

TRIANA: a control strategy for Smart Grids

Forecasting, planning & real-time control

Vincent Bakker



TRIANA
a control strategy for Smart Grids

Forecasting, planning and real-time control

Vincent Bakker

Members of the dissertation committee:

Prof. dr. ir.	G.J.M. Smit	Universiteit Twente (promotor)
Prof. dr.	J.L. Hurink	Universiteit Twente (promotor)
Dr.	M.J. Arentsen	Universiteit Twente
Prof. dr. ir.	T.H. van der Meer	Universiteit Twente
Prof. dr. ir.	R. Belmans	Katholieke Universiteit Leuven
Dr. ing.	S. Klous	KPMG
Prof. dr. ir.	J.G. Slootweg	Technische Universiteit Eindhoven
Prof. dr. ir.	A.J. Mouthaan	Universiteit Twente (chairman and secretary)



This research have been funded by Essent, GasTerra and Technology Foundation STW.

CTIT Ph.D. thesis Series No. 11-215
Centre for Telematics and Information Technology
University of Twente, P.O.Box 217, NL-7500 AE Enschede

Copyright © 2011 by Vincent Bakker, 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 Drukkerijen, The Netherlands.

ISBN 978-90-365-3314-0
ISSN 1381-3617 (CTIT Ph.D. Thesis Series No. 11-215)
DOI 10.3990/1.9789036533140

TRIANA
A CONTROL STRATEGY FOR SMART GRIDS

FORECASTING, PLANNING AND REAL-TIME CONTROL

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 vrijdag 13 januari 2012 om 14.45 uur

door

Vincent Bakker

geboren op 26 april 1982
te Wageningen

Dit proefschrift is goedgekeurd door:

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

Voor Ellen en Iris

ABSTRACT

Technological advances of the last centuries has given the western world a high quality of life, a good health system and for most people a comfortable home to live in. As a result, we live in a high-tech society, and this technology has to be fueled by energy, where electricity is commonly used as a primary energy source. Currently, a number of trends can be identified in the electricity supply chain. Firstly, the electricity demand is still increasing and will become more fluctuating due to increasing prosperity and the electrification of many activities/devices, e.g. electrical vehicles. Secondly, the stochastic behavior of the electricity demand and the lack of flexibility on the demand side requires constant adjustment on the production side of electricity, decreasing the efficiency of power plants. The increasing demand and extra fluctuation will put more stress on the whole supply chain.

On the other side, society's desire to reduce our environmental footprint requires a reduction in CO₂. To achieve this, also changes are required in the electricity supply chain. More distributed generation based on renewables or (bio-based) fossil fueled generation with a higher efficiency is desired. Unfortunately, these renewable sources are often based on (partly) uncontrollable and very fluctuating sun-, water- and wind power, requiring an introduction of flexibility in the grid elsewhere.

These trends result in more renewable, distributed energy production and higher peaks in the electricity demand. Consumers should become more active to exploit their flexibility to cope for the inflexibility of renewable generation. Techniques like controllable distributed generation, distributed storage, and smart appliances can introduce the required flexibility in the grid. Smart appliances with more intelligence can adjust their demand profile dependent on the current situation in the grid, without any loss of comfort of the end-users. For example, freezers and fridges may advance or postpone their cooling cycles to better match their electricity demand to the production capacity available. Appliances like a washing machine and/or dishwasher can postpone their start-time, but may also alter their electricity consumption pattern by reducing the amount of power they consume for heating the water.

To exploit the newly available flexibility in the grid and maintain a proper functioning, affordable electricity supply, the grid has to become more intelligent. In the more intelligent grids, called Smart Grids, production, transmission, distribution and consumption are continuously monitored, managed and coordinated to maintain grid stability and reliability. This requires that technical, economical and

legislative challenges need to be tackled. It is generally agreed that ICT plays an essential role within such Smart Grids to control the whole system.

In this thesis TRIANA, a the control methodology for the Smart Grid consisting of forecasting, planning and real-time control is presented. The goal of this energy control methodology is to manage the energy profiles of individual devices in buildings to support the transition towards an energy supply chain which can provide all the required energy in a sustainable way.

TRIANA has been developed at the University of Twente. The focus of this thesis is on the first two steps of the method. In the first step of TRIANA, the forecasting step, the scheduling freedom (flexibility) of a device is determined. This scheduling freedom is dependent on the type of device, the device specific restrictions and the environment of the device. For example, the scheduling freedom of the already mentioned smart freezer is determined by the state of the freezer (based on the current temperate and the allowed temperature range), the insulation quality of the freezer, the number of freezing modes and the interactions with the residents, etc. For each individual device a forecast is made, since device specific information and restrictions need to be known in order to control the device.

Due to the enormous amount of devices in the grid, each with individual information, restrictions and environment, forecasting is performed for each individual device by a local controller present in the building. This results in a scalable system, since no information about the device and the environment needs to be communicated and the required computational power required for forecasting is distributed. By performing the forecasting locally, local building/resident specific characteristics can be taken into account. This can improve the prediction quality. Furthermore, since the environment of a device may not be static due to the stochastic nature of the residents, local information can be used to adapt to changes. By using locally harvested data, a fully autonomous forecasting system without direct interaction with the residents can be built. As an example of the forecasting system, the forecasting for a micro-Combined Heat and Power (CHP) appliance in combination with a heat store is researched in more depth. The freedom in using a micro-CHP appliance is determined by the size of the heat store, the storage level of the buffer and the expected heat demand of the building. Here, local information like historical heat demand and weather information are considered as good candidates to be used as input data for the heat demand forecasting. Simulated annealing is used to determine good forecast model parameters. Using the simulated annealing algorithm, forecasts with a Mean Percentage Error of around 10% can be achieved.

In the second step of TRIANA, a planner tries to exploit the freedom of the devices determined in the first step for his objective. The objective of the planner could be to reach a flat profile, or only to consume electricity at certain hours of the day. Using this information, the planner generates a desired profile for a fleet of buildings by generating a planning for each building/device individually. Planning is performed by a hierarchical system consisting of a top planner, multiple levels with intermediate grid planners and at the bottom of the structure the individual building controllers. By subsequently dividing the overall planning problem into smaller subproblems which are solved at lower levels, a scalable system is achieved. By aggregating

information at each level in the structure, the amount of communication upwards is reduced. Using steering signals, the building controllers generate a planning for each device based on the cost functions of the devices, the steering signals and the locally generated forecasts. By iteratively adjusting the steering signals, the profiles of the individual houses are reshaped to reach the global objective.

There are multiple ways to determine the steering signals and they can be determined at multiple levels within the hierarchical structure. The best results are achieved by using different steering signals for each building, determined by the lowest grid planner. Since the schedules generated in the planning phase are based on forecasts, and forecasts often are not perfect, deviations from the planning can occur. In a replanning phase, a new planning can be generated, based on the real situation and improved forecasts based on more recent information. Enabling replanning shows a significant improvement in reaching the desired objective.

The third and final step of TRIANA is the real-time control step performed by the local controller. Based on criteria set by the planning, the local controller controls the devices to achieve this planning in the best possible way. The real-time control should be able to work around forecast errors, or if this is not possible, signal the planner that the planning cannot be reached and a new planning, based on the current situation, is required. Essential in this approach is that the comfort level of the residents must be maintained.

In order to analyze the impact of the control methodologies introduced for the smart grid, a simulator has been built. The simulator is developed based on an energy model. The basic elements of the model are individual devices and between devices energy streams (electricity, heat, gas etc) are defined. Devices can consume, buffer, convert and exchange energy, resulting in four categories of devices. The energy streams are connected via so called pools, which represent the physical connections between the devices. These pools are used to ensure the energy balance. The model has been implemented in the simulator and a controller based on cost functions is used to control the devices. The cost functions provide a generic and flexible, but still powerful method to control current and future devices.

Furthermore, frameworks for configuration of the model, the addition of stochastic variations and tools for result analysis are provided. Since the model is computationally intensive, the simulator can be organized in a distributed fashion to allow simulations using multiple computers. The underlying communication framework for distributed simulation can also be used for distributed control of different controllers present in the smart grid. The simulator evenly distributes the load over the available computers involved in a simulation. The simulation speeds up linearly to the number of computers in the simulation, but the speed up is limited by the slowest computer.

To study the effectiveness of the control methodology, to find the best parameters of the control methodology and to study the most economic use of the flexibility of devices, multiple scenarios have been simulated. The simulations show that the control methodology can optimize the energy flows and can control the operation of the domestic devices in an economic manner without discomfort for the residents. Via TRIANA, different objectives can be reached and the optimization potential can

be exploited.

Based on the simulations we conclude that TRIANA is a control methodology capable of monitoring and adjusting the energy profiles and electricity streams. It can autonomously determine the optimization potential of a large group of buildings and exploit this potential to work towards global objectives. The combination of forecasting, planning and realtime control is a promising direction for control methodologies for Smart Grids: it is scalable, generic and reliable.

x

SAMENVATTING

Technologische ontwikkelingen van de afgelopen eeuwen hebben de westerse wereld een goede kwaliteit van leven, een zorgstelsel en voor de meesten een comfortabel plek om te wonen gebracht. Daardoor leven we in een 'high-tech' maatschappij en deze technologie moet gevoed worden met energie, waarbij elektriciteit een veelgebruikte energiebron is. In de huidige elektriciteitsvoorziening zijn een aantal trends te herkennen.

Ten eerste zien we dat de elektriciteitsvraag nog steeds toeneemt en meer fluctuerend wordt als gevolg van de toenemende welvaart en het toenemen van elektriciteit verbruikende activiteiten, zoals elektrisch koken en elektrisch rijden.

Ten tweede vraagt het stochastische gedrag van de elektriciteitsvraag, als gevolg van menselijk handelen, en het gebrek aan flexibiliteit aan de consumptiekant van de energieketen dat de elektriciteitsproductie constant aangepast moet worden om vraag en aanbod in balans te houden. Deze constante aanpassingen verlagen de efficiëntie van de elektriciteitscentrales. Door de toenemende vraag en extra fluctuatie zal in de toekomst de belasting op de energieketen alleen maar toenemen.

Als laatste is er een sterke maatschappelijke wens om onze belasting op het milieu te verminderen door minder CO₂ uit te stoten. Een dergelijke vermindering vereist veranderingen in de huidige elektriciteitsvoorziening. Met name elektriciteitsproductie gebaseerd op hernieuwbare bronnen of het gebruik van biobrandstoffen met een hogere efficiëntie is wenselijk. Helaas is de productie van hernieuwbare bronnen slecht of zelfs niet te besturen vanwege de afhankelijkheid van zon-, water- of windenergie, met als gevolg dat elders in de keten flexibiliteit zal moeten worden toegevoegd.

Bovengenoemde trends zullen leiden tot een toename van op hernieuwbare bronnen gebaseerde gedistribueerde energieproductie en hogere pieken in de elektriciteitsvraag. De benodigde extra flexibiliteit nodig voor een goed functionerend elektriciteitsnetwerk zal deels moeten komen van de consumptiekant. Consumenten of consumerende apparaten zouden actiever bezig moeten gaan en een deel van hun flexibiliteit moeten benutten om inflexibiliteit van generatie op basis van hernieuwbare bronnen te compenseren. Technieken als bestuurbare gedistribueerde productie, gedistribueerde energieopslag en slimme apparaten kunnen de benodigde flexibiliteit introduceren in de huidige keten. Slimme apparaten met meer intelligentie kunnen hun vraagprofiel aanpassen aan de situatie in het elektriciteitsnetwerk, zonder enig verlies van comfort voor de eindgebruikers. Apparaten

als diepvrieskisten en koelkasten kunnen hun koelcycli naar voren of naar achter schuiven om zo bijvoorbeeld de vraag en aanbod van energie eenvoudiger in balans te houden. Apparaten als wasmachines en vaatwassers kunnen naast het verschuiven van hun startmoment ook hun energieverbruik aanpassen door bijvoorbeeld wat sneller of langzamer het water te verwarmen.

Om gebruik te maken van deze mogelijke nieuwe flexibiliteit zal het netwerk zelf ook intelligenter moeten worden. Deze nieuwe, meer intelligente netten worden ook wel 'Smart Grids' genoemd. In een Smart Grid worden productie, transport en consumptie continue gemonitord, geregeld en gecoördineerd om het systeem stabiel en betrouwbaar te houden. Dit vereist dat zowel technische, economische als wetgevende uitdagingen aangepakt worden. Men verwacht dat ICT hierbij in de toekomst een essentiële rol zal gaan spelen.

In dit proefschrift wordt TRIANA, een besturingsmethode voor Smart Grids bestaande uit voorspelling, planning en real-time aansturing gepresenteerd. Het doel van deze besturingsmethode is het beheren van energiprofielen van individuele apparaten in gebouwen om zo de energietransitie naar een duurzame energievoorziening mogelijk te maken. TRIANA is ontwikkeld op de Universiteit Twente. Dit proefschrift richt zich vooral op de eerste twee stappen van TRIANA. In de eerste stap, de voorspelstap, wordt de flexibiliteit in de aansturing van een apparaat bepaald. Deze flexibiliteit is afhankelijk van het type apparaat, de apparaatspecifieke beperkingen en de omgeving waarin het apparaat zich bevindt. De flexibiliteit van bijvoorbeeld de al eerder genoemde diepvriezer wordt bepaald door de toestand van de diepvriezer (huidige interne temperatuur, de toegestane temperatuurlimieten), de isolatiekwaliteit van de diepvriezer, het aantal vriesstanden, hoe vaak bewoners etenswaren erin/eruit halen etc. De voorspellingen worden voor ieder individueel apparaat gemaakt, aangezien apparaatspecifieke informatie en beperkingen bekend moeten zijn om ze correct aan te sturen.

Door het enorme aantal apparaten aanwezig in het netwerk, elk met eigen informatie, beperkingen en omgeving, worden de voorspellingen bepaald door een ingebed systeem aanwezig in het gebouw. Dit ingebedde systeem, in dit werk een lokale controller genoemd, is een systeem dat kan communiceren met de slimme apparaten en informatie over (de bewoners van) een gebouw kan verzamelen. Door de voorspellingen lokaal uit te voeren is het systeem schaalbaar, omdat er geen informatie over elk apparaat gecommuniceerd hoeft te worden naar een centrale plek. Bovendien wordt de vereiste rekenkracht nodig voor de voorspelling verspreid. Door de voorspellingen lokaal uit te voeren kunnen lokale gebouw- en bewonersspecifieke karakteristieken meegenomen worden in de voorspelmethoden. Dit kan de kwaliteit van de voorspellingen verbeteren. Bovendien kan de omgeving van een apparaat veranderen door het natuurlijk gedrag van de bewoners. Door gebruik te maken van lokale observaties kunnen deze veranderingen gebruikt worden om het voorspellingssysteem aan te passen aan nieuwe situaties. Het doel is om, door gebruik te maken van lokaal verzamelde data, een volledig autonoom voorspellingssysteem te bouwen dat geen interactie met de bewoners nodig heeft. Als voorbeeld voor een dergelijk systeem worden de voorspelmogelijkheden van een HRe ketel verder onderzocht in dit proefschrift. De aansturingvrijheid van

een HRe ketel (in combinatie met warmte-opslag) wordt bepaald door de grootte van het warmtevat, het huidige opslagniveau en de verwachte warmtevraag van het gebouw en de bewoners. In dit geval worden historische warmtevraaggegevens en weerinformatie beschouwd als goede kandidaten om als invoer voor het voorspelsysteem te gebruiken. Via een 'simulated annealing' algoritme worden de juiste invoergegevens en andere voorspellingsparameters bepaald. Door gebruik te maken van simulated annealing kan de voorspelfout beperkt worden.

In de tweede stap van TRIANA probeert een planner de besturingsvrijheden van apparaten, zoals bepaald in de eerste stap, te gebruiken voor een bepaald doel. Het doel van de planner kan bijvoorbeeld het bereiken van een geheel vlak verbruiksprofiel zijn, of zoveel mogelijk energievraag te verschuiven naar bepaalde perioden van de dag. Gebaseerd op de informatie van de eerste stap genereert de planner een gewenst en uitvoerbaar profiel voor een grote groep gebouwen, bestaande uit profielen voor elk individueel gebouw/apparaat.

Het planproces wordt uitgevoerd door een hiërarchisch planning systeem, bestaande uit een centrale planner boven in de hiërarchie, meerdere tussenliggende planners verantwoordelijk voor een bepaald gebied en onderaan in de hiërarchie de controllers aanwezig in de gebouwen. Door het opeenvolgend opdelen van het planningsprobleem in kleinere problemen, die opgelost worden binnen de lagere niveaus, wordt wederom een schaalbaar systeem bereikt. Door het samenvoegen van informatie op ieder niveau in de hiërarchie wordt de hoeveelheid communicatie naar boven telkens verminderd. Gebaseerd op stuursignalen ontvangen van de centrale planner en de lokaal voorspelde besturingsvrijheid genereert de planner een planning voor elk apparaat via apparaat specifieke kostenfuncties. Door het iteratief aanpassen van de stuursignaal worden de profielen van de gebouwen zodanig aangepast totdat een doelprofiel bereikt is.

Het bepalen van de stuursignalen kan op verschillende manieren en op verschillende plekken binnen de hiërarchie. De beste resultaten worden bereikt door voor elk gebouw unieke stuursignalen te gebruiken, bepaald door de planner op het laagste niveau in de hiërarchie.

Omdat de gemaakte planningen gebaseerd zijn op voorspellingen en deze voorspellingen niet foutloos zijn, kan het voorkomen dat er afwijkingen ontstaan ten opzichte van de planning. In een herplanningsfase wordt een nieuwe planning gemaakt, gebaseerd op de huidige situatie en verbeterde voorspellingen gebaseerd op meer actuele informatie. Het gebruik van herplanning leidt tot een significantie verbetering in het behalen van het gewenste profiel.

In de derde en laatste stap van TRIANA worden alle apparaten bestuurd door de lokale controller. Op basis van de wensen van de planner probeert de controller de apparaten zodanig aan te sturen dat de gewenste planning zo goed mogelijk bereikt wordt. De aansturing moet hierbij rekening houden met eventuele voorspelfouten, de huidige situatie en het comfort van de bewoners. Indien een planning niet haalbaar blijkt te zien kan de controller de planner waarschuwen. Daarna kan, indien nodig, een nieuwe, op de huidige kennis gebaseerde planning gemaakt worden. Belangrijk hierbij is dat ten alle tijde het comfort van de bewoners behouden blijft.

Om de invloed van verschillende besturingsmethoden voor een Smart Grid te analyseren is er een simulator ontwikkeld. Als basis voor deze simulator is een energiemodel ontwikkeld. Met behulp van het energiemodel kunnen individuele apparaten en de energiestromen (elektriciteit, warmte, gas enz.) tussen deze apparaten uitgedrukt worden. Apparaten kunnen energie consumeren, opslaan, converteren of uitwisselen. De apparaten worden met elkaar verbonden via 'pools', die representaties zijn van de fysieke verbindingen tussen apparaten. Via deze pools wordt er altijd gezorgd voor een energiebalans door te eisen dat de som van de energiestromen altijd nul moet zijn. Het model is geïmplementeerd in de simulator en de lokale controller bestuurt de apparaten via apparaatspecifieke kostenfuncties. Deze kostenfuncties bieden een algemene, flexibele maar toch krachtige manier om de huidige, maar ook toekomstige apparaten te besturen.

Naast een implementatie van het energiemodel biedt de simulator ook raamwerken voor de configuratie van instanties van het model, de mogelijkheid om stochastische variaties toe te voegen en middelen om de resultaten van een simulatie te analyseren. Doordat het energiemodel rekenintensief is, biedt de simulator de mogelijkheid om een simulatie te distribueren over meerdere computers. Het onderliggende benodigde communicatieraamwerk voor de gedistribueerde simulatie kan ook gebruikt worden door de verschillende besturingsmethoden. De simulator verdeelt de last over de beschikbare computers in het netwerk en bereikt daarbij een snelheidswinst lineair aan het aantal computers dat meewerkt aan een simulatie. De snelheidswinst wordt wel beperkt door de traagste computer.

Om de effectiviteit van TRIANA te toetsen, de beste parameters van de besturingsmethode te vinden en de rendabiliteit van slimme apparaten te bepalen zijn er verschillende scenario's gesimuleerd. Simulaties van deze scenario's tonen aan dat TRIANA in staat is de energiestromen te beheren zonder enig comfortverlies voor de bewoners. Met behulp van TRIANA is men in staat om uiteenlopende doelfuncties te bereiken en verschillende soorten apparaten aan te sturen. Met behulp van TRIANA kunnen dus de energieprofielen gemonitord en aangepast worden.

TRIANA kan autonoom de besturingsvrijheden van een grote groep apparaten bepalen en met behulp van planning en aansturing deze vrijheden gebruiken. De combinatie van voorspellen, plannen een realtime aansturing is een veelbelovende richting voor besturingsmethoden voor Smart Grids; de aanpak is schaalbaar, generiek en betrouwbaar.

DANKWOORD

Dit is hem dan, mijn proefschrift. Dit werk dat nu voor je ligt is het resultaat van iets meer dan vier jaar werk, wat ik zeker niet in isolement bereikt hebt. Daarom wil ik graag via deze weg een aantal mensen bedanken die direct en indirect het behalen van deze prestatie mogelijk hebben gemaakt.

Ten eerste wil ik graag mijn promotoren/begeleiders Gerard Smit, Johann Huring en Simon Kolin bedanken. Toen ik eenmaal besloten had om bij CAES, toentertijd nog CADTES, mijn afstuderen te doen is Gerard vanaf de eerste dag mijn begeleider geweest. Ondertussen is Gerard leerstoelhouder geworden van CAES, met als gevolg een enorm volgeboekte agenda. Toch weet Gerard altijd tijd voor je te maken als dat nodig is, iets wat ik enorm waardeer. Tijdens sommige van onze voortgangsgesprekken opperde Gerard of ik interesse had om te promoveren binnen zijn vakgroep. Uiteraard had ik daar wel eens over nagedacht en was ik vereerd dat mij een promotieplek werd aangeboden. Het onderzoeksgebied was ‘Virtual PowerPlants’, voor mij en de rest van de leerstoel nog een geheel nieuw gebied. Daardoor hebben de promovendi in dit gebied veel vrijheid gekregen en ik ben dan ook blij dat we er toen op tijd bij waren. Het doet mij dan ook deugd dat de Universiteit Twente ondertussen het belang van energie en duurzaamheid herkend heeft en één van de speerpunten van deze universiteit gemaakt heeft.

Al vrij vlot na aanvang van mijn onderzoekstraject leerde ik Johann kennen. Johann weet altijd de juiste kritische vragen te stellen en, waar nodig, de juiste, tactische opmerking te plaatsen. Zeker op momenten waarop het allemaal tegen lijkt te zitten krijgt Johann het voor elkaar dat je weer met een goed gevoel naar huis gaat.

Simon heb ik leren kennen tijdens mijn stage bij HOMA Software BV. Ik heb toen met heel veel plezier bij hem op kantoor gewerkt en enorm genoten van de gesprekken/discussies die we hadden over allerlei zaken en met name over ondernemerschap. Zijn vooruitziende blik en innovatieve idee brachten hem uiteindelijk bij de leerstoel van Gerard Smit, wat heeft geleid tot de onderzoeksrichting waarin ik actief ben geweest.

Naast dit proefschrift is een belangrijk en spannend onderdeel de openbare verdediging, waarin ik ondersteund wordt door mijn paranimfen Jan Willem van Houwelingen en Albert Molderink. Jan Willem ken ik al vanaf de middelbare school, waarbij wij eigenlijk vanaf de dag dat wij elkaar ontmoetten meteen goede vrienden waren. Jan Willem is, netzo als ik, lekker eigenwijs en heeft overal een

beetje verstand van en heeft vooral overal een mening over. Dat zijn allemaal eigenschappen die ik enorm waardeer.

Albert heb ik leren kennen tijdens mijn afstuderen. Ook met Albert kon ik van begin af aan goed opschieten. Albert is wellicht nog eigenwijzer dan ik, enorm gemotiveerd en kan hard werken. Daarnaast houdt hij zijn collega's scherp en dat heeft er zeker aan bijgedragen dat ik nu de eindstreep behaald heb. Uit de vele, pittige discussies die we hadden zijn we altijd gekomen en deze discussies hebben volgens mij altijd zeer vruchtbare resultaten gebracht.

Naast Albert heb ik mijn kamer mogen delen met Maurice, Karel en later ook Hermen. Tussendoor kwamen daar ook nog de nodige afstudeerders bij, wat resulteerde in een drukke, maar zeer gezellige kamer, waar wij tussen het werk door voor genoeg afleiding hebben gezorgd. Maurice kwam als AIO een jaartje later in onze onderzoeksgroep en was als TW-er meteen een goede aanvulling, zowel op wetenschappelijk als persoonlijk vlak. Met zijn tweetjes zaten wij in hetzelfde project, waarin we denk ik leuke resultaten hebben geboekt. Karel zorgt er met zijn cynische opmerkingen en droge humor voor dat het nooit saai in onze kamer was. Maar ook voor een goed gesprek kan je altijd bij Karel terecht. De laatste lading droge humor kwam van Hermen, die naast zijn schitterende 'is nooit grappig' opmerkingen ook een mooie bijdrage heeft mogen leveren aan mijn resultatenhoofdstuk.

Ook wil ik graag alle andere vakgroepgenoten bedanken voor de goede sfeer en gezelligheid tijdens de koffiepauzes en andere activiteiten die door iedereen georganiseerd worden. Uiteraard verdienen Marlous, Nicole en Thelma een speciaal bedankje. Bedankt dat jullie altijd voor me klaar hebben gestaan, al was het maar voor een gezellig gesprekje!

Als laatste wil ik mijn vrienden en familie bedanken. Het gaat te ver al om mijn vrienden individueel te bedanken, maar ik denk dat de juiste personen zich wel aangesproken voelen.

Uiteraard wil ik mijn ouders bedanken. Zij hebben mij altijd gemotiveerd om goed je best te doen en hebben mij geleerd dat als je iets wilt bereiken je er voor moet werken. Bedankt voor jullie steun en dat jullie het mogelijk hebben gemaakt dat ik zo fijn heb kunnen studeren in Enschede. Ook Remon, Elise, Pascal en Ingeborg bedankt voor jullie interesse in mijn, toch wat ver weg staande, werk. Ondanks dat we tegenwoordig wat verder van elkaar wonen hebben we nog steeds een uitstekende band en reis ik graag wat kilometers om van onze gezinnetjes te genieten. Ook mijn schoonfamilie, Jos, Toos en Harm wil ik graag bedanken voor hun steun. Al vanaf de eerste keer dat ik in Bergen kwam voelde het als een tweede thuis.

Tot slot wil ik mijn vrouw Ellen bedanken. Jij haalt het beste in mij naar boven en zonder jou was het afronden van mijn promotieonderzoek niet mogelijk geweest. Je bent mijn steun en toeverlaat en je hebt altijd achter mijn beslissingen gestaan. Bedankt voor onze fijne tijd samen en het schenken van onze dochter Iris. Hopelijk mogen wij nog een lang, gelukkig leven met zijn allen hebben.

Vincent Bakker
December 2011

CONTENTS

Abstract	vii
Samenvatting	xi
Dankwoord	xv
1 Introduction	1
1.1 Expected changes in the energy supply chain	2
1.2 Transition to a smart grid	4
1.2.1 ICT and energy consumption	6
1.2.2 Privacy	7
1.3 Problem statement	8
1.4 Approach and contribution of this thesis	9
1.5 Outline of this thesis	11
2 Background and Related Work	13
2.1 Current energy supply chain	14
2.1.1 Electricity markets	16
2.2 Smart grids	18
2.2.1 The Drivers towards Smart Grids	18
2.2.2 Smart Grid	22
2.2.3 Technical challenges	24
2.2.4 Smart Grid control	26
2.3 TRIANA: a three step control methodology for smart grids	28
2.3.1 Requirements	29
2.3.2 Related work	30
2.3.3 Position of TRIANA	32
2.3.4 Three-step approach	33
2.4 Emergence	34
2.4.1 Programs and alliances	35
2.4.2 Test sites	36
2.5 Conclusion	36
3 Forecasting	37

3.1	Requirements	39
3.2	Related work	40
3.3	Approach	42
3.4	Forecasting model	43
3.4.1	Multi-layer feed-forward networks	46
3.5	Results: Heat Demand Forecasting	47
3.5.1	Forecast model	47
3.5.2	Input selection	48
3.5.3	Forecasting Quality	51
3.5.4	Results	52
3.6	Searching for adequate model parameters	56
3.7	Conclusions	60
4	Distributed Control	63
4.1	Goal and requirements of the planning methodology	64
4.2	Related work	67
4.3	Approach	68
4.3.1	Distributed Iterative planning	72
4.3.2	Example	75
4.3.3	Impact of different price vectors	76
4.3.4	Communication	81
4.4	Replanning	82
4.5	Conclusions	85
5	Energy flow simulation	87
5.1	Requirements	88
5.2	Related work	90
5.3	Energy Model	92
5.3.1	Devices	93
5.3.2	Limitations and options	94
5.3.3	Streams Between Devices	95
5.3.4	Complete Model	95
5.3.5	Formal model description	97
5.4	Simulator	98
5.4.1	Discrete simulation	98
5.4.2	Configuration	99
5.4.3	Control	101
5.4.4	Smart grid control	104
5.4.5	Logging and results analysis	105
5.4.6	Simulator architecture and implementation	106
5.5	Distributed simulation	109
5.5.1	Protocol	111
5.5.2	Software architecture	113
5.5.3	Validation	113
5.5.4	Performance	113

5.6	Conclusions	114
6	Results	117
6.1	Virtual Power Plant	118
6.1.1	Determining the scheduling freedom of a micro-CHP ap- pliance	119
6.1.2	Use case description	121
6.1.3	Results	123
6.1.4	Conclusions	126
6.2	Heat pump use case	126
6.2.1	Heat pump model	127
6.2.2	Optimal solution	129
6.2.3	Effectiveness TRIANA approach	130
6.2.4	Using forecasted heat demand values	131
6.2.5	Conclusions	133
6.3	Multiple devices use case	134
6.3.1	Results	134
6.3.2	Conclusions	135
7	Conclusions	137
7.1	Conclusions	140
7.2	Recommendations for future work	141
	Acronyms	143
A	Details use cases	145
A.1	Freezer use case Chapter 4	145
A.1.1	Pseudorandom forecasting errors	146
A.2	Heat pump use case	146
	List of publications	149
	Refereed	149
	Non-refereed	151
	Bibliography	153

INTRODUCTION

Technological advances of the last centuries has given the western world a high quality of life, a good health system and for most people a comfortable home to live in. As a result, we live in a high-tech society, and this technology has to be fueled with energy. Unfortunately, most of the energy nowadays is provided, directly or indirectly, by using fossil fuels. Furthermore, looking at our past energy consumption, a continuing trend of increased energy consumption can be seen and it is expected that this increase will continue for the coming years.

Besides our continuous increasing demand for energy, developing countries like China and Brazil also need a lot energy to feed their growth. The increasing demand for energy, and thus the demand for fossil fuels, puts stress on the fossil fuel production. The production capacity and the limited availability of easily accessible resources cannot keep up with the increasing demand, leading to shortage and higher prices. Furthermore, fossil fuels are mostly harvested in political less stable regions, on which most countries do not want to be dependent on.

Next to the increasing price, environmental concerns are becoming more and more important. The consumption of fossil fuels leads to a major increase of greenhouse gases, with all the negative impacts on the environment as a result. Therefore, alternative sustainable methods to provide society's energy are important topics.

One way to decrease the dependence on fossil fuels is to change to bio-fuels, which have the same properties as conventional fossil fuels, but can be created using crops. Therefore, the amount of greenhouse gases does not increase. Another advantage is that the currently used system (based on fossil fuels) can be used, since they can be fueled with bio-fuels.

Other options are to change to an all-electric energy system, and produce this electricity as sustainable as possible. Obviously, these devices should be efficient in energy consumption. An advantage of this approach is that it is flexible to be implemented. The electricity production can change to a more sustainable

production process without any change on the consuming devices. This, however, assumes that all the required electricity can be produced and transported efficiently. Both are nowadays not the case, and additional solutions in the grid are necessary.

On the production side, sustainable electricity production can be achieved using production methods based on renewables sources like wind and solar. A problem with wind and solar is that their production can only be decreased, and is thus only partly controllable. The production capacity is determined by nature, and other mechanisms are required to keep the production and consumption of electricity in balance. Electricity buffers might be a solution, but most of the currently available electricity buffer technology has a low efficiency and problems with wearing, which makes it an economically infeasible solution.

Next to an predictable production, also the transportation and distribution network has to be able to transport all the demand. The current grid is designed and built decades ago. If the current design and control philosophy is continued, the rising energy demand requires a significant increase in grid capacity and operation costs.

To avoid this, a more intelligent grid should be created, in which Information and Communication Technology (ICT) systems help to better match demand and supply and increase the amount of possible renewables in the grid while maintaining a safe, dependable grid. The focus of this thesis is on the requirements, possibilities and algorithms of such ICT systems.

In the remainder of this chapter first the current energy supply chain and its expected changes are described. These expected changes result in a transition to a new, improved energy supply chain which is described next. Based on the imposed challenges, the problem description and the research focus is given in Section 1.3. The approach and the contribution of this thesis is given in Section 1.4. Finally, the outline of this thesis is given in Section 1.5.

1.1 EXPECTED CHANGES IN THE ENERGY SUPPLY CHAIN

Traditionally, most western countries supply domestic electricity demand through generation in large central power stations, with subsequent transmission and distribution through networks. The generation efficiency of the power stations varies between around 35% (older coal stations) to over 50% (modern combined cycle stations), averaging to about 39%. When transmission and distribution losses are considered, the average overall efficiency of the system drops to 35% [26].

Besides the inefficiency of the various power stations, more inefficiency is added to the grid due to the stochastic nature of the consumption. Since electricity storage cannot be achieved efficiently and economically, the whole electricity supply chain is based around balance. Balance is achieved by constantly adjusting the production of electricity to the demand. If the demand increases, management and control systems in the grid ensure that production of the power plants is increased as well. Similarly, the production capacity is lowered when necessary.




Figure 1.1: Electricity load of the Dutch transmission grid of March 9–17, 2011




Figure 1.2: Energy demand of Dutch households

An example of a typical load profile, gathered from TenneT, is depicted in Figure 1.1, in which the load on the Dutch transmission grid in MW per hour of March 9 2011 up to March 17 2011 is shown. Here a typical periodic behavior of the consumption can be seen. During the night the demand drops to the base load, which results from all continuously switched on devices. In the morning, when people start their day, we see the first peak of the day. A second peak can be seen in the evening, when people are at home.

