production side is visible, since the consumption equals the generation at all time intervals.

6.5 CONCLUSION


This chapter treats the general energy planning problem. This problem is an extension of the Unit Commitment Problem (UCP), which treats the commitment and economic dispatch of electricity generators. We add distributed generation, distributed storage and demand side management possibilities to this problem, thereby shifting the focus of this optimization problem towards the decentralization within the Smart Grid.

The general energy planning problem differs from the UCP in size and in objective. The amount of appliances that are operated within the problem increases significantly. The inclusion of a Virtual Power Plant shows that the problem is not only focusing on cost minimization for the central generators in the traditional UCP, but it also concentrates on other objectives (e.g. profit maximization for power plants or the cost effective use of local storage possibilities and demand side management).

We propose a hierarchical structure to solve the general energy planning prob-



(a) The rough planning



(b) The final planning

Figure 6.10: Comparison of rough planning and final found solution for the planning using 10 small power plants, 3000 microCHPs, 2000 heat pumps, 1000 electrical cars, 5000 freezers and 5000 batteries

lem, in which the different elements are solved by formulating different sub problems in different planning levels. The general framework consists of creating patterns for single entities/appliances, combining patterns for such appliances on higher levels into so-called aggregated patterns, and using these aggregated patterns to solve a global planning problem.

Two different case studies show the applicability of the method. The first case study treats a group of 5000 microCHPs in combination with 10 power plants, to show the direct combination of optimizing a VPP and solving a traditional UCP. The second case study adds heat pumps, electrical cars, freezers and inhome batteries as examples of other types of distributed generation, distributed storage and demand side management.

CONCLUSION

The electricity supply chain is changing, due to increasing awareness for sustainability and an improved energy efficiency. The traditional infrastructure where demand is supplied by centralized generation is subject to a transition towards a Smart Grid. In this Smart Grid, sustainable generation from renewable sources is accompanied by controllable distributed generation, distributed storage and demand side load management for intelligent electricity consumption. The transmission and distribution grid have to deal with increasing fluctuations in demand and supply. Since realtime balance between demand and supply is crucial in the electricity network, this increasing variability is undesirable.

Monitoring and controlling/managing this infrastructure increasingly depends on the ability to control distributed appliances for generation, consumption and storage. In the development of control methodologies, mathematical support, which consists of predicting demand, solving planning problems and controlling the Smart Grid in realtime, is of importance. In this thesis we study planning problems which are related to the Unit Commitment Problem: for a set of generators it has to be decided when and how much electricity to produce to match a certain demand over a time horizon. The planning problems that we investigate in this thesis are part of a control methodology for Smart Grids, called TRIANA, that is developed at the University of Twente.

7.1 CONTRIBUTION OF THIS THESIS

This thesis presents new planning problems in the emerging Smart Grid. It treats these planning problems from two perspectives. First and foremost, the case of a Virtual Power Plant, consisting of a collection of domestic combined heat and power generators (microCHPs), is proposed. For this Virtual Power Plant case

mathematical planning problems are formulated and solution approaches are developed.

The microCHP planning problem

These planning problems deal with the decisions that have to be made for the operation of microCHPs in households. These appliances need to fulfill the heat demand in the household, while globally the aggregated electricity output of the microCHPs needs to fulfill a desired production pattern. The operation of the microCHP itself is restricted to binary decisions to switch the appliance on or off. In the microCHP planning problem, the profit of the Virtual Power Plant on an electricity market is maximized and/or the total deviation from a desired aggregated electricity output is minimized.

Exact methods and heuristic methods

For the considered planning problem a mathematical description is given, which is shown to be \mathcal{NP} -complete in the strong sense, by reducing 3-PARTITION to the microCHP planning problem. Exact formulations by modelling the problem as an Integer Linear Programming or a dynamic programming model show that practical instances are indeed difficult to solve in limited computational time. We explore different heuristic methods to solve the microCHP planning problem. A local search method, based on the dynamic programming formulation, shows a large improvement in computational time; the deviation from the desired bounds however is not satisfying. A basic version of an Approximate Dynamic Programming method is used to estimate the outcome of the large dynamic programming structure. Finally, a column generation technique offers a nice framework to minimize the deviation from the desired aggregated electricity output. For simplified instances it is shown, based on a lower bound calculation, that this method can solve this deviation (close) to optimality. The difficulty in the microCHP planning problem is shown in the feasibility aspects of finding solutions that respect requirements on the energy efficient operation of the individual appliances, on the fulfillment of local heat demand, and on the total electricity output simultaneously. These heuristic methods are appropriate for a practical implementation in the context of scalability. The division in global aggregation/optimization problems and local optimization problems offers a framework that is scalable, in case of the local search method and the column generation approach. In case of the Approximate Dynamic Programming approach, simplifications in the treatment of the DP tree lead to scalability.

Uncertainty in a practical setting

Next to solving the microCHP planning problem, we shortly sketch the influence of demand uncertainty in the broader context of controlling this Virtual Power Plant in a real world setting. The TRIANA methodology is able to capture such uncertainty in realtime. Moreover we show that we can anticipate on this uncertainty

by reserving part of the total heat capacity, such that the original planned operation can be completely followed during realtime operation.

Acting on electricity markets

In a real world setting, the insight that we can derive from the lower bound calculation can help us in the practical situation of the Virtual Power Plant, in which we would like to act on an electricity market. In this case we want to guarantee that a desired aggregated pattern can be reached by the individual generators. In an exploratory phase, a sketch of the aggregated output can be found, using the lower bound calculation as a guideline. The actual planning of the individual microCHPs can be postponed until a rough sketch is found that satisfies the (profit maximization) objective of the owner of the Virtual Power Plant, and that has a promising lower bound.

To be able to sell the production of a Virtual Power Plant, we want to offer the planned production to a day ahead electricity market. We show methods to construct bids for two auction mechanisms on the day ahead electricity market. In comparison with existing approaches, our bid construction has the special form of having limited flexibility in the variation of the quantity-to-offer combined with the requirement of a minimum probability of winning the auction; bids are constructed in the absence of a cost function for the VPP.