If we look at the energy demand of the Dutch households over the last decade¹, as shown in Figure 1.2, we see two trends. First, the natural gas demand is decreasing, caused by better insulation quality of modern buildings and other, alternative means for heating. Furthermore, the overall energy demand roughly maintained the same, caused by the increasing electricity demand. In other words, electricity will become a more dominant energy carrier in the future.

It is expected that this growth for electricity will continue. Emerging technologies, like electrical transportation, increase the electricity demand significantly. Where most domestic devices have a relative low demand, charging power will require in the order of 10 up to 20 kW to charge the battery within a short amount of

¹from <http://www.compendiumvoordeleefomgeving.nl/indicatoren/nl0035-Energieverbruik-door-huishoudens.html?i=6-40>

time. As shown in Figure 1.1, the load is quite dynamic already and the introduction of electric transportation can make the demand even more dynamic. Not all power plants can adjust their production capacity fast enough to maintain the delicate balance. Therefore, peak power plants with a very dynamic production capacity are used within the grid. These peak power plants are, due to their flexibility, very inefficient. It would be better to maintain a steady production pattern, for which the production plant could be designed.

The dynamic behavior does not only determine the required production capacity, it also determines the transmission and distribution capacity of the grid. The grid must be able to handle the highest peak, resulting in high investments in the grid and thus, for a large part of the day, low utilization of the system.

Another trend in the coming decades is that a larger part of electricity demand will be generated in a distributed setting. Photovoltaics (PV) solar cell and wind turbines offer sustainable methods to generate electricity. Especially PV cells are good distributed generators which can be easily installed on top of buildings and houses. Also other kind of micro-generators are emerging, like micro-Combined Heat and Power (CHP) appliances, micro gas turbines, micro windmills, heat pumps, etc. Although some of these micro-generators are still fueled by fossil fuels, their efficiency often is very high. For some of these generators both electricity and heat are produced while consuming fossil fuels. In case of a conventional power plant, this heat is often wasted, while in the building the heat can be used for central heating or for hot tap water. Due to the increased efficiency of these micro-generators, the overall efficiency can increase.

The last important trend in the energy supply chain is the introduction of so called smart appliances. Like described above, currently demand follows consumption in the total supply chain. However, generators based on renewables lack proper ways of control, requiring adjustment in the consumption if these technologies are introduced on a larger scale. Therefore, appliances with more intelligence are needed to adjust their demand profile dependent on the current situation in the grid. For example, freezers and fridges may advance or postpone their cooling cycles to better match their electricity demand to the production capacity available. Appliances like a washing machine and/or dishwasher can advance/postpone their start-time, but may also alter their electricity consumption pattern by reducing the amount of power they consume for heating the water. All these appliances introduce a lot of flexibility. In total up to 50% of the electricity consumption of a household can be shifted in time [13].

1.2 TRANSITION TO A SMART GRID

The trends for the energy supply chains, as described in the previous sections, introduce challenges that need to be tackled in the near future. Social pressure to reduce our CO₂ footprint and keep our energy supply affordable requires a paradigm shift in the way we produce and consume energy. ICT can play an important role to facilitate the required changes in the energy supply chain, preferably without any

loss of comfort for the end users. On the production side of the supply chain, many control systems already monitor and manage the production process. Since the amount of production has to follow the energy consumption, an advanced control and management system is required. Furthermore, advanced computer models to forecast the electricity demand are used to optimize the fuel purchase and control strategy of big power plants.

Although on many locations the electricity flow is measured, the amount of intelligence in the distribution/transportation network is limited. Currently, operators continuously monitor and manage the grid and changes in the operation often have to be performed manually. One of the problems with the current grid is that it is designed and built decades ago. The hardware in the grid endures for a long period of time, slowing down the innovation and renewal of the grid. Especially in the field of communication a lot of improvement have been achieved only in the last years. The addition of communication within the grid is essential when automating decision making within the distribution and transportation network, since information about subsequent and surrounding networks is required. Institutes like National Institute for Standardization and Technology (NIST) and Institute of Electrical and Electronics Engineers (IEEE) are working hard to determine Smart Grid communication standards. By continuously monitoring and managing power flow and power quality parameters, a more robust and fault-tolerant grid can be achieved.

Using ICT in buildings for energy management is a quite new concept. For large buildings with complex Heating, Ventilating, and Air Conditioning (HVAC) and facade systems, advanced management and control systems do exist, but mainly operate to minimize the energy requirements while maintaining a comfortable temperature for the residents. These systems work independently, without cooperation with other parties in the grid. Within houses, most devices also work independently. Devices are switched on by the residents and continuously working appliances like freezers and fridges solely use local circumstances like the internal temperature in their control. As mentioned earlier, quite some load of a household can be shifted in time, and a smart control system in the house can coordinate this task. By continuously monitoring the status and electricity load of appliances, the control system can adjust the energy profile. The goal of such a control system is to exploit the flexibility of devices, while maintaining the comfort level of residents. Perhaps a certain level of comfort may be sacrificed, for which a resident may be compensated and/or rewarded.

Using control systems in individual buildings, controlling different appliances, the overall energy profile of the building is adjusted. By adjusting the profiles, the objective of adjusting the demand to the production based on renewables can be reached. Another objective can be adjusting the demand to reach a flattened demand profile which can be provided more efficiently by power plants. The objective of the control system can thus differ, and perhaps cooperation and coordination with other control systems in the electricity supply chain may be needed/desirable, i.e. optimization and control can occur on multiple levels in the grid and with a different scope:

In this work, three possible control levels are used:

Local Control A local controller placed in a building might alter the energy profile of a building without any coordination with other energy management controllers in other buildings. All optimization is on a local, in-building level. Example objectives might be to shift as many load as possible to periods with low energy tariffs, or peak shaving to flatten the overall energy profile.

Micro-grid control A different local energy management controller may cooperate within a so called micro-grid. By using a separate controller for this micro-grid, this micro-grid controller coordinates with the different local energy management controller present in each building. The micro-grid can be a part of the distribution network, for example within the neighborhood. Using this approach, assets spread in the neighborhood may be utilized more efficiently. For example, if in a neighborhood multiple local generators are present, like PV cells or micro windturbines, the electricity produced by these generators can be utilized more efficiently by shifting consumption of devices within the neighborhood to match this production. The electricity then stays within the neighborhood and thus does not have to be transported elsewhere.

Virtual Power Plant An option on an even broader level is the creation of a Virtual Power Plant (VPP). The basic principle of a VPP is to emulate the production capacity of a conventional power plant by controlling a (very) large fleet of (controllable) micro-generators. By using a proper control scheme, the large fleet of micro-generators can be used for commercial exploitation, i.e. using the production capacity on an electricity market.

This thesis is carried out as part of the SFEER project, funded by Essent², GasTerra³, and technology foundation STW⁴, in which the creation of a VPP using micro-CHP appliances is researched. A micro-CHP appliance is a system that produces heat and — as a by-product during the heat production — electricity. By adding a heat store to the heating system, the production and consumption of heat can be decoupled (within the boundaries set by the heat demand and the buffer dimensions).

In the VPP scenario, the local controller is responsible for the control of the micro-CHP appliance. By exploiting the flexibility added by the heat buffer, the production of the micro-CHP appliance can be changed from a heat demand-driven control to a electricity-demand control. The produced electricity can than, for example, be traded on an electricity market. In Section 6.1 results of the creation of a VPP based on real life data is presented.

1.2.1 ICT AND ENERGY CONSUMPTION

Adding ICT systems to the grid enables the possibilities mentioned above. The smart controllers present in each house can continuously cooperate and coordinate

²<http://www.essent.nl>

³<http://www.gasterra.nl>

⁴<http://www.stw.nl>




Figure 1.3: Energy consumption and saving of ICT (Source: Simon Mingay of Gartner UK and Ministry of Economy, Trade & Industry, Japan)

with each other to achieve the desired objectives. However, ICT systems themselves consume energy as well. The saving and improvements achieved by incorporating ICT into the grid should clearly be bigger than the extra energy consumed by these systems. These conflicting issues are illustrated in Figure 1.3, using data from the Ministry of Economy, Trade & Industry in Japan. The figure shows that, although the energy consumption of ICT as a whole is still increasing, this increase in energy consumption is superseded by the impact of smart ICT for energy management.

1.2.2 PRIVACY

Although the efficiency of the energy supply chain may be increased by using ICT, with all the beneficial effect on the environment, still there are some non-technical concerns. Especially privacy is a sensitive and important issue. When in 2008 the Dutch government tried to make a smart meter compulsory, many citizens and digital rights organizations started to protest against the new proposal. Smart meter data contains a lot of information about whether or not people are at home, what their habits are etc. Therefore, when building technical solutions for a smarter grid, these privacy concerns should be taken into account. The information communicated within the grid should contain the least possible information about specific household. If some privacy sensitive information has to be sent, it should be sent as anonymously as possible. Without taken the privacy issues into account when designing the smart grid, the technology may not be accepted by society. People must clearly see the added value of the smart grid, without any or with the least possible loss of comfort and privacy.

1.3 PROBLEM STATEMENT

The goal of this research is to investigate the possibilities of using ICT control systems to change the energy profiles of buildings. By properly changing the energy profile of buildings, the current electricity supply chain can be more efficient. Newly introduced technology enhances the grid and enables the possibility for better demand and supply matching and allows a larger part of the electricity production via renewable energy sources. By improving the efficiency and enabling the possibility to enlarge the share of renewables in the chain, a more sustainable supply of electricity can be achieved.

To change energy profiles of a building, flexibility is needed in the way devices can be used, preferable without any loss of comfort. Therefore, an ICT system should be able to determine the flexibility of a device. To determine the flexibility, information about the device and the environment where it is located are important factors. A crucial part of a devices' environment are the residents of the building where the devices are located. People might be willing to put some effort into energy saving, but the system should not require too much intervention from the residents. Preferably, a system should be able to determine the flexibility of the devices completely autonomously. Therefore, a system should be generic and flexible to be able to work with different kind of devices.

Another important aspect of the system is that it consists of many elements. There are quite some buildings/houses in the grid, each with many devices. The whole system may consist of up to millions of devices which need to be managed. This requires a fast and scalable system. When cooperating within such a large group of systems, an efficiently organized command structure is required, since some decisions have to be made really fast (in real-time) to maintain proper and stable functioning of the electricity network. It might be acceptable that when a device like a washing machine is switched on, it takes a couple of seconds before the devices actually start. But for a system which regulates power quality, decisions have to be made much quicker.

Since cooperation is required, communication within the system is essential. When designing the smart grid, opportunities and limitations set by the communication system should be taken into account. For example, it may not be feasible to communicate with millions of buildings within a short period of time from one central place. Effects of latency and limited bandwidth must be dealt within the system design.

In this research, the above mentioned goals with the accompanying challenges are considered. More concrete, in this work the following research questions are stated:

1. What is the optimization potential of devices located in buildings/houses?
2. How can this optimization potential be exploited?
3. How can a control system autonomously determine the optimization potential of devices?

4. What is a proper control system and methodology to utilize the optimization potential, taking the size and timing constraints of the system into account?

1.4 APPROACH AND CONTRIBUTION OF THIS THESIS

As mentioned earlier, the goal of this research is to investigate the possibilities of using ICT control systems to change the energy profiles of buildings. At the University of Twente, TRIANA, a three step control methodology has been developed. The control methodology consists of multiple control systems, located at different levels within the grid. More precisely, it consists of the following three steps:

1. **Forecasting** In order to adjust or alter the consumption pattern of building by controlling the devices, the device status and their possibilities to shift/adjust their consumption pattern must be known in advance. For certain classes of devices, possibilities can arise for flexibility of control. Using information about the device and the environment in which a device is located, the control freedom (flexibility) of that device can be determined. For example, in case of the VPP in the SFEER project, the heat production (and thus also the electricity production) of a micro-CHP is determined by the heat demand of a household. Furthermore, if there is a heat buffer present, other aspects like the state of charge of the buffer, the insulation quality of the buffer and the size of the buffer determine the production capacity.
2. **Planning** Using the information about the flexibility of the devices, a planner located in the network tries to use this flexibility to achieve a certain objective. The planner generates schedules for a group of devices for a given time in the future. The planning is dependent on the objective, and while planning, device specific constraints should be taken into account. In the VPP example, based on the expected heat demand and the restrictions of the heat buffer and the micro-CHP appliance, a planning of the runtime of the micro-CHP can be determined. The objective can be to maximize the earnings on the produced electricity by producing the most electricity during high price periods. The heat demand adds a constraint on the amount of heat that can be generated. The heat buffer loosens the heat demand constrains a bit by allowing extra filling or emptying of the heat buffer, but this limited by the buffer size. The micro-CHP appliance itself also adds constraints on the runtimes of the appliance. Once it is started, it takes a while before the appliance produces electricity at the maximum rate. Furthermore, there may be a minimal off-time before a consecutive start of the appliance is allowed.
3. **Real-time control** The planner generates a certain energy profile of a building by determining the run-times of the controllable devices within that building for a certain period. The local controller, located in the building, is responsible for reaching the planning. Based on criteria set by the planning, the local controller controls the devices to achieve this planning to the best possible way. Since the planning is made in advance based on forecasts, the actual

situation might differ from what was forecasted. Therefore, the real-time control should be able to work around these forecast errors, or perhaps signal to the planner that the planning cannot be reached and a new planning, based on the current situation, is required. Essential in this approach is that the comfort levels of the residents must be maintained. If there is a conflict between what the planner wants and what the residents require, the residents must have precedence.

In this work, a number of elements of the three step approach are researched. First, the forecasting step in the three step approach is investigated (Chapter 3). In contrast to former research forecasting energy demand on large systems, in our approach forecasting is performed on an individual device level. We mainly focus on individual heat demand prediction, which is the first contribution.

Furthermore, the ways in which this forecasted data is used in the other two steps is researched. In this part, multiple (hierarchical) control structures are possible (Chapter 4). The second contribution is the analysis of different control structures between the planner and real-time controllers.

Before the effects of the control system and algorithms can be analyzed, first models of the energy supply chain are required. Since the focus is on controlling individual devices, energy flows between devices needs to be modeled. The third contribution of this is the development of such an energy flow model, powerful and generic enough to model energy flows of different energy carriers on a device level (Chapter 5). Based on this model and real demand data, realistic computer models of the energy flows of a group of households can be simulated.

On top of this model, control strategies and algorithms can be implemented. The energy flow model, control strategies and algorithms all have been bundled into an energy flow simulator. Due to the architecture of the simulator, different large scale scenarios comprising different combinations of technology, use cases and control algorithms can be quickly simulated and analyzed. This simulator is the fourth contribution of this work (Chapter 5).

Using the developed models within a simulator, the effects of the three step control methodology are analyzed. Based on acquired real life data, the opportunities of the smart grid are analyzed via multiple use cases. In these use cases, the forecasting, planning and control algorithms have been investigated for scenarios with mostly heat demand driven devices. Since real life data has been used, this gives a good indication of the power and possibilities of the control methodology. The brief analysis of the possibilities of the control methodology is considered as the last contribution of this work (Chapter 6).

Besides technical and economical aspects of using the domestic potential, also the cooperation of residents is important. As mentioned above, the system should be running autonomously. It may be that some residents are willing to put some effort in energy saving, but this may differ per household. Furthermore, autonomous systems taking decisions about when or how devices are controlled in your own house can introduce a certain feeling of loss of control and privacy, and interference with their personal lives. However, in our research, we approached the optimization

potential from a purely technical view, the social acceptance and economical analysis are left out of scope. This is left for future work in the Route 14 Energy track of the University of Twente.

1.5 OUTLINE OF THIS THESIS

In this chapter a brief introduction of the current and expected energy supply with the corresponding challenges has been given. In the following chapter, some more background information on the current energy supply chain and corresponding markets are given. Furthermore, the drivers towards the smart grid and its technical challenges are discussed. After given related work on smart grids research and control system, our developed three step approach TRIANA is described. The first step of TRIANA is forecasting, which is topic of Chapter 3. The ways in which these forecasts are used in the other two steps of the distributed control are discussed next in Chapter 4. As described above, a energy model has been made. This model has been implemented in a energy stream simulator, which is discussed in Chapter 5. In Chapter 6 some results of test scenarios using the three step approach TRIANA are discussed. We end this thesis with conclusions and future work in Chapter 7.

BACKGROUND AND RELATED WORK

ABSTRACT – This chapter describes the current electricity supply chain, markets and the foreseen changes. Current grids, designed decades ago, are demand driven and electricity is mainly produced centrally. Due to difficulties storing electricity, balance within the grid is maintained by continuously adjusting the production to the demand. The required physical flow of electricity, both for balancing and consumption, is traded via many energy markets. Although this system has worked very well for more than a century, the rising energy prices and environmental concerns require a more sustainable process of providing electricity. Furthermore, the electricity demand is still increasing and is expected to keep increasing in the future. These trends result in more renewable, distributed energy production and higher peaks in the electricity demand. To facilitate the distributed renewable generation and the increasing demand, the grid has to become a Smart Grid. In the Smart Grid production, transportation and consumption has to be continuously monitored, managed and coordinated to maintain grid stability and reliability. Consumers have to become more active to exploit their flexibility to cope for the inflexibility of renewables. This requires that technical, economical and legislative challenges need to be tackled. ICT plays an essential role within this Smart Grid to control the whole system. A proposed control methodology for the Smart Grid is the three step control methodology TRIANA, consisting of forecasting, planning and real-time control.

The Smart Grid can be seen as an evolution of the current energy supply chain, which has provided us with energy for more than a century. In this chapter, some background of the current energy supply chain, its infrastructure and the corresponding markets are given. Since this thesis mostly focuses on reshaping the electricity profiles in the grid, this chapter only describes the electricity part of the energy supply chain.

Parts if this chapter have been presented at [VB:16] and [VB:18] .

Although the current grid works very stable, issues like improving the overall energy efficiency, the introduction of large scale renewables and the reduction in CO₂ require a different paradigm. Therefore, the expected challenges and driving factors of the transition towards a Smart Grid are given next. After giving the definition of Smart Grids in Section 2.2.2, the opportunities and possibilities when using a Smart Grid are described.

New Information and Communication Technology (ICT) and smart control systems and distributed generation in the grid enable new (business) opportunities and will change the way electricity is produced and consumed. The introduction of communication, distributed generation and smarter appliances can lead to a system with constant coordination and cooperation between multiple parties in the grid. TRIANA, a three step control methodology, exploiting the new opportunities introduced by the new ICT technology in the grid, has been developed to manage a large fleet of devices present in the Smart Grid. In Section 2.3 more information about each the three steps of this methodology is given.

TRIANA is the result of a cooperation within a team of three PhD students. Due to this cooperation, there is a large overlap in the background and related work of the three projects. As a consequence, some parts of this chapter are reused from ‘On the three-step control methodology for Smart Grids’ by Albert Molderink [62]. Furthermore, due to the collaboration with Dutch partners, parts of this chapter are motivated from a Dutch perspective as it was during the start of these projects in 2007.

2.1 CURRENT ENERGY SUPPLY CHAIN

Today most buildings in the western world are constantly provided with energy required by the devices in the building. For example, when residents of a house switch on their washing machine or television, the device is switched on instantly. Automatically, somewhere in the grid, the power required for this device is generated.

Besides a electricity connection, other grid connections might be present. For example, many households in the Netherlands have a natural gas connection. These households use gas for cooking, and most importantly for providing heat for central heating and hot tap water via an installed boiler. Often, in (large) flats, heat is provided via a district heating plant or with a mini-Combined Heat and Power (CHP) appliance.

These connections to a building provide the residents of these buildings all the energy they demand. However, the supply of the energy-carriers (electricity, gas, etc) is handled via different supply chains. Gas is harvested on many different locations throughout the world and transported via many gas pipelines to storage facilities and/or end users. A similar supply chain exists for electricity. However, the difference between electricity and other energy-carriers is that electricity can be transported very efficiently, but not be stored efficiently. Gas and heat on the other hand can be stored easily, but transportation requires more effort. The characteristics

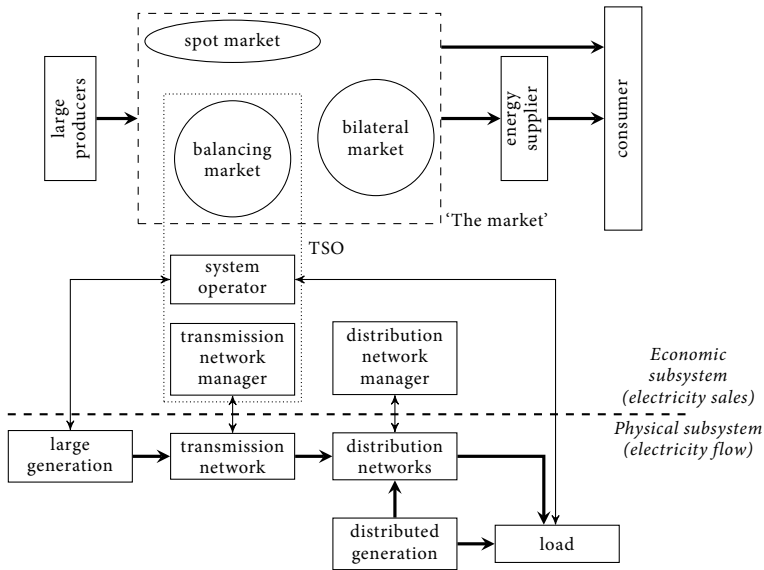


Figure 2.2: Simplified overview of the electricity network and markets

the highest peak is called the *peak load*.

On the supply side, fluctuating demand is supplied by different kinds of power plants. The base load is supplied with large power plants, often fired by coal or nuclear power. These power plants are, due to their nature, inflexible in adjusting their production output in short term and therefore more suitable for supplying the base load. Other power plants, with a more flexible production pattern can be used to handle the short term fluctuations in the demand. Unfortunately, the flexibility of these plants often comes with a decrease of efficiency. For the really short term fluctuations, spinning reserves are used. These spinning reserves are generators which can be started and stopped within a couple of minutes and are often used to stabilize the grid. In general their efficiency is low, but they are required to keep the grid functioning properly and ensure a stable power supply.

2.1.1 ELECTRICITY MARKETS

Generating electricity requires resources like power plants, fuel, operators etc. Ensuring balance and a properly functioning grid requires, due to the fluctuating demand, even more resources. Therefore, electricity has a value, which is traded via different electricity markets.

The current (European) electricity markets and the companies trading energy on these markets are a result of the liberalization and privatization of the energy markets in the late nineties. Although the liberalization has been implemented differently in different countries, roughly the companies were separated into produc-

tion companies, distribution companies and energy suppliers. In the Netherlands, the production companies and energy suppliers became privately owned. The Transmission System Operator (TSO) responsible for the high voltage networks was nationalized via a new grid company TenneT. To operate the distribution network (medium and low voltages), multiple Distribution System Operator (DSO) companies were founded, each responsible for a certain area. Since these DSOs and the TSO still are monopolies, they are regulated by the government.

In the lower part of Figure 2.2 an simplified overview of the electricity network, as described above, is depicted. In the bottom of this figure, the physical flow from supplier ('large generator') via the subsequent networks to the end users ('load' in the figure). New in the (smart) grid is the upcoming share of distributed generation, feeding the locally produced electricity directly into the distribution network. Although the physical flow is straightforward, the economical flow has a totally different structure.

As already mentioned, balance is essential to maintain a properly functioning grid. Achieving this balance has a cost, both in terms of purchase costs for the fuel, as the costs of offering reserve production capacity, since there is a high uncertainty in when spinning reserves are required. Since demand and supply always have to be in balance, multiple electricity markets exist. On these market, many players are present, each offering, purchasing and trading energy.

The base load which solely depends on long term changes (like the seasonal changes) is most often traded using long term contracts (the bilateral market). In the Netherlands, 85% of the electricity is traded via these long term contracts [50]. The short term fluctuations are traded on a day-ahead market, which trades electricity with a granularity of one hour (the spot market). Since these short term fluctuations are very dynamic, a production facility which can handle this dynamic behavior is required. Since this is more difficult to achieve, the prices on the day-ahead market are in general more volatile and can be much higher compared to the long term contracts. The most expensive electricity is the electricity required for the real-time stabilization of the grid, which is traded per quarter of an hour on very short terms (balancing market). As can be seen in the upper part of Figure 2.2, small consumers buy their energy from energy suppliers, which in return buy their energy on different markets. Very big consumers can buy their electricity directly on the market.

To ensure balance in the grid, suppliers (producers) and consumers (retailers) of electricity have to specify one day in advance what their electricity profile is going to be for each quarter of an hour for the next day. For producers this means to specify the amount of electricity they are going to produce and put into the network. Retailers specify the amount of electricity they, i.e. their costumers, are going to consume. The profiles are based on forecasts of the electricity demand and the long term contracts. The difference between what has been traded on before hand (based on long term forecasts) and the expected electricity profile for the next day (short term forecast), the day-ahead market can be used to close this mismatch by trading this difference on the day-ahead market.

Furthermore, deviation from this specification will result in an imbalance and is

penalized by a central authority (the TSO). If a deviation occurs, it has to be compensated elsewhere in the network. The network operators continuously monitor the stability of the grid. Electricity producers with fast available production capacity can offer ancillary services to the TSO. The TSO can order the right amount of balancing power to keep the grid functioning properly.

2.2 SMART GRIDS

The current grid has evolved into a very stable and reliable system. For example, Enexis (one of the Dutch DSOs) has an average downtime of 20.6 minutes in 2009. This is an average uptime of 99.996%. Although the grid is working so well as it is now, it was designed decades ago with different design principles, and environmental and societal circumstances. Back then, fossil fuels were cheap and abundant. Electricity was produced at central places and transported one-way towards to the customers [8, 71]. Although nowadays renewables have an increasing share in the energy mix, fossil fuels are still dominant. For example, the energy mix of the Netherlands in 2004 [33] shows that 89.4% of the electricity production was fueled by fossil fuels. But the circumstances are changing: fossil fuels are becoming expensive and are produced by political less stable countries.

Besides the expected problems in harvesting these fossil fuels, most of the fossil fuels are consumed with a very low efficiency. The generation efficiency of power stations varies between around 35% (older coal stations) to over 50% (modern combined cycle stations), averaging to about 39%. When transmission and distribution losses are considered, the average overall efficiency of the system drops to 35% [26]. These fossil fueled power plants exhaust a lot of CO₂, with all the resulting environmental problems. Therefore, groups of countries (e.g. the G8 countries) made agreements about CO₂ emission reduction [8, 71], for example in the Kyoto agreement. CO₂ reduction is in principle possible, but requires (radical) changes in the way we currently generate, transport and consume electricity. These required changes and the need for other, sustainable energy sources drive a transition towards an electricity grid that is monitored and managed. A change towards another supply chain with more sustainable energy production via continuous management of production, transportation and consumption requires a so called 'Smart Grid'. In the next section, the driving factors towards this Smart Grid are discussed. After this section, a more formal definition of the Smart Grid is given. We end this section by discussing the technical challenges and possible smart grid control.

2.2.1 THE DRIVERS TOWARDS SMART GRIDS

The European Technology Platform Smart Grids [71] has identified three groups of the driving factors towards a smart grid, which are depicted in Figure 2.3. To ensure stability and security of supply, with reduces environmental effects for a affordable price, the grid needs to be updated to keep up with changes in demand and supply. Now is the right moment since the lifetime of a lot of grid elements comes to an end and need to be replaced [28].




Figure 2.3: Schematic of the driving factors for a Smart Grid [71]

In all three parts of the supply chain (generation, transport and distribution/-consumption) driving factors exist for the transition towards a smart grid.

In this subsection the driving factors for the three parts of the supply chain and the driving factor due to liberalization are discussed.

Generation

Today, coal is the main source of electricity generation. A future without electricity generation using coal is almost unthinkable since coal is cheap, is still almost abundant and can be harvested in more stable countries [71]. However, coal is one of the most polluting fossil fuels concerning the amount of CO_2 emission. One of the solutions is to capture and store the CO_2 , so called Carbon Capture and Storage (CCS). At the moment, a couple of CCS installations are in use (for example [21]). However, it is not (yet) a broad applicable option and still based on fossil fuels. Another, better option is sustainable electricity generation using renewable sources (sun, wind, tides, etc). However, this requires thorough changes and improvements of the electricity grid.

First, large scale sustainable electricity generation is often only possible on remote places with a low density of population and therefore a low electricity demand (e.g. large wind power farms offshore or solar panels in the desert). Therefore, electricity needs to be transported to the customers, requiring a large transmission capacity. It is expected that the renewable potential in Europe is large enough to supply all electricity [8]. Mainly in the southern parts of Europe there is a large (solar energy) potential and when even the Northern part of Africa is taken into account, the potential is huge. In the Desertec project [34] a concept is proposed for making use of sustainable energy worldwide. In Figure 2.4 the proposed solution is shown. The dark squares in the desert depict the area required to generate enough



Figure 2.4: Sketch of possible infrastructure for a sustainable supply of power to Europe, the Middle East and North Africa (EU-MENA) (Euro-Super grid with a EU-MENA-Connection proposed by TREC) [34].

electricity using thermal solar power plants for the world, Europe, the Middle East and Northern Africa (MENA) and the proposed combination of Europe and MENA.

Second, large scale sustainable electricity generation is quite different from conventional power plants, both in generation capacity and controllability. It is, in general, agreed that it is both desirable and necessary to manage this new type of generation and adapt the rest of the grid infrastructure to facilitate the sustainable, unmanageable generation. Also on a domestic level more and more electricity is generated using micro-generators. Lower capacity generation on various sites, i.e. generation with lower capacity than conventional large power plants (e.g. sustainable and domestic generation), is called Distributed Generation (DG). Scott et al. [73] state that a fit-and-forget introduction (just install DG without any control) of domestic DG will cause stability problems, amongst others by large flows from lower to higher voltage levels. A study of the International Energy Agency (IEA) concludes that, although DG has higher capital costs than power plants, it has a huge potential and that it is possible with DG to supply all demand with the same reliability, but with lower capacity margins [35]. The study foresees that the supply can change to decentralized generation in three steps: 1) accommodation in the current grid, 2) introduction of a decentralized system cooperating with the central system and 3) supplying most demand by DG.

So, next to high capacity lines for long-distance transportation of electricity, a

sustainable electricity supply also requires more and better monitoring and control capabilities of all types of generation on different levels of the grid.

Consumers

Next to the supply of electricity, also the demand of electricity changes, especially when the trend towards a mainly electricity based energy supply continues. The overall electricity demand increases every year and is expected to keep increasing in the coming years. Furthermore, when more electricity consuming technologies are introduced (e.g. electrical cars) the demand will increase and become more fluctuating. A fit-and-forget introduction will have a severe impact on the grid and generation. More capacity and flexibility is required to ensure the expected reliability and stability of supply. A naive introduction will lead to large investments and decreased generation and transportation efficiency.

New domestic appliances lead to some freedom in the electricity consumption patterns of these devices. They can be monitored and managed to change their consumption profile. So, monitoring and control on the lowest level, on a device level, is desired. With monitoring the expected consumption and production of domestic devices can be forecasted and control can enable the possibility to exploit scheduling freedom of domestic devices to work towards (global) objectives.

Transport

For a transition towards a sustainable energy supply with electricity as the main energy carrier of the future, both more renewable generation and more flexible electricity consuming devices are required. To merge these two tendencies, generation and consumption need to be matched. To make this possible, significant improvements in the grid infrastructure and more intelligence in the grid are required.

The foreseen changes in production and consumption as described in the previous paragraphs will increase the stress on the grid while at the same time stability, reliability and self-healence of the grid becomes more important due to the increasing importance of electricity for society. Therefore, the streams through the grid should be monitored and managed. Geidl et al. [38] propose an alternative transport medium by combining multiple energy carriers in one “cable”. These new cables interconnect a set of so called ‘energy hubs’, where an energy hub is considered a unit where multiple energy carriers can be converted, conditioned, and stored. This leads to a more flexible supply of combinations of energy carriers and to a synergy of energy carriers. For example, natural gas can cool the electricity transportation cables resulting in less transport losses. At the same time the transport losses of the electricity transport can be compensated at the energy hub by converting natural gas to electricity via a Combined Heat and Power (CHP) plant.

Another issue is the large distance between areas with high potential for generating renewable electricity and areas where the electricity is consumed. To transport the sustainable electricity from the generation site towards the customers, an European wide interconnected high capacity electricity grid is required, in combination

with a European wide electricity market. One of the technologies for the high capacity backbone for this European network is High Voltage Direct Current (HVDC). This technology is already used for transport of electricity from offshore wind parks to the coast. Battaglini et al. [8] propose a super Smart Grid: a combination of a European wide HVDC backbone (super grid) to transport the unevenly distributed renewable potential through Europe, partly from the African continent. This is combined with clusters of Smart Grids, using decentralized generation and demand side management in combination with electricity supplied by the super grid to supply all demand.

The earlier mentioned Desertec project also proposes a super grid, not only covering Europe but also Northern Africa and parts of the Middle East (see Figure 2.4). They propose a 17,000 km² solar power system in the Sahara desert producing the main part of the electricity consumption in Northern Africa and 15% of the electricity needs of Europe.

Liberalization

A third driving factor for improving grid capabilities, next to environment and changes in demand/supply, is the liberalization of the electricity market. Due to competition between (distribution) companies they aim for an affordable electricity supply and a stable and reliable grid. Furthermore, an innovative and 'green image' is important for companies.

2.2.2 SMART GRID

The next generation of the grid is often called a Smart Grid. It is hard to give a definition of a Smart Grid. Different parties have their own definition, ranging from a grid capable of charging electrical cars up to a completely controlled grid, including producers, transmission and consumption. The authors of 'A vision for the smart grid' [55] state that the Smart Grid is not a "thing" but rather a "vision": "*The Smart Grid vision generally describes a power system that is more intelligent, more decentralized and resilient, more controllable, and better protected than today's grid*". Another definition of a smart grid given by Scott et al. [73]. This definition is rather common, so we choose for this definition:

"A Smart Grid generates and distributes electricity more effectively, economically, securely, and sustainably. It integrates innovative tools and technologies, products and services, from generation, transmission and distribution all the way to customer devices and equipment using advanced sensing, communication, and control technologies. It enables a two-way exchange with customers, providing greater information and choice, power export capability, demand participation and enhanced energy efficiency."

To develop a Smart Grid, it is important to incorporate the complete grid including supply and demand [54]. Within the grid itself monitoring and switching possibilities are added, generation and consuming devices are (partially) extended with a monitoring and managing interface [71]. Optionally this can be extended with

(electricity) storage. A (central) monitoring and control system matches generation and consumption in the most efficient way (the definition of efficiency can differ depending on the stake holder: reduced CO₂ emission, less fossil fuel usage, more profit, etc). In [54] it is stated that the U.S. electrical infrastructure will evolve as a highly automated and interconnected network much in the fashion of the Internet; one where information and knowledge will flow through intelligent systems to serve the entire grid community; one where a dynamic network of smart devices enables the realtime balance of generation and delivery of electricity with the highest reliability and lowest cost.