For the auction mechanism uniform pricing, the bid construction is given by a unique bid for positive market prices and (possibly) an additional bid for negative prices, in case the probability of winning the auction cannot be satisfied with the first bid.

For the auction mechanism pricing as bid, the bid construction is given by successive bids (p_t, q_t) , for which the quantity q_t increases with the minimum required difference of 0.1 MWh and the price p_t is based on the predicted values for the market clearing price μ (mean price), σ (standard deviation of the price) and a coefficient a_t , such that $p_t = \mu + a_t\sigma$. The values of the different coefficients a_t are optimized for a given range of the fraction $\frac{\mu}{\sigma}$. Application of this form of bid construction to real world data shows that 88% of the market clearing price can be reached as average settlement price, when at most 5 different bids are used.

The general energy planning problem

The second problem treated in this thesis concerns the general energy planning problem. This problem is an extension of the Unit Commitment Problem (UCP), which treats the commitment and economic dispatch of electricity generators. We add distributed generation, distributed storage and demand side management possibilities to this problem, thereby shifting the focus of this optimization problem towards the decentralization within the Smart Grid.

The general energy planning problem differs from the UCP in size and in objective. The amount of appliances that are operated within the problem increases significantly. The inclusion of a Virtual Power Plant shows that the problem is not

only focusing on cost minimization for the central generators in the traditional UCP, but it also concentrates on other objectives (e.g. profit maximization for power plants or the cost effective use of local storage possibilities and demand side management).

We propose a hierarchical structure to solve the general energy planning problem, in which the different elements are solved by using different sub problems in different planning levels. The general framework consists of creating patterns for single entities/appliances, combining patterns for such appliances on higher levels into so-called aggregated patterns, and using these aggregated patterns to solve a global planning problem.

Two different case studies show the applicability of the method. The first case study treats a group of 5000 microCHPs in combination with 10 power plants, to show the direct combination of optimizing a VPP and solving a traditional UCP. The second case study adds heat pumps, electrical cars, freezers and inhome batteries as examples of other types of distributed generation, distributed storage and demand side management.

7.2 POSSIBILITIES FOR FUTURE RESEARCH

For the Virtual Power Plant case we propose a planning problem and explore corresponding solution techniques, which give an indication of the real world problem instances that we should be able to solve. To accomplish the step towards solving a practical realization of the electricity supply chain, we need to find production profiles for a VPP by using a framework that is open to different kinds of interaction between households, the operator of a VPP, electricity markets, etcetera. The base for this framework is given by the microCHP planning problem, which could be solved iteratively to include additional requirements from the real world. An example of such an additional requirement treats the coupling between different days, in contrast with the time horizon of one day that we use in this work (in this context we may think of a variable granularity of the time intervals).

From a mathematical point of view, the solution methods of the microCHP planning problem may be studied in more detail. The different approaches that are studied in this thesis might be open for improvements. An extra challenge may be introduced by including stochastic aspects in the problem formulation, which requires advanced methods to cope with demand and price uncertainty.

Regarding the general energy planning problem, we have sketched a framework, by which combined problems are solved in the energy supply chain. Of course this problem is open for different combinations of objectives and to include different elements of a Smart Grid. A real world implementation requires cooperation between different end-users. The general energy planning problem can be a helpful tool in such situations, as long as we continue to model the possible extensions to this problem by small-scale, tractable problems.

CREATION OF HEAT DEMAND DATA

To represent the heat demand, sets of heat vectors $S^s = \{heat_1^s, \dots, heat_N^s\}$, are created, with each vector $heat_i^s$ characterizing the heat demand for the 24 hours of the day. If in some scenario a higher resolution of the time intervals is used, we simply downscale these hourly values (no interpolation between hours is done). The hourly vectors are generated as in Algorithm 5, with $w = 4$ and $s = 0$. In a winter day the average daily heat demand is assumed to be 54 kWh, which is typical for a cold day in winter in The Netherlands. This equals the average total heat demand, if the hourly demand (in Wh) is picked uniformly from the interval $I = [500, 4000]$. The function $create_heat(p, I)$ creates a heat demand based on the random number p and the interval I (using a uniform distribution over the interval I). Using the decimal development of π in samples of four digits length, semi-random numbers are used to produce 24 hourly heat demands. These 24 values are then ordered semi-randomly. The goal of this ordering is to achieve heat profiles which are somehow close to a real profile. We specify this by creating a profile with two peaks, one in the morning hours (between 7 a.m. and 11 a.m.) and one in the evening hours (between 6 p.m. and 10 p.m.). To create these peaks, two hours are selected semi-randomly from the corresponding sets $\{7, 8, 9, 10\}$ and $\{18, 19, 20, 21\}$ and the two highest of the 24 values are assigned to these selected hours. The function $match(p, F)$ selects the hour from the set F , where the probability p is used to pick the hour in the set F following a uniform distribution. Then, the remaining values $T_{unassigned}^i$ are, in decreasing order, assigned to an hour, where this hour can only be chosen from the set of unassigned hours that differs one from an already assigned hour (i.e. the set $freeplaces$ that is created by the function $freeplaces(heat_i^s)$, which assigns immediate neighbouring hours of already assigned hours to a set). The semi-random numbers and the function $match(p, F)$ are used again to decide which of these hours in $freeplaces$ is assigned the given value.