The Smart Grid should be accessible for distributed generation and renewable energy sources, comply with different forms of generation, enable local energy demand management (optionally through smart metering systems) and facilitate dynamic control techniques [71]. Furthermore, it should facilitate high levels of power security, quality, reliability and availability with minimum negative side-effects on the environment and the society [71]. The required functionality is summarized in seven important characteristics of a Smart Grid [55]:

1. enable active participation by consumers,
2. accommodate all generation and storage options,
3. enable innovation, new products and services by market forces,
4. provide high quality power,
5. optimize asset utilization and operate efficiently,
6. anticipate and respond to system disturbances (self-healing),
7. operates resiliently against attacks and natural disasters.

An important issue is the large number of stake holders involved in the transition towards a Smart Grid: governments, regulators, consumers, generators, traders, power exchanges, transmission companies, distribution companies, power equipment manufactures and ICT providers [71]. These stake holders need an incentive to cooperate while in first instance it seems to be unattractive for companies. However, distribution companies can decrease operating and maintenance costs and reduce capital costs. Production companies can introduce new types of generation and increase generation by relatively cheap base-load plants [28]. The consumers can reduce their costs and increase power quality and finally society will benefit from a stimulated economy and improved environmental conditions [28].

Both Scott et al. [73] and Fraser [35] indicate that commercial attainability and legislation are important issues for the success of the introduction of DG. The opinions on the investments and profits differ strongly. On the one hand, the European Climate Forum states that large investments are required while it is unknown what the actual benefits and profits are [8]. On the other hand, the U.S. department of energy states that the transition towards a Smart Grid already started and that profits are higher than the investments [28]. They even claim that due

to all benefits (e.g. improve safety and efficiency, better use of existing assets) the transition towards a Smart Grid will be market driven.

2.2.3 TECHNICAL CHALLENGES

For a successful introduction of a Smart Grid we face a number of technical challenges. As a result, there is a lot of research ongoing on different fields for a more efficient and sustainable electricity supply. New power plants themselves are much more efficient, CCS is developed to decrease emissions, transmission mediums with lower losses are developed and domestic devices become more efficient and controllable. In [54] five key technologies required for the Smart Grid are identified:

1. *Sensing and measurement*

Since one can only manage what one can measure, sensing and measuring are an important part of the Smart Grid. The health parameters of the transmission lines and substations should be monitored to prevent the grid from outages. Monitoring and forecasting of the weather can be used for forecasting load and potential output of renewable sources. This can subsequently be correlated with transmission line capacity. Next to the grid, also the generation, storage and consumptions sites and devices need to be monitored to be capable of balancing generation and usage and respecting transmission limitations. An Advanced Metering Infrastructure (AMI) is not only used for billing, but also for monitoring domestic usage, voltage and power quality. Furthermore, the Smart Meter can be used as a gateway to the domestic devices and to determine the optimization potential

2. *Integrated communications*

To transport all information, a high speed communication infrastructure is required. This Integrated Communications (IC) infrastructure moves the information from sensing and measurements devices towards the operators and management information back to the actuators. Creating a homogeneous communication infrastructure requires standards respected by all stake holders, from home networks and all devices connected to it via the smart meters and the distribution companies to the overall network operators. The National Institute for Standardization and Technology (NIST) addressed this problem and is working together with Institute of Electrical and Electronics Engineers (IEEE) to create Smart Grid standards [65]. The IC infrastructure should be designed with future requirements in mind. The capacity, security and performance should be sufficient to facilitate also future applications. A fast, reliable and well designed IC infrastructure glues all the parts of the Smart Grid together.

3. *Advanced components*

A Smart Grid is build up by a network of advanced components. The grid itself should consist of efficient transmission elements connected by advanced flow con-

trol devices, e.g. HVDC lines and solid state transformers. On domestic level a lot of technologies are in development. These technologies range from Photovoltaics (PV) panels on roofs and micro-CHP [76] up to controllable devices [39]. The technologies can be subdivided in three groups:

Distributed Generation (DG) In contrast to a few years ago where electricity was generated in a few large power plants, nowadays and in the future a growing share of the electricity is generated in smaller, geographically distributed generators. This DG ranges from wind turbine parks with a capacity on a MW level up to domestic DG on a kW level.

Distributed Storage (DS) Especially with a growing amount of renewable sources in the electricity supply chain there is a growing demand for electricity storage [36]. Electricity can be produced more efficiently (e.g. at daytime) or at certain periods (e.g. wind, sun) when it is not consumed and thus needs to be stored [71].

Storage can take many forms, can be spread across a large geographic area and can be connected to any voltage level [71]. Especially with the large scale introduction of electrical cars huge distributed controllable storage capacity becomes available. Furthermore, in multiple projects hardware is developed to manage domestic electricity streams and store electricity within buildings, e.g. within the PowerRouter project¹. On the other hand, also larger scale electricity storage is developed, for example in the Smart Substation project of a Dutch consortium [48], which can be seen as a large version of the PowerRouter. An important research area for storage is the development of batteries supplying high requirements concerning capacity, charge/discharge currents and lifespan of many charge/discharge cycles.

Demand Side (Load) Management (DSM) has the goal to modify the consumption pattern of consumers. About 50% of the load in houses is dedicated to controllable devices such as refrigerators, freezers, heaters, washing machines and dryers [13]. These devices can be managed with only a little discomfort for the residents in contrast to lights and a television, which cannot be switched off or shifted without discomfort. Field tests in the US have shown that optimizations using these manageable devices already can lead to significant peak reductions [39]. Furthermore, when residents allow a certain level of discomfort, e.g. a deviation of 0.5°C from the settled room temperature, even more scheduling freedom is gained. Of course there has to be an incentive for the residents to accept the reduction in comfort.

4. Interfaces and decision support

The job of a grid operators became much more challenging in recent years. New tools are required to assist the grid operators. To have enough information in

¹www.powerrouter.com

order to take correct decisions, data mining is very important. Data is produced by measurement devices, transported via the communication infrastructure and gathered and presented by advanced visualization tools. This improved interface is required to visualize the large amount of data on such a way it can be understood at a glance. Furthermore, decision support tools help taking decision, for example using fast simulations to forecast consequences of decisions.

5. *Advanced control*

To make use of all control capabilities and to exploit all optimization potential, advanced control systems need to be developed. Advanced protection systems can adjust relay settings in time for better protection of the grid and even increased power flows in some cases [54]. Controlling flows can for example increase stability, increase damping of oscillations, operate transmission networks as efficiently as possible and assure maximum utilization of transmission assets. The growing share of technologies on a lower voltage level that can influence real and reactive flow, can enhance operators' ability to influence grid conditions significantly. Furthermore, the coordination of (renewable) generation, storage and consumption is fundamental to reach the targets of a Smart Grid.

2.2.4 SMART GRID CONTROL

To create a successful Smart Grid solution and exploit all optimization potential, the introduced technologies need to be monitored and synchronized to each other. On the production side of the electricity chain already a lot of control is available. Controllable production sites, e.g. central power plants, adjust their production to the demand, i.e. a feedback loop adjusts the settings based on the grid frequency to react on small variations. This monitoring and controlling system is a so called Supervisory Control And Data Acquisition (SCADA) control system. Next to adjusting the amount of production in a production site, grid operators can decide to start or stop a (peak) production site completely based on the demand forecasts and the current power plant states and energy demand. Production sites that are not controllable need to be monitored and their production needs to be forecasted.

At the moment, the electricity flows within the grid are mainly only monitored. In a Smart Grid these flows also need to be managed by (automatically) adjusting transformer settings to maintain stability and prevent blackouts caused by overburdening.

However, production and transportation control are also important issues. The biggest challenge is managing the technologies connected to the medium and lowest voltage levels of the grid, the medium sized Direct Current (DC) and domestic technologies. By managing the electricity production, storage and consumption of these technologies a lot of electricity flows in the net can already be managed. The combined flexibility of these technologies is high, but to exploit this potential a lot of devices need to be monitored and managed. Therefore, scalability, communication and uniformity issues need to be solved.

To overcome the scalability and communication issues the structure of the control system is important. A hierarchical structure with data aggregation on the different levels is an often proposed scheme. Such a structure is scalable while the amount of communication can be limited. However, when data is aggregated, information gets lost, i.e. there is a trade-off between precision and the amount of communication.

The goal of a low voltage control system is to manage the cooperation between the domestic technologies to use the maximum optimization potential. The primary functionality of the system is to control the domestic generation and buffering technologies in such a way that they are used properly and efficiently. Furthermore, the required heat and electricity supply and the comfort for the residents should be guaranteed. Some devices have scheduling freedom in how to meet these requirements. This scheduling freedom of the domestic devices is limited by the comfort and technical constraints and can be used for optimizations.

The optimization objective can differ, depending on the stake holder of the control systems, the system state and the rest of the electricity infrastructure. The objective for residents or utilities can be earning/saving money and therefore the goal is to generate electricity when prices are high and consume electricity when prices are low. For network operators the goal can be to maintain grid stability and decrease the required capacity while an environmental goal can be to improve the efficiency of power plants. Therefore, a control methodology should be able to work towards different objectives.

Next to different objectives, control methodologies can have different scopes for optimization: a local scope (within the building), a scope of a group of buildings e.g. a neighborhood (micro-grid) or a global scope (Virtual Power Plant). Every scope again might result in different optimization objectives.

Local scope

Also on a local scope the import from and export into the grid can be optimized, without cooperation with other buildings. Possible optimization objectives are shifting electricity demand to more beneficial periods (e.g. nights) and peak shaving. The ultimate goal can be to create an independent building. This can be done in two forms: *energy neutral* or *islanded*. Energy neutral implies that there is no net import from or net export into the grid. A building that is physically isolated from the grid is called an islanded building.

The advantages of a local scope is that it is relatively easy to realize; it needs no communication with others (less privacy intrusion) and there is no external entity deciding which devices are switched on or off (less privacy issues and hence a better social acceptance). Objectives like peak shaving can easily be achieved by preventing the simultaneous usage of many devices. However, an islanded situation required more investments and is harder to achieve. In this case, higher investment costs, e.g. in storage capacity and micro-generation, are required in order to supply all the locally required energy. Some assets, like local storage, might be shared

with neighbors, reducing the individual costs. This brings us to the next scope, the micro-grid.

Micro-grid

In a micro-grid a group of buildings together optimize their combined import from and export into the grid, optionally combined with larger scale DG (e.g. wind turbines). The objectives of a micro-grid can be shifting loads and shaving peaks such that demand and supply can be matched better internally. The ultimate goal can be perfect matching within the micro-grid, resulting in a neutral or islanded micro-grid. The advantage of a group of buildings is that their joint optimization potential is higher than that of individual buildings since the load profile is less dynamic (e.g. startup peaks of devices disappear in the combined load). Furthermore, multiple micro generators working together can match more demand than individual micro-generators since better distribution in time of the production is possible [1]. Finally, within a micro-grid the locally produced electricity can be used locally, saving transmission costs and preventing streams from lower to higher voltage levels. However, for a micro-grid a more complex control methodology is required.

Virtual Power Plant (VPP)

The original Virtual Power Plant (VPP) concept is to manage a large group of micro-generators with a total capacity comparable to a conventional power plant. Such a VPP can replace a power plant while having a higher efficiency, and moreover, a greater flexibility than a normal power plant. Especially this last point is interesting since it expresses the possibility to react on fluctuations. This original idea of a VPP can of course be extended to other domestic technologies. Again, for a VPP a complex control methodology is required. Furthermore, communication with every individual building is required and privacy and acceptance issues may occur.

2.3 TRIANA: A THREE STEP CONTROL METHODOLOGY FOR SMART GRIDS

As mentioned in the Section 2.2.4, a successful Smart Grid solution requires new control methodologies to monitor and coordinate the large amount of newly added technology, like distributed generation and -storage, smart appliances etc. In this section TRIANA, a three-step control methodology for smart grids is introduced. The goal of this energy control methodology is to manage the energy profiles of individual devices in buildings to support the transition towards an energy supply chain which can provide all the required energy in a sustainable way. Therefore, the objectives of the control methodology are mainly based on electricity streams. However, since these streams are tightly interweaved with the other domestic energy streams, all domestic energy streams are incorporated in the control methodology.

As mentioned earlier, TRIANA is the result of a cooperation within a team of three PhD students. Each researcher developed a part of the control methodology. The first step of the three steps, the forecasting, is topic of this work. The other

two steps are mainly performed by the other two PhD students. Therefore, these two steps are described to the extent which is necessary to understand the overall control methodology. For more detailed information on the planning step, see the papers of Bosman [15, 16, 17, 18]. Details about the real time control can be found in the PhD thesis of Molderink [62].

The remainder of this section is organized as follows: the next section outlines the requirements of the control methodology followed by a section describing related work and giving a motivation for the chosen control methodology. In Section 2.3.4 the overall control methodology is explained and the last two steps are described. Results achieved with the three-step control methodology are presented and discussed in Chapter 6.

2.3.1 REQUIREMENTS

The goal of the control methodology is to monitor, control and optimize the domestic import/export pattern of electricity and to reach objectives which may incorporate local but also global goals. In this context, local objectives concern energy streams within the building, e.g. lowering electricity import peaks and using locally (in or around the building) produced electricity within the building. Global objectives concern energy streams of multiple buildings, e.g. in a neighborhood, city or even (parts of) a country. To work towards local and global objectives, the control methodology optimizes individual devices on a domestic level.

Since there are a lot of different (future) domestic technologies and building configurations, the control methodology should be able to work independently of the actual configurations. Furthermore, the methodology should be flexible such that new technologies can be added in the future. Consequently, the control methodology needs to be very flexible and generic.

The control methodology should be able to optimize for a single building up to a large group of buildings. Thus, the algorithms used in the control system should be scalable and the amount of required communication should be limited. The control methodology should exploit the potential of the devices as much as possible while respecting the comfort constraints of the residents and the technical constraints of the devices. Furthermore, the control system should consume significant less electricity than it saves.

Using TRIANA, the energy profiles of buildings can be reshaped. Based on the stakeholder, many applications are possible. One possible application of the control methodology is to act actively on an electricity market for a group of buildings. To trade on such a market, an electricity profile must be specified one day in advance. Therefore, it should be possible to determine a forecast of the net electricity profile of the managed group of buildings one day in advance.

Another application can be to react on fluctuations in the grid, for example caused by renewable generation, asking for a realtime management. Reacting on fluctuations requires a realtime control and the availability of sufficient generation capacity at every moment in time to be able to increase or decrease the consumption or generation. To achieve sufficient capacity, again forecasts and a planning must

be determined in advance, in combination with realtime control to react on the fluctuations. Thus, a combination of forecasting of consumption and generation of devices and a planning of the use of these devices is needed.

Once a planning has been made, the methodology should try to achieve this planning as good as possible. Deviations from the planning are often caused by forecasting errors. Therefore, the control methodology should not exploit all the forecasting scheduling freedom during planning to cope with these forecasting errors. This can prevent oscillating behavior caused by over-steering and large fluctuations and peaks in the energy profile, e.g. when a large group of buildings simultaneously react the same on the steering signals.

Furthermore, limitations of the communication links and power lines should be taken into account. Due to the latency of communication links, sending information from the local controllers to global controllers and sending the decision from the global controller back to the local controller requires a certain amount of time. However, deciding whether it is profitable to switch on a large consuming device (e.g. a washing machine) or reacting on fluctuation in generation need to be done virtually instantaneously. Thus, the local controller also has to be able to make these realtime decision independent of the global controller or these decision need to be taken on beforehand. More information about the steering signals and network communication can be found in Chapter 4.

Summarizing, the requirements for the control of a smart grid are:

1. local as well as global control and optimizations,
2. specify and strictly obtain a desired profile up to 24 hours in advance,
3. generic and flexible,
4. scalable,
5. respect the comfort of the residents,
6. limited requirements on the communication links,
7. local controller must be able to work independently.

2.3.2 RELATED WORK

Several control methodologies for distributed generation, energy storage, demand side load management or a combination of these can be found in literature. Roughly these control methodologies can be divided into two groups: 1) agent-based market mechanisms and 2) discrete mathematical optimizations. The advantage of agent-based market mechanisms is that no knowledge of the local situation is required on higher levels, only (aggregated) biddings for generation/consumption are communicated. The advantage of mathematical optimizations is that the steering is more direct and transparent, the effect of steering signals is better predictable. Another important difference is that in an agent based approach often every buildings works

towards its own objectives where in a mathematical approach the buildings can work together to reach a global objective.

Most of the research considers agent based control methodologies. These agent based control methodologies propose an agent per device [66]. The agents give their price for energy production (switching an appliance off is seen as production); via a market principle it is decided which agents are allowed to produce. Since there are a lot of agents, the information is aggregated on different levels in a hierarchical way. The research described in [13] combines domestic generation, consumption and buffering of both heat and electricity. They propose an agent based system where buildings are divided into groups (microgrids) which are loosely connected to the conventional large-scale power grid. In first instance the goal is to maintain balance within the microgrid without using the large-scale power grid. Furthermore, agents use predictions to determine their cost function. Field studies show that 50% of the domestic electricity demand can potentially follow a planned schedule (within certain boundaries). To reach this potential, there have to be incentives for the residents to allow some discomfort.

The PowerMatcher described in [52] and [45] additionally takes the network capacities into account. This control methodology is rather mature; it is a product capable of being used in field tests [79]. In these field tests, a peak reduction of 30% is reached when a temperature deviation of one degree of the thermostat in the buildings is allowed. To be able to reach objectives, business agents can be added that influence the biddings in the auction market.

Dimeas and Hatziargyriou [27] compare the results of individual (local) and overall (global) optimizations. They conclude that global optimizations lead to better results. Next, they claim that agent based control methodologies outperform non-agent based control methodologies since agent based control methodologies take more (domestic) information into account.

In literature, also some mathematical control methodologies are proposed. The research described in [20] proposes a methodology that is capable to aim for different objectives. For every device a cost function is determined for both heat and electricity. Using a Non Linear Problem definition the optimal on/off switch pattern is found. The authors of [40] address the problems of both agent and non-agent based solutions: non-agent based solution are less scalable and agent based solutions need local intelligence and are not transparent. Therefore, they propose a combination: aggregate data on multiple levels, while these levels contain some intelligence. The aggregation is done with a database, the control methodology is rule based. In [24] a control methodology is proposed using stochastic dynamic programming. The stochastic part of the control methodology considers the uncertainty in predictions and the stochastic nature of (renewable) production and demand. The authors of [32] propose a control methodology based on Time Of Use (TOU) pricing, where electricity is cheaper during off-peak periods. They combine this approach with a domestic wireless sensor network: when a Smart Appliance would like to switch on, it has to send a request to a controller. This controller decides, based on the electricity price and the status of the other devices, whether the appliance is allowed to switch on. The TOU pricing can be seen as global steering

signals, however, it is a rough steering signal which is equal for a large group of buildings. Furthermore, it is not known in advance what the impact of the steering signals is.

In [11] a combination of existing tools together with a new developed platform is used. The electricity consumption and production per device is forecasted and using a genetic algorithm the best runtime for every device is determined. The platform exists of two levels, a global level for global optimizations sends steering signals to the local level and a local level control which uses the global steering signals as input and determines the runtimes based on the steering signals while respecting local constraints.

2.3.3 POSITION OF TRIANA

As described above, there are many research projects investigating the optimization of energy efficiency. From the mentioned research, simulations and field tests it can be concluded that the efficiency can be improved significantly, especially when all three types of technologies (consuming, buffering and generating) are combined. All control methodologies split the control into a local and a global part, most of them using a hierarchical structure for scalability. Furthermore, most control methodologies use an online algorithm deciding on device level and some control methodologies use forecasts to adapt the production and demand patterns. However, this forecast data is only used on a local level and, therefore, on a global level hardly any forecasted knowledge is available. The control methodology proposed in this thesis adds the predictability on a global level. This is e.g. required for electricity market trading, insight in the effect of choices and, therefore, is related to dependability.

We chose to use mathematical optimization techniques and a combination of forecasting, offline global planning based on the forecasts and online realtime control based on the global planning. We use forecasts on a device level to be able to predict the overall result, planning to estimate and steer the energy profile of the buildings and the grid, and realtime control to respond to changes (e.g. fluctuations in renewable generation) and work around forecasting errors.

Based on the above considerations, the control methodology consists of three steps and is split up into a local and a global part: 1) local offline forecasting, 2) global offline planning and 3) local online control. Because of scalability reasons, the global planning has a hierarchical structure and can aggregate data and plannings on different levels. More information about the hierarchical structure is given in Chapter 4.

Especially the three steps and the global planning differ from the control methodologies described in literature. Furthermore, the control methodology is not agent based and uses other mathematical optimization methods or heuristics than the control methodologies described above.

Due to the forecasts and planning in advance, the predictability of the global electricity streams is improved. The combination of planning (aggregated knowledge on higher levels) and mathematical optimization result in better dependability

and the combination of planning and realtime control improves the damage control. Furthermore, the amount of communication can be limited due to the hierarchical structure of the planning. Finally, the requirements on the communication medium is low since the local controller can work independently and a lot of information can be sent on beforehand without high latency requirements. Therefore, our three-step control methodology fulfills all requirements mentioned in Section 2.3.2.

2.3.4 THREE-STEP APPROACH

As described in the previous section, TRIANA consists of three steps, which are executed on different levels within the grid. A schematic overview of the three steps is given in Figure 2.5. As shown in Figure 2.5(b), at the lowest level we envision a controller present in the building: the local controller. The higher level nodes are planning nodes and are placed through the grid (see Section 4.3).

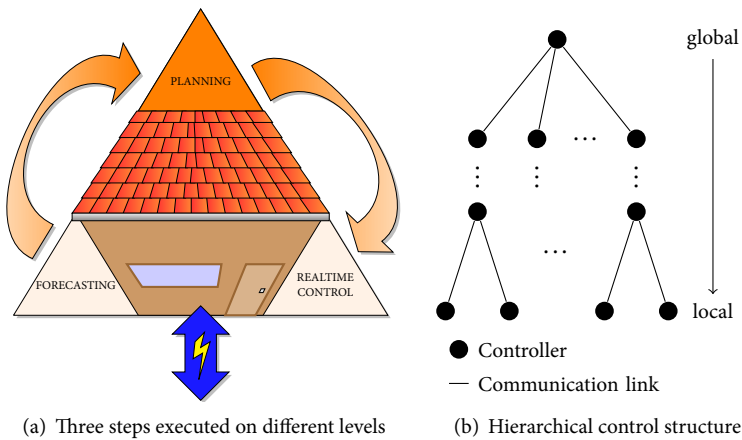


Figure 2.5: Overview of TRIANA

In the first step of the three-step approach, which is described in more detail in Chapter 3, the local controller learns the behavior of the residents and the influence of external factor like the weather. Combining this knowledge with the knowledge about the devices present in the house, the expected energy profile and corresponding possibilities to change this profile, i.e. the scheduling freedom, is forecasted.

In this context all devices which are not steerable are integrated to one device without optimization potential. Information about these non controllable devices are required in the control methodology to be able determine the overall energy profile of the building.

The result of the first step is a forecasting and the corresponding scheduling freedom on a building level. Note that in this step the devices are not controlled yet.

In the second step, the output of the first step is used as input. The forecasted scheduling freedom can be used by a global planner to exploit the optimization potential to work towards a global objective. The global controller can make use of the hierarchical tree structure of control nodes, local or global algorithms and heuristics to optimize the electricity streams in the grid given the objective. In this way, it can give a prediction of expected electricity streams for the forecast period.

To achieve this, the planner tries to reshape the energy profile of the buildings. Based on the forecasts determined in the first step and the desired objective, the planner determines steering signals for its children nodes to ask them to work towards the global objective. Each intermediate node in the tree determines steering signals for its children based on the received steering signals from its parents. Finally, the building controllers (the leaf nodes) can determine an adjusted profile, taking the steering signals into account. This profile is sent upwards in the tree and when necessary the root node can adjust the steering signals again. The planning is thus an iterative, distributed methodology led by the global controller. More information about this iterative, distributed control can be found in Chapter 4. The result of the second step is a planning for each building for the upcoming day (which can be described by the used steering signals for this planning) and a prediction of the resulting global electricity streams.

The third step uses the steering signals of the planning step as input, in combination with information about the status of the devices in the building and the grid. In this final step, a realtime control algorithm decides at which times devices are switched on/off, when and how much energy flows from or to the buffers and when and which generators are switched on. The local control algorithm has to take three different inputs into account while making these decisions. First, it uses the offline steering signals from the global planning. Furthermore, it can use realtime steering signals of the global controller, based on the status of the grid, to respond on fluctuations in the grid, e.g. caused by wind farms. Finally, the local controller has to work around forecasting errors and preserve the comfort of the residents in conflict situations. To reach its goals, the local controller can use, next to realtime optimization algorithms, also local (short-term) forecasts and planning. More details can be found in [62].

The combination of forecasting, planning and realtime control exploits the potential of the overall system at the most beneficial times. The hierarchical structure with intelligence on the different levels ensures scalability, reduces the amount of communication and decreases the computation time of the planning.

As described above, the control methodology is very flexible and the proposed infrastructure supports different algorithms for the global and local controllers.

2.4 EMERGENCE

TRIANA is only one of the many initiatives control smart grids. The emergence of energy-efficient electrification of the society and a sustainable supply leads to a lot of activities from academia, business and authorities. Many projects are started,

conferences are organized and test sites are built up. In this section a short survey of a number of programs, alliances and projects is given.

2.4.1 PROGRAMS AND ALLIANCES

On different levels programs and support initiatives are started. The Institute of Electrical and Electronics Engineers (IEEE), the world's largest professional association for the advancement of technology, organizes many conferences covering all aspects of the Smart Grid with a focus on both academia and industry. The IEEE Power & Energy Society moves its focus more and more towards Smart Grid technologies and IEEE even started a web portal focussed on Smart Grids². Furthermore, a new journal is initiated focussing on Smart Grids, IEEE Transactions on Smart Grids.

A lot of governments provide extra funds for research on Smart Grids and energy departments more and more focus on Smart Grid. For example, the US invests nearly \$ 100 million³ on Smart Grid technology. Other countries like the UK⁴ and the Netherlands⁵ have similar programs. On a more global level, the European Union, the emergence of Smart Grids is also realized. They state that for a successful transition to a future sustainable energy system all the relevant stakeholders must become involved and coordination at regional, national and European levels is essential. Therefore, the Smart Grids Technology Platform [71] has been designed to facilitate this process. Other institutes, like The International Energy Agency also published a roadmap, stating that demonstrations of Smart Grid technology is urgently needed [3].

Probably the largest industrial initiative is the Gridwise Alliance [39], a collaboration of companies to realize a Smart Grid, with as mission 'to transform the electric grid to achieve a sustainable energy future'. But also on a smaller scale there are a lot of industrial initiatives, for example Renqi [47], a collaboration of three Dutch research facilities, the Smart Energy Collective⁶ and the Smart Power Foundations⁷.

Furthermore, a lot of projects concerning Smart Grid are ongoing, both on the technical side as well as on the control side. A lot of projects in Europe are funded by the European programs FP5–FP7 [23, 75]. Next to these European fundings, also on a national levels a lot of research programs started to fund this research. Example are the earlier mentioned \$100 million of the USA and natural gas profits in the Netherlands used for Smart Grid research⁸.

²smartgrid.ieee.org

³www.energy.gov/news/8842.htm

⁴www.decc.gov.uk/en/content/cms/what_we_do/uk_supply/network/smart_grid/smart_grid.aspx

⁵www.twanetwerk.nl/default.ashx?DocumentId=13256

⁶www.smartenergycollective.com

⁷www.smartpowerfoundation.nl

⁸nlenergieenklimaat.nl

2.4.2 TEST SITES

As mentioned in the road map of the IEA, demonstration projects are very important to promote developed technologies. Examples of such demonstration projects are Boulder Colorado, Mannheim-Wallstadt, Meltemi and Powermatching City. In Boulder Colorado a Smart Grid City [31] is implemented with as goal to incorporate 1000 households in the Smart Grid project. The Smart House project[82] is a FP7 funded research project with as goal to demonstrate how ICT can help to achieve maximum energy efficiency. In the project 100 houses in Mannheim, Germany are connected to each other with as main goal to supply washing machines with electricity produced by PV panels. Within the same project, on the Greece island Meltemi, on a camping site PV panels and a diesel generator are installed to reach islanded operation of the camping site. The Powermatching City project [12] in Groningen, the Netherlands, is a testbed for the PowerMatcher[45], an optimization algorithm developed by the Energy Centrum Nederland (ECN). In this project houses are equipped with smart devices and a local controller to test the optimization abilities of the PowerMatcher algorithm. In [43] the results of creating a real VPP with micro-CHP appliances to reduce the load on the distribution network using the PowerMatcher are given.

2.5 CONCLUSION

Changing circumstances ask for a renewed electricity grid to maintain an affordable and reliable supply, to shift towards more sustainable generation and to keep up with the electrification of the energy supply. Since the lifetime of a lot of elements in the electricity grid comes to an end, now is the time to implement a smarter grid. The renewed electricity grid should support distributed (sustainable) generation and should be able to supply the growing demand of electricity. To reach this, consumers have to change from passive consumers to active prosumers, cooperating with each other. Furthermore, plant/grid operators need to maintain grid stability and reliability under the changing circumstances. However, to reach this, a number of technical, economical, legislative and ethical challenges have to be tackled. For the technical challenges, ICT is one of the key technologies. Essential in a Smart Grid is a monitoring and management system that monitors and manages all parts of the grid, from central generation and large scale renewable generation, via transportation up to consumption/generation at the consumers on device level, in a cooperative way. One of the possible smart grid management methodologies is TRIANA. TRIANA consists of three steps and is split up into a local and a global part: 1) local offline forecasting, 2) global offline planning and 3) local online control. TRIANA is only one of the many initiatives to investigate the management of smart grids. The emergence of smartening the grid and updating the electricity supply chain is emphasized by the numerous initiatives worldwide from the European Union, governments, industry as well as from the academic world.

FORECASTING

ABSTRACT – In this chapter the first step of the three step methodology, the forecasting step, is described. Forecasts are used to determine the scheduling freedom of a device. Therefore, for each individual device a forecast is made, since device specific information and restrictions are required in the planning and control of the device. Due to the enormous amount of devices in the grid, each with individual information, restrictions and environment, the forecasting is performed for each individual building by a local controller. This results in a scalable system, since no information about the device and the environment needs to be communicated and the required computational power required for forecasting is distributed. By performing the forecasting locally, local building/resident specific characteristics can be taken into account. This can improve the prediction quality. Furthermore, since the environment of a device may not be static due to the stochastic nature of the residents, local information can be used to adapt to changes. By using locally harvested data, a fully autonomous forecasting system without direct interaction with the residents can be built. As an example of the forecasting system, the forecasting for a micro-CHP appliance is researched in more depth in this chapter. Here, local information like historical heat demand and weather information are considered as good candidates to be used as input data for the heat demand forecasting. Multiple representations of these influence factors have been analyzed. In the last part of this chapter, a simulated annealing method is presented, which has been used to determine good representation of the influence factors and other forecast model parameters.

As mentioned in Chapter 2, the electricity demand can roughly be divided into three parts. The largest part is the more or less regular base load, mainly influenced by the season. On top of the regular base load there are specific fluctuations caused

Parts if this chapter have been presented at [VB:1] and [VB:2] .

by the day to day behavior of the consumers. The smallest part of the electricity demand is the electricity needed to keep the grid stable and functioning properly.

In many European countries, for each of these parts of the total electricity markets exist on which electricity can be traded. In long term contracts the base load is traded, which is roughly 80% of the total electricity demand. Since this base load is more or less fixed and easy to predict, the prices of electricity in these contracts are low.

The other 20% of the electricity demand is traded on short-term markets like the day-ahead market and the balancing market. On these markets, electricity demand and production of the upcoming day / quarter of an hour are traded. For this, forecasts are made of the expected energy consumption. Electricity retailers try to make good forecasts about the expected electricity demand of their customers and power suppliers use forecasts to optimize the control of the power plants, and to purchase their fuel. On the day ahead market, energy is traded 24 hours in advance, based on forecasts of the expected electricity consumption. Energy traded on the day-ahead has to be delivered/consumed exactly as specified in the contract made on this market. Deviating from the agreed electricity production/consumption results in a shortage/surplus of electricity, which can cause instabilities in the grid. Therefore, every deviation from the agreed electricity pattern is penalized by a central authority.

The agreed electricity patterns are based on forecasts, which are in general imperfect. As a result, deviations from the agreed energy pattern are normal. To stabilize the grid, balancing power is required, which is traded on the balancing market. This small part of electricity demand needed for balancing is only known a short time in advance or in the worst case when the customers have already switched on their devices. Due to the short time between purchase and delivery of electricity on this market, and thus necessary capacity reservations, prices are very high. Once electricity is purchased on this market, the agreed electricity must be delivered within a couple of minutes.

The introduction of Smart Grid technology enables the possibility to achieve a better matching between demand and supply. Electricity producers and consumers can cooperate together to adjust their production/consumption pattern to keep the electricity network stable. For example, some devices of consumers might have adjustable or shiftable load, e.g. a freezer might actually be turned on a little bit earlier to consume a surplus of electricity. Or a washing machine might use a little bit less electricity to heat up the water. If a large group of devices located at consumers can be controlled, the energy consumption of a part of the total energy consumption can be controlled. Considering that households use 27% of the total electricity consumption in the Netherlands [78] and that 50% of this consumption can be shifted in time [13], controlling a large group of devices has a significant potential. This flexibility in the usage of the devices can be exploited on the short-term energy markets (balancing/day ahead), giving this control quite some value. If devices in other kinds of buildings, like in offices, can also be shifted/alterd, even a bigger part of the total electricity profile can be altered.

In order to adjust the consumption pattern of households by controlling the

devices, the device status and their possibilities to shift/adjust their consumption pattern, i.e. their flexibility, must be known in advance. Reusing the freezer example, it might be beneficial to reduce the electricity consumption, by switching off the freezer for a certain period time. This is only possible if the temperature inside the freezer allows this, more precisely by keeping the freezer off the temperature must be kept below the required maximum temperature during the whole period. Based on the characteristics of the freezer, like insulation and power consumption, and the expected interaction of the residents with the freezer, the possibilities of the freezer for altering the consumption pattern can be determined. This freedom in controlling the freezer is called the *scheduling freedom*.

To determine the scheduling freedom, device specific information and the expected interaction with the device's environment are required. For this reason, forecasts for a controllable appliance with scheduling freedom are made. The rest of this chapter is about the forecast methods to determine the scheduling freedom. First, the requirements of the forecasts are given. Then, the related work is given in Section 3.2. In Section 3.3 the chosen approach is given. A general description of the forecasting model is given in Section 3.4. In Section 3.5 multiple approaches for individual heat demand forecasting to determine the scheduling freedom of a micro-Combined Heat and Power (CHP) are given. We finalize this chapter by describing a search method to determine adequate parameters for the forecasting model.