Algorithm 5 Heat demand creation for N microCHPs

Input: decimal development of π , starting from a position s , a width w , and an interval I

$pirest \leftarrow \pi \times 10^s - \lfloor \pi \times 10^s \rfloor$

for $i=1$ to N **do**

for $j=1$ to 24 **do**

$tempheat_{i,j} \leftarrow create_heat(\lfloor pirest \times 10^w \rfloor, I)$

$pirest \leftarrow pirest \times 10^w - \lfloor pirest \times 10^w \rfloor$

end for

$tempheat_i \leftarrow sort_non_increasing(tempheat_i)$

end for

for $i=1$ to N **do**

for $j = 1$ to 24 **do**

$p_{i,j} \leftarrow \frac{\lfloor pirest \times 10^w \rfloor}{10^w}$

$pirest \leftarrow pirest \times 10^w - \lfloor pirest \times 10^w \rfloor$

end for

end for

for $i=1$ to N **do**

$T_{unassigned}^i \leftarrow tempheat_i$

$freeplaces \leftarrow \{7, 8, 9, 10\}$

$heat_{i,match(p_{i,1},freeplaces)}^s \leftarrow first(T_{unassigned}^i)$

$freeplaces \leftarrow \{18, 19, 20, 21\}$

$heat_{i,match(p_{i,2},freeplaces)}^s \leftarrow first(T_{unassigned}^i)$

$k \leftarrow 3$

while $T_{unassigned}^i \neq \emptyset$ **do**

$freeplaces \leftarrow freeplaces(heat_i^s)$

$heat_{i,match(p_{i,k},freeplaces)}^s \leftarrow first(T_{unassigned}^i)$

$k \leftarrow k + 1$

end while

end for

BIBLIOGRAPHY

- [1] <http://www.aimms.com>.
- [2] <http://www.apxendex.com>.
- [3] <http://www.cbs.nl>.
- [4] <http://www.tsp.gatech.edu/concorde/index.html>.
- [5] <http://www.fchea.org/>.
- [6] <http://www.akira.ruc.dk/~keld/research/LKH/>.
- [7] Personal communication with Matthieu Bosman.
- [8] <http://www.claymath.org/millennium/>.
- [9] Specifications of the Nedap PowerRouter Battery Manager.
- [10] <http://www.netbeheernederland.nl>.
- [11] Personal communication with Stefan Nykamp of RWE Westfalen-Weser-Ems Netzservice GmbH.
- [12] Personal communication with Rob Bosman.
- [13] <http://www.tennet.org/>.
- [14] <http://comopt.ifi.uni-heidelberg.de/software/TSPLIB95/>.
- [15] <http://www.whispergen-europe.com/>.
- [16] Kari Alanne and Arto Saari. Sustainable small-scale chp technologies for buildings: the basis for multi-perspective decision-making. *Renewable and Sustainable Energy Reviews*, 8(5):401–431, 2004. ISSN 1364-0321.
- [17] N. Amjady and F. Keynia. Day-ahead price forecasting of electricity markets by mutual information technique and cascaded neuro-evolutionary algorithm. In *IEEE Transactions on Power Systems*, volume 24, pages 306 – 318, February 2009.
- [18] E. J. Anderson and A. B. Philpott. Optimal offer construction in electricity markets. *Mathematics of Operations Research*, 27(1):82–100, February 2002.

- [19] D. Applegate, R. Bixby, V. Chvátal, and W. Cook. Implementing the dantzig-fulkerson-johnson algorithm for large traveling salesman problems. *Mathematical Programming*, 97:91–153, 2003.
- [20] APXINDEX. Apx power nl. In *APX Power NL data sheet 2011*, pages 1 – 2, 2011.
- [21] C. L. Archer and M. Z. Jacobson. Spatial and temporal distributions of u.s. winds and wind power at 80m derived from measurements. *Journal of Geophysical Research*, 108(D9), 2003.
- [22] I. Arsie, V. Marano, G. Rizzo, and M. Moran. Integration of wind turbines with compressed air energy storage. In *Proceedings of the Second Global Conference on Power Control and Optimization (PCO)*, pages 11–18, 2009. Keynote lecture.
- [23] L.M. Ausubel and P. Cramton. Using forward markets to improve electricity market design. *Utilities Policy*, 18(4):195 – 200, 2010.
- [24] L. M. Ayompe, A. Duffy, S. J. McCormack, and M. Conlon. Validated real-time energy models for small-scale grid-connected pv-systems. *Energy*, In Press, Corrected Proof, 2010. ISSN 0360-5442.
- [25] A. Azadeh, M.R. Skandari, and B. Maleki-Shoja. An integrated ant colony optimization approach to compare strategies of clearing market in electricity markets: Agent-based simulation. *Energy Policy*, 38(10):6307 – 6319, 2010.
- [26] B. Baarsma, S. Bremer, M. De Nooij, and J. Poort. Toetredingsdrempels kleinverbruikersmarkt energie. In *SEO Economisch Onderzoek*, pages 1 – 60, 2007.
- [27] R. J. Bagby. Calculating normal probabilities. *The American Mathematical Monthly*, 102(1):46–49, January 1995.
- [28] A. Bakirtzis, V. Petridis, and S. Kazarlis. Genetic algorithm solution to the economic dispatch problem. *Generation, Transmission and Distribution, IEE Proceedings-*, 141(4):377 –382, July 1994. ISSN 1350-2360.
- [29] V. Bakker. *Triana: a control strategy for Smart Grids: Forecasting, planning & real-time control*. PhD thesis, University of Twente, Enschede, January 2012.
- [30] V. Bakker, A. Molderink, M. G. C. Bosman, J. L. Hurink, and G. J. M. Smit. On simulating the effect on the energy efficiency of smart grid technologies. In *Proceedings of the 2010 Winter Simulation Conference, Baltimore, MD, USA*, pages 393–404, USA, December 2010. IEEE Computer Society.
- [31] M. Beccali, M. Cellura, V. Lo Brano, and A. Marvuglia. Short-term prediction of household electricity consumption: Assessing weather sensitivity in a mediterranean area. *Renewable and Sustainable Energy Reviews*, 12:2040 – 2065, 2008.
- [32] R. Bellman. Dynamic programming treatment of the travelling salesman problem. *J. ACM*, 9:61–63, January 1962.

- [33] G. Berdichevsky, K. Kelty, J.B. Straubel, and E. Toomre. The tesla roadster battery system. In *Tesla Motors*, page 5, 2007.
- [34] D. Bienstock and S. Mattia. Using mixed-integer programming to solve power grid blackout problems. *Discrete Optimization*, 4(1):115 – 141, 2007. ISSN 1572-5286.
- [35] C. Block, D. Neumann, and C. Weinhardt. A market mechanism for energy allocation in microchp grids. In *Proceedings of the 41st Hawaii International Conference on System Sciences*, pages 1 – 11, 2008.
- [36] F. Boffa, V. Pingali, and D. Vannoni. Increasing market interconnection: An analysis of the italian electricity spot market. *International Journal of Industrial Organization*, 28(3):311 – 322, 2010.
- [37] M. G. C. Bosman, V. Bakker, A. Molderink, J. L. Hurink, and G. J. M. Smit. Load control in low voltage level of the electricity grid using microchp appliances. In *Proceedings of the Nineteenth Annual Workshop on Circuits, Systems and Signal Processing (ProRISC)*, pages 25–29. STW Technology Foundation, 2008. ISBN 978-90-73461-56-7.
- [38] M. G. C. Bosman, A. Molderink, V. Bakker, G. J. M. Smit, and J. L. Hurink. Multilevel unit commitment in smart grids. In *Proceedings of the First International Conference on Operations Research and Enterprise Systems*, pages 361–370, February 2012.
- [39] P. Carlsson. *Algorithms for Electronic Power Markets*. PhD thesis, Uppsala University, 2004.
- [40] Claus C. Caroe and Ruediger Schultz. A two-stage stochastic program for unit commitment under uncertainty in a hydro-thermal power system. In *Preprint SC 98-11, Konrad-Zuse-Zentrum fuer Informationstechnik*, pages 1–17, 1998.
- [41] G. Casolino, G. Liuzzi, and A. Losi. Unit commitment in oligopolistic markets by nonlinear mized variable programming. *Optimization and Engineering*, 11(2): 213 – 245, 2010.
- [42] Santiago Cerisola, Alvaro Baillo, Jose M. Fernandez-Lopez, Andres Ramos, and Ralf Gollmer. Stochastic power generation unit commitment in electricity markets: A novel formulation and a comparison of solution methods. *Operations Research*, 57(1):32–46, 2009.
- [43] D.M. Clucas and J.K. Raine. Development of a hermetically sealed stirling engine battery charger. In *Proceedings of the Institution of Mechanical Engineers*, pages 357 – 366, 1994.
- [44] A.J. Conejo, J. Contreras, R. Espínola, and M.A. Plazas. Forecasting electricity prices for a day-ahead pool-based electric energy market. *International Journal of Forecasting*, 21(3):435 – 462, 2005. ISSN 0169-2070.

- [45] S. A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the third annual ACM symposium on Theory of computing*, pages 151–158. ACM, 1971.
- [46] L.M. Costa and G. Kariniotakis. A stochastic dynamic programming model for optimal use of local energy resources in a market environment. In *Power Tech, 2007 IEEE Lausanne*, pages 449–454, July 2007.
- [47] P. Cramton and S. Stoft. Why we need to stick with uniform-price auctions in electricity markets. *The Electricity Journal*, 20(1):26–37, 2006.
- [48] G.S. Crawford, J. Crespo, and H. Tauchen. Bidding asymmetries in multi-unit auctions: Implications of bid function equilibria in the british spot market for electricity. *International Journal of Industrial Organization*, 25(6):1233–1268, 2007.
- [49] G. Dantzig, R. Fulkerson, and S. Johnson. Solution of a large-scale traveling salesman problem. *Journal of the Operations Research Society of America*, 2:393–410, November 1954.
- [50] A. de Jong, E.J. Bakker, J. Dam, and H. van Wolferen. Technisch energie- en CO_2 -besparingspotentieel van micro-wkk in nederland (2010-2030). In *Technical report, Werkgroep Decentraal*, pages 1–45, July 2006.
- [51] J.F. DeCarolis and D.W. Keith. The economics of large-scale wind power in a carbon constrained world. *Energy Policy*, 34:395–410, 2006.
- [52] DESERTEC. Clean power from deserts. In *Whitebook - 4th edition*, pages 1–65, 2009.
- [53] Dienst uitvoering en toezicht Energie. Informatie- en consultatiedocument ten behoeve van het advies van de aan het ministerie van economische zaken (ez) inzake virtual power plant veilingen voor de nederlandse groothandelsmarkt. In *Informatie- en Consultatiedocument VPP*, pages 1–38, 2005.
- [54] DISPOWER. Dispower - distributed generation with high penetration of renewable energy sources. pages 1–98. DISPOWER, 2006.
- [55] S.-O. Fridolfsson and T.P. Tangerås. Market power in the nordic electricity wholesale market: A survey of the empirical evidence. *Energy Policy*, 37(9):3681–3692, 2009.
- [56] M. R. Garey and D. S. Johnson. *Computers and intractability: a guide to the theory of NP-completeness*. W.H. Freeman and company, San Francisco, 1979.
- [57] P. Giesbertz and M. Mulder. Economics of interconnection: the case of the northwest european electricity market. In *The International Association for Energy Economics Energy Forum*, volume 17, Second Quarter, pages 17–21, 2008.