3.1 REQUIREMENTS

Based on forecasts, which is the first step in the three step approach described in Chapter 2, the scheduling freedom of individual devices is determined. In the second step, this scheduling freedom of devices is used to plan the runtime/control of the appliances to reach a certain objective. In the last step, based on the planning made in step two and, when necessary new forecasts using more recent information, the devices are controlled to achieve the objective.

Key in this approach is the autonomy of the whole system. The system should just work, without any intervention of the users. Furthermore, the system should be scalable. Since eventually each device is controlled individually, a forecast per device is required. There are millions of devices connected to the grid, each with their own device-specific constraints and placement within their own environment. The system should be able to make forecasts for this large group of devices to exploit the available scheduling freedom of this large group of appliances. Preferably, the device-specific constraints and environment should be taken into account while forecasting, assuming this improves the quality of the forecasts.

Another important aspect of the forecast is the timespan of the forecast and the time the forecast must be available. When the forecast is required in the last step, the real time control of the device for the next hour, the forecast only has to be available just before making a control decision. However, when using the system to

control a large group of devices to act on the day-ahead market, the forecast of the whole next day must be available one day ahead.

When making a forecast for a device, the environment in which the device is placed is very important. However, this environment may not be static. When a device is used intensively by the residents of a building, information about this usage is important. However, human behavior might change, and this change should be taken into account when generating a forecast. These changes can be caused by many factors, e.g. the season of the year or changes in habits of the residents. The forecasting system should be able to adapt to these changes. It should periodically evaluate the quality of the forecasts and perhaps adjust the forecast model to incorporate recent changes, e.g. when a family is on holidays.

Summarizing, the system responsible for the forecasts should fulfill the following five requirements:

1. Generate forecasts for each individual device: device specific information is required in order to control the device.
2. Be scalable: the Smart Grid contains a very large group of devices.
3. Work for different timescales: the system must be able to make forecasts at least one day ahead, but also a few hours ahead.
4. Run autonomously: the system must be running on its own, without requiring interaction with the residents of the building.
5. Be adaptable to changes: the environment of the device may not be static and relevant changes in the environment should be dealt with.

3.2 RELATED WORK

As described above, forecasting is very common in the energy supply chain. Accurate load forecasting has a great potential for electricity producers. Bunn and Farmer [19] claimed in 1985 that one percent increase in load forecasting error in the British system resulted in a loss of 10 million pounds per year. Due to the economical impact of good electricity demand prediction, a large amount of research has been performed in this topic. During the years, a lot of methods for demand forecasting have been analyzed. According to an overview of Alfares and Nazeeruddin [4] the techniques can be categorized into nine categories: (1) multiple regression, (2) exponential smoothing, (3) iterative reweighted least-squares, (4) adaptive load forecasting, (5) stochastic time series, (6) Autoregressive moving average with exogenous variable models based on genetic algorithms, (7) fuzzy logic, (8) neural networks and (9) expert systems. Furthermore, a trend towards fuzzy logic, genetic algorithms, expert systems and neural networks can be seen. Comparative studies showed that fuzzy logic and neural networks outperform the autoregressive models. Based on these reasons, and the requirements set in the previous section, neural network techniques are used in this work.

In the use case of the Virtual Power Plant (VPP), the scheduling freedom of micro-CHP appliances needs to be determined. This scheduling freedom is determined by the heat demand of the household and the heat buffer connected to the micro-CHP. Therefore, heat demand forecasting on an individual level is required. Although electricity load forecasting is a well studied topic, heat demand forecasting research has been limited. Dotzauer [29] describes a heat demand prediction for large scale systems. The presented model forecasts the load by combining two functions. The first function is a (piecewise) linear function used to determine the influence of the weather and uses outdoor temperatures as input. The second function is used to describe the social factor and is modeled as a constant (total load – weather dependent part). During training, the best combination of the two functions is searched.

Serban and Popescu [74] describe a heat demand prediction for district heating systems using times series analysis. A time series can be described as output of a system that has as input a white noise signal. The output consists of a (linear) combination of observed data (history) and the present input.

In the work of Nielsen and Madsen [63] the heat consumption in a district heating system is forecasted. In this work, the theoretical knowledge of the system is combined with measurements performed on the system to determine a mathematical description of the system. The heat load is forecasted using meteorological forecasts and the developed mathematical description.

Similar to the forecasting of heat demand is the forecasting of cooling load. Ben-Nakhi and Mahmoud [10] describe the forecasting of the cooling load of a building using neural networks techniques by using hourly temperatures as input for the network. Yao et al. [86] forecast the hourly cooling load using a combination of multiple forecasting techniques.

Common in related work mentioned is that the forecast of load is made for large systems, not individual households. Forecasting heat demand for individual buildings still is a unexplored field of research. In contrast to forecasting individual heat demand, in many studies models of houses and buildings have been made to analyze the impact of the energy demand on the whole system. In the work of Heller [41] the heat-load is modeled for large system, in which the heat demand is divided into multiple elements: space heating, domestic hot-water preparation, distribution loss and additional work-day loads. By combining models of each element, an overall model for the total heat load of a district-heating is generated.

A method of formulating the energy load profile of a domestic building in the UK is described by Yao and Steemers [85]. This energy profile comprises both heat and electricity, where 56% of the energy-consumption was used for space heating and 24% for domestic hot water. The profile is mainly determined by behavioral and physical components. The behavioral component is defined as the habits of the occupants of a building, which is slightly influenced by season. The physical component comprise the climate, the physical properties of the building like size and design. Here, multiple types of houses and occupants habits are used to model the heat load using a thermal resistant model.

The potential saving of using a micro-CHP appliance in the house is analyzed by

Pearce et al. [69]. Here, the heat load is modeled using a ‘house thermal equivalent circuit’. Using this circuit, consisting of thermal constants of a house (information about how heat flow of and within the house), the outside temperature and the desired room temperature, the heat load of the building is determined.

In the work of Pearce et al. [69], no behavior pattern is taken into account for modeling the heat load of a building. This is in contrast to Lampropoulos et al. [57], in which the importance of including behavior of small prosumers in power system planning is stressed. In their modeling methodology a combination of deterministic (devices operation), probabilistic (user groups, user behavior) and stochastic models (weather and external parameters) are used. Studies in the US, The Netherlands and the UK show that 26-36% of the domestic energy end-use variations is due to the behavior of the residents [83].

3.3 APPROACH

As described in Section 3.1, the forecast system should be scalable and capable of forecasting data for each individual device. One may chose to do the forecasting for each device centrally to be able to use similarities between consumers. However, such an approach is not really scalable. Above a certain fleet size a central system may get too inefficient. Furthermore, in most of the cases there are consumer specific characteristics which may be worthwhile to integrated in a forecasting scheme. In a central approach this requires a lot of communication, again resulting in scalability and privacy problems.

If we integrate the forecasting system into a local control system located in a house, local individual information can be incorporated. The information for adequate forecasting can be harvested locally. For example, human behavior has a big impact on the energy usage of a household. Although there are some similarities between households, still each household has its own habits and characteristics. External factors like insulation of the house, required comfort levels of the residents etc, can be different for each household. Assuming that the local control system continuously measures the electricity flow of each device, historical energy demand patterns containing information about this specific household can be generated. Other relevant information, like for example weather information, can be collected locally or perhaps downloaded from a central authority or the internet. By continuously collecting the required information locally, the quality of the information increases.

Another advantage of continuously collecting the required information locally is that changes in the environment are present in the collected data. When necessary, the forecast scheme can reevaluate the forecasts and can, when necessary, adjust the forecasting scheme to fit better to recent changes in the household.

Furthermore, when using the local approach the information does not need to be transmitted to a foreign system, which improves both scalability and privacy issues of the smart grid.

Another strong point of using individual forecasts and the three step optimization methodology is that when the overall system breaks down, a fallback mechanism can be build in. For example, if the local controller is part of a Virtual Power Plant or a micro-grid, where the local controller cooperates with other local-controllers in other buildings and perhaps a grid/fleet controller, and for some reason there is (temporarily) no communication possible, the local controller can still use the local forecast for a local optimization objective, such as peak shaving.

Summarized, a local forecast improves both privacy, scalability and the quality of the information required to make a good forecast. The processing power available in a local controller can be exploited since the forecasting software runs on each the local system instead of centrally. Since most required information is gathered locally, less information needs to be transmitted by the local control systems. Furthermore, the information contains household specific details which may be required for certain devices. In case of communication failure within the smart grid, a fallback mechanism can be enabled, using the local control system for local objectives.

3.4 FORECASTING MODEL

As described in the previous section, for each device forecasts are made to determine the scheduling freedom of the device. What needs to be forecasted is dependent on the device (requirement 1). The function of the forecasting model is to generate expected data of future events, based on the information currently available and information from the past. The forecasting model of a device can be seen as a function $\mathcal{F}(\mathcal{I}, \mathcal{F}_p) \rightarrow \mathcal{O}$, where \mathcal{I} is the input for the forecast model, \mathcal{F}_p are the forecast model parameters and \mathcal{O} the forecasted output values.

Function \mathcal{F} should map the input \mathcal{I} to the output \mathcal{O} . However, in general this mapping is unknown and must be estimated. Furthermore, as described in the previous section, the mapping might change over time since the device or the environment of a device can change as well (requirement 5). On the one hand, the changes can be present in the input \mathcal{I} . But on the other hand, the relation between \mathcal{I} and \mathcal{O} can change as well. Since the system should be running autonomously (requirement 4), the forecasting system should be able to learn this relation without any intervention of the residents. And since the relation between the input and output might change over time, the system should periodically evaluate the performance of the predictions and perhaps relearn, incorporating recent changes. While relearning, a good trade off between focusing too much on recent (incidental) changes and learning real consistent changes in the environment must be made. Furthermore, it may be necessary that the model has to relearn often. Since the (re)learning process is executed on a local controller, which is expected to be a power efficient device which little computational power, the learning process cannot be too computational intensive.

In this work neural network techniques are used as a mapping function. Neural networks are simple, robust, and fast models which can learn, based on given training examples, (non-linear) relations between the input and the output. The learning

process can be self-supervised, i.e. without human interaction. By periodically relearning the neural network, the forecasting model is adaptive to changes.




Figure 3.1: A single processing unit of a neural network

Neural networks, as described in [53], are computational models based on biological neurons. They are able to learn, to generalize, and to cluster data. Their operation is based on parallel processing. A neural network consists of a pool of simple processing units, which communicate by sending signals to each other over a large number of weighted connections. An example of a processing unit, called a neuron, is depicted in Figure 3.1. Each neuron basically performs one task. It receives inputs from neighbors and computes an output signal, based on the received inputs, which is propagated to other neurons.

Within the neural network three types of neurons exist: input neurons which receive their input from outside the network, output neurons which send data out of the network and hidden neurons whose in- and output remain inside the network. Neurons are connected to each other via weights w_{jk} , which determines the effect of a signal of neuron j on neuron k . The total input of neuron k normally is simply the weighted sum of the separate outputs of the neurons connected to k plus a bias or offset θ_k , but other propagation rules exist [53].

The activation function \mathcal{A}_k of neuron k determines the new level of activation (the output) based on the effective input s_k . There are many possible activation functions, but generally some sort of threshold function is used. Examples are hard threshold functions, like shown in Figure 3.2(a), a semi-linear function (Figure 3.2(b)) or a commonly used smoothing limiting function (Figure 3.2(c)). The goal of the threshold function is that the artificial neuron, similar to biological neurons, only fires after receiving sufficient stimulus from other neurons.

The ‘knowledge’ of a neural network is present in the weights w in combination with the activation functions. Using the weights and the activation functions, the neural network approximates a desired function. A neural network is configured such that the application of the neural network to a set of given inputs produces the desired outputs (which are also given), i.e. the right weights in the neural network are set. If a priori knowledge is available, this can be used to pre-specify the weights and choose the appropriate activation functions. If this is not the case, the desired function must be determined in the form of learning.




Figure 3.2: Possible activation function

The learning process consists of given examples of input and output pairs for the neural network. The idea is that the required mapping function is present in the example data and that by providing sufficient examples to the neural network the neural network itself is capable of learning this correct mapping. There are many learning algorithms, but the basic principle is to adjust the weights of the hidden neurons. The weights in the hidden neurons must be adjusted such that the application of the neural network to the inputs gives the correct response, i.e. the output of the neural network is equal or as close as possible to the example outputs. By adjusting the weights according a certain training rule, it is tried to minimize the error between the network output and the expected output.

A neural network can only properly learn if the example input/output pairs during the training phase contain enough information of the function to be learned. If the examples during the training phase are not representative for the problem, a wrong relation between the input/output can be learned. Generally, when training a neural network, the whole set of example is divided into three subsets: a training set, a validation set and a test set. The examples are often divided randomly to get a good distribution of the examples over the set. The training set is used during the training phase, which can consists of multiple iterations of altering the weights. During training a general mapping applicable to all the examples should be found. However, it may be possible that instead of finding a generally applicable mapping, overfitting occurs. In the case of overfitting, a mapping too specific for the given examples is learned. Therefore, the validation set is used after each learning iteration to detect overfitting. After each iteration, the neural network is tested against the validation set. If after a training iteration the performance of the neural network on the validation set decreases, the training process gets too much focussed on the examples in the training set. In other words, the training algorithm is applied as long as the performance of the neural networks on both the training set as the validation sets keeps improving. After training, the test is used to analyze how general the trained neural network is, i.e. how well the neural network performs on samples which where not using during training. Therefore, the test set contains examples which were not used during training.




Figure 3.3: Example of used network structure in the prediction model using 15 hidden neurons.

In our approach we use a multi-layer feed-forward network (see Figure 3.3). Each layer consists of neurons which receive their input from a layer directly in front (left in the figure) and send their output to a layer directly behind (right in the figure). The layers in between the input and output layers are called hidden layers, where the neurons in a hidden layer are referred to as hidden neurons. There are no connections between neurons within the same layer. The first, most left layer, performs no computation but only distributes input, and this layer is therefore not counted as a layer. The example given in Figure 3.3 is thus a two-layer feed-forward neural network. Important parameters when using feed-forwarding neural networks is the number of layers used and the amount of neurons present in each layer. The amount of layers/neurons required can be very specific for the problem and is often determined by starting with a relatively small network, i.e. with only two layers and a low number of hidden neurons. The neural network is then trained and the performance of the neural network is determined. This process is repeated

after adding one or more hidden neurons to the neural network. This continues as long as the overall performance of the new neural networks keeps improving. To prevent getting stuck on a local minima while determining the proper network size, one can choose to test a limited amount of larger networks and see whether these larger networks all perform less than the current neural network.

3.5 RESULTS: HEAT DEMAND FORECASTING

One of the use cases of the three step approach is the creation of a Virtual Power Plant consisting of a large group of micro-CHP appliances. A micro-CHP appliance is a system that consumes natural gas and produces heat and — as a by-product during the heat production — electricity. It generates electricity at the kilowatt level, allowing these units to be installed in an individual house. They can be connected directly to the domestic heating and electrical systems, which leads to a very high efficiency (up to 90%) in usage of primary energy. The heat is used for the heat demand in the home such as central heating, showering, hot water taps etc. The electricity can be used in the home or, when not needed, be exported to the electricity distribution network. When a micro-CHP appliance is connected to a heat buffer, the periods in which heat, and thus electricity, is produced are decoupled from the moment the heat is consumed. Ultimately, the control of a micro-CHP can be shifted from a heat demand driven to electricity demand driven control scheme, within the limits set by the heat demand and the heat buffer.

For a Virtual Power Plant consisting of many micro-CHP appliances, the heat demand of the households determines the amount of electricity that can be produced. Considering that on average in the European Union 69% of the household energy consumption goes to heating and 14% to hot tap water [2], quite some heat — and thus electricity — can be generated using a micro-CHP appliance. When a heat buffer is installed in the house, a lot of scheduling freedom is available. This makes a micro-CHP appliance a good candidate to be used in a VPP. Based on the expected heat demand, the buffer characteristics like dimension, loss and State of Charge, a schedule for the runtimes of a micro-CHP can be derived. There are in general many valid schedules of a micro-CHP providing all the heat demand. However, some may be more beneficial than others and, therefore, a most beneficial runtime for a micro-CHP can be searched for. In order to determine the scheduling freedom of a micro-CHP, the expected heat demand for the coming day must be forecasted.

3.5.1 FORECAST MODEL

In first instance, we are interested in the heat demand for the upcoming day, since in the second step of TRIANA a planning for the day ahead market is needed. For a good planning it is preferable to forecast the heat demand as accurately as possible (in the order of minutes). However, the amount of information available to us is not enough to give such an accurate forecast. For this reason, we forecast the heat demand on an hourly basis. Furthermore, it is questionable if a reasonable forecast for smaller time units is possible 24 hours ahead.

To incorporate the differences in human behavior on different days in the week, we use a neural network for each weekday. It might be the case that for some weekdays the important influence factors and the amount of their influence differs from other weekdays. As a consequence, the amount of required neurons might be different as well. By using a different neural network for each weekday, these differences can be dealt with.

The neural networks used consists of two layers (the input layer, one hidden layer and the output layer). The reason for this is to keep the neural networks simple. For the hidden layer, the best network size (number of hidden nodes) is determined as described in Section 3.4.

3.5.2 INPUT SELECTION

To set up a neural network to predict the expected heat demand, we must analyze the factors that may have influence on this heat demand of a household. The forecasting model should learn the relation between these influence factors and the heat demand. Examples of such influence factors are the size, insulation and location of a house. The combination of these tree factors determine how much heat is needed to reach and keep a certain temperature in a house for central heating. Other important factors are the thermostat program, the required temperature in the house, the amount of residents and their behavior. Residents staying at home during the day means more central heating demand and a larger number of residents also means more required hot tap water demand.

Another important factor is the weather. During winter more central heating is required. When the outdoor temperature is low and there is a strong wind, especially old houses without double glazing and bad insulation loose a lot of their heat. There may be many other influences, but the most relevant factors can be roughly divided into three categories: house, human behavior and weather [81].

Since houses do not change that often, we may consider the characteristics of the house static. Because of this, the neural network should be able to learn the characteristics since they are present in the data used.

For the behavior factor, information about the thermostat program and user overwritten settings might give important data. However, we assume that the thermostat program is fixed and thus can be learned by the system. If the thermostat program is not fixed, it is dependent on the behavior of the residents and cannot be determined on before hand. Other information about the user, for example holidays, is not used since this requires interaction with the user. Although it might be possible to use information, such as electronic agendas, it also introduces additional issues like privacy. Therefore, this addition is left for future work.

We aim for a system running autonomously. Furthermore, people have different behavior on different days on the week and their behavior changes in time. Changes in behavior should be learned quickly in order to cope with changes like holidays.

To learn the behavior of the residents, historical heat demand data is used. In the test set-ups used in this research, four households were equipped with a

WhisperGen¹ and a 220 liter heat store. This heat store is used for both central heating demand and hot tap water. To determine the heat demand of these houses, status information of the heat stores and installed micro-CHP appliances of these houses, kindly made available by Essent and GasTerra (two local energy companies), were used. For each household, the status of the heat store and the micro-CHP appliance were monitored on a minute basis for roughly one year, starting around the beginning of the year 2007. From this information, the heat demand for these four households is derived by combining the micro-CHP appliance status and the changes of the tank levels. Since the micro-CHP units were used for testing, some gaps in the measurements data occurred. All days with less than 1200 of the total 1440 measurements are filtered out and not used as input for our model.

For the weather, outdoor temperatures and information about wind and sun can be used. Information about temperatures and wind speeds can be obtained from weather stations nearby the households. These weather stations output Meteorological Aerodrome Reports (METAR) every half hour, from which the temperatures and wind speeds are extracted. When deploying the prediction system, the half hourly weather information is of course not available yet and forecasted values must be used.

The output of the prediction model is the expected heat demand for the next day on an hourly basis. More fine grained prediction of the expected heat demand might be wanted, but given the available input data we think this hourly heat demand prediction is a good tradeoff between accuracy and the quality of the information. Therefore, the neural network consists of 24 output neurons, one for each hour of the day.

The possible inputs \mathcal{I} for our prediction model have to be selected from all possible inputs \mathcal{I}_{pos} , which consists of historical heat demand data H , outdoor temperatures T and wind speed information W , i.e. $\mathcal{I}_{\text{pos}} = H \cup T \cup W$.

The historical heat demand data set H consists of hourly heat demand data of recent days H_n ($n \in \{-1, -2, \dots, -7, -14\}$), where n is the number of days relative to the the day that is being forecasted. For example, H_{-1} is the hourly heat demand data the day before the day this is being forecasted.

For information about the weather only very recent information is used. It is expected that for example the weather information a week earlier is not related to the expected heat demand. Therefore, only the (forecasted) outdoor temperatures for the to be forecasted day and the outdoor temperatures one day before the to be forecasted day are used. These temperatures can be represented in different ways, i.e. per half hour, per hour etc. Therefore, three different representation of outdoor temperatures are used. Therefore, the above mentioned set T with the possible outdoor temperatures consists of $T_{n,t}$, where n ($n \in \{0, -1\}$) is the number of days relative to the day to be forecasted and t ($t \in \{mm, 30, 60\}$) is the chosen representation of the weather. If $t = mm$, only the minimum and the maximum outdoor temperatures are used. For $t = 30$ and $t = 60$ the temperatures per half

¹<http://www.whispergen.com>

Table 3.1: Possible inputs for a day

Input	Description
H_{-1}	Hourly heat demand 1 day earlier
H_{-2}	Hourly heat demand 2 days earlier
H_{-3}	Hourly heat demand 3 days earlier
H_{-4}	Hourly heat demand 4 days earlier
H_{-5}	Hourly heat demand 5 days earlier
H_{-6}	Hourly heat demand 6 days earlier
H_{-7}	Hourly heat demand 7 days earlier
H_{-14}	Hourly heat demand 14 days before D_n
$W_{0,mm}$	The forecasted minimum and maximum windspeed
$W_{0,30}$	The average forecasted windspeed (per half hour)
$W_{0,60}$	The average forecasted windspeed (per hour)
$W_{-1,mm}$	The minimum and maximum windspeeds 1 day earlier
$W_{-1,30}$	The average windspeed (per half hour) 1 day earlier
$W_{-1,60}$	The average windspeeds (per hour) 1 day earlier
$T_{0,mm}$	The forecasted minimum and maximum temperatures
$T_{0,30}$	The average forecasted temperatures (per half hour)
$T_{0,60}$	The average forecasted temperatures (per hour)
$T_{-1,mm}$	The minimum and maximum temperatures 1 day earlier
$T_{-1,30}$	The average temperatures (per half hour) 1 day earlier
$T_{-1,60}$	The average temperatures (per hour) 1 day earlier

hour and per hour respectively are used. The set of possible windspeed inputs W is constructed in a similar way as T . A complete overview of \mathcal{I}_{pos} is given in Table 3.1.

As described in Section 3.4, a neural network has to be trained and the available training data is separated into multiple sets to function as training and validation data. Often this separation is done randomly.

In case of the heat demand prediction, the reason for choosing a random set for training is to make the prediction more general and to find as much behavior as possible. However, in case of human behavior and seasonal weather data, this might not be the best choice.

Instead of using random weeks of the whole data set, only a limited subset could be used for a training set. Using this limited subset, a good tradeoff between finding general behavior and recent behavior can be made. This implies that each day, (re)learning is required using information of the last weeks. Therefore, a ‘sliding window’ [7] is introduced, where only data from the last weeks is used while training the neural network. By using this sliding window, more recent behavior is used to learn the behavior of the residents. This makes the prediction model more flexible and adaptable. Furthermore, we expect that behavior may change during the year, and that current behavior resembles more to the recent behavior instead of general behavior during the whole year.

One important decision to make while using the sliding window approach is how many weeks of history should be used for training. Adding more weeks might give more information about the general behavior, but this can cause ‘too general’ or outdated behavior. Using too few weeks can cause overfitting to the recent behavior, missing the general behavior people might have on a certain week day.

The addition of a sliding windows also introduces the possibility to use different input for each day that has to be forecasted. As described above, it might be the case that due to changes in the environment, other influence factors become more dominant while forecasting the expected heat demand. In this work, the same input set is used for all the evaluated days. Using different input sets for each individual day that is evaluated is left for future work.

3.5.3 FORECASTING QUALITY

Once the networks are trained, we have to determine the performance of the neural network. As the heat demand forecast is used to optimize the runtime of a micro-CHP, two factors are important. First, the forecasted heat demand profile for the day should be as close to the actual heat demand as possible. In other words, the amount of heat demand and when the heat demand peaks occur should be predicted accurately, since these are important influence factors on the runtime of the micro-CHP. For this reason, the Mean Absolute Percentage Error (MAPE) is used as a quality measure for the forecast. The MAPE, as given in (3.1), gives the mean absolute deviation of the forecasted values from the actual values over all days for a given house. This error can be used to express how ‘well’ the forecasted heat demand is matched with the real heat demand. In (3.1) D_h denotes the set of days that are forecasted for a house h , $\mathcal{O} = \mathcal{F}(\mathcal{I}, \mathcal{F}_p)$ denotes the forecasted values and \mathcal{A} the actual values of the heat demand.

$$MAPE = \frac{1}{24|D_h|} \sum_{d \in D_h} \sum_{t=1}^{24} \frac{|\mathcal{O}_{d,t} - A_{d,t}|}{F_{d,t}} \quad (3.1)$$

$$F_{d,t} = \begin{cases} A_{d,t} & , \text{ if } A_{d,t} \neq 0 \\ \frac{1}{24} \sum_{i=1}^{24} A_{d,i} & , \text{ otherwise} \end{cases} \quad (3.2)$$

The second important factor is the deviation from the actual total heat demand, which should be predicted accurately. If for some reason a heat demand peak is forecasted an hour too early, it is still important that the volume of the forecasted profile is correct since we should be able to produce enough electricity. To express this, the Mean Percentage Error (MPE) is used, which is defined as

$$MPE = \frac{1}{24|D_h|} \sum_{d \in D_h} \sum_{t=1}^{24} \frac{\mathcal{O}_{d,t} - A_{d,t}}{F_{d,t}},$$

where $F_{d,t}$ is given by (3.2). For example, a positive MPE means that the prediction model forecasted a higher total heat demand than the actual heat demand.

Since both errors are important, they are both used to determine the performance of a network. We have chosen to determine the error for a network via

$$E_{\text{total}} = MAPE + |MPE|.$$

3.5.4 RESULTS

For the first computational test we used as input for our prediction model an input vector which consist of three groups of data: (a) the heat demand of the previous day (24 values), (b) the heat demand of the same day one week earlier (24 values) and c) the average (predicted) temperatures per hour of the day (24 values) (see Figure 3.3 on page 46) [6]. More formally, the input set of specified by $\mathcal{I} = \{H_{-7}, H_{-1}, T_{0,60}\}$. H_{-7} was selected since it is expected that there is a regularity in the behavior of the residents on each weekday. For example, more people are at home on a Wednesday afternoon since then the schools are free. H_{-1} is chosen since there is a strong relation between the heat demand of the current day and a day earlier. The temperature was added to \mathcal{I} since it is expected it is an important factor that influences the heat demand.

To determine the best model parameters, i.e. the the amount of hidden neurons, given the selected \mathcal{I} , all combinations of network sizes, weekdays and households in our database have been trained using the Fast Artificial Neural Network library [64]. While training, the mean squared error was minimized using the RPROP [72] training function, where a sigmoid activation function was used.

Table 3.2: The number of days present in the validation sets


House	Sun	Mon	Tue	Wed	Thu	Fri	Sat	Total
1	21	21	18	20	19	21	23	143
2	19	18	17	18	18	18	18	126
3	17	18	17	18	17	18	18	123
4	13	13	11	12	13	14	14	90

As mentioned above, the database with heat demand data contained some gaps. The number of days of the validation sets after filtering are given in Table 3.2.

The results of the initial approach are given in Table 3.3. We can see that the heat demand of some households is more difficult to forecast than the demand of other households. For example, household 3 is the easiest to forecast. This household probably has a very regular heat demand, probably due to a very fixed thermostat program and good insulation. Another interesting observation is that for this household Saturdays show significantly less performance than the other days. Analyzing the MPE column for each weekday in Table 3.3, all the prediction models have a positive MPE value. The prediction model thus always predict a higher total heat demand with respect to the actual values. Again, household 3 shows the best performance.

Table 3.3: The Mean Absolute Percentage Error and Mean Percentage Error of the initial results

H	Sun		Mon		Tue		Wed		Thu		Fri		Sat	
	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE
1	1.96	1.16	2.59	1.81	6.04	5.43	2.13	1.22	2.64	1.88	1.18	0.26	2.75	1.97
2	1.57	0.73	2.10	1.16	2.30	1.57	3.95	3.27	2.39	1.61	3.48	2.65	4.62	3.98
3	0.54	0.03	0.56	0.16	0.46	0.10	0.55	0.10	0.50	0.08	0.53	0.13	5.40	5.17
4	1.09	0.17	3.42	2.84	1.22	0.35	2.36	1.75	1.54	0.63	2.49	1.68	1.39	0.60

Figure 3.4: Heat demand forecasts and actual values for household 1 in July using $\mathcal{I} = \{H_{-7}, H_{-1}, T_{0,60}\}$

To give an impression of the forecast quality, the forecasted and real heat demand of the first week of July of 2007 of household 1 is depicted. As can be seen in the picture, the actual heat demand consist of many high peaks, with moments of no heat demand between the peaks. The trained neural networks is unable to properly forecast these high peaks, resulting in a MAPE is 17.02 and the MPE is 16.56 for this week.

Although the initial results look promising, a better heat demand forecast is required. Especially for certain households there is a large space for improvement. Furthermore, by dividing the input data randomly into the training and validation sets, like in the initial results, a large data set is required before training and forecasting can start. In other words, when such a forecasting system is used in real life, it takes a very long period to collect the information to train the forecasting model.

A possibility to solve this problem is to add more information about the environment and human behavior to the input set. By adding more information to the training data, less training data might be required to properly train the neural network. As mentioned earlier, the user-interaction should be kept to a minimum, thus information about human behavior is not a good candidate to add to the input

Table 3.4: The Mean Absolute Percentage Error and Mean Percentage Error with wind, no sliding window

H	Sun		Mon		Tue		Wed		Thu		Fri		Sat	
	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE
1	2.10	1.35	2.26	1.46	6.32	5.61	3.94	3.13	3.48	2.74	1.65	0.63	2.09	1.21
2	2.48	1.65	4.74	3.91	2.43	1.72	2.75	1.94	1.92	1.06	2.19	1.39	3.24	2.64
3	0.47	0.05	0.55	0.16	0.53	0.07	0.48	0.02	0.61	0.10	0.57	-0.06	2.61	2.14
4	1.91	1.14	1.99	1.40	1.37	0.41	2.69	1.89	1.34	0.57	2.97	2.13	1.52	0.78

Table 3.5: The Mean Absolute Percentage Error and Mean Percentage Error without wind, 4 weeks sliding window

H	Sun		Mon		Tue		Wed		Thu		Fri		Sat	
	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE
1	2.23	1.51	1.47	0.65	1.37	0.74	2.51	1.72	1.29	0.56	4.52	3.69	1.48	0.70
2	1.71	1.11	1.11	0.33	1.90	1.25	1.66	0.96	1.65	0.89	0.73	-0.03	1.38	0.80
3	0.38	0.02	0.33	0.01	0.44	0.13	0.45	0.13	0.38	0.04	0.72	0.32	3.79	3.42
4	0.75	0.11	1.35	0.94	2.55	1.92	1.12	0.40	6.07	5.44	1.14	0.48	1.22	0.42

set. Since the METAR report also provides wind speed information, we have analyzed the influence of adding wind speed information [7].

Table 3.4 shows the result of using input set $\mathcal{I} = \{H_{-7}, H_{-1}, T_{0,60}, W_{0,60}\}$, thus adding wind speed information to the initial approach. For households 1 and 4 only two out of the seven weekdays show an improvement, where the improvement is only minor. For households 2 and 3 three respectively four weekdays show a significant improvement, where the less performing models for the other weekdays show more or less similar results. Wind speed information can thus be a valid addition to the input set.

As described in Section 3.5.2, another possibility is to use a sliding window approach. Therefore, the effect of a sliding window is analyzed as well [7].

The results in Table 3.5 up to Table 3.7 show the results of adding a sliding window to the initial approach. We see that by using the sliding window an improvement in the prediction quality can be seen for many households/weekdays. In general, using a sliding window size of 4 weeks, which is the smallest window size analyzed, often gives the best results, especially for household 2, where for six of the seven weekdays the best results are obtained using a sliding window size of 4 weeks. Only for Sundays the results are worse, but still close to the original values.

In Table 3.8 up to Table 3.10 the results of both adding windspeed information and using a sliding window are depicted. Again, compared to the non-sliding window approach using wind information, an improvement can be seen for many

Table 3.6: The Mean Absolute Percentage Error and Mean Percentage Error without wind, 5 weeks sliding window

H	Sun		Mon		Tue		Wed		Thu		Fri		Sat	
	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE
1	2.38	1.70	1.50	0.73	2.00	1.42	2.69	1.94	1.62	0.87	6.29	5.48	1.46	0.69
2	2.08	1.35	1.36	0.59	2.01	1.32	2.23	1.59	2.25	1.47	1.10	0.34	1.62	1.04
3	0.50	0.17	0.38	0.01	0.49	0.14	0.48	0.07	0.45	0.13	0.94	0.58	0.86	0.44
4	1.26	0.59	1.94	1.51	3.28	2.54	1.32	0.56	4.95	4.39	1.58	0.85	1.31	0.61

Table 3.7: The Mean Absolute Percentage Error and Mean Percentage Error without wind, 6 weeks sliding window

H	Sun		Mon		Tue		Wed		Thu		Fri		Sat	
	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE
1	2.64	1.91	1.60	0.77	2.08	1.49	2.80	2.01	2.12	1.37	6.61	5.82	1.42	0.60
2	1.69	1.03	1.63	0.85	2.18	1.46	2.11	1.47	2.49	1.73	1.77	1.03	1.71	1.14
3	0.52	0.18	0.37	0.01	0.57	0.21	0.57	0.16	0.48	0.09	0.66	0.28	1.65	1.33
4	1.51	0.81	2.30	1.83	2.40	1.51	1.72	1.00	3.95	3.35	3.82	3.14	1.83	1.13

Table 3.8: The Mean Absolute Percentage Error and Mean Percentage Error with wind, 4 weeks sliding window

H	Sun		Mon		Tue		Wed		Thu		Fri		Sat	
	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE
1	1.83	1.15	1.16	0.36	1.33	0.73	2.91	2.14	1.58	0.86	2.16	1.18	1.24	0.39
2	1.84	1.15	1.22	0.45	1.35	0.70	1.50	0.79	2.21	1.51	0.89	0.12	1.69	1.13
3	0.42	0.09	0.30	-0.00	0.44	0.13	0.44	0.11	0.41	0.05	0.80	0.41	1.50	1.12
4	0.81	0.18	1.18	0.77	2.18	1.38	0.97	0.27	3.61	2.85	1.08	0.35	1.12	0.36

households/weekdays. Again, a window size of 4 weeks often gives the best results.

To determine which approach is the best choice, Table 3.11 depicts the inputs sets which gave the best results. The number in the table represents the sliding window size, where a value of -1 represents to not use a sliding window. The + or - after the number represents whether or not windspeed information is used. Surprisingly, the prediction models for household 1 and 4 perform best when windspeed information and a sliding window is used. This is in contrast to the results shown in Table 3.4, where only windspeed information was added to the input set. For household 2 not adding the windspeed information seems to be the best choice. For household 3

Table 3.9: The Mean Absolute Percentage Error and Mean Percentage Error with wind, 5 weeks sliding window

H	Sun		Mon		Tue		Wed		Thu		Fri		Sat	
	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE
1	1.31	0.59	1.57	0.81	1.55	0.91	2.48	1.69	2.00	1.33	9.85	9.05	1.21	0.38
2	2.10	1.38	1.31	0.52	1.81	1.17	2.05	1.35	2.50	1.86	1.31	0.56	1.87	1.29
3	0.56	0.22	0.30	-0.01	0.53	0.20	0.49	0.06	0.41	-0.03	0.69	0.35	1.37	1.00
4	1.32	0.75	1.67	1.27	2.50	1.81	1.13	0.37	2.99	2.23	2.24	1.49	1.43	0.63

Table 3.10: The Mean Absolute Percentage Error and Mean Percentage Error with wind, 6 weeks sliding window

H	Sun		Mon		Tue		Wed		Thu		Fri		Sat	
	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE	MAPE	MPE
1	2.25	1.54	1.59	0.79	2.05	1.44	3.20	2.42	1.96	1.20	7.67	6.83	1.25	0.36
2	1.82	1.12	1.91	1.15	2.58	1.90	2.23	1.51	2.27	1.49	1.85	1.10	1.85	1.28
3	0.51	0.14	0.35	0.02	0.46	0.09	0.58	0.15	0.49	0.14	0.62	0.28	1.77	1.39
4	1.33	0.64	1.81	1.33	2.35	1.56	1.40	0.65	4.08	3.41	5.17	4.48	1.60	0.82

Table 3.11: Best input sets (window size/use windspeeds)

House	Sun	Mon	Tue	Wed	Thu	Fri	Sat
1	5+	4+	4+	5+	4-	-1+	5+
2	6-	4-	4+	4+	4-	4-	4-
3	4-	4+	6+	-1+	4-	-1+	5-
4	4-	4+	-1+	4+	-1+	4+	4+

the results are not decisive, where using the wind speed information has a slight advantage over not using the wind speed information.

Overall, the optimal window size seems to be 4 weeks. For households 2 and 4 a sliding window size of 4 week gives 6 respectively 5 of the 7 weekdays the best result. For the other households, the results are not that decisive. It seems that each household is unique, and that the best approach is very dependent on the household.

3.6 SEARCHING FOR ADEQUATE MODEL PARAMETERS

As shown in the previous section, there is no single input set selection which always gives the best results. Each house and each household is unique, so the inputs in

the input set should be selected such that it gives the best results for this particular combination of household and house. Furthermore, for the results of the previous section, only historical heat demand data the day before and a week before is used. Perhaps other combinations of historical heat demand can give better results.

For this reason, a general search method is required to find a suitable input set, network size and window size. Similar to the work of Pai and Hong [67], a Simulated Annealing (SA) algorithm is used to optimize the input set selection, windows size and number of neurons.

SA is an heuristic optimization technique analogous to the annealing process of material physics. Annealing is a process that produces conditions by first heating a metal to above the recrystallization temperature. Then the metal is slowly cooled, causing changes in its properties such as strength and hardness. The heat causes the atoms to become unstuck from their initial positions and wander randomly through states of higher energy. By slowly cooling the material, the atoms can find a configuration with a more orderly state and therefore a lower internal energy than the initial one.

SA is first described by Kirkpatrick et al. [49] and is based on the The Metropolis algorithm [59]. In SA each point in the search space S is analogous to the state of the physical system and has a certain energy $E(s)$, which has to be minimized. The goal is to bring the system from an arbitrary initial state to the state with the minimal possible energy.

At each iteration, the SA heuristic considers some 'neighboring' state s' of the current state s . For the state s' the energy $\mathcal{E}(s')$ is calculated. If the change in energy $\Delta\mathcal{E} = \mathcal{E}(s') - \mathcal{E}(s)$ is not positive, the s' is accepted and considered as the best solution. If $\Delta\mathcal{E}$ is positive, the state is accepted with a probability p determined by the Boltzmann distribution. More precisely,

$$p = e^{\frac{-(\mathcal{E}(s') - \mathcal{E}(s))}{kT}},$$

where k represents the Boltzmann constant and T the temperature. If a new state s' is not accepted, another neighboring state is analyzed.

In other words, a new state s' is accepted when the new energy is lower and the SA algorithm moves to the considered state s' . If the new energy is higher, a transaction can occur, and the likelihood is influenced by the temperature T and the energy difference $\Delta\mathcal{E}$. During each step of the algorithm, the temperature is lowered according to a cooling schedule.

The probability of accepting a state with a higher value is what allows simulated annealing to get out of local minima. In the beginning, when the temperature is high, states with higher energy values are more likely to be accepted and the SA algorithm keeps exploring the search space. However, when the temperature drops, the probability of accepting a state with higher energy reduces significantly, and thus mainly only states with lower energy values are accepted.

Considering our forecasting model as a function $\mathcal{F}(\mathcal{I}, \mathcal{F}_p) \rightarrow \mathcal{O}$, the SA algorithm is used to determine proper values for \mathcal{I} and \mathcal{F}_p . In case of the heat demand prediction, finding a proper value for \mathcal{I} reduces to determining a good subset of

Table 3.12: The Mean Absolute Percentage Error and Mean Percentage Error using Simulated Annealing

H	Sun		Mon		Tue		Wed		Thu		Fri		Sat	
	MAPE	MPE	MAPE	MPE	MAPE	MPE	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	0.71	-0.01	0.53	-0.16	0.47	-0.10
2	0.85	0.30	0.69	0.00	0.97	0.51	0.46	-0.12	0.68	0.03	0.47	-0.19	0.86	0.35
3	0.20	0.01	0.16	-0.01	0.16	-0.08	0.23	-0.06	0.19	-0.06	0.21	-0.06	0.22	0.01
4	0.39	-0.14	0.31	-0.03	0.50	-0.02	0.50	-0.08	0.40	-0.22	0.50	-0.04	0.45	-0.05

\mathcal{I}_{pos} (see Section 3.5.2). The \mathcal{F} is a neural network, thus the model parameters \mathcal{F}_p are the proper sliding window size and the number of neurons in the hidden layer.

Translating this into a state s for the SA approach, a state consists of the number of neurons to use s_{nn} , the sliding window size s_{sw} and the input $S_{\mathcal{I}}$. In this setting, the error E_{total} is used as the energy value for a state s , which should be minimized. Based on the results in the previous section, the number of allowed neurons is kept according to the network sizes in already analyzed networks, i.e. $4 \leq s_{nn} \leq 35$. Since a small window size showed the best result, the allowed minimal window size is reduced to 2 weeks, i.e. $2 \leq s_{sw} \leq 6$. For the input $S_{\mathcal{I}}$ a (sub)set from \mathcal{I}_{pos} is allowed, with a restriction on W and T . If one variation of information about windspeed/temperature is used for a day, the other two are not allowed any more for that day. For example, if $T_{-1,mm} \in \mathcal{I}$, then $T_{-1,30}$ and $T_{-1,60}$ are not allowed. Furthermore, $S_{\mathcal{I}}$ is not allowed to be an empty set.

A neighboring state s' is determined by either altering s_n , s_{sw} or s_{is} , where the probability of changing one of these fields is uniformly distributed, i.e. all three possibilities have a probability of $\frac{1}{3}$.

In Table 3.12 the MAPE and MPE after searching the optimal parameters is depicted. Our initial temperature for our SA algorithm was set to 15. The temperature was lowered after 5 steps with a damping factor of 1.04 until the temperature dropped under 0.5. The MAPE and MPE values in Table 3.12 show a significant improvement of both error values.

In Table 3.13 the parameters found by the SA algorithm are depicted. We see that for almost all households and weekdays a very small window size performs best, which is similar to the results shown in previous sections.

In Figure 3.5 the histogram of the input options is depicted. Here we see that the heat demand data a week earlier provides the most useful information in our prediction model, followed by the heat demand a day earlier. Thus our initial attempt, using $\mathcal{I} = \{H_{-7}, H_{-1}, T_{0,60}\}$, was a good starting point. Furthermore, we see that the SA often finds input sets where temperature information is used. Again, using wind speed information in the prediction model is dependent on the household.

In Figure 3.6 the forecast of a week in July of household 1 after using the SA

Table 3.13: Model parameters after running the Simulated Annealing algorithm

H	Weekday	N	WS	Inputs
1	Sunday	26	2	$H_{-4}H_{-5}H_{-7}T_{0,60}T_{-1,mm}$
1	Monday	24	3	$H_{-1}H_{-4}H_{-7}H_{-14}W_{0,60}W_{-1,mm}T_{0,60}T_{-1,30}$
1	Tuesday	28	2	$H_{-1}H_{-3}H_{-5}H_{-7}H_{-14}W_{-1,60}T_{0,60}$
1	Wednesday	31	2	$H_{-1}H_{-4}H_{-7}H_{-14}W_{-1,30}T_{-1,mm}$
1	Thursday	18	2	$H_{-6}H_{-7}W_{0,30}T_{0,30}$
1	Friday	32	2	$H_{-1}H_{-2}H_{-4}W_{0,30}W_{-1,30}T_{0,30}T_{-1,30}$
1	Saturday	31	2	$H_{-1}H_{-3}H_{-6}H_{-7}H_{-14}W_{-1,30}$
2	Sunday	23	2	$H_{-1}H_{-2}H_{-4}H_{-7}W_{0,mm}T_{0,mm}T_{-1,mm}$
2	Monday	12	2	$H_{-1}H_{-2}H_{-3}H_{-5}H_{-7}W_{-1,mm}$
2	Tuesday	31	2	$H_{-7}W_{-1,60}T_{0,30}$
2	Wednesday	24	2	$H_{-2}H_{-4}H_{-7}H_{-14}W_{-1,30}T_{-1,mm}$
2	Thursday	19	2	$H_{-1}H_{-4}H_{-5}H_{-7}H_{-14}W_{-1,mm}T_{0,60}T_{-1,30}$
2	Friday	12	2	$H_{-7}W_{-1,60}T_{0,60}T_{-1,30}$
2	Saturday	27	3	$H_{-1}H_{-2}H_{-6}H_{-7}$
3	Sunday	28	2	$H_{-1}H_{-2}H_{-3}H_{-4}W_{0,60}T_{0,mm}T_{-1,60}$
3	Monday	32	2	$H_{-2}H_{-3}H_{-5}H_{-7}W_{-1,30}$
3	Tuesday	33	2	$H_{-2}H_{-6}H_{-14}W_{0,30}W_{-1,mm}T_{0,60}T_{-1,mm}$
3	Wednesday	26	2	$H_{-1}H_{-3}H_{-5}H_{-6}W_{0,30}W_{-1,30}T_{0,mm}T_{-1,60}$
3	Thursday	27	2	$H_{-1}H_{-4}H_{-5}H_{-7}H_{-14}W_{0,mm}W_{-1,30}T_{0,30}$
3	Friday	26	2	$H_{-1}H_{-4}H_{-5}H_{-6}H_{-7}W_{0,mm}T_{0,30}T_{-1,30}$
3	Saturday	29	2	$H_{-1}H_{-3}H_{-4}H_{-5}H_{-7}H_{-14}W_{0,30}W_{-1,30}T_{0,mm}T_{-1,30}$
4	Sunday	22	2	$H_{-6}H_{-7}W_{0,mm}W_{-1,mm}T_{0,60}T_{-1,60}$
4	Monday	32	2	$H_{-1}H_{-5}H_{-7}H_{-14}T_{0,30}$
4	Tuesday	26	2	$H_{-6}H_{-7}W_{0,60}T_{-1,mm}$
4	Wednesday	25	2	$H_{-6}H_{-7}H_{-14}W_{0,mm}T_{0,mm}T_{-1,mm}$
4	Thursday	32	2	$H_{-1}H_{-6}H_{-7}W_{-1,60}T_{0,mm}T_{-1,mm}$
4	Friday	22	2	$H_{-1}H_{-3}H_{-5}H_{-7}T_{0,60}T_{-1,30}$
4	Saturday	24	2	$H_{-1}H_{-7}T_{0,mm}$

algorithm is depicted. In this week, the MAPE is 5.84 and the MPE is 5.75. Again, sometimes a high demand peak is forecasted, while in the real heat demand no such peak is present. However, compared to the figure in Figure 3.4, a significant improvement in the quality of the forecasts can be seen.

To give more insight in the quality improvements of the different approaches used in this chapter, the error values (E_{total}) of four approaches for household 1 are depicted in Figure 3.7. The initial approach is depicted to show how well the first attempt performed. The two suggested improvements based on the initial results, adding wind speeds information and adding the sliding window approach, shows that no single approach always performs best. The usage of the SA algorithm finds a proper approach, yielding a significantly increased forecast quality.




Figure 3.5: The histogram of the used input options




Figure 3.6: Heat demand forecasts and actual values for household 1 in July

3.7 CONCLUSIONS

During the first step of TRIANA, the forecasting step, forecasts are used to determine the scheduling freedom of a device. Therefore, for each individual device a forecast is made, since device specific information and restrictions are required in the planning and control of the device. In order to incorporate device specific information in a forecasting system, without the necessity of communicating all required information of each device to a central location, forecasting is performed by the local house controller in each building. Using such an approach, the requirement of a scalable forecasting system can be met. The local controller can use locally harvested data and can be programmed to forecast device specific information, resulting in a flexible system. In the use case of individual heat demand prediction to determine the scheduling freedom of a micro-CHP appliance, neural network techniques are used. The possibility of autonomous learning of (non-linear) relations between the input- and outputdata makes neural network techniques a good candidate. By adjusting the neural network structure, forecast on different timescales can be generated. Furthermore, periodically evaluating the forecasting quality and, when




Figure 3.7: Performance of the different forecast approaches of household 1

necessary, adjusting the input for the neural network, results in a system adaptable to changes. The usage of a SA searching algorithm results in an automated search for a proper input set and neural network parameters.

DISTRIBUTED CONTROL

ABSTRACT – This chapter describes how the forecasts are used in the second step, the planning. Based on the forecasts and the scheduling freedom, the planner tries to exploit this scheduling freedom of the devices for his objective. This leads to a desired profile for a fleet of buildings, which needs to be achieved, by generating a planning for each building/device individually. Planning is performed in a hierarchical way consisting of a top planner, multiple levels with intermediate grid planners and at the bottom of the structure the individual building controllers. By subsequently dividing the overall planning problem into smaller subproblems which are solved at lower levels, a more scalable system is achieved. By aggregating information at each level in the structure, communication is reduced. Using artificial cost price vectors, the building controllers generate a planning for each device based on the costs functions of the devices, the artificial price vector and the locally generated forecasts. By iteratively adjusting the price vectors, the profiles of the individual houses are reshaped to reach the global objective. There are multiple ways to determine the price vectors. They can be determined at multiple levels within the hierarchical structure. The best results are achieved by using different price vectors for each building, determined by the grid planner at the lowest level. Since the schedules generated in the planning phase are based on forecasts, and forecasts often are not perfect, deviations from the planning can occur. In a replanning phase, a new planning may be generated, based on the real situation and improved forecasts based on more recent information. Enabling replanning shows a significant improvement in reaching the desired objective.

Continuous communication and cooperation between all the parties in the energy supply chain is one of the key features of smart grid technology. Constant cooperation between producers, grid operators and consumers allows for optimizations to

Major parts if this chapter have been presented at [VB:3] .

improve the overall energy efficiency, reduce CO₂ and allow more renewables in the supply chain. In the previous chapter the first step of the three step approach, the forecasting step, has been introduced. In this chapter, the cooperation between the second step, the planner, and the third step, the real time control, is discussed in more detail.

Using the forecasting step, the scheduling freedom of a device can be determined. The task of the planner is to exploit this scheduling freedom to achieve a certain objective. Example objectives are peak shaving, reducing overall energy purchase costs or supplying the demand with the highest possible share of renewable energy sources. Certain global objectives require cooperation of multiple buildings. If the group of buildings becomes large, a distributed approach with one or multiple planners present in the grid may be required. In this case each grid controller cooperates with a group of other grid controllers or buildings controllers.

Based on the objective and the scheduling freedom, in our approach a planning for each building or each device is determined. This planning is transmitted to the individual building controllers (also called house controllers), which should follow this planning as good as possible. Since the planning is based on forecasts and forecasts often are not perfect, the controller must cope with differences between the forecasted and real situation. When there is a large deviation between the forecast and the real situation, the planning may become infeasible and replanning is required. Using the recently observed information, a new short term forecast can be made. Based on this newly determined scheduling freedom, the overall planner can create a new, more feasible planning.

In the remainder of this chapter, first the requirements and goals of the planner and the cooperation between the planner and the building controller is introduced. Then, some related work on smart grid control systems is given. In Section 4.3, the chosen approach is described. The algorithms used in this approach are given in Section 4.3.1. After presenting some preliminary results in this section, the process of rescheduling is described in Section 4.4. This chapter finalizes with conclusions in Section 4.5.

4.1 GOAL AND REQUIREMENTS OF THE PLANNING METHODOLOGY

The goal of the three step control methodology is to exploit the optimization potential within a building in a generic way to achieve a certain objective. As mentioned earlier, this objective can differ and is dependent on the stakeholders of the control system. For example, a building controller might cooperate with a planner owned and controlled by a grid operator. Such a grid operator might prefer peak shaving by shifting load to minimize the required grid capacity. If on the other hand the buildings controller cooperates with an operator of a Virtual Power Plant (VPP), the operator might try to maximize the profit by steering the electricity production of a fleet of micro-CHP appliances towards high price periods.

Based on the above considerations, the proposed methodology has to be flexible in both the optimization objective and the technologies available. After all, the

objectives may differ over time and in different buildings different technologies may be installed.

Within the three step methodology, cooperation and communication between a house controller and a planner is possible. Cooperation implies that decisions are made on different levels within the system, and these decisions have impact on the overall functioning of the system. In the remainder of this section, an exploration of the used communication and decision structure of the the overall control system is given.

As mentioned before, the control strategy consists of three steps. In the first step, a system located in the buildings forecasts the production and consumption pattern for all devices for the upcoming day. For example, in a normal household multiple devices like a TV, washing machine and central heating are present. For each devices, based on the historical usage pattern of the residents and external factors like the weather, a forecasted energy profile is generated. The combination of these energy profiles determines the optimization potential of all devices located in the house. The optimization potential of a device is dependent on the type of the device.

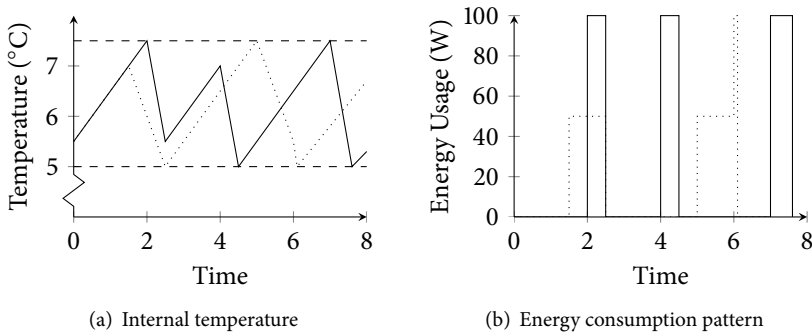


Figure 4.1: Scheduling freedom of a fridge

For example, a fridge has to maintain its internal temperature within a certain limit (see Figure 4.1). How and when to achieve this can be varied. For a fridge, the freedom is dependent on the internal temperature of the fridge, which is mainly influenced by the interaction of the residents with the freezer, by adding and removing goods.

In the second step of the control strategy, these optimization potentials can be used by a central planner to exploit the potential to reach a global objective. The result of the second step is a planning for each building or device for the upcoming day.

In the final step, a realtime control algorithm decides at which times devices are switched on/off, when and how much energy flows from or to the buffers and when and which generators are switched on. This realtime control algorithm uses steering signals from the global planning as input, but preserves the comfort of

the residents in conflict situations. In this way, the algorithm can also deal with forecasting errors.

The first step has been described in Chapter 3 and for the third step we refer to [62]. The mathematical principles of the second step are for a micro-CHP appliance is investigated by Bosman et al. [17]. In this chapter a more general analysis of the second step, applicable to different types of devices, is given.

Key in cooperative optimization methodologies is communication between all the different parties involved. This communication can take place at different levels, from short-range communication in buildings between devices and the building controller to wide-range communication between neighborhood control nodes and a central planner. For example, the simplest case of communication in a house is the wire connecting the thermostat and the (High Efficiency) boiler. However, smart devices and Demand Side Load Management require a connection between the building controller and these smart devices. For this communication, Power Line Control (PLC) or a wireless standard (e.g. Zigbee) can be used since their installment is relatively cheap. New smart meters with some computing capacity are a very good candidate to act as building controller.

On top of communicating within the building, the optimization methodologies need communication links with parties in the grid to work towards a global objective. The most common connection is a connection to the distribution company, but other schemes are possible as well. Especially in future smart grids topologies it is very likely that even more communication takes place. Therefore the control strategy should take communication restrictions into account.

Besides limited communication, also only a limited amount of computational resources may be available. Since finding an (near) optimal utilization of the available resources in the smart grid may require a lot of computational power [17], a control scheme must exploit the computational power spread in the smart grid as much as possible. Furthermore, the control system should be dependable and must be able to cope with failures within the system.

Summarized, the system should:

Be device agnostic The system should work with different kinds of devices. In a building many kinds of devices are present, each with different characteristics and limitations. Still a generic control scheme is required.

Respect communication restriction Communication links have limited bandwidth and, more importantly, introduce latency effects.

Respect limited computational power Finding a (optimal) solution for a large group of devices may take too much time. The system should exploit the available computational power in the grid.

Dependable The system should be able to cope with (forecasting) errors.

4.2 RELATED WORK

Most approaches about smart grid control strategies found in literature are hierarchically structured, agent based solutions, where the hierarchical structure ensures the scalability of the solution. Although a lot of approaches claim to be distributed without a central algorithm, all approaches found have one decision-making element (node, agent, etc.).

Hines et al. [42] describe a decentralized, agent-based, realtime control methodology. On different levels in the network agent-based nodes are installed controlling and/or monitoring a network element (e.g. house appliances, capacitor banks, substation voltage). These nodes communicate with each other to exchange information, optionally also with human operators. The communication is structured in a sort of hierarchical way: the nodes are divided in groups exchanging information with each other. When all information is available, the group leader makes a decision, optionally in cooperation with other group leaders.

The PowerMatcher described by Hommelberg et al. [44, 45] is an agent-based, hierarchical optimization methodology. Every device in the house can be controlled by an agent. This agent sends a bid (amount of demand/generation and a price) to the agent one level higher (house agent), which aggregates all bids and sends the aggregated bids one level higher, etc. The root agent decides, based on the bids and the objective, the market clearing price. This price is distributed and each agent knows what to do based on his bid made and the market clearing price. The agents coupled to a device can use predictions to optimize their bidding strategy. However, this is on a local (device) level, only leading to profit optimization of the device-agent itself.

For the GridWise project [22] a decentralized control methodology with dynamic pricing is used. In this approach, no centralized algorithm is used. Field tests showed that the dynamic pricing can reduce peaks up to 15%.

Koch et al. [51] describe a methodology to manage thermal household devices (fridge, freezer, boiler), which can have a buffering property: cooling a freezer during low demand periods may prevent having to cool during peak demand periods. Their approach uses one central controller, all houses are directly connected to this controller. To prevent sending too much detailed information (privacy), only the costs to switch a device on or off are sent. The central controller decides how much devices should be switched on and determines the switch-on-price and switch-off-price.

The similarities between the described approaches and our approach is reflected in the control up to a device level and the hierarchical structure with aggregation on each level. The main differences are the prediction/planning and the lack of agents. Although the PowerMatcher approach uses prediction and planning on a device level, this is utilized for profit raising of the agent itself. The latter is also the main difference between our approach and an agent-based approach: agents are greedy and try to optimize their own profit where our optimization methodology tries to reach a global objective for the whole fleet. Furthermore, our approach uses

variable steering signals instead of the same price/signal for everyone and has a larger planning horizon (up to one day in advance).

4.3 APPROACH

There is a large group of buildings present in the grid, each again with a multiple devices. Some of these devices might have some scheduling freedom, while others have none. Preferably, a single control system present in the building should be able to exploit the scheduling freedom of the devices, without knowing the individual restrictions and characteristics of each individual device.

When multiple buildings cooperate to reach a more global objective, for example an energy neutral or energy autonomous neighborhood, communication between the building controllers is required. The envisioned global objective influences the requirements on the overall control system in both network and computational requirements. For instance, if a centralized system has to be used for realtime power quality control, latencies of only a few milliseconds are allowed considering the current power supply uses a 50 or 60 Hz alternating current. Modern communication technology like ADSL or cable have latencies of 1 up to 10 ms. Furthermore, the controllers need to process the message and afterward communicate with the devices present in the house to give a control signal. Again, these devices need to process these commands as well. Furthermore, reaching controllers spread over the grid for power quality control most like requires multiple communication hops, making the latency requirement very hard to reach or only with substantial costs.

Based on this, a better approach is to make the building controller more autonomous. The controller should be able to take many decisions locally and only swiftly communicate with other controllers. The coordination between the controllers is then more for decisions on the long term, in the order of seconds or even minutes. Using such an approach, network latencies of modern communication technologies are sufficient and cause no problems. Still global objectives like peak shaving, virtual power planting and demand and supply matching are possible, but now in a distributed way.

For this reason, our planning and control methodology is organized in a tree structure. The root node of the tree contains the global planner. This global planner tries to optimize the energy profile of the whole fleet, based on a given objective. Taking into account the information on the forecasted energy profiles of all consumers, this given objective leads a desired energy profile of the fleet. However, since matching this desired energy profile exactly is very difficult or even impossible, the desired energy profile is described by a lower and upper bound on this profile (see Figure 4.2). The aggregated load profile of all buildings steered by the planner should fall within these bounds.

The structure of our proposed control system, as depicted in Figure 4.3, allows different strategies to achieve the desired profile. At multiple levels in the hierarchical structure the overall profile can be divided into smaller subprofiles, which need to be reached by lower grid controllers. As depicted in Figure 4.3(b), somewhere in

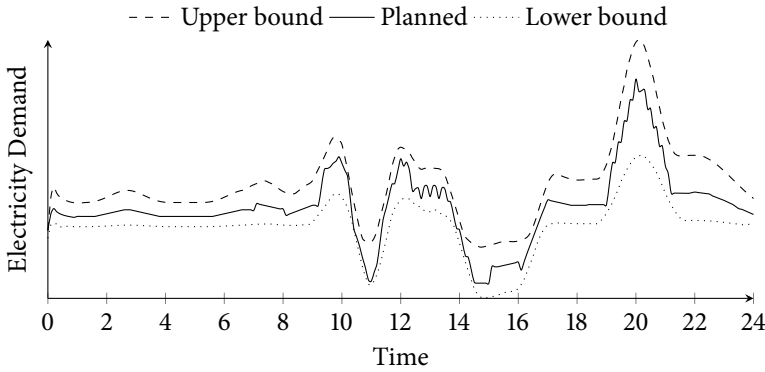


Figure 4.2: Desired profile with upper and lower bounds





Figure 4.3: Organization and information flow of the proposed control scheme

the hierarchical structure a controller switches from subprofiles to steering signals, in the form of price vectors, to reach the desired profile requested from the parent grid controller. The location where this switch is made is an important design choice and determines which decisions are taken where. This also has influence on the autonomy of each controller in the structure. In the following sections, multiple possibilities of switching points and steering vectors are explored.

Determining steering vectors at the top of the structure

As mentioned above, the root planner tries to achieve a desired profile. An approach to reach this desired profile is by directly using steering signals. This planning strategy is depicted in Figure 4.4. The steering signals are sent to all the connected

subgridcontrollers. The subcontrollers then distributes the steering signals to the controllers below. At the bottom of the structure, the local controllers generate their planning based on the received steering signal and send these upwards to the subgridcontrollers. The subgridcontrollers aggregate the received information and submit this information upwards in the tree. Based on the aggregated profiles and the corresponding mismatch with the given lower and upper bounds, the top controller determines new steering signals and sends these adjusted steering signals downwards in the tree. This process is repeated until the top planner is satisfied with the achieved results.

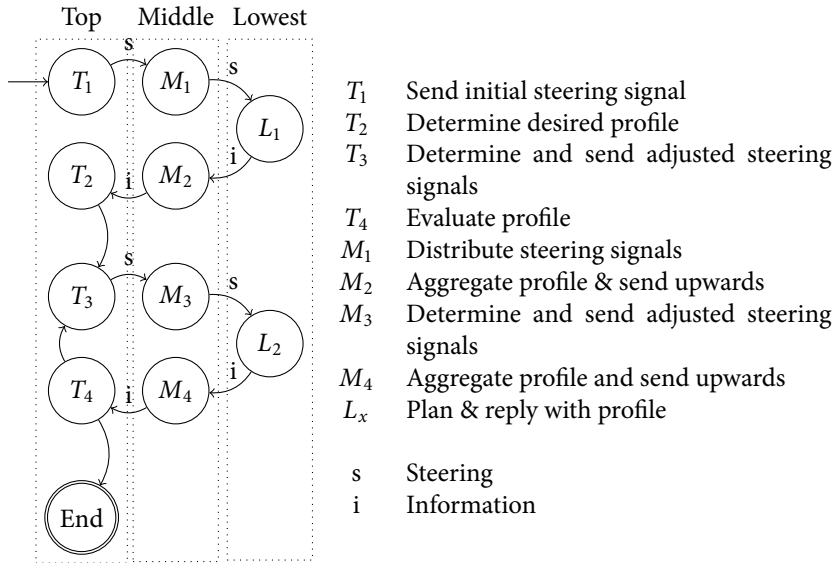


Figure 4.4: Distributed planning strategy workflow, steering at top

Determining steering vectors at the bottom of the structure

Another possible control approach is determining the steering vectors at the bottom of the tree structure. This planning strategy is depicted in Figure 4.5. In this approach, the root planner tries to achieve the desired profile by decomposing the profile into subparts and delegating these parts over the planners located in the nodes of the hierarchical structure directly below him (see Figure 4.3(a)). Each planner directly below the root planner is responsible for planning its part of the tree such that his desired share of the global profile is reached. Again, these planners try to achieve this goal by delegating subparts of their desired profile to their children. The planners on the bottom of the tree are directly connected to the controllers located in the buildings. They try to achieve their given energy profiles by sending

steering signals to the building controllers. Based on these steering signals, these controllers generate a planning for the coming day and send their planning to their planner.

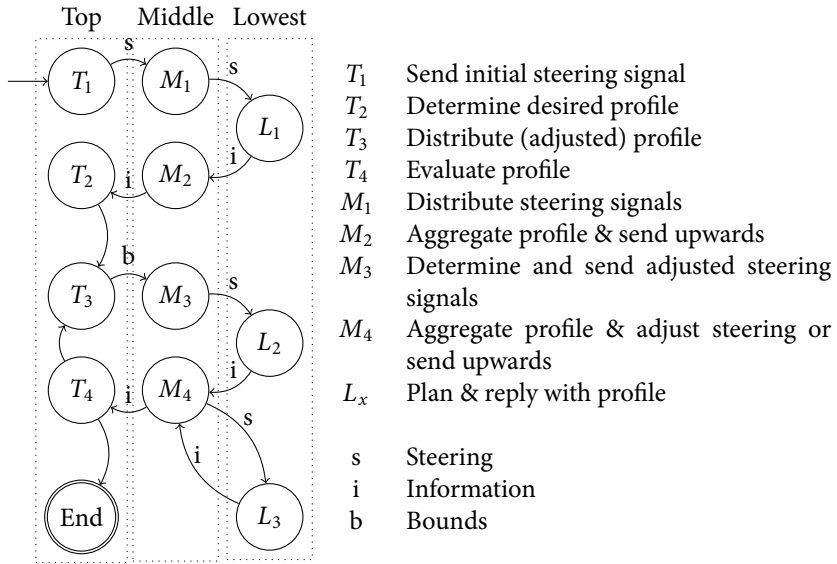


Figure 4.5: Distributed planning strategy workflow, steering at bottom

On every level, the received data from the level below is aggregated and sent further upward in the tree. Based on the mismatch between the planning and the desired profile, the root node adjusts the partition of the profile, sends this to the planners directly below him and then the process starts again. This iterative process is repeated until the resulting profile falls between the lower and upper bound (if possible).

Differences between the planning strategies

In the first approach, all the autonomy is with the top planner. The decisions are made at the top level, where the subgridcontrollers are only used to distribute the price signals and aggregate the data. This keeps the planning process more comprehensible, since only at one part of the planning structure decisions are made. However, changes in the planning require a traversal over the whole tree. Furthermore, this approach can result into an oscillating effect of the planners. Since a lot of subgrid- and local controllers receive the same steering signal, they may all react on a similar way, all shifting their load to the same time periods. The central planner will adjust the steering signal accordingly, again shifting all the load to another period.

In the second approach, the planning is more distributed over the tree structure, as the (sub)gridcontroller is responsible for solving its own subproblem. This leads to a more sophisticated but also more scalable solution, since each subgridcontroller can autonomously plan a large group of houses. Only when problems occur during planning the desired (sub)profile, like the infeasibility of a profile, more coordination and communication with other grid controllers may be required.

Common in all approaches is the exploitation of the computational power available within the grid. The actual device planning, which required the most computation power, is executed at the local controllers. Due to the tree structure and continuous aggregation, the required amount of information to be communicated is not very large. This fits well within the expected structure of smart grids, where faster communication links between buildings and the nearby planner may be available and higher level planners may be further away. Dependent on the used communication links, and their required investment costs, probably less advanced communication links can be used to communicate to higher levels. It is therefore desired to reduce the amount of data exchanged between the higher gridcontrollers, which is achieved by the aggregation of information at each level.