- [58] P. C. Gilmore and R. E. Gomory. A linear programming approach to the cutting-stock problem. *Operations Research*, 9(6):849–859, Nov. 1961.
- [59] J. P. Green, S. A. Smith, and G. Strbac. Evaluation of electricity distribution system design strategies. *IEEE Proceedings of Generation, Transmission and Distribution*, 146(1):53–60, 1999. ISSN 1350-2360.
- [60] Nicole Groewe-Kuska and Werner Roemisch. *Stochastic Unit Commitment in Hydrothermal Power Production Planning*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2005. ISBN 9780898718799.
- [61] A. D. Hawkes and M. A. Leach. Cost-effective operating strategy for residential micro-combined heat and power. *Energy*, 32(5):711 – 723, 2007.
- [62] A.D. Hawkes and M.A. Leach. On policy instruments for support of micro combined heat and power. *Energy Policy*, 36(8):2973 – 2982, 2008.
- [63] M. Held and R. M. Karp. A dynamic programming approach to sequencing problems. *Journal of the Society for Industrial and Applied Mathematics*, 10(1): 196–210, March 1962.
- [64] K. Helsgaun. An effective implementation of the lin-kernighan traveling salesman heuristic. *European Journal of Operational Research*, 126(1):106 – 130, 2000.
- [65] K. Helsgaun. General k-opt submoves for the lin-kernighan tsp heuristic. *Mathematical Programming Computation*, 1:119–163, 2009.
- [66] B.F. Hobbs, M.H. Rothkopf, R.P. O’Neill, and H. Chao, editors. *The Next Generation of Electric Power Unit Commitment Models*. Kluwer Academic Publishers, New York, Boston, Dordrecht, London, Moscow, 2001. ISBN 0-306-47663-0.
- [67] M.P.F. Hommelberg, B. Roossien, C.J. Warmer, J.K. Kok, F.J. Kuijper, and J.W. Turkstra. Geïntegreerde micro-wkk’s als virtuele centrale - first trial smart power system. In *openbare eindrapportage*, pages 1 – 27, 2007.
- [68] M. Houwing. *Smart Heat and Power - Utilizing the Flexibility of Micro Generation*. PhD thesis, Technische Universiteit Delft, 2010.
- [69] M. Houwing and I. Bouwmans. Agent-based modelling of residential energy generation with micro-chp. In *Proceedings of the 2nd international conference on integration of renewable and distributed energy resources*, pages 1 – 27. NAPA, 2006.
- [70] M. Houwing, P.W. Heijnen, and I. Bouwmans. Deciding on micro-chp; a multi-level decision-making approach. In *IEEE International Conference on Networking, Sensing and Control*, pages 302 – 307, 2006.
- [71] N. Ingvar and C. Persson. Vehicle to grid state of the art system design - plug in hybrid electric vehicles as source of power for the control markets in sweden and germany. In *Master thesis, KTH Stockholm*, pages 1 – 128, 2010.

- [72] International Energy Agency. Advanced local energy planning (alep): a guide-book. In Reinhard Jank, editor, *Technical report*. International Energy Agency, 2000.
- [73] A.G. Kagiannas, D.T. Askounis, and J. Psarras. Power generation planning: a survey from monopoly to competition. *International Journal of Electrical Power and Energy Systems*, 26(6):413 – 421, 2004. ISSN 0142-0615.
- [74] R.G. Karandikar, S.A. Khaparde, and S.V. Kulkarni. Strategic evaluation of bilateral contract for electricity retailer in restructured power market. *International Journal of Electrical Power and Energy Systems*, 32(5):457 – 463, 2010.
- [75] F. Karlsson. *Capacity Control of Residential Heat Pump Heating Systems*. PhD thesis, Chalmers University of Technology, 2007.
- [76] W. Kempton, J. Tomic, S. Letendre, A. Brooks, and T. Lipman. Vehicle-to-grid power: Battery, hybrid, and fuel cell vehicles as resources for distributed electric power in california. In *Institute of Transportation Studies*, pages 1 – 95, 2001.
- [77] R.H. Kerr, J.L. Scheidt, A.J. Fontanna, and J.K. Wiley. Unit commitment. *Power Apparatus and Systems, IEEE Transactions on*, PAS-85(5):417 – 421, May 1966. ISSN 0018-9510.
- [78] J. C. P. Kester, P. J. M. Heskes, J. J. Kaandorp, J. F. G. Cobben, G. Schoonenberg, D. Malyna, E. C. W. De Jong, B. J. Wargers, and A. J. F. Dalmeijer. A smart mv/lv-station that improves power quality, reliability and substation load profile. In *20th International Conference on Electricity Distribution*, 2009.
- [79] A. Kets, P.G.M. Boonekamp, and J. Jelsma. Afstemmen van vraag en aanbod door middel van gedragsmatig verschuiven. In *ECN report*, pages 1 – 52, 2002.
- [80] A.R. Kian, J.B. Cruz, and R.J. Thomas. Bidding strategies in oligopolistic dynamic electricity double-sided auctions. In *IEEE Transactions on Power Systems*, volume 20, pages 50 – 58, February 2005.
- [81] J. Kiviluoma and P. Meibom. Influence of wind power, plug-in electric vehicles, and heat storages on power system investments. *Energy*, 35:1244 – 1255, 2010.
- [82] J. K. Kok. Short-term economics of virtual power plants. In *20th International Conference on Electricity Distribution*, 2009.
- [83] J. K. Kok, C. J. Warmer, and I. G. Kamphuis. Powermatcher: multiagent control in the electricity infrastructure. In *AAMAS '05: Proceedings of the fourth international joint conference on Autonomous Agents and MultiAgent Systems*, pages 75–82, New York, NY, USA, 2005. ACM. ISBN 1-59593-093-0.
- [84] J. Lagorse, D. Paire, and A. Miraoui. A multi-agent system for energy management of distributed power sources. *Renewable Energy*, 35(1):174 – 182, 2010.