4.3.1 DISTRIBUTED ITERATIVE PLANNING

As mentioned in the previous section a desired profile has to be achieved. Somewhere in the hierarchical approach, via steering signals the devices in the buildings are steered to reach this profile in the best possible way. By using steering signals, which in our approach always are price vectors, the electricity profiles of the devices that are steered are reshaped. In this section we introduce a method to determine these price vectors. An example use case and objective is given to illustrate the working of the chosen approach.

First, the planning horizon (the time period we want to plan) is discretized into NR intervals of equal lengths. Given a fleet of H buildings, we first need to determine the total amount of electricity TEC , which all buildings together would consume based on their individual planning, is calculated. The total electricity consumption is determined by first sending a fixed, flat price vector with a fixed price for all the time intervals (see the first round trip in Figures 4.4 and 4.5). All local planners in response send back their expected electricity profile, based on the forecasts made in the first step of TRIANA. So TEC denotes the total planned consumption of electricity, if each building optimizes the use of the devices for his own benefit. Based in this information and the global objective, a desired pattern P for the fleet is determined. Example objective are a profiles with a shifted load to certain time periods, or a flattened profile. Via steering towards this desired pattern we try to achieve the desired objective.

The principle behind the approach is that we iteratively adjust a (virtual) cost price vector, which contains a artificial cost price for the electricity prices per time period. This vector is distributed to and used by the individual building controller to optimize their devices for maximal profitability. By adjusting the prices, we try to reshape the energy profile of a (group of) building(s). In case there are consumers

present in a building, a building controller tries to shift as much consumption as possible to low price periods. When producers are present in a building, production capacity is shifted to high price periods.

As mentioned above, the energy profile of a fleet of buildings should meet a desired pattern P . However, since reaching exactly P might be impossible, we allow the expected consumption in each interval to be in a certain interval band, for example $[0.8P, 1.2P]$. As a measure to analyze the quality of a planning, we define the following. Let TPL_j be the planned consumption in interval j ($TPL_j \in P$). Then $M_j := \max\{TPL_j - 1.2P_j, 0.8P_j - TPL_j, 0\}$ denotes the measure of quality for interval j , and the total mismatch $M = \sum_{j=1}^{NR} M_j$ can be used as measure of quality for the planning approach.

The base of the heuristic for a building i is a local program DP^i that determines the usage of each device in the building, only regarding local constraints (the operational characteristics of the devices determine the possible states of the devices) and a price vector p^i :

$$DP^i(p^i) \rightarrow PL^i, \quad (4.1)$$

where PL^i is the resulting profile.

The planned consumption TPL of the whole fleet is an aggregation of the local profiles PL^i resulting from the programs DP^i :

$$TPL = \sum_{i=1}^H PL^i = \sum_{i=1}^H DP^i(p^i). \quad (4.2)$$

Solving towards a global optimum (minimal mismatch) means that an exploration of the price space of possible price vectors for all building is needed. The power of the local program is that it produces a fast solution for a local problem, instead of applying an optimization algorithm to the complete set of buildings simultaneously. The proposed heuristic uses this fast local approach iteratively to find a solution for the complete set of buildings. However, for practical use, the number of iterations needs to be limited to gain as much as possible from this advantage of having a fast local procedure.

The used price vectors are used to steer the profile to the desired profile, but can also be related to the energy cost price. As a consequence, the costs can be determined by two parts: electricity costs ep^i and artificial costs ap^i . In the use case the electricity costs remains constant during the different iterations of the planning process. The artificial costs are used to move the consumption of the devices to specific periods. Summarizing, we use as 'interface' to the local planning performed by the local program the artificial costs ap^i .

In general the iterative search works as follows. The initial price vector for all building, used to determine the TEC , is set to the electricity prices ep .

$$p^{i,1} = ep, \quad (4.3)$$

$$(4.4)$$

If after the planning in iteration n the consumption of building i is too high for a certain interval, the price for this interval is enlarged based on the deviation from

the target planning and the artificial price ap^i . If on the other hand the consumption of the building is too low for a certain interval, the price for the interval is similarly lowered. To ensure stabilization of the algorithm, the magnitude of the artificial prices is reduced after each iteration. More formally, the prices for iteration n and building i are changed via:

$$p^{i,n+1} = p^{i,n} + \text{addPrice}(ap^{i,n}, TPL^n), \quad (4.5)$$

$$ap^{i,n} = ap^{i,n-1} - \frac{p^{i,1}}{\max_{it}}, \quad (4.6)$$

where \max_{it} is the maximum number of iterations and TPL^n is the planned total consumption in iteration n . The *addPrice* function is described in the following section.

Price distribution

The structure in Figure 4.3 allows different strategies to partition the overall planning problem over the tree. At multiple levels in the hierarchical structure the overall profile can be divided into smaller subprofiles, which need to be reached by lower grid controllers. As depicted in Figure 4.3(b), somewhere in the hierarchical structure a controller switches to steering signals, in the form of price vectors, to reach the desired profile requested from the parent grid controller. In this section, the effect on how and where these price vectors are constructed is discussed.

The simplest possibility to construct the price vectors is to use an uniform price vector. In this case, all the buildings receive the same price vector. Here the function $\text{addPrice}(ap^{i,n}, TPL^n)$ returns a vector (a_1, \dots, a_{NR}) , where:

$$a_j = \begin{cases} ap_j^{i,n} & \text{if } TPL_j^n > 1.2P \\ -ap_j^{i,n} & \text{if } TPL_j^n < 0.8P \\ 0 & \text{otherwise.} \end{cases} \quad (4.7)$$

The function *addPrice* only results into steering via a_j in case the result of the last planning iteration is above or below the desired profile. In this approach the search space is limited since we use an equal price for all buildings.

In a second variant we create diversity between buildings by allowing only a fraction of the buildings to change its price. The vector (a_1, \dots, a_{NR}) resulting from *addPrice* is then defined as follows:

$$a_j = \begin{cases} ap_j^{i,n} & \text{if } (TPL_j^n > 1.2P) \text{ AND } (\text{rand}(0,1)^i < \frac{M_j^n}{TPL_j^n}) \\ -ap_j^{i,n} & \text{if } (TPL_j^n < 0.8P) \text{ AND } (\text{rand}(0,1)^i < \frac{M_j^n}{TPL_j^n}) \\ 0 & \text{otherwise,} \end{cases} \quad (4.8)$$

where $\text{rand}(0,1)^i$ is a random number between 0 and 1 for building i , where higher mismatch results in more buildings receiving an adjusted price vector. Using

this variant of the *addPrice* function, the artificial price is still dependent on the mismatch, but due to the random not every building receives the same price vector.

Within the hierarchical structure we can implement the second variant (the usage of different price vectors) in two different ways. These two ways differ where the stochastic choice of changing a price is made. On the one hand, the random choice to change an additional artificial cost in a certain interval can be picked at the bottom of the hierarchical structure, i.e. just above the building controller. On the other hand, we may also pick this choice at higher branch in the structure, meaning that all underlying buildings get the same price vectors. This can be seen as a compromise between the idea of the first and second variant, allowing only different prices on a higher level in the structure.

Using the heuristic, the planner tries to reshape the energy profile to the desired profile. It might need a number of iterations, which must be limited in order to find a solution on time. It might be possible that the desired profile is unfeasible to reach. The heuristic can use different stop criteria. In the current (initial) implementation the heuristic ends when the maximum number of iterations \max_{it} is reached to analyze the impact of the algorithm in each iteration. In future work, each planner within the structure can end the search in its substructure whenever the aggregated planning of this substructure confirms the desired planning for this planner or if there has been no significant improvement in the last iterations.

4.3.2 EXAMPLE

To illustrate the working of the distributed iterative approach we consider a use case in which a large group of freezers has to be steered. The internal temperature of freezers has to maintain between certain limits. Since the environment of a freezer is relatively static, freezers show a very regular and predictable pattern. When the cooling element of a fridge is switched off the temperature rises slowly. Once a certain upper temperature is reached, the cooling element is switched on to cool to a certain lower temperature threshold. By advancing and postponing the switching points of the cooling element, the energy profile of the freezer can be altered, as long as the temperature boundaries are respected. The regular pattern of the freezer simplifies the planning of the device.

The local planning program, of which some states are depicted in Figure 4.6, generates a profile for the freezer. Each state change requires a certain amount of electricity and leads to a certain cost, similar to the example of the fridge in Figure 4.1 (on 65). As will be explained in Chapter 5, each individual device is controlled using a set of costs function. Based of the state change costs, which are determined by the costs functions of the device and the price vector, the local program tries to minimize the costs. The path with the lowest costs determines the planned behavior and the corresponding usage of the freezer.

The global objective of the use case is to spread the electricity consumption of the whole fleet equally over the planning horizon, which is one day. The general idea is that, if we are able to flatten the consumption, this electricity can be produced




Figure 4.6: Planning of a freezer

more efficiently, since constant generation is preferred over adaptive (fluctuating) generation.


To analyze the performance of our algorithm, it has been implemented in our simulator, which is described in Chapter 5. Two scenarios are simulated. The first scenario is a simulation consisting of 50 houses (with the above described freezer). This scenario can be used to analyze the performance of actual planner cooperating with the house controllers. Due to the small amount of houses, only one central planner and no intermediate planners are used. Using only one planner, the effects of the different approached determining the the steering signals can be used.

In the second scenario, 200 houses are controlled. For controlling this a larger group, a hierarchical approach is used. Again, different steering signals can be used to study the effect on the whole group, e.g. multiple ways to divide the whole planning can be tested. Using this use case, different strategies on where to switch from (sub)profiles to steering signals can be analyzed. The strategies depicted in Figures 4.4 and 4.5 are used in this scenario. Furthermore, like in the 50 houses use case, different approaches on how to determine the steering signals can be applied.

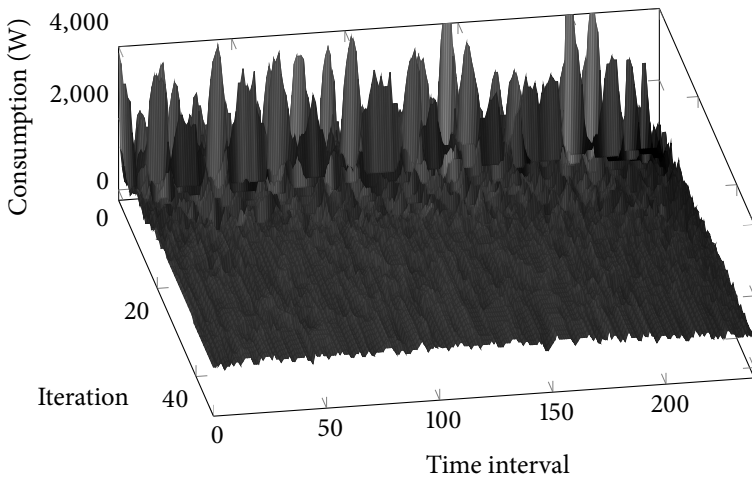
In the next subsections, first the different approaches in distributing the price vectors are described. Another important aspect is the required communication, which is analyzed in subsection 4.3.4.

4.3.3 IMPACT OF DIFFERENT PRICE VECTORS

As described above, there are different possibilities a) how price vectors are determined and b) how they are distributed over different branches of the tree. In the simplest case, there is only one global planner and the house controllers directly cooperate with this global planner. This is the case in the first simulation, where 50 houses are simulated. The freezers in the houses all start with different starting temperatures and have different characteristics, i.e. there are differences in cooling capacity and the insulation quality. For example, the freezer in Figure 4.6 has a cooling capacity of 0.5°C and cold loss of 0.1°C per time interval. Appendix A describes how these different starting temperatures and characteristics are determined.



(a) Using uniform price vectors



(b) Using different price vectors

Figure 4.7: Results of planning for 50 houses

Determining the prices vectors via (4.5), (4.6) and (4.7) result in identical price vectors for all houses. The results of the planning using this approach is depicted in Figure 4.7(a). In this figure, the evolution of the energy profile can be seen. In the first iteration, the price vector contains only a single price for all time intervals. This results in a typical cyclic load, which aggregated leads to large peaks in demand. Since all house controllers also receive the same price vector in the next iterations, they all try to shift their cooling cycles in the time interval with low costs, resulting in

alternating pattern of consumption profiles. Only after reducing the price changes with the increasing number of iterations, the load pattern get more flattened, but still has a peaky structure.

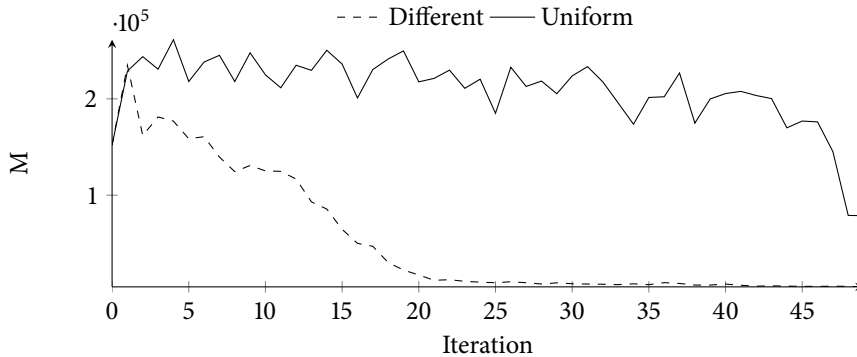


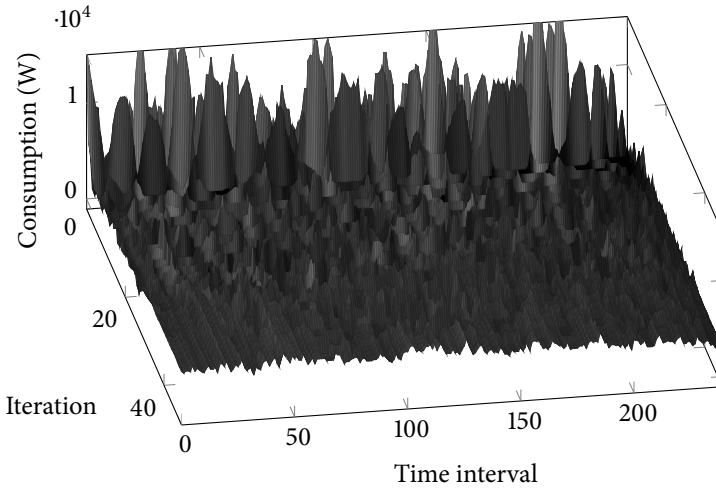
Figure 4.8: Evolution of M for planning with 50 houses

When using different prices for each house controller, as described by (4.8), the quality of the planning improves significantly. The result of this scheme is depicted in Figure 4.7(b). The production pattern is more flattened (as asked for by the global objective).


In Figure 4.8 the evolution of the mismatch M is depicted. As can be seen in the figure, the usage of different price vectors for the individual houses clearly outperforms the usage of a uniform price vector. Not only is the mismatch lower using different prices, the mismatch is also reduced significantly earlier. Although the production pattern is quite flat, the global objective is not reached for every time interval, since there is still some mismatch after the last iteration. This is caused by the limited amount of freezers. Although each freezer has a different state, the number of state combinations of the group is limited.

When 200 houses are simulated, the number of possibilities increases. It is then expected that a more flattened production pattern can be reached. The planning of the 200 houses is organized in a tree structure with one global planner, four intermediate planners, each planning 50 houses. One possibility to alter the price vectors is to use different price vector for a whole branch of the tree. In this case, the prices are adjusted with a certain probability only at the top of the tree. Each of the intermediate planners receive a different price vector, but distributes this price vector unaltered. In this case, since the group of 200 houses is divided over four intermediate planner, again a group of 50 houses receive the same price vector.

The most advanced scheme is to distribute the load evenly over the tree and let the bottom planning determine the chances to change the prices vector. The top planner here determines the global objective, i.e. the demand profile for each subtree. Each planning below then distributes its part of the profile to its children.



(a) Determining the vectors at the top, resulting in a vector per branch



(b) Determining the vectors at the bottom, resulting in a vector per house

Figure 4.9: Results of planning for 200 houses

The bottom planners, communicating with the building controllers, try to reach this profile using the different price vectors per house as described by (4.8). Once the bottom planners are satisfied with the result for their subproblem, the resulting load profile is communicated upwards the tree. If the global objective is not reached, the global planner may determine another distribution of the production pattern over the tree and repeat this process until the global objective is met. Using such an

approach, the most variation in price vector is obtained.



Figure 4.10: Evolution of M for planning with 200 houses

The results using determining the price vector at the top of the hierarchical structure are depicted in Figure 4.9(a). Similar to the case with 50 houses with a uniform price vector, still an alternating pattern can be seen in the figure. Again, all the houses in one group receive the same price, but since there are four different groups, enough variations between each group is present to flatten the overall profile.

Figure 4.9(b) depicts the results when determining the price vectors at the bottom of the structure. Since each house receives a different price vector, the steering is more flexible and leads to the best results in both in the amount of required iterations as in the end result of the planning.

In Figure 4.10 the evolution of the mismatch M when planning the 200 houses is depicted. When using the scheme where the prices are changed at the top of the tree, at the global level, a price change has an impact on a whole subtree. As expected, due to the different pricing scheme for subtrees within the whole tree, a more flattened production profile is obtained. However, since still a large group of house controllers obtain the same price vector, still a relatively large mismatch remains after the last iteration. The case where each individual house has a different price vector has the best result, as was expected. This scheme gives a very fine grained approach to steer an individual house.

Another advantage of this approach is that the optimization problem is very distributed, resulting in a fast planning. Each planner on the bottom of the hierarchical structure execute in parallel. Most of the computation and communication occurs nearby the house, only the results (the load profiles) are feed back to the global planner. Using such an approach, parts of the three can work autonomously when necessary. This results in a more dependable and scalable system.

4.3.4 COMMUNICATION

Due to the tree like structure, the amount of communication is reduced by aggregating the data at the intermediate levels. When dividing all the buildings over the tree structure, the tree should be as wide as possible in the bottom of the tree and as narrow as possible in the top, since this leads to the highest amount of aggregation. Furthermore, using this structure, the algorithms can be distributed to the bottom of the tree as much as possible, exploiting the available computation power. Furthermore, it is assumed that building controllers and the bottom planners are close to each other with good communication links, resulting in fast communication. Going further up the tree the amount of communication is reduced, which unifies with the fact that communication links between higher planners can have less bandwidth and higher latency.

Since in our approach only price vectors and energy profiles are used, the amount of information that needs to be exchanged is limited. Furthermore, no other privacy sensitive information is transmitted. For each time interval to be scheduled, only a single number needs to be transmitted. Since a production or a price can be represented by a single number, the message size between two entities is determined by the planning horizon (the amount of time intervals to be planned) and the representation of the numbers. In our simulation of 200 houses, a planning is made for each six minutes for a whole day, leading to 240 time intervals.

Since only small numbers are exchanged at the lower level planners, fewer bits are required at these levels. This is beneficial, since most data is transmitted at these levels. Going upwards in the tree, the numbers increase and the amount of bits can be increased accordingly. Another approach is to switch the scale, for examples switching from Watts to kilowatts. Although some accuracy might be lost after switching, it is expected that higher tree the accuracy at the watt level is not required.

Message Size (4 Bytes)	Message Type (1 Byte)	Message content
---------------------------	--------------------------	-----------------

Figure 4.11: Used packet format for sending price/production vectors

The prototype implemented in the simulator uses an own developed protocol, of which the packet format is depicted in Figure 4.11. Using this prototype protocol, 32-bit integer values for prices and profiles are used at all levels.

In Table 4.1 the amount of network traffic used for the planning of a single day for each of the proposed methods using 20 iterations is given. The best planning method, using different price vectors per house, determined at the bottom of the tree structure, required a total amount of 16.4 MB, of which all data was transmitted at the bottom level. Although this seems quite a lot of data, this is only roughly 82 kB per building. Note that a very fine grained planning is obtained for a whole day. Furthermore, many iterations were executed to analyze the effect of the algorithms. By planning for a time interval of fifteen minutes (similar to the electricity

Table 4.1: Network traffic different planning methods using 20 iterations

Houses	Steering vector	Traffic	% at bottom
50	Uniform	4.1 MB	-
50	Different	4.1 MB	-
200	Top	16.7 MB	98.24%
200	Bottom	16.4 MB	99.99%

markets time intervals), the amount of data can already be reduced by a factor of two and a half. Another improvement can easily be made to use less bits at the lower levels. For example, the demand profile of a house is limited to a couple of kilowatts, thus a 16-bit value is more than sufficient to specify the demand in each time interval. The same holds the prices, allowing 16-bits values to be used between the house controller and the fleet controller. Since 99.99% of the data transmitted at the bottom of the tree structure, another factor of roughly two can be achieved. Furthermore, since the algorithm reaches a good result after only a few iterations, another reduction can be achieved by reducing the number of iterations at the bottom planners. These numbers can be reduced even further by adding compression and optimizations in information encoding, which is left for future work.

4.4 REPLANNING

As shown in the previous section, the iterative planning algorithm is able to reshape the demand profile of a group of houses. The overall planner distributes the planning over the hierarchical structure to different intermediate planners and eventually all the different house controller execute the planning of the devices. Based on forecasts of the environment, together with the cost functions, the house controller uses the price vector to determine demand profile for the planning horizon. As a consequence, the planning is based on forecasts.

If a forecasting error has been made, it might be possible that the device does not behave as initially planned and that a deviation from the planning occurs. Perhaps such a deviation is compensated elsewhere in the grid, but it might also happen that the planning is no longer feasible.

As a consequence, when the real situation differs too much from the forecasted situation, an evaluation of the generated planning may be necessary. Using new, improved short(er) term forecasts the feasibility of the current planning can be determined and, when required, new iterations of the planning process can be executed. The decision to perform a replanning is dependent on the new information obtained from the new forecasted, the desired objective function and the general forecasting quality. Therefore, a suitable *replanning threshold* is required. If, for example, the forecasted data has an average error of 25% and the local controller has difficulties coping with this prediction error locally, starting a replanning session

when a deviation of 5% occurs will result in a lot over replanning sessions. Furthermore, replanning requires a certain amount of processing and network traffic, which should be included when determining the replanning threshold.

Based on the objective of the planner, a new overall energy profile of a group can be desired, or it might be necessary to stick closest possible to the initial desired profile, but with a different planning. For example, in case of a VPP it is beneficial to stick the closest possible to the original energy profile, since this energy profile has been sold on the electricity market and deviations from this profile results in a penalty. Another case can be that demand side load management is used to balance the imbalance caused by renewables sources, and that the production of these renewables is different than forecasted, requiring a different demand profile.

Again, based on the objective a (possible new) desired energy profile P needs to be reached. In the initial planning stage, a planning is generated for a certain time period T , which was discretized into NR time intervals, resulting in a desired consumption TPL_j for each time period $j \in [t_0, t_f]$. Based on the TPL_j , a (different) price vector for the buildings present in the grids are generated using the iterative approach. The planner, present in the grid, can continuously monitor the real demand profile and compare this real demand with the planned demand. When in a certain time interval t_p the deviation is unacceptable, the planner can execute another planning session. In this replanning session, a new improved planning is generated. Instead of generating a new planning for the whole interval $[t_0, t_f]$, only a planning for $[t_{p+1}, t_f]$ is generated.

The used approach is the same as used in the initial planning session, but for a smaller time horizon and thus a shorter price vector. Furthermore, the building controllers may use new forecasts based on the more recent information and for a shorter time period. These forecasts are expected to have lower forecast errors, probably resulting in a better achievable planning. When generating the schedules for the devices, the current actual situation is used, compensating the differences between the forecasted situation and the actual situations.

To analyze the effect of replanning, again the use case with 200 freezers are planned with the same objective using the simulator. In the simulations in the previous section, a perfect forecast of the behavior of the freezer was assumed, resulting in no deviation from the planning. Now, an artificial forecasting error is introduced. For each freezer, a single human interaction with the freezer between 7:12 and 21:36h (30% and 90% of the time intervals) is simulated. The effect of the interaction with the freezer is that the internal temperature of the freezer increases. These internal temperature increase is chosen to be between 0°C and 1.8°C . The temperature change in the freezers results in another demand profile, since the offset temperatures are reached earlier.

When creating the forecasts for the freezer, this single human interaction is also forecasted. During forecasting, a pseudorandom timestamp of this interaction and corresponding temperature increase is used. These values, as described in Appendix A.1.1, are used in the planning and replanning phases. For the real situation, another pseudorandom timestamp and temperature increase are picked. The differences between the forecasts and the real situation results in a deviation

from the planning. For each time interval the deviation from the planning is determined with

$$Dev_j = \left| \frac{E_j - TPL_j}{TPL_j} \right|,$$

where E_j is the actual electricity demand for time interval j . Since only one forecasting error occurs, with only a mild deviation from the planning, the replanning threshold is chosen in such a way that when the actual situation differs 10% or more from the planning a replanning is executed, i.e. $Dev_j \geq 0.1$.

As a performance measure of the replanning, the deviation of the real electricity flow from the planned profile is used. The goal of the replanning is to minimize this deviation, thus a lower value is better.

During the (re)planning phases, the variant using different price vectors for each house is used, since this approach yields the best results. The price vectors are dependent on the deviation from the objective and the probability a price in the vector is changed. Therefore, the overall results are dependent the values drawn from the probability distribution and thus different in each planning. For this reason, both the scenario with and without replanning are simulated 50 times.

The mean deviation of the 50 simulations without replanning was 33588, with a standard deviation of 843. The mean deviation of the 50 simulations with the replanning enabled was 21351, with a standard deviation of 2568. As can be seen, a significant reduction of the mismatch is achieved. Since the forecasts in this case are static and thus do not improve, a zero mismatch cannot be achieved. The improvement comes purely from the adjustment of the schedule of each freezer during a replanning session. In the 50 simulations with replanning enabled, on average 3.2 replanning sessions were performed, with a standard deviation of 0.6.




Figure 4.12: The best electricity profiles of 50 simulations with 200 freezers

In Figure 4.12 the two best of the 100 simulations are depicted. The dotted graph is the result of the best simulation without replanning enabled, i.e. the one with the lowest deviation. The best results of the simulation with replanning is depicted via the solid graph. The dashed line in this figure indicate the times the replanning threshold was exceeded and a replanning session was executed. Up to the 72nd time interval, the differences in the graphs are caused by differences in the steering

signals caused by the probabilistic planning. Since there are no forecast errors in this period, no replanning is required. Just before the first replanning session a sharp drop in the demand can be seen. In this time interval the deviation is too big and a replanning is initiated. Initially, the profile is nicely flattened, but the real situation of the freezers is continuously deviating more from the predicted situation, requiring in more fluctuation and eventually another replanning session. Overall, the profile is more flattened, especially after the second replanning session.

4.5 CONCLUSIONS

The result of the freezer use case shows that the iterative approach is able to apply global optimization techniques for a large group of houses. Although only a freezer is scheduled using this approach, it may be clear that the use of generic price vectors can also steer different kind of devices. Since both the forecasts and planning are performed at each building, all the required information for planning a (group of) devices is locally available. The building controller can use this information, and use device specific constrains to generate a planning for the device. In this approach, the distributed computational power available in the grid is exploited. Since less information about the device and its environment needs to be sent outside the building, less privacy sensitive data is exposed. Due to the subsequent division of the large optimization problems into subproblems via a tree structure, a fast scalable system is achieved.

Using uniform prices for all buildings leads to a sort of ‘worst case’ scenario. Since all individual building controllers try to minimize their own cost, they all optimize to periods with low costs, leading to a shift of demand creating new peaks instead of the desired profile. Addressing each building individually by using different steering signals gives the best results. However, independent on the way the price vectors are determined, a proper tree structure helps that communication requirements can be kept low. The use of simple price and production patterns lead to small messages to be sent.

Since the planned schedules are based on forecasts, forecasting errors can lead to deviations from the initial planning. By detecting these deviations and executing a replanning based on the actual situation and an improved short term prediction, deviations can be minimized. However, replanning requires a certain amount of processing and network traffic. Therefore, when determining the replanning threshold, issues like communication, processing power, forecasting quality, allowed deviation and the ability of the local controller to cope with deviations should be taken into account. Optimizations in the communication protocol can reduce the amount of traffic to the order of tens of kilobytes per building. With modern communication standards, these amount of network traffic can easily be processed.

Although the initial version presented in this chapter already shows promising results, improvements are still possible and needed. On multiple levels in the tree, better approaches to determine when to stop the planning are needed to reduce the amount of iterations required during planning. Furthermore, by optimizing the

way information is encoded during communication, bandwidth requirements can further be reduced.

ENERGY FLOW SIMULATION

ABSTRACT – Smart grid technology will change the demand profiles of buildings. To analyze the impact of the control methodologies introduced by the smart grid, a simulator has been built. The simulator is developed based on an energy model. The basic elements of the model are individual devices and between devices, energy streams (electricity, heat, gas etc) are defined. Devices can consume, buffer, convert and exchange energy, resulting in four categories of devices. The energy streams are modeled via so called pools, which represent the physical connections between the devices. Using the pools, the energy balance can be ensured, by enforcing the sum of the energy flow to and from a pool to be zero. The model has been implemented in the simulator and a controller based on cost functions is used to control the devices. The cost functions provide a generic and flexible, but still powerful method to control current and future devices. Furthermore, frameworks for configuration of the model, addition of stochastic variations and results analysis are provided. Since the model is computationally intensive, the simulator can be organized in a distributed fashion to allow simulation using multiple computers. The underlying communication framework for distributed simulation can also be used for distributed control between grid- and housecontrollers. The simulator evenly distributes the load over the computers involved in a simulation. The speed up is limited by the slowest computer, but speeds up linearly to the number of computers in the simulation.

Using the three step methodology, the energy profile of individual devices and their corresponding buildings are altered. To analyze the effect of the methodology, a model of the energy flows has been developed. Using this energy model, all devices present in the grid/gas network and their interconnection can be modeled. With such a model, the effects of decisions made by the overall control system

Parts of this chapter have been presented at [VB:14] and [VB:4] .

can be modeled and analyzed. Furthermore, by modeling future changes in the energy infrastructure, like the introduction of large scale electric cars or distributed electricity buffers at/in each building, the requirements of future grids can be analyzed.

Based on this model, a simulator has been developed. In order to enable the simulation of a large group of buildings, in spite of the computational complexity of the model, the simulator has been designed such that it can be distributed over multiple computers within a network. Besides analyzing the energy flows within the whole energy supply chain, the simulator is also able to cope with the cooperation between the local and global controllers. Therefore, different cooperation schemes and their corresponding communication requirements can be studied.

In the remainder of this chapter, first the requirements of the model and the corresponding simulator are given. Then, related work on other (energy) simulation software/frameworks is given. In the related work, the motivation of designing and implementing a simulator is also given. Next, in Section 5.3, the energy model is described. In Section 5.4 the design and implementations details of the simulator are given. Information about the distributed simulation is provided in Section 5.5. This chapter finalizes by some results and conclusions about the developed simulator.

5.1 REQUIREMENTS

The goal of the simulator is to provide a tool to analyze the effect of optimization methodologies, residential generation, storage technologies and smart devices for a large fleet of buildings. These buildings can be houses, schools, offices or even small factories. However, this work is mostly focussed in houses and their corresponding devices.

Since the control methodologies steers individual (future) devices, these new technologies need to be modeled and simulated at an appropriate level. Since residential generators, buffers and consumers are devices located in or near a building, individual buildings need to be modeled in detail, i.e. on a device level. The simulated situation should be an accurate model of the actual situation. Therefore, every individual device present in buildings has to be modeled. To make the model more accurate, measurement data (e.g. electricity usage) should be incorporated into the simulated devices.

Smart grid technology enables the possibility to develop new, smarter devices with more controllable options. The simulator should be able to simulate these to incorporate future devices and technologies. Since it is not yet known what future devices, technologies and scenarios look like, the model of a building needs to be very generic and flexible.

Due to the differences in the behavior of the residents of a building, the devices they own and use, the number of residents in a building etc. almost all buildings are different and have different energy profiles. Therefore, the modeled building should be able to represent different types of buildings (family/single-person houses,

big/small houses, a small office, etc). We thus need a realistic, generic and flexible model of a building.

Within a building often a controller is present. This can be a simple controller like a thermostat and a controller which decides when the boiler has to fill the heat buffer, but a more sophisticated controller which takes for example electricity prices into account is also possible. The simulator should provide a framework to emulate such a controller and analyze changes in the control system.

For the simulation of a massive introduction of micro-generation, multiple buildings are combined in a grid. The basis for this simulation is a realistic, generic and flexible model of a building since each building should be individually addressed due to its individual characteristics and internal state. The buildings need to be grouped together to form a grid, e.g. a city with a realistic mix of buildings. There are a lot of possible scenarios (different combinations of buildings, local controllers, etc), thus this part of the model needs to be flexible as well. Especially when local optimizations are coordinated on a global level, i.e. the local controllers are able to communicate and cooperate with a coordinating global controller present elsewhere in the grid, a lot of scenarios are possible. The global controller can have different objectives, for example the creation of a Virtual Power Plant [73] or peak shaving. The local and global controllers need to cooperate and communicate with each other to achieve the desired objective. This cooperation introduces requirements on the communication between the local and global controllers. The simulation framework should provide a means to analyze, implement and evaluate the communication between the controllers.

Summarizing, the base of the simulator needs to be a realistic, generic and flexible model of a building on a device level. Multiple buildings are combined in a grid to analyze the effect of a large group of buildings. The simulator should be easily adaptable to new types of micro-generators, controllers and other supported elements. It should be easy to simulate different scenarios, combinations of different buildings, local controllers and global objectives.

Using the simulator, the quality of the controllers and their control algorithms can be studied. This requires that sufficient information about the devices and the grid/gas connection etc. must be available. However, dependent on the objective and the control algorithms, the amount of detail of the simulation must be limited to keep the overview in the otherwise massive amounts of simulation data, especially when the amount of simulated devices/building becomes very large. Therefore, the logging of data should be flexible.

Another important requirement concerns the speed and memory usage of the simulator. The tool has to be able to simulate a large fleet of buildings in detail. For example, an average windmill park produces around 50 MW. In order to have an Virtual Power Plant (VPP) that is comparable to such a windmill park, a generation potential of 50 MW is necessary. Therefore, the number of buildings with for example a 1 kW micro-Combined Heat and Power (CHP) that should be simulated within a reasonable time (hours) should be at least 50.000, leading to requirements on processing power and memory usage. Again the amount of data is dependent on how fine-grained the simulation has to be.

The requirements described above can be briefly summarized as:

1. Simulation of **realistic** settings and devices.
2. Simulation of both a **single house** and a grid with a **large amount of buildings**.
3. **Flexible**, both in adding new elements and in the supported scenarios.
4. Simulation of the **whole energy chain**.
5. **Adjustable logging/precision**.
6. **Simulation speed and memory usage** sufficient for ‘normal’ computers.
7. It should support a **network communication framework** for cooperation between local and global controller.