- [85] R. Lanzafame and M. Messina. Power curve control in micro wind turbine design. *Energy*, 35(2):556–561, 2010. ISSN 0360-5442. ECOS 2008, 21st International Conference, on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems.
- [86] J. Levine, G. Martin, R. Moutoux, and F. Barnes. Large scale electrical energy storage in colorado. In *CERI Research Report*, pages 1 – 49, 2007.
- [87] Lichtblick. Zuhausekraftwerk: die intelligenz des schwarms. Lichtblick ZuhauseKraftwerk GmbH, 2009.
- [88] M.G. Lijesen. The real-time price elasticity of electricity. *Energy Economics*, 29(2):249 – 258, 2007. ISSN 0140-9883.
- [89] S. Lin and B. W. Kernighan. An effective heuristic algorithm for the traveling-salesman problem. *Operations Research*, 21(2):498–516, March/April 1973.
- [90] R. Madlener, B. Glensk, and G. Westner. Optimization of e.on’s power generation with a special focus on renewables. In *E.ON Energy Research Center Series*, pages 1 – 27, 2010.
- [91] C. E. Miller, A. W. Tucker, and R. A. Zemlin. Integer programming formulation of traveling salesman problems. *J. ACM*, 7:326–329, October 1960.
- [92] B. Mills and J. Schleich. What’s driving energy efficient appliance label awareness and purchase propensity? *Energy Policy*, 38(2):814–825, 2010. ISSN 0301-4215.
- [93] D. Moest and M. Genoese. Market power in the german wholesale electricity market. *The Journal of Energy Markets*, 2(2):47 – 74, 2009.
- [94] A. Molderink. *On the three-step control methodology for Smart Grids*. PhD thesis, University of Twente, May 2011.
- [95] A. Molderink, M. G. C. Bosman, V. Bakker, J. L. Hurink, and G. J. M. Smit. Simulating the effect on the energy efficiency of smart grid technologies. In *Proceedings of the 2009 Winter Simulation Conference, Austin, Texas, USA*, pages 1530–1541, Los Alamitos, December 2009. IEEE Computer Society.
- [96] A. Molderink, V. Bakker, M. G. C. Bosman, J. L. Hurink, and G. J. M. Smit. Management and control of domestic smart grid technology. *IEEE Transactions on Smart Grid*, pages 109–119, 2010.
- [97] National Renewable Energy Laboratory. Getting started guide for homer version 2.1. In *Technical report*. National Renewable Energy Laboratory, 2005.
- [98] P.J. Neame, A.B. Philpott, and G. Pritchard. Offer stack optimization in electricity pool markets. *Operations Research*, 51(3):397–408, 2003.

- [99] V. Oikonomou, F. Becchis, L. Steg, and D. Russolillo. Energy saving and energy efficiency concepts for policy making. *Energy Policy*, 37(11):4787–4796, 2009. ISSN 0301-4215.
- [100] S. Oren. When is a pay-as bid preferable to uniform price in electricity markets. In *Power Systems Conference and Exposition, 2004. IEEE PES*, pages 1618 – 1620, October 2004.
- [101] Z. Ouyang and S.M. Shahidehpour. Heuristic multi-area unit commitment with economic dispatch. In *IEE Proceedings of Generation, Transmission and Distribution*, pages 242 – 252, 1991.
- [102] N.P. Padhy. Unit commitment-a bibliographical survey. *Power Systems, IEEE Transactions on*, 19(2):1196 – 1205, May 2004. ISSN 0885-8950.
- [103] A.D. Peacock and M. Newborough. Impact of micro-chp systems on domestic sector co2 emissions. *Applied Thermal Engineering*, 25(17):2653 – 2676, 2005.
- [104] N. Peerbocus and A. Melino. Export demand response in the ontario electricity market. *The Electricity Journal*, 20(9):55 – 64, 2007.
- [105] A. B. Philpott, M. Craddock, and H. Waterer. Hydro-electric unit commitment subject to uncertain demand. *European Journal of Operational Research*, 125(2): 410 – 424, 2000. ISSN 0377-2217.
- [106] M. Pomp and V. Shestalova. Switching costs in the netherlands energy markets: can liberalisation bring benefits to small customers? In *De Economist*, volume 155, pages 305 – 321, 2007.
- [107] W. B. Powell. Approximate dynamic programming: Lessons from the field. In *Winter Simulation Conference, 2008*, pages 205 –214, December 2008.
- [108] W.B. Powell. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. John Wiley & Sons, New Jersey, 2007.
- [109] S. Quoilin. Experimental study and modeling of a low temperature rankine cycle for small scale cogeneration. In *Master thesis, University of Liege*, pages 1 – 129, 2007.
- [110] M.H. Rothkopf. Control of market power in electricity auctions. *The Electricity Journal*, pages 15 – 24, 2002.
- [111] T. Schalm, J. Cordeweners, M. Noorlander, and E. Kahl. Monitor consumentenmarkt elektriciteit en gas 2010. In *Nederlandse Mededingingsautoriteit - Energiekamer*, pages 1 – 65, 2011.
- [112] C. Schulz, G. Roeder, and M. Kurrat. Virtual power plants with combined heat and power micro-units. In *International Conference on Future Power Systems*, pages 1 – 5, 2005.