5.2 RELATED WORK

Simulation solutions already exist in a lot of different areas, e.g. optimizations for logistics, 3D modeling or process management. However, most simulation software is domain specific and is not easily portable to different application areas.

We have chosen not to use simulation frameworks like Tortuga [80] or SimPy (Simulation in Python). Although these frameworks may provide some generic functionality required for our simulation, they still require to create our own model within the limitations of the framework. We considered it easier and more promising to create an own model in a familiar environment and still have the flexibility to reuse work in the literature. Furthermore, our aim is to have a simulator that is very fast and memory efficient. Preferable the simulator should be usable on multiple platforms, which makes distributing the simulations easier. For this reason, we chose C++ as programming language, using Nokia’s QT library. This library provides efficient, cross platform libraries for data storage, network communication and user interfaces. Furthermore, the C++ programming language enables us to use other C and C++ libraries.

The focus of our research is on simulating the effect of (domestic) energy streams on the system as a whole. Commercial software is available where heat- and electricity load of large buildings can be simulated. These systems are often used to optimize Heating, Ventilating, and Air Conditioning (HVAC) systems and facade control systems or to build more energy efficient buildings and take into account the structure of the building and the materials used [5, 68]. However, our goal is not to simulate the expected heat and electricity load, but to simulate the control methodologies to supply the loads resulting from the structure of the building. Other available energy simulation software focuses on a specific technology, for example wind parks or solar cells.

The Advanced Local Energy Planning (ALEP) simulation framework [46] is an initiative of the International Energy Agency and is developed in cooperation

with multiple countries. This simulation framework focusses on the local energy supply. It can analyze which mix of generators can best be used for certain (remote) areas. It takes, next to technical constraints, also management and social issues into account. The goal of the tool is to analyze the impact of different generation components in an area. Another free tool is HOMER, an optimization model for distributed power [56]. It is a model that simplifies the task of evaluating design options for both off-grid and grid-connected power systems for remote, stand-alone, and distributed generation applications. It can also take buffering and deferrable loads into account. The goal of the tool is to find the best combination of supply components and parameters for these components. Thus, these two tools have a different focus and do not meet all of our requirements.

In related work regarding energy-optimizing control strategies the focus is mostly on agent based approaches [25, 52]. An example of an agent-based system is the PowerMatcher [52], which creates a virtual market to determine who can produce/consume energy and for which price. Important in such an approach is the stability and reliability of the bidding system, which is often simulated. The simulator described in this chapter has a more generic approach, where the focus is on control of the system and the influence of that control system on the whole system. Due to the flexible design, the PowerMatcher bidding system could be embedded in our simulator.

In [40] a custom simulation system for the coordination of decentralized energy conversion is described. Similar to our approach, a custom simulator is developed. However, little detail about the underlying design is given and only one device specific example is given. Our approach has a more flexible design. Different control strategies can easily be added and the design is more flexible to future technologies.




Figure 5.1: Concept of a Hybrid Energy Hub

The energy model described in the next section is inspired by the work described in [37], where the concept of a Hybrid Energy Hub is introduced (see Figure 5.1). The basic principle is that multiple energy carriers — for example electricity, gas, but also heat — are connected to the loads via energy hubs. Each hub has a certain

conversion matrix, where elements in the hub convert one or multiple energy carriers into other carriers. After conversion, the converted energy streams supply the required load. In the hub, both conversion and storage might be present. The conversion matrix defines the ratio between the input streams, the storage and the output streams. In this conversion matrix, efficiencies of the conversion, storage and transportation can be incorporated. The problem with this approach is that the load is inelastic, while in our approach we want to analyze the possibilities to change the load, since for certain devices a certain freedom exists in controlling the device.

5.3 ENERGY MODEL




Figure 5.2: Overview of a expected house

As described in Section 5.1, the simulator is used to analyze the effect of control methodologies for residential generation, storage technologies and smart appliances. Therefore, the basis of the simulator is a model of a building. As an example of a building, an example house as depicted in Figure 5.2. In a building, multiple devices are present, each with its own functionality. In this figure, for example the television and lighting are electricity consuming appliances. Electricity is imported from the grid and consumed by the devices, after being distributed throughout the house.

Besides electricity, heat is also consumed in the house for hot tap water and central heating. In the figure, the heat is stored (buffered) in a heat store for more comfort.

Heat can be, like electricity, imported from a district heating station or, more common in the Netherlands, converted from gas using a (High Efficiency) boiler. In general, energy can be imported and exported, consumed, converted and buffered. The energy model should therefore be able to support this. Like in [37], the energy model is based on a set of energy-types. All energy streams within a building are seen as streams of a certain energy-type: heat, electricity and gas, but next to that also sunlight, wind, etc. Which energy-types are taken into account can be defined per instance of the model. This set can for example be electricity, gas, sun and wind, but it is also possible to add reactive power to the set to be able to model phase shifting.

Within a building energy is converted, (temporarily) stored and consumed by devices. E.g. a hot water tap is a heat consuming device just as a television is an electricity consuming device. The model distinguishes four different types of devices within the house: 1) exchanging devices, 2) converting devices, 3) buffering devices and 4) consuming devices.

5.3.1 DEVICES

The basis of the model are devices and energy streams between these devices: energy of a specific type flows from device to device. Every device does something with the energy-types (exchange, convert, buffer, consume). So, a device is an entity where energy flows in and/or out and where for each in/out port the type of the energy streams is specified.

All devices within the building are modeled individually, since the optimization algorithms optimize the behavior of individual devices. Such behavior can for instance be the decision when to run a converting device (e.g. starting a micro-CHP).

As mentioned earlier, the model distinguishes four kinds of devices:

Exchanging devices exchange energy with the outside world. Concerning a building as an entity that is modeled, a building exchanges energy with its environment. For most conventional houses only electricity can be imported and exported and gas that can be imported. But some building, like flats, also import heat from district heating. Furthermore, also sunlight and wind are modeled as energy imports when a building is equipped with a solar panel or micro wind turbine. In our model an exchanging device can only exchange one energy-type with the outside world.

Converting devices convert one or more energy-types into one or more other energy-types. In our model, the amount of energy streaming into these device is equal to the amount of energy streaming out of these devices. Although the energy-types are different, no energy is lost during conversion. However, energy conversions (often) have a certain amount of loss during conversion.

This is modeled as a separate energy stream out of the device. So, a running micro-CHP for example has a gas stream in (100%) and a heat stream (80%), an electricity stream (15%) and loss stream (5%) out.

Buffering devices can temporarily store an energy-type. These devices have an energy-type stream in and the same energy-type stream out. This separation of the stream in and stream out is necessary since these streams are not always shared, e.g. most currently installed hot-water buffers have separate in and out flow circulations. When the in and out stream are shared, the streams can be combined when implementing of the device. Next to the in and out stream, a separate energy-type stream out can be used for modeling loss.

Consuming devices consume one or more energy-types in a certain ratio. For most devices, the amount of energy consumed in a certain time interval (the consumption profile) is a characteristic of the device and is therefore defined on beforehand. A loss device is also modeled as a consuming device. For this device it is not defined how much energy it consumes, it simply consumes all loss (since loss streams are connected to this device).

5.3.2 LIMITATIONS AND OPTIONS

For every device certain limitations are given, e.g. the amount of energy it can import, the amount of energy it can convert, etc. Furthermore, within the limitations of the device there are often multiple options possible on how to use the device. For example, a high efficiency boiler might have a modulating burner, allowing different heat production levels. Every device has a set of possible options O (micro-CHP on/off, buffer charge or discharge and how much) and based on the state s of the device (e.g. State of Charge (SoC) of the buffer) a subset O_s of the set of options are valid options. A control algorithm needs to decide which option to choose.

The limitations concerning the exchanging devices are the amount of energy which can flow in and out. They are characteristics of the exchanging devices and are modeled within the devices. The decision is how much energy is exchanged, within the bounds of these limitations. The amount of every energy-type flowing in and out, as well as the limits, can be defined in watts.

For converters, the amount of energy that can be converted is often limited. Furthermore, often there are only a few possible levels of conversion (e.g. 0 kW, 30 kW and 70 kW) or conversion regions (e.g. 0–20 kW and 65–90 kW) and the efficiency can differ per level or region. The decision in this case is to choose a valid conversion level out of the set of options. Next, there are fixed ratios between the different input streams (optionally different per level). To calculate the amount of every energy-type flowing out based on the amount of energy flowing in, conversion matrices are used (again optionally different per level). The conversion matrices can depend on the status of the device. For example, a micro-CHP consumes electricity during startup and produces electricity when it is running at full speed (so it is possible to have energy streams of the same energy-type in and out).

For buffering devices the maximum amount of energy which can flow to or from the buffer depends on the State of Charge (SoC) of the buffer and the buffer size. The loss can also depend on the SoC, so a buffering device needs to keep track of its status too. The decision is how much energy flows in or out the buffer.

The amount of energy consumed by consuming devices is normally fixed once it has been switched on. However, when the energy streams are optimized the load of (some) consuming devices can be shifted in time, or it can be decided to temporarily switch off a device.

5.3.3 STREAMS BETWEEN DEVICES

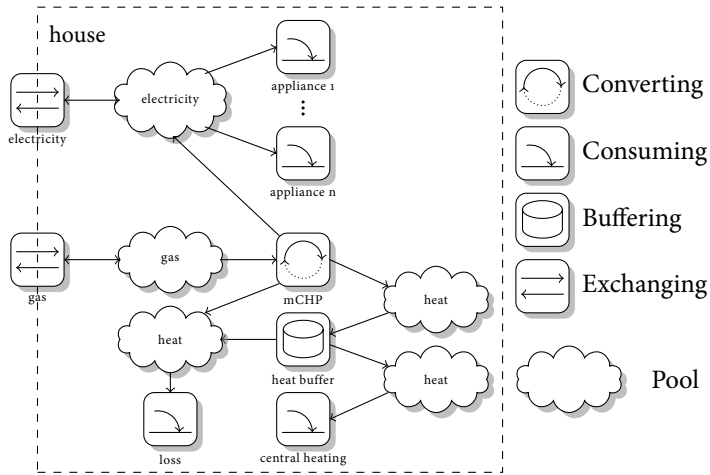
As mentioned before, the basis of the model are devices and energy streams between these devices. Every device has certain streams in and/or certain streams out. Each stream consists of a flow of one energy-type. Input and output streams of devices are coupled, so energy can flow between devices. To manage these flows in a proper way pools of energy-types are introduced. Each pool is of a certain energy-type and combines streams of this energy-type. More precisely, a set of input and output streams of the given energy-type are combined to the pool and this pool has to be in balance, i.e. it has no loss. This means that the amount of energy flowing into the pool is equal to the amount of energy flowing out of the pool.

For example, in most houses all electricity producing and consuming devices are connected to one grid in the house. Electricity can flow from every electricity output stream to every electricity input stream. On the other hand, hot water flows from the boiler via a pipe to the hot-water buffer and via another pipe to the consuming devices. This leads to two separate “hot water pools” in a house, as depicted in Figure 5.3(a). Note that also a third ‘artificial’ heat pool is present, which models the loss of the heat buffer. Summarizing, within every pool one energy-type is transported and every stream is connected to one pool. The amount of energy flowing from and to a pool can be limited due to limits in the transportation medium. This introduces a lot of expression power. For example, we can model the situation that (a part of) the house is protected with an Uninterruptible Power Supply (UPS) system.


5.3.4 COMPLETE MODEL

The complete model of a house combines the four different types of devices and the pools. This enables the possibility to model all energy streams of a complete house. An example of such a model is shown in Figure 5.3(a). This house consists of multiple electricity consuming appliances, central heating, a heat buffer and a micro-CHP. This model of the house can easily be extended with other energy-types, e.g. the water stream, the reactive power or Photovoltaics (PV) on the roof.

The concept of devices and pools can also be used to model the whole energy supply chain, as depicted in Figure 5.3(b). For example, electricity can be generated by a conventional power plant and/or by a wind turbine, which is then transported via the grid (in the form of a pool) to the houses.



(a) Model of domestic energy streams.



(b) Model of energy streams in the network(s).

Figure 5.3: Example instances of the given energy models

When using the model, certain constraints have to be respected. First of all, all streams should be within the bounds set by the devices. Furthermore, the energy flowing in and out of a device has to be in balance. The sum of the flows in and out of a converting device has to be zero. The sum of the flows in and the flows out of a buffering device has to be equal to the change of S_oC . In other words, for every device a valid option from its set of possible options O_s has to be chosen. Finally, the energy streams within a pool have to be in balance, i.e. the sum of flows in and out should be zero. When all the pools and appliances are in balance, the complete energy chain in the house is in balance.

5.3.5 FORMAL MODEL DESCRIPTION

The proposed model describes the energy streams within a house on a device level in a generic way. It is based on devices and energy streams between these devices. The energy streams are separated in flows of different energy-types, from heat and electricity to sunlight and windenergy. Devices exchange, buffer, convert or consume energy streams. Energy flows between devices via pools, where every stream of every device is connected to a pool and there has to be as much energy flowing into the pool as out.

For every device limits for the exchanging devices, conversion matrices for the different levels of converting devices, etc. are defined. Furthermore, the model defines a framework to model the different devices. Each different device needs to be modeled using this framework, for example using an internal state machine.

Formally, as described in [62], this means that for each building a set of energy-types EC is defined. For every energy-type $ec \in EC$ at least one pool is defined, resulting in a set of pools P :

$$P = \cup p_{ec},$$

where $p_{ec} = \{p_{ec_1}, \dots, p_{ec_{N_{ec}}}\}$. Furthermore, a set of devices is present in every building. These devices are split up in exchanging, converting, buffering and consuming devices:

$$Dev = Dev_{ex} \cup Dev_{conv} \cup Dev_{buf} \cup Dev_{cons},$$

where $Dev_{ex} = \{d_{ex_1}, \dots, d_{ex_{N_{ex}}}\}$, $Dev_{conv} = \{d_{conv_1}, \dots, d_{conv_{N_{conv}}}\}$, $Dev_{buf} = \{d_{buf_1}, \dots, d_{buf_{N_{buf}}}\}$ and $Dev_{cons} = \{d_{cons_1}, \dots, d_{cons_{N_{cons}}}\}$. These devices are connected to the pools via streams. Streams are uni-directional, there are streams from the device to the pool and vice-versa:

$$str = (p, d), str \in Str, p \in P, d \in Dev,$$

$$Str = Str^p \cup Str^d.$$

The amount of energy flowing through a streams is defined by $x_{(p,d)}$ and $x_{(d,p)}$. Since the pools are abstract devices introduced for modeling purpose they cannot contain energy, the sum of the energy flowing inside the pool should be zero at all time:

$$\sum_{(p,d) \in Str^p} x_{(p,d)} = \sum_{(d,p) \in Str^d} x_{(d,p)} \quad \forall p \in P.$$

For every device $d \in Dev$ an internal energy stream x_d is defined. Not all values for x_d are valid, therefore a set of option O^d is defined for every device. Since only one option can be chosen, variable $c_o \in \{0, 1\}$ is introduced:

$$\sum_{o \in O_{state}^d} c_o = 1.$$

An option exists of a valid interval for x_d :

$$c_o \times F_o \leq x_o \leq c_o \times T_o.$$

The internal value x_d is the sum of all values of x_o (note that only one can be non-zero):

$$x_d = \sum_{o \in O^d} x_o. \quad (5.1)$$

The flow through every stream in and every stream out of a device is defined by the internal energy stream x_d and a multiplication factor $M_{(str)}^o$:

$$x_{(p,d)} = \sum_{o \in O^d} M_{(p,d)}^o x_o,$$

$$x_{(d,p)} = \sum_{o \in O^d} M_{(d,p)}^o x_o.$$

Since the state of a device changes, only a subset O_s^d of all options O^d is valid in a certain state s . All options that are not valid in a state s should not be chosen:

$$c_o = 0, \quad \forall o \in O^d \setminus O_s^d.$$

5.4 SIMULATOR

The presented model forms the foundation for simulations of the energy streams within buildings over a certain period. The model describes the devices present in the building and how the energy flows between the devices. To be able to use the simulation model we need to include the notion of time and the corresponding changes in production, transmission and demand, which is described in Section 5.4.1. Furthermore, the model needs to be configured such that multiple types of buildings can be modeled, for example representing different households with different devices and usage. Section 5.4.2 described how the simulation model can be configured and setup. Once the simulation model is configured and setup, a control system present in the building needs to properly control the devices, which is described in Section 5.4.3. After simulation a certain period of time, the results of the decisions made during simulation need to be analyzed. The provided framework for storing and analyzing results is described in Section 5.4.5. In Section 5.4.6 some more detail about the simulator architecture and implementation is given.

5.4.1 DISCRETE SIMULATION

As mentioned above, the presented model in the previous section has no notion of time, which needs to be added to simulate the dynamic behavior of the energy supply chain. Possible options for this dynamic behavior is to create a continuous, event-based or a discrete simulation. In a continuous simulation, all streams need to be described using a continuous function and requires a continuous analysis and control. Furthermore, within the optimization methods also continuous predictions of energy demand and production have to be used.

For a discrete simulation the simulation horizon is divided into a set of consecutive time intervals. The number of intervals depends on the length of the planning

horizon and the chosen length of the time intervals. For example, a whole day (24 hours) can be divided into fixed time intervals of five minutes, resulting in 288 consecutive time intervals. During each time interval, the energy flows between the devices can be determined. Thus, for each time interval for each device a new set of options O_s^d is determined, from which the control system present in the house takes a decision.

An extension to a discrete simulation is an event based simulation, where entities in the simulation model can trigger events. Events can for example be reaching the minimum level of a buffer, or switching on a consuming device. Events need to be triggered when something changes in a device, requiring a new decision to be made by the control system. The advantage of an event based simulation that only decisions need to be made when something has changed. However, events can be highly correlated, resulting in continuing chain of events. One can argue that there can be timed events, which basically results in a regular discrete simulation.

We have chosen to use discrete simulation since it allows a less complex analysis and control. Only at the beginning of each time interval, the control decision has to be made, which remains fixed during the whole time period. At the end of each time interval, the internal states of the devices are updated once and the new set of options O_s^d for each device $d \in Dev$ is determined. In other words, in each time interval the simulation model describes the status of the building in that time interval. The control system chooses an option for each device, resulting in the amount of energy flowing between the devices during this time interval.

When using a discrete simulation, one has to choose a proper value for the length of a time interval. On the one hand, a larger time interval makes the simulation less accurate. When using a large time interval, very dynamic behavior is not visible since the energy flow is only determined for whole time interval. This makes the simulation less complex and faster, since less decisions have to be made. However, if one wants to perform an accurate simulation, a short time interval length is preferable. A five minute time interval is a good tradeoff between accuracy and simulation speed [84].

5.4.2 CONFIGURATION

The simulation model describes the energy flow between the devices present in the building, and by using a discrete simulation the dynamic behavior of the devices and their usage can be simulated. Using the simulation model we should be able to model a range of buildings and devices. During simulation we must be able to create different instances of the model, representing different mixtures of buildings and devices.

In the simulator, an instance of a model can be configured. The workflow in the simulator is to first configure the different entities present in the grid. First, all the different kinds of devices implemented in the simulator, which are described in more detail in Section 5.4.6, can be configured. A device might have changeable parameters, e.g. the capacity of a heat store, or the heat production capacity of a

micro-CHP. These parameters are stored in configuration files, describing a certain version of that specific device.

After configuring each device, the devices are grouped together in a building. Since this work is mostly focussed on households, the devices are grouped in a *house*. Within a house, energy pools can be added and the streams of the devices present in the house can be connected to each other. Grid connections can be represented using exchanging devices. After specifying the devices present in the house, and determining their interconnections, the controller (see Section 5.4.3) can be selected. The group of devices, the energy pools and the controller are stored in a house configuration file.

Similar to the grouping of devices in a house configuration, a grid is constructed by grouping a set of houses. In a group, a mixture of different house configurations can be combined to create a representable mixture of buildings present in the grid. By simply adjusting a parameter in the configuration file, multiple scenarios can be easily simulated.

After configuring the grid, global parameters of a simulation can be configured. In a simulation configuration, the grid that has to be simulated can be selected and parameters determining the time span and the length of a time interval can be configured. If parameters of an individual device, like for example the demand, is defined with a different time interval length, the simulator automatically converts these parameters to the configured time interval length setup in the simulation configuration.

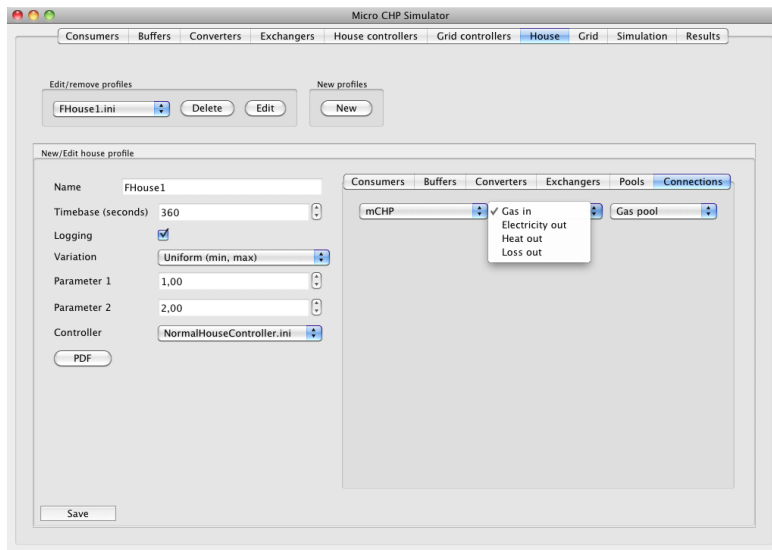


Figure 5.4: Configuration of a house

Since every entity (devices, houses, controllers) in the simulator can be config-

ured and has different configurable parameters, each entity available for simulation is responsible for providing a Graphical User Interface (GUI) where that entity can be configured. An example of the GUI to configure a house is depicted in Figure 5.4. In the figure, you can see tabs to add the four device types, the pools and how to connect the streams.

Stochastic variation

Using the structured approach as presented above of creating a configuration of the devices, houses, controllers, grids and simulations the simulator provides a lot of flexibility to simulate different scenarios. As stated in requirement 1 on page 90, the simulator must be able to simulate a realistic mix of buildings and devices. By creating a realistic mix of buildings in a grid, this can be partly achieved. By properly assigning the available devices to the houses and configuring realistic demand profiles for a house, a realistic demand profile for the grid can be simulated. However, not every building is equal and the behavior of the residents is different as well. Therefore, the simulator provides a framework to add stochastic variations to the demand profile. The simulator provides several distributions (uniform, exponential, Weibull, normal and Poisson), which gives enough flexibility to create a realistic mix of devices.

In the simulator, the demand profile of each (consuming) device can be varied. The start time of device can be shifted and/or the total load profile might be higher or lower. The shifting of the start time can be used to simulate human behavior. The adjustment of the load profile simulates the differences between individual devices of the same type.

On top of the individual device changes, a variation can be added to an individual house. Using this variation, the overall energy profile of a house can be adjusted. Using this approach, a good variation is made, preventing that only exactly identical houses are simulated.

5.4.3 CONTROL




Figure 5.5: Artificial cost function for three options of a device

After instantiating a simulation instance, based on the simulation configuration, the simulator will determine the energy flows between the devices for each time


interval. In every time interval a valid option for the devices has to be chosen such that all the pools are in balance. In a house, a set devices Dev exists, and the controller has to control all devices in this set. As described earlier, each device has as set of possible options O^d on how to control the device. Based on the state s of a device, only a subset O_s^d of O might be valid options to chose from. As mentioned above, a controller present in each modeled house is responsible for determining the right set of control signals to all the devices, i.e. choosing the right options from the sets O_s^d for each device.

As described in Section 5.3.5, each option o of a device ($o \in O^d$) describes the possible amount of energy flowing in and out of the device using the variable x_d . An example of a x_d with three valid intervals is given in Figure 5.5. The value of x_d in this example should be chosen on one of the intervals $F_1 \leq x_d \leq T_1$, $F_2 \leq x_d \leq T_2$ or $F_3 \leq x_d \leq T_3$.

The introduced constraints force all technical and non-technical constraints to be satisfied (e.g. balance in the pools and supply all demand). However, within these constraints often multiple sets of values for x_d for every device can be chosen. For example, all electricity can be imported from the grid or it can be (partly) drawn from a battery. Therefore, costs are assigned to every possible value of x_d . In other words, a cost function is defined for every device expressing the preferences of the residents, wearing of the devices, SoC of the buffers, etc. The cost functions should express the ‘quality’ of the decision for a certain value of x_d . Some decisions are more preferable than others for the residents, e.g. temporarily switching off a television is less desirable than temporarily switching off the freezer. Furthermore, switching on and off a device too often may lead to wearing. Finally, the amount of electricity imported or exported is topic of desirability, depending on the objective. These preferences can be expressed using cost functions. The costs exist of costs for picking an option (e.g. switch a device off) and the costs for the energy stream (e.g. flows from/to buffers). Therefore, the costs for every option exist of a part A depending on the internal energy stream x_o of the device and a fixed part B for choosing the option: $A_o \times x_o + B_o$. For example, if a device is on and it is preferable to keep it running, a high value of B can prevent it from being switched off. Or, if a device is not even allowed to switched off, the option of switching the device off is not added to set of possible options O_s for that specific device.

In Figure 5.6(a), an example cost function of a freezer is given. Goal of the freezer is to maintain within certain temperature ranges, preferable without as less runs as possible to minimize wearing. Therefor, there must be costs associated with switching on or off a freezer, dependent on the temperate. If the temperate is higher, the freezer not allowed to stay off and therefore the cost of switching on becomes lower. A similar rule exists for switching off. In the freezer example, no costs A dependent on the amount of energy flowing is used since the costs of electricity consumption is represented by grid import. The cost function of the grid import for example can solely depend on the amount of energy imported/exported, as shown in Figure 5.6(b).

Using the artificial costs defined for every device, the controller has to chose the best option for each device. The controller has to solve an optimization problem, in



(a) Costs for switching on/off the freezer as function of the temperature (via B) (b) Costs for importing/exporting electricity via a grid connection (via A)

Figure 5.6: Example costs functions

which for all the devices a proper value for x_d has to be chosen. Due to the artificial costs the optimization problem is reduced to a cost minimization problem with constraints:

$$\text{minimize } \sum_{d \in Dev} tc_d,$$

where

$$tc_d = \sum_{o \in O_s^d} A_o \times x_o + B_o \times c_o.$$

The big advantage of using this approach with generic cost functions is the added flexibility. Devices and their behavior can be presented to the controller in an abstract way, without knowing the internals of that specific device. This allows the integration of new, future devices in the same control methodology.

Enhanced control

The controller present in the house is responsible for choosing the right set of options for all devices present in the house. In each time interval the devices update their valid option set O_s^d and the corresponding costs functions, enforcing proper control of the devices. Energy balance constraints are added to the cost minimization problem to ensure the correctness of the model and the energy flow.

If we consider a normal house without intelligent control as described in Figure 5.3(a), the heat consumption just appears, i.e. no information on the demand is known in advance to the controller. Whenever there is a central heat or hot tap water demand e.g. a resident taking a shower, the heat demand appears to the controller. In the given situation, the controller just supplies the heat demand, both central heating and hot tap water, by the heat buffer. The only decision the controller has to take is determine when to start the micro-CHP, and this is done when the heat buffer level drops below a certain threshold level. In a normal house the electricity

consumption also just appears and together with the use of the micro-CHP this gives a total electricity flow in the house. The electricity surplus or shortage is simply exchanged with the grid. In other words, the only decision to be taken is when to start the micro-CHP and that decision is based on the level of the heat buffer.

More sophisticated control algorithms may be used to optimize runtimes of converting and consuming appliances and make smart use of the buffers. For example, some (expected) limitations about future states of the device, i.e. expected set O_s^d and their corresponding costs functions, can be taken into account when controlling a device.


Furthermore, controllers might use predictions about future energy demand to optimize the runtime of devices. As described in Chapter 3, information about historical demand data and harvested local data can be used in a forecasting scheme, determining the scheduling freedom of a device. Using the predictions, the controller is able to determine what the effect of a decision is on the future. A controller can be extended using Model Predictive Control (MPC) [9] to not only take the current time interval into account, but also future time intervals using short term predictions. More details about MPC in the local controller can be found in [62].

In the simulator, a software interface is implemented to handle devices that are able to perform prediction and give information about their possible future states. Although it is expected a real controller present in a real house will be executing the forecasting, the forecasting has been abstracted to the devices in the simulator. By shifting the forecasting intelligence into the devices, the implemented controller stays device agnostic and generic, allowing easy addition of new devices.

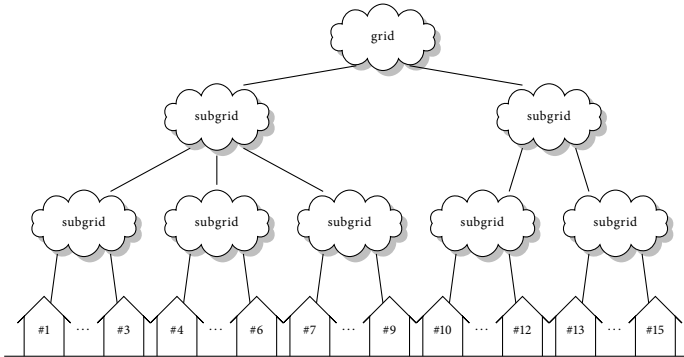
5.4.4 SMART GRID CONTROL

The house controller is responsible for properly controlling the devices in a house. However, one of the possibilities of Smart Grid technology is to incorporate cooperation between different houses and/or a global controller cooperating with multiple houses into the control scheme. The house controller, also called the local controller, can be extended to be able to communicate with other controllers, and can be influenced by using steering signals received from external grid controllers. The grid controller has its own objective and is able to cooperate with/steer a (large) group of house controllers. For example, information about the individual houses can be used to determine a planning for a group of houses and their corresponding devices (step two of the three step methodology). Dependent on the planning methodology and the complexity of the planning problem, it may be necessary to limit the amount of houses that a certain grid controller is responsible for. The deviation of the problem into multiple subproblems, as described in Chapter 4, leads to a scalable solution.

The limitation on the number of houses a grid controller can steer may lead to the use of many grid controllers. Furthermore, the grid controllers too may cooperate with other grid controllers, possibly in a hierarchical structure, as depicted in Figure 5.7. Using this approach, the overall grid is split up into different so called



(a) Steer maximal 10 controllers



(b) Steer maximal 3 controllers

Figure 5.7: Hierarchical structure of the grid infrastructure for 15 houses

subgrids, each with their own planner. Again, subgrids can be divided in the same way into a number of subgrids.

The controllers present in the grid must be able to cooperate with each other, requiring communication protocols on top of communication links. The simulator provides a generic framework that is able to divide a grid into a number of subgrids, creating and setting up grid controllers for each subgrid and creating communication links between all the controllers present in the (sub)grids. The simulator can automatically (recursively) divide the group of houses, as depicted in Figure 5.7(b), over multiple subgrids. When such a tree approach is used, the grid controller can cooperate with either a group of houses, or a group of other grid controllers. Based on the limitation of the number of houses/controllers a grid controller can plan for, the framework can automatically instantiate and interconnect extra grid controllers.

5.4.5 LOGGING AND RESULTS ANALYSIS

During simulation, the behavior of devices and the energy flows between the devices is simulated. However, a simulation is only useful if the results of a simulation can

be analyzed. In this section, the logging and analysis framework of the simulator is discussed.

The simulator provides a generic logging framework which each entity in the simulation can use to store (intermediate) results. For example, the electricity demand of a device can be logged per time interval. Using this information, the quality of the control algorithm can be analyzed.

Dependent on the simulation, not all data generated by all devices has to be analyzed and thus logged. Therefore, per entity logging can be enabled and disabled when configuring the entity. Furthermore, if for an individual house logging has been disabled, all the corresponding devices in that house configurations will not log their data as well.

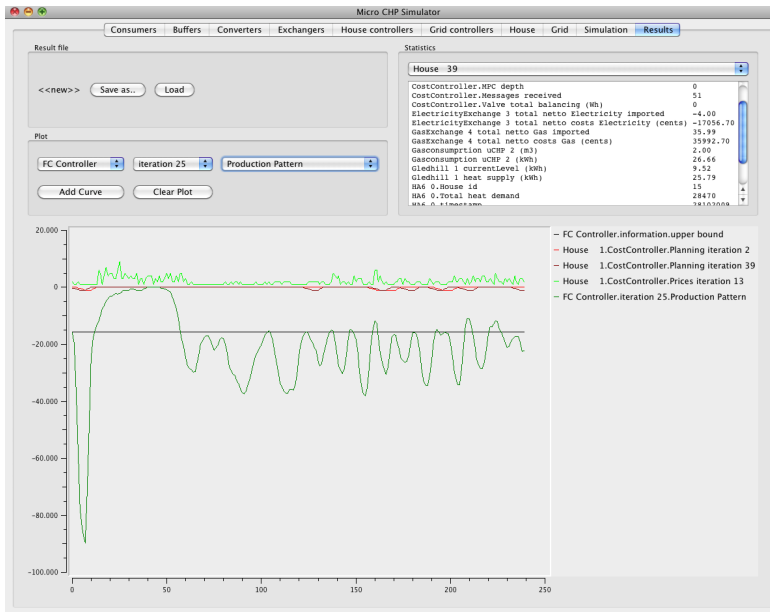


Figure 5.8: Results analysis window in the simulator

Using the logging framework, the simulator provides GUI elements to analyze the results (see Figure 5.8). In the results GUI, data varying over time can be displayed as graphs. Furthermore, overall information about the whole time period can be displayed. The simulation results can also be stored to disk and can be loaded later to analyze differences between multiple simulation configurations.

5.4.6 SIMULATOR ARCHITECTURE AND IMPLEMENTATION

Based on the energy stream model and the requirements set in Section 5.1, a software design and implementation of the simulator has been designed and implemented.

The simulator has been designed using an object-oriented approach, where the C++ programming language is used to implement the software.

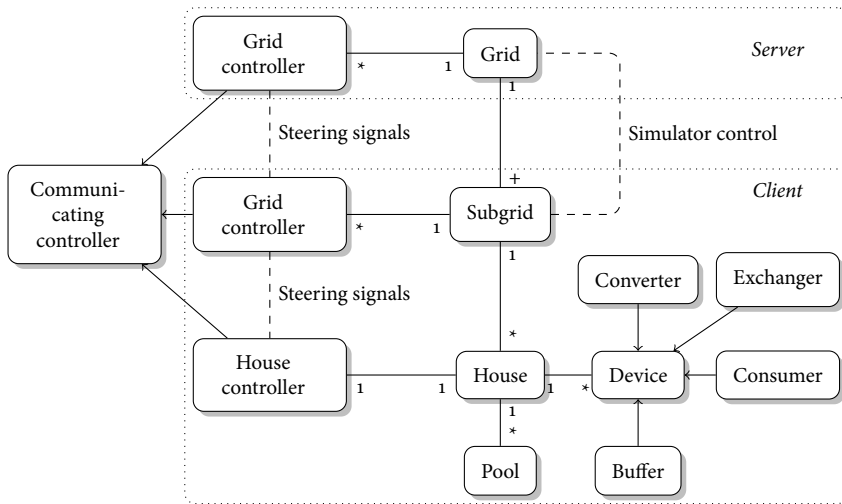


Figure 5.9: Design of the model and the network architecture.