- [113] J. Scott, P. Vaessen, and F. Verheij. Reflections on smart grids for the future. In *KEMA consultancy document for the Dutch Ministry of Economic Affairs*, pages 1 – 71, March 2008.
- [114] E.A. Setiawan. *Concept and Controllability of Virtual Power Plant*. PhD thesis, University of Kassel, 2007.
- [115] G.B. Sheble and G.N. Fahd. Unit commitment literature synopsis. *Power Systems, IEEE Transactions on*, 9(1):128 –135, February 1994. ISSN 0885-8950.
- [116] S. Soares, C. Lyra, and H. Tavares. Optimal generation scheduling of hydrothermal power systems. *IEEE Transactions on Power Apparatus and Systems*, 99(3): 1107 – 1118, 1980.
- [117] L.C. Spencer. A comprehensive review of small solar-powered heat engines: Part ii. research since 1950 - conventional engines up to 100 kw. *Solar Energy*, 43(4): 197 – 210, 1989. ISSN 0038-092X.
- [118] P. Stephenson and M. Paun. Electricity market trading. *Power Engineering*, pages 277 – 288, 2001.
- [119] S. Takriti, B. Krasenbrink, and L.S.-Y. Wu. Incorporating fuel constraints and electricity spot prices into the stochastic unit commitment problem. *Operations Research*, 48(2):268–280, 2000.
- [120] J.W. Taylor, L.M. De Menezes, and P.E. McSharry. A comparison of univariate methods for forecasting electricity demand up to a day ahead. *International Journal of Forecasting*, 22:1 – 16, 2006.
- [121] CSIRO National Research Flagships Energy Transformed. Intelligent grid: a value proposition for distributed energy in australia. pages 1 – 592. CSIRO, 2009.
- [122] United States Department of Energy. The micro-CHP technologies roadmap. *Results of the Micro-CHP Technologies Roadmap Workshop*, 2003.
- [123] M.A. Uytterlinde, M. Junginger, H.J. De Vries, A.P.C. Faaij, and W.C. Turkenburg. Implications of technological learning on the prospects for renewable energy technologies in europe. *Energy Policy*, 35:4072 – 4087, 2007.
- [124] M. Van den Akker, G. Bloemhof, J. Bosman, D. Crommelin, J. Frank, and G. Yang. Optimal distributed power generation under network-load constraints. In *Proceedings of the 72nd European Study Group Mathematics with Industry*, pages 25 – 38, 2010.
- [125] M. Ventosa, A. Baillo, A. Ramos, and M. Rivier. Electricity market modeling trends. *Energy Policy*, 33(7):897 – 913, 2005.

- [126] M. Vermeulen, M. Mulder, W. Van den Reek, G. Thomeer, and M. De Kleijn. Marktmonitor, ontwikkeling van de groothandelsmarkt voor elektriciteit in 2006. In *Nederlandse Mededingingsautoriteit - Directie Toezicht Energie*, pages 1 – 50, 2007.
- [127] M.A. Vinke. An approximate dynamic programming approach to the microchp planning problem. In *Master thesis (to appear), Universiteit Twente*, 2012.
- [128] A.P. Wemhoff and M.V. Frank. Predictions of energy savings in hvac systems by lumped models. *Energy and Buildings*, 42(10):1807 – 1814, 2010. ISSN 0378-7788.
- [129] S. Wharton and J.K. Lundquist. Atmospheric stability affects wind turbine power collection. *Environmental Research Letters*, 7:1 – 9, 2012.
- [130] G. J. Woeginger. Combinatorial optimization - eureka, you shrink! chapter Exact algorithms for NP-hard problems: a survey, pages 185–207. Springer-Verlag New York, Inc., New York, NY, USA, 2003.
- [131] A. Wright and S. Firth. The nature of domestic electricity-loads and effects of time averaging on statistics and on-site generation calculations. *Applied Energy*, 84(4):389–403, 2007.
- [132] Z. Xiaoyan. Double-sided auction mechanism design in electricity based on maximizing social welfare. *Energy Policy*, 37(11):4231 – 4239, 2009.
- [133] J.H. Yan and G.A. Stern. Simultaneous optimal auction and unit commitment for deregulated electricity markets. *The Electricity Journal*, pages 72 – 80, 2002.

LIST OF PUBLICATIONS

REFEREED

- [MB:1] V. Bakker, M. G. C. Bosman, A. Molderink, J. L. Hurink, and G. J. M. Smit. Demand side load management using a three step optimization methodology. In *First IEEE International Conference on Smart Grid Communications (SmartGridComm 2010)*, Gaithersburg, MD, USA, pages 431–436, USA, October 2010. IEEE Communications Society.
- [MB:2] V. Bakker, M. G. C. Bosman, A. Molderink, J. L. Hurink, and G. J. M. Smit. Improved heat demand prediction of individual households. In *Proceedings of the first Conference on Control Methodologies and Technology for Energy Efficiency*, Vilamoura, Portugal, pages 1–6, Oxford, March 2010. Elsevier Ltd.
- [MB:3] V. Bakker, A. Molderink, M. G. C. Bosman, J. L. Hurink, and G. J. M. Smit. On simulating the effect on the energy efficiency of smart grid technologies. In *Proceedings of the 2010 Winter Simulation Conference, Baltimore, MD, USA*, pages 393–404, USA, December 2010. IEEE Computer Society.
- [MB:4] M. G. C. Bosman, V. Bakker, A. Molderink, J. L. Hurink, and G. J. M. Smit. The microchp scheduling problem. In *Proceedings of the Second Global Conference on Power Control and Optimization, PCO 2009, Bali, Indonesia*, volume 1159 of *AIP Conference Proceedings*, pages 268–275, London, June 2009. Springer Verlag.
- [MB:5] M. G. C. Bosman, V. Bakker, A. Molderink, J. L. Hurink, and G. J. M. Smit. Benchmarking set for domestic smart grid management. In *Innovative Smart Grid Technologies Conference Europe, ISGT Europe 2010, Gothenburg, Sweden*, pages 1–8, USA, October 2010. IEEE Power & Energy Society.
- [MB:6] M. G. C. Bosman, V. Bakker, A. Molderink, J. L. Hurink, and G. J. M. Smit. On the microchp scheduling problem. In *Proceedings of the 3rd Global Conference on Power Control and Optimization PCO, 2010, Gold Coast, Australia*, volume 1239 of *AIP Conference Proceedings*, pages 367–374, Australia, February 2010. PCO.
- [MB:7] M. G. C. Bosman, V. Bakker, A. Molderink, J. L. Hurink, and G. J. M. Smit. Scheduling microchps in a group of houses. *Global Journal on Technology and Optimization*, 1(1): 30–37, June 2010. ISSN 0925-5001.
- [MB:8] M. G. C. Bosman, V. Bakker, A. Molderink, J. L. Hurink, and G. J. M. Smit. Controlling a group of microchps: planning and realization. In *First International Conference on*

Smart Grids, Green Communications and IT Energy-aware Technologies, ENERGY 2011, Venice, Italy, pages 179–184, Canada, May 2011. International Academy, Research, and Industry Association.

196

- [MB:9] M. G. C. Bosman, V. Bakker, A. Molderink, J. L. Hurink, and G. J. M. Smit. Planning the production of a fleet of domestic combined heat and power generators. *European Journal of Operational Research*, 216(1):140–151, January 2012. ISSN 0377-2217.
- [MB:10] M. G. C. Bosman, A. Molderink, V. Bakker, G. J. M. Smit, and J. L. Hurink. Multilevel unit commitment in smart grids. In *Proceedings of the First International Conference on Operations Research and Enterprise Systems*, pages 361–370, February 2012.
- [MB:11] A. Molderink, V. Bakker, M. G. C. Bosman, J. L. Hurink, and G. J. M. Smit. Domestic energy management methodology for optimizing efficiency in smart grids. In *IEEE Bucharest Power Tech Conference, Bucharest, Romania*, pages 1–7. IEEE, October 2009.
- [MB:12] A. Molderink, M. G. C. Bosman, V. Bakker, J. L. Hurink, and G. J. M. Smit. Hard- and software implementation and verification of an islanded house prototype. In K. J. Brunham and O. C. L. Haas, editors, *Proceedings of the 2009 International Conference on Systems Engineering, Coventry, UK*, pages 327–332, UK, September 2009. Coventry University.
- [MB:13] A. Molderink, M. G. C. Bosman, V. Bakker, J. L. Hurink, and G. J. M. Smit. Simulating the effect on the energy efficiency of smart grid technologies. In *Proceedings of the 2009 Winter Simulation Conference, Austin, Texas, USA*, pages 1530–1541, Los Alamitos, December 2009. IEEE Computer Society.
- [MB:14] 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. In *Proceedings of the 2010 IEEE Innovative Smart Grid Technologies Conference, Gathersburg, USA*, pages 1–8. IEEE, January 2010.
- [MB:15] A. Molderink, V. Bakker, M. G. C. Bosman, J. L. Hurink, and G. J. M. Smit. Improving stability and utilization of the electricity infrastructure of a neighbourhood. In *Proceedings of the 2010 IEEE Conference on Innovative Technologies for an Efficient and Reliable Electricity Supply, CITRES 2010, Waltham, Boston, USA*, pages 233–239, USA, September 2010. IEEE Power & Energy Society.
- [MB:16] A. Molderink, V. Bakker, M. G. C. Bosman, J. L. Hurink, and G. J. M. Smit. Management and control of domestic smart grid technology. *IEEE transactions on Smart Grid*, 1(2):109–119, September 2010. ISSN 1949-3053.
- [MB:17] A. Molderink, V. Bakker, M. G. C. Bosman, J. L. Hurink, and G. J. M. Smit. On the effects of mpc on a domestic energy efficiency optimization methodology. In *Proceedings of the 2010 IEEE International Energy Conference (EnergyCon 2010), Al Manamah, Bahrain*, page 1569327927, USA, December 2010. IEEE Power & Energy Society.

NON-REFEREED

- [MB:18] V. Bakker, A. Molderink, M. G. C. Bosman, J. L. Hurink, and G. J. M. Smit. Improved simulator to analyse the impact of distributed generation on the electricity grid. In *Proceedings of the Twentieth Annual Workshop on Circuits, Systems and Signal Processing (ProRISC), Veldhoven, The Netherlands*, pages 197–201, Utrecht, November 2009. Technology Foundation STW.
- [MB:19] M. G. C. Bosman, V. Bakker, A. Molderink, J. L. Hurink, and G. J. M. Smit. Load control in low voltage level of the electricity grid using μ chp appliances. In *Proceedings of the Nineteenth Annual Workshop on Circuits, Systems and Signal Processing (ProRISC), Veldhoven, The Netherlands*, pages 25–29, Utrecht, November 2008. Technology Foundation STW.
- [MB:20] M. G. C. Bosman, V. Bakker, A. Molderink, J. L. Hurink, and G. J. M. Smit. Production planning in a virtual power plant. In *Proceedings of the 20th annual workshop on Program for Research on Integrated Systems and Circuits, Veldhoven, Netherlands*, page 6, Utrecht, November 2009. Technology Foundation STW.
- [MB:21] M. G. C. Bosman, V. Bakker, A. Molderink, J. L. Hurink, and G. J. M. Smit. Applying column generation to the discrete fleet planning problem. In *Proceedings of the STW.ICT Conference 2010, Veldhoven, The Netherlands*, pages 200–204, Utrecht, The Netherlands, November 2010. Technology Foundation STW.
- [MB:22] A. Molderink, V. Bakker, M. G. C. Bosman, J. L. Hurink, and G. J. M. Smit. Simulation of the effect of introducing micro-generation, energy buffers and accompanied optimization algorithms on the energy efficiency. In *Proceedings of the Nineteenth Annual Workshop on Circuits, Systems and Signal Processing (ProRISC), Veldhoven, The Netherlands*, pages 72–76, Utrecht, November 2008. Technology Foundation STW.
- [MB:23] A. Molderink, V. Bakker, M. G. C. Bosman, J. L. Hurink, and G. J. M. Smit. Domestic energy efficiency improving algorithms. In *Proceedings of the 2009 ProRISC Workshop, Veldhoven, Netherlands*, pages 229–236. Technology Foundation STW, November 2009.
- [MB:24] A. Molderink, V. Bakker, M. G. C. Bosman, J. L. Hurink, and G. J. M. Smit. Steering the smart grid. In *2010 STW.ICT Conference, Veldhoven, the Netherlands*, pages 1–7, Utrecht, November 2010. Technology Foundation STW.
- [MB:25] G. J. M. Smit, A. B. J. Kokkeler, V. Bakker, M. G. C. Bosman, and A. Molderink. Wat maakt een slimme meter echt slim? In *Duurzame ICT*, pages 35–50. Academic Service, Den Haag, August 2010.