For each model-entity in Figure 5.9, a separate class is built. Flexibility is obtained by defining the devices and controllers as abstract classes. In the abstract classes minimum required functionality is implemented such that all requirements of the model are met. An implementation of an actual element consists of a class that extends the abstract class and implements the abstract defined functionality. Optionally, standard behavior can be adapted by overriding the corresponding functions. In this way the class implements the specific behavior of the element.

Although all devices have the same interface in order to be treated equally by the controller, still four pre-defined subclasses of devices are specified. All device-category specific basic behavior is implemented in these four abstract classes representing *Converters*, *Exchangers*, *Consumers* and *Buffers*. In each of these four subclasses, basic elementary features required by the corresponding device-category are implemented. All buffers for example have a certain State of Charge, while all consuming devices have a certain demand profile. Using such an approach, only limited amount of code has to be written to add a new device type to the model. More precisely, only roughly 100 lines of codes have to be written to add a new device implementation.

For each of the four devices categories, a number of devices has been implemented, which are discussed briefly:

Converter Up to now, a High Efficiency Boiler, a WhisperGen micro-CHP appliance and a GenericConverter class are the converters available in the simulator. The first two classes represent real life-appliances with their device

specific characteristics. For example, a WhisperGen micro-CHP appliance has limitations like minimum runtime and minimum cool-down periods. The GenericConverter class is a basic converting device which can convert one energy stream into another with a configurable efficiency, and can be used to model converters without state.

Consumer The available consuming devices within the current simulator are a standard appliance and a freezer. A freezer is a device with keeps tracks of the internal temperature. Based on the internal temperature, the set of valid options and costs functions is determined each time interval to ensure the internal temperature stays within the specified bound. Thus, based on this temperature the electricity demand is determined.

A standard appliance can model appliances which have a certain predefined demand profile. A demand profile is determined by the time interval the device is switched on and a sequence of electricity demand values. For example, a lamp has a fixed pattern once it is switched on. Its corresponding demand pattern is for example that it switches on at 6pm and has a constant electricity demand for 5 hours of 40 W. However, other demand sequences are also possible, like the changing electricity demand of a washing machine during the different phases of a washing cycle. A standard appliance can be configured to have a start time and a pre-defined demand profile, consuming one or multiple energy-types. To create different instances of the same house type, on the start time and runtime of these appliances a stochastic variation can be added to model more realistic user profiles.

Buffer A Gledhill heat store and an implementation of the Kinetic Battery Model (KiBaM) [58] are available buffer devices. The heat store can be configured to be of different sizes. Using this size parameter, the effect of the size of the heat store can be analyzed. The KiBaM models a realistic battery, emulating the characteristics of a real battery.

Exchanger Only a single exchanging device type is implemented. Each exchanging device can exchange one energy carrier. For each time interval of the simulation, a limitation on the amount of energy that can be exchanged can be configured. For example, a electricity connection to a house can be limited to 35 A or the amount of wind imported to the house can be set to a predefined pattern.

The devices are connected with each other using pools, which are also present as a separate class. A pool object keeps track which devices are connected to that pool, allowing validation of the required energy balance in each pool.

Using the house class, all the devices and pools are connected and grouped. Furthermore, a house controller, is assigned to a house here. The house controller is delegated, allowing multiple implementations of a house controller. For example, more advanced house controllers with different control strategies can be implemented and used to analyze the effect of the new control algorithms. Due to the

abstract definitions of the device class, the same interface to communicate with and control devices can be used.

The grid class is responsible for grouping and creation of all the houses during a simulation. Although the model support pools between each house, a simplification has been made that the grid has no limitations on transportation and 'production'. Limitation for each household, for example the maximum amount of flow to each household, can be enforced using the exchanging devices. The grid also handles the limited amount of houses a grid controllers can steer and creates and configures a hierarchical tree structure of multiple grid controllers, as depicted in Figure 5.7(b) on page 105. Although depicted in this figure as two separate classes, in the actual implementation only one class is used which can function as the grid and as a subgrid.

Since the house controllers can cooperate, and thus communicate, with other controllers, a basic *CommunicatingController* class has been implemented to handle the communication. It provides a framework in which a controller can become a server (for example the top grid controller), a client (for example a house controller) or both (all the intermediate grid controllers).

5.5 DISTRIBUTED SIMULATION

Up to now, the simulation model and the translation of the model into a simulation design and implementation have been given. The simulation model is flexible and versatile, but it also comes with a cost of a higher complexity to determine the energy flows. To speed up the simulator and enable the possibility of simulating a large group of house, a distributed version of the simulator has been developed, which is discussed in more detail in this section.

As discussed in the previous sections, all the energy pools must be in balance, which can be reached in many different ways with different associated costs. Determining balance for all energy pools with minimal costs for all time intervals makes a simulation with many devices and many houses a computational intensive job. To keep the simulation time within reasonable limits, the computation can be partitioned in smaller parts and distributed over multiple computers via a network. Considering the entities in the model, the model has a lot of opportunities for parallelization. Houses are independent of one another, allowing these entities to be simulated on different machines. Furthermore, within a house, a house controller has to decide which options are chosen for each device. Once this decision is made, devices may need to update their internal state. These updates can all be performed in parallel.

For this reason, the simulation model is extended with a server-client model (see Figure 5.9). The server has three responsibilities. The first task is the configuration of a simulation. A configuration defines a grid and a possible grid controller. A grid consists of a group of (different kinds of) houses. A house consists of multiple devices, where each device is connected to one or multiple pools. As mentioned earlier, each house has a house controller.

Each entity, i.e. the grid, house controllers, devices etc; can be configured with the GUI at the server. Once a configuration is built, it can be simulated. When a simulation is started at the server, it automatically distributes the houses configured in the grid over so called subgrids and sends the configuration for each subgrid to each connected client.

During simulation, information is exchanged between the server and the clients. The server is responsible for synchronizing the clients, since some clients may be faster than others.

When all clients are finished, the server is responsible for aggregating all local simulation results into a global simulation result. When all data is aggregated, it can be analyzed on the server. Since the amount of data can be quite substantial, simulation results are kept at the clients as much as possible. Only the relevant information required to aggregate (sub)grids and information about the simulation results are requested by the server. For example, all houses, controllers and devices can generate simulation results. On the client, these results can be displayed in the GUI via text or plots. The GUI needs to know which information can be displayed and on which client this information is available. When certain data is requested, the server looks up the origin of the information and sends a request for the information to the client. This minimizes network traffic and the amount of required memory at the server. When a simulation has to be saved (to disk) for later analysis, all received information is stored directly to disk, minimizing the amount of data stored in memory.

The division of the grid into multiple subgrids enables the possibility to create a hierarchical structure. For example, multiple neighborhoods can be simulated. Each neighborhood consists of a mixture of different houses with different usage profiles. This mixture of neighborhoods can be divided into different subgrids, which can be simulated at different clients. The grid controller at the server cooperates with the grid controller of each subgrid on the clients. Each grid controller of a subgrid cooperates with both the main grid controller and each house controller. An example of such a hierarchical approach is the the creation of a Virtual Power Plant [60]. The goal here is to optimally control a large group of micro-generators, like micro-CHP appliances, to generate a certain electricity profile. Objectives can be to minimize purchase costs of energy retailers or to ensure stability on the electricity network. Due to the large size of the fleet, it is impossible to optimally control the fleet centrally [17]. By dividing the whole control system in smaller subsystems, the subgrid controllers and house controllers can optimize the runtime of the micro-generators within the subgrid. The approach leads to a faster, more scalable system.

The network stack provides an interface to facilitate the communication between the controllers. Another advantage of this approach is that also the communication requirements between the controllers in a real life setting can be analyzed. For example, characteristics of a certain communication medium like GPRS can be emulated. All data sent between the controllers can be sent via this emulated communication channel. Communication properties like delay, packet loss or limited bandwidth can be simulated. Using such an approach, the amount of

required bandwidth or the fault tolerance of the controllers can be determined.

5.5.1 PROTOCOL

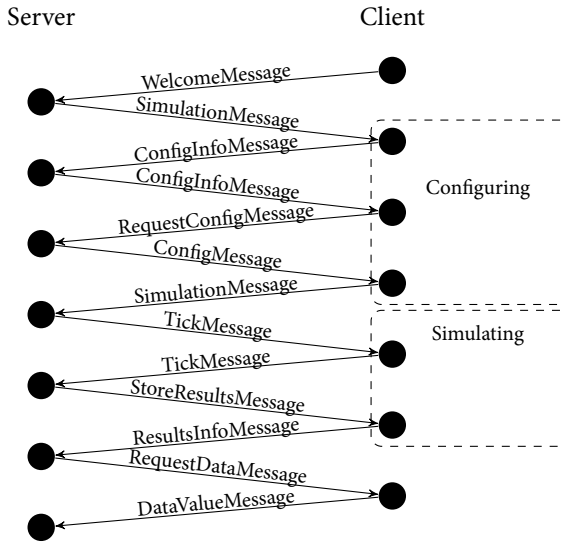


Figure 5.10: The developed network protocol

One of the requirements of the improved simulator was the ability to simulate larger groups of houses. By distributing the houses over multiple computers, the memory and computational requirements of the simulation are spread over multiple computers. To make the approach scalable the network overhead has to be limited, requiring a fast and efficient protocol.

Since a simulation is distributed over multiple computers, all information required for a simulation has to be available at each client. As a consequence, the protocol has to thus be able to distribute all required information prior to the construction of the (sub)grids, houses etc.

Based on the flow of a simulation, a protocol between the server and the client is developed. After connecting, a welcome message is sent to confirm that a simulation client is connected and not some random process.

When the user selects a simulation configuration to be simulated, the server automatically distributes the required houses over the connected clients. Simulation configurations are simple files, which can be easily transferred. For each client, a new configuration file of the subgrid is created, and information about this file is sent via a SimulationMessage. When the client retrieves this SimulationMessage, it first looks if it has a cached version of the configuration file. If it has a cached version, it has to check whether this file has been changed since it was obtained. This is done by sending a ConfigInfoMessage, which contains the filename of the configuration and

a hash of its contents. When the server retrieves a `ConfigInfoMessage`, it constructs a `ConfigInfoMessage` from its own configuration and sends it back to the client. When the hashes are equal, the file is up-to-date.

If the hashes are not equal, or when the configuration file did not exist at the client, a `RequestConfigMessage` is sent. When the server receives such a message, the file is read from disk, compressed and sent to the client. Once the configuration file is up-to-date at the client, it is scanned for dependencies. For example, a grid consists of houses and houses have their own configuration (files). For each dependency, their configuration files are exchanged in a similar way. When all configuration files are up-to-date, the client confirms it has finished the configuration phase by sending a `SimulationMessage` to the server.

When all clients are configured, the simulation can start. Since a discrete simulation is used, the simulation starts by simulating the first time interval. In each time interval, all the entities in the model receive a so called 'tick', signaling the start of a new time interval. Therefore, after configuring the server sends a `TickMessage` to simulate the first tick.

Each client simulates the time interval and when it is done it confirms replying with a `TickMessage`. The server waits for all clients to confirm their tick, and then sends a new `TickMessage` to each client. This way, all the clients are synchronized with each other, which is required when simulating a global optimization algorithm.

When all time intervals are ticked, a `StoreResultsMessage` is sent. During this phase, statistics and information about the whole simulation are calculated at the clients. The completion of this phase is confirmed by each client with a `ResultsInfoMessage`. This `ResultsInfoMessage` contains information about which simulation results are available at the client. Note that all the results are still stored distributed over the clients. The server only collects where the information is stored.

Since the grid is split up into subgrids, the server combines the subgrids into a global grid by aggregating the data. When data of a specific house or grid is required, the server requests the data by sending a `RequestDataMessage` to the client which has this data. The client sends the required information to the server with a `DataValueMessage`. After aggregating all subgrid data into a global grid, the GUI on the server is informed that the simulation has completed.

When the simulation has to be saved to disk, or when information about the simulation has to be displayed in the GUI, the server requests the required data from the clients on demand. This way, the amount of memory required on the server and the amount of data transferred over the network is limited.

The simulation results are discarded at the client locally when the client is disconnected. Since it might be possible another simulation has to be executed, the client automatically reconnects to the server, and the process starts again from the beginning. Using this approach, a client can be installed on many machines, creating a large cluster of simulation clients. Since the discovery of the server is completely automatic, no user interaction at the client is required.




Figure 5.11: The simulation software architecture

5.5.2 SOFTWARE ARCHITECTURE

Based on the model and the developed protocol, the simulation is split up into multiple parts. The basis of the simulator (the presented model) has been written into a library. Since the server and the stand-alone GUI share user interface elements to configure all the entities in the model, these elements are added to the library.

A ‘stand-alone’ simulation without a server/client model can be used for smaller simulation instances. For small simulation instances which can be performed on a single machine fast enough, it is not required to add the network overhead.

Added to the code base is the simulation protocol and the required c++ classes for a server and the client. The server is similar to the stand-alone version of the simulator. With this program, the entities can be configured and simulation configurations can be created. When one or more clients are connected to the server, simulations can be performed and their results can be displayed. Note that the server cannot simulate a configuration, this always has to be delegated to a client. The client is responsible for the real simulation. Based on the configuration received from the server, the real model is calculated at the client.

5.5.3 VALIDATION

As mentioned in previous sections, the energy flow during a time interval is determined by the energy demand, the amount of available (stored) energy and the production capacity of the producers. The house controllers are responsible for selecting the right set of options during simulation. As described in [61], the model can be expressed as a Mixed Integer Problem (MIP). Using this approach, the validity of the model can be ensured by adding the proper constraints to the MIP. For example, the balance within the pools can be easily expressed by requiring the sum of the energy flow in the pool to be zero. Furthermore, assertions in the code are used to check that the assumed conditions remain valid. The correctness of the implementation of the other classes is checked using unit testing.

5.5.4 PERFORMANCE

The developed simulator is used to analyze many use cases and scenarios, of which some of the results are given in Chapter 6. To simulate a large fleet of houses a lot of computational power is required. This is especially caused by the local control algorithm that is executed every time interval for every house.

Table 5.1: Simulation speed (seconds).

	1 house	100 houses	500 houses	1000 houses
standalone	1	47	234	490
network, one client	1	50	246	491
network, two clients	-	28	123	245
network, three clients	-	20	88	170

To measure the speed and the influence of the network layer the same scenario is simulated for 1, 100, 500 and 1000 houses. In this scenario, each house consists of multiple electricity consuming appliances, a heat consuming appliance (central heating), a heat buffer and a micro-CHP. The objective is to act as a Virtual Power Plant: an on beforehand determined planning of the runtimes of the micro-CHPs should be met. The planning is based on prediction of the heat usage.

The scenario is simulated using the standalone version, the network version with one client (measuring network overhead) and the network version with multiple clients with similar specifications to analyze the speedup.

The simulation times are shown in Table 5.1. It can be seen that the network overhead is negligible for large simulation instances and the speedup using the network version is significant. Simulating larger instances leads to less network overhead, since the aggregation of the subgrids is determined by the amount of clients, not the amount of houses. The speedup is almost linear with the number of clients (only a bit network overhead), the merging of the results at the end of the simulation is done by the server and decreases the speedup slightly. Due to the synchronization in the protocol, the speedup is limited by the slowest client.

5.6 CONCLUSIONS

The energy model and its translation to the simulator provide useful tools to analyze current and future smart grid technology and their impact on the grid. The model's flexibility enables the possibility to simulate the whole energy supply chain and even future (smart) devices (requirements 3 and 4). By combining a realistic set of devices in a house and a realistic mix of house configurations to a grid, a realistic simulation of energy flow of houses and the whole grid can be simulated (requirements 1 and 2). This is further improved by using the stochastic variation provided by the simulator. By distributing the simulation over multiple PCs, both the simulation runtime and memory requirements can be distributed over multiple machines (requirements 2 and 6). By distributing the logged data over multiple machines, the amount of data is still feasible to store in Random Access Memory (RAM). The amount of data can even be configured on a device level, allowing flexible analysis of simulation (requirement 5). The speed up reached using the distributed simulation is determined by the slowest machine in the network, but scales almost linearly. The required communication framework of the distributed simulation can also

be exploited to analyze the network requirements of cooperating optimization methodologies (requirement 7)

Although all requirements are met, still the simulator can be improved. As mentioned above, the grid is assumed to provide and transport all the required energy. Although the model provides the possibility to model the grid with real devices for the production, storage and transport of energy, this is not implemented yet and left for future work. Furthermore, in the current implementation the smallest allowed time period is one second. For analysis of issues like power quality (for example frequency deviation) and communication latency a shorter time period may be necessary.



RESULTS

ABSTRACT – In this chapter the results of different use cases are presented. Common in all use cases is that the heat demand forecasts, as determined in Chapter 3, are used in real situations. The first use case, creating a VPP using a fleet of micro-CHP appliances, shows that TRIANA is capable of reshaping the production profile to a predetermined profile. Results show that a good tradeoff during planning between buffer size, forecast quality and allowed exploitation of the scheduling freedom must be made.

In the second use case, a fleet of heat pumps is steered towards a flattened profile. Although both a micro-CHP and a heat pump can be used for providing heat to a building, their internal working and restrictions on how to operate the devices are different. An extra addition in this use case is the usage of replanning. Like in the VPP use case, the planning program should reserve some capacity to deal with forecasting errors locally, reducing the deviation from the planning.

In the last use case, again the objective is to steer the electricity profile towards a flattened profile. In this use case, a mixture of micro-CHP appliances and heat pumps are simulated. Where a micro-CHP appliance generates electricity while generating heat, consumes a heat pump electricity while generating heat. Although the devices have different characteristics, they both respond to the shared steering vector. However, changes in the steering vector have bigger effects on the micro-CHP appliance than on the heat pump. Therefore, the costs functions, used during planning, should be defined such that they are evenly sensitive to the steering signal. Another possibility is to create a planner which generates a planning, respecting the characteristics and limitations of the device, and extending the cost function by adding additional costs for deviating from this planning.

In this chapter different use cases are presented. Common in all use cases is that the heat demand forecasts, as determined in Chapter 3, are used in real situations.

Parts of this chapter have been presented at [VB:2], [VB:6], [VB:8], [VB:9], and [VB:19]

In the first use case, described in the next section, we analyze the possibility of creating of a Virtual Power Plant (VPP) using micro-Combined Heat and Power (CHP) appliances using TRIANA. Section 6.2 describes a use case where TRIANA is used to flatten the consumption profile of a fleet of heat pumps. In the last use case, described in Section 6.3, a fleet containing both micro-CHP and heat pumps are steered. In this use case, we try to steer the mixture of devices to also reach a flattened profile by steering the electricity consumption of the heat pumps toward the periods where the micro-CHP appliance generates electricity.

6.1 VIRTUAL POWER PLANT

This work, as mentioned in the introduction, was part of the SFEER project. The goal of this project was the creation of a VPP using a large fleet of micro-CHP appliances. A micro-CHP appliance is a system that produces heat and — as a by-product during the heat production — electricity. Current generation micro-CHP appliances are fueled by natural gas. They can generate electricity at the kilowatt level which allows these units to be installed in an individual home. They are connected directly to the domestic heating and electrical systems, which leads to a very high efficiency (up to 97%) in usage of primary energy. The heat is used for the heat demand in the home such as central heating, showering, hot water taps etc. The electricity can be used in the house or, when not needed, be exported to the electricity distribution network.

The (electricity) production of a micro-CHP appliance is heat driven, since it only produces electricity while producing heat. Adding a heat buffer (hot water tank) decouples the demand and production of heat, within the limits of the heat demand and the buffer size. This gives flexibility in the electricity production, allowing the production of electricity on more beneficial periods. For example, a peak in the electricity demand can be seen when people get home. During this period, electricity can be generated by the micro-CHP system and used within the home by the appliances switched on. The heat can be used for central heating or to fill the hot water tank. The stored heat can be used the next morning for showering.

It is expected that micro-CHP appliances will replace the current high efficiency boilers [76]. For example, in the Netherlands, the target market in 2002 for micro-CHP appliances was 4.4 million households [77]. It is expected that in 2020 between 0.9 and 1.4 million households will have a micro-CHP appliance and is expected to grow up to 2 million in 2030. When the number of micro-CHP appliances becomes high enough, generators can be virtually grouped together and become a VPP. By controlling and smart scheduling such a fleet of generators a VPP may replace a conventional (less-efficient) power plant. Using a VPP instead of a conventional one will result in a significant reduction in costs and CO₂ emissions due to a more optimal use of primary energy sources.

The goal of the control system is to autonomously determine the scheduling freedom introduced by the heat buffer and use this scheduling freedom of the VPP for commercial exploitation. As mentioned in Section 2.1.1 on page 16, the electricity

infrastructure always has to be in balance. The transportation and distribution network operators use different control systems and corresponding electricity markets to ensure this balance. As a result, balance has a value.

Since a VPP consists of many small generators, which can start and stop within a couple of minutes, a VPP has the potential to be used on the short-term markets. Dependent on the stakeholder of the VPP, the VPP can be used for balancing the grid by a network company or to reduce purchase costs/penalties by a utility. Important when using a large fleet of small generators is the available production capacity of such a large fleet and the guarantee that this production capacity can be exploited. As a consequence, to use a VPP, the production capacity of the fleet has to be forecasted at reasonable accuracy. This will ensure that the promised production capacity is really available.

For micro-CHP appliances the electricity production capacity is based on the heat demand. Thus an accurate heat demand forecast is required. In Chapter 3 the used approach in determining the heat demand for a household is described. To test and analyze the possibilities of creating a VPP using TRIANA, in which these forecasts are used, a use case consisting of 50 houses equipped with a micro-CHP appliances is used. In the remainder of this section first the determination of the scheduling freedom and the exploitation of this scheduling freedom is explained. Then the defined use case is described in more detail. In the last section, results and discussion of the simulations are given.

6.1.1 DETERMINING THE SCHEDULING FREEDOM OF A MICRO-CHP APPLIANCE

The micro-CHP appliance together with the heat store are responsible for providing heat to the house. This heat store can be used only for hot tap water to provide a generous supply, while the central heating is provided directly by a boiler. Another possibility is using the heat store for both central heating and tap water. In this approach, the micro-CHP appliance can be used to produce all the heat required in the house, maximizing the electricity production of the micro-CHP appliance. In this use case, a heat store providing both the hot tap water and the central heating demand is used.

In order to provide the required heat demand, the buffer should always contain enough heat to supply the requested demand. The state of charge of the buffer should thus remain above a certain threshold level. Once the state of charge drops below this lower level, the micro-CHP should be started. After running the micro-CHP for a while, the micro-CHP should be switched off when the buffer is (almost) full. This process is illustrated in Figure 6.1.

However, the micro-CHP has some limitations on how the appliance can be used. For example, after switching on a micro-CHP appliance, it takes a while before it starts to produce heat at maximum level. In this startup period, the heat and electricity production increases up to the maximum production capacity (see interval S in the figure). A lower level in the heat store is used to provide any heat demand request while the micro-CHP is not producing heat yet.

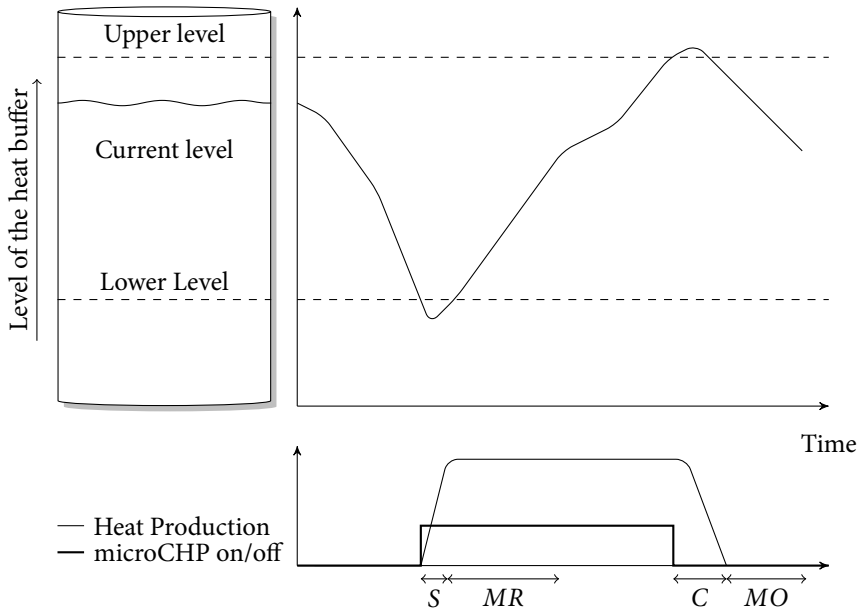


Figure 6.1: The use and limitations of the heat buffer in combination with a micro-CHP appliance

Once the appliance is running at full capacity, it has to keep running for at least a minimum amount of time. This minimal runtime (MR) ensures the appliance makes long runs instead of running in series of small runs. After the minimal runtime has been reached, the device is allowed to switch off, which normally is signaled when the heat store is almost full. Once the micro-CHP appliance has received a signal to switch off, the Stirling engine still contains enough heat to continue the heat production for a short period of time. During this cooldown period C the heat production slowly decreases. To ensure that the heat being produced during the cooldown period can be stored, the upper level switch point must be chosen such that the heat store can absorb this heat at all times.

The final constraint of the micro-CHP appliance is the requirement that once the appliance has been shut down, it has to maintain off a while before it can be used for the next run. During this minimum offtime (MO) the device is not allowed to start.

Within the minimum and maximum level, the micro-CHP appliance is allowed to start whenever is required, as long as the device constraints of the minimum runtime and minimum offtime are respected. This flexibility is what determines the scheduling freedom. Preferably, the micro-CHP appliance should be running when the electricity prices are high, since then the purchase costs are reduced the

most, or in case of exportation to the grid the electricity gives you the most profit.

The problem with determining and exploiting the scheduling freedom of the micro-CHP appliance is that once a decision on the runtime of the micro-CHP is taken, it restricts the possibilities for a long period of time. Therefore, the chosen time intervals to start a micro-CHP must be chosen carefully. Due to these restriction, the planning of a fleet of micro-CHP appliances is NP-complete [17].

In this use case a central planner described by Bosman et al. [14] is used to determine the required production profile of the fleet, optimizing the profitability of the runtime of the machine. Based on the prices of the day-ahead market (see Section 2.2 for information about the available energy markets), the electricity production of the micro-CHP is shifted towards the high-price periods. Using this approach, less expensive energy has to be bought by an energy supplier, which in this case is the operator of the fleet of micro-CHP appliances. During the planning process other constraints, for example imposed by a grid operator, can be added.

To ensure that a local controller is able to work around forecast errors, not the full optimization potential should be exploited by the planner. The above described upper and lower levels, together with capacity required by the real time controller to cope with forecast errors, determine the (artificial) upper and lower bounds the planner has to respect while planning.

6.1.2 USE CASE DESCRIPTION

In this use case, a VPP consisting of 50 houses equipped with a micro-CHP and a heat store is used. The currently commercially available micro-CHP devices based on a Stirling engine have a 1 kW electrical and 8 kW thermal production capacity. Therefore, in our simulation model, a micro-CHP appliance with this electrical and thermal production capacity has been used.

For the thermal store, a 10 kWh and 20 kWh heat buffer size has been used. Using the 10 kWh buffer size, two different planning settings have been used. In the first variant, the planner uses the 1 kWh and 9 kWh as the lower respectively upper bound on the allowed heat store level during planning. In the second variant, the bounds are set to 2 kWh and 8 kWh as bounds, introduced more reserve capacity for the real-time controller to handle forecast errors. Using the 20 kWh buffer size, the bounds during planning are set to 4 kWh and 16 kWh. It is expected that the increased buffer size introduces more scheduling freedom, allowing better steering towards the desired periods of production.

To simulate the heat demand of 50 houses, data extracted from our heat demand database is used. This database contains heat demand data of four households from the beginning of January up to December 2006, and of six houses from October 2009 up to February 2010. Reusing the forecasting results from Chapter 3, 50 days with a predicted heat demand between 50 kWh and 80 kWh have been selected to represent cold days. The average forecasted heat demand (per day) was 63.7 kWh, with a standard deviation of 9.2 kWh. The corresponding average real heat demand of these days was 57.8 kWh, with a standard deviation of 17.4 kWh. The forecasted heat demand was thus on average 5.9 kWh higher than the actual heat demand,




Figure 6.2: The APX electricity prices used

which is a substantial part of the heat store capacity. This makes it very hard for the real-time controller to cope with the forecast errors.

Planning the micro-CHP appliances

The objective of the planner is to maximize the revenue on the produced electricity, using the prices of the APX market. Therefore, the planner steers the micro-CHP appliances to produce electricity during the periods with a high electricity price. The used electricity prices, which are market prices for electricity of Nov 2, 2010, used in use case are depicted in Figure 6.2.

As can be seen in the table, during the morning period and at the beginning of the evening, the prices are the highest. It is thus beneficial to produce electricity during these periods.

Since the planner tries to solve a maximization problem, the planner could fill the buffers as much as possible towards to end of the day, since this can produce more electricity and may thus improves the desired objective. Therefore, we restrict the amount of flexibility the planning can exploit. At the beginning of the day, all the heat buffers are modeled to be half full. At the end of the day, the total amount of energy stored should be less than or equal to the amount of energy stored at the beginning of the day. This way we ensure that only the actual heat demand is generated, and enough flexibility remains for the next day.

6.1.3 RESULTS

Using the simulator, the 50 houses equipped with the micro-CHP appliance and the heat stores have been simulated. First, these houses are simulated without any steering to determine the reference case, i.e. see what happens without using TRIANA. The results of these reference simulations are depicted in Figure 6.3. In the figure, clearly the morning peak can be detected. With both buffer sizes, the buffers are half full in the beginning. Due to the heat demand during the night, these buffers are emptied at the start of the day, resulting in the first filling cycle. When in the morning these buffers are emptied by the morning heat demand peak, all the micro-CHP appliances again start to produce heat and thus electricity. As also shown in the figure, the production peaks are very dependent on the buffer size.

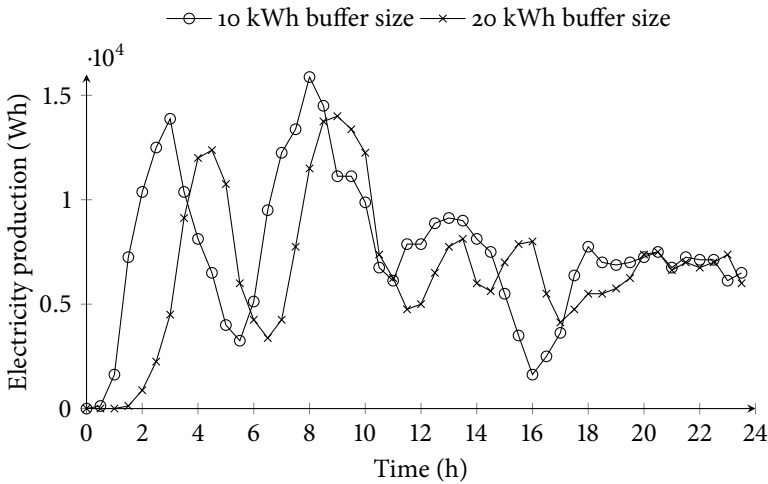


Figure 6.3: Electricity production profile without any steering for the two buffer sizes

The results after generating a planning for houses equipped with the 10 kWh buffer size and the 1-9 kWh planning bounds are depicted in Figure 6.4. Here the influence of the planning can clearly be seen. As expected, the production of electricity is shifted towards the periods with the highest price. However, since the planning is generated based on the heat demand forecasts, the planning cannot be reached for all time intervals. Furthermore, since the planning uses a very big part of the scheduling freedom for planning, the real-time controller does not have enough reserve capacity to handle the forecast errors. The desired peaks at 7h and 11h therefore cannot be reached.

In Figure 6.5 the results after introducing more reserve capacity for the real-time controller during planning, i.e. set the bounds to 2-8 kWh, is depicted. Due to the restrictions added to the planner, less optimization potential is exploited, as




Figure 6.4: The target and realized electricity production of the vpp (10 kWh/1-9kWh bounds)

can be seen in the figure by the more flattened desired profile. As a result, the real-time controller is better able to follow the desired planning, resulting in a lower mismatch between the planning and the realization of the planner.

By doubling the buffer sizes and the corresponding bounds, more flexibility for

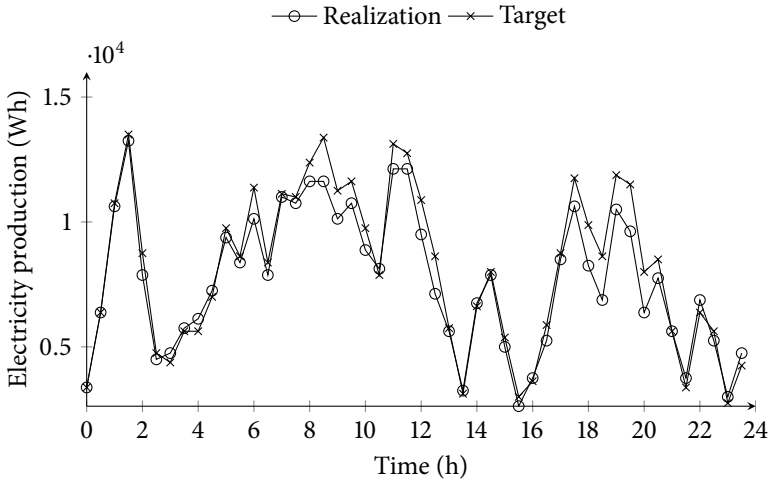


Figure 6.5: The target and realized electricity production of the vpp (10 kWh/2-8kWh bounds)




Figure 6.6: The target and realized electricity production of the vPP (20 kWh/4-16kWh bounds)

Table 6.1: Quality figures of the vPP use case

Buffer size	Field	Realization	Planning (Forecast)	Planning (Actual)
10 kWh	Total production (Wh)	364000	389625	350875
1-9 kWh	Average price (Euro/MWh)	30.39	30.35	31.30
10 kWh	Total production (Wh)	367250	390250	354625
2-8 kWh	Average price (Euro/MWh)	29.90	29.50	30.70
20 kWh	Total production (Wh)	364625	392125	354375
4-16 kWh	Average price (Euro/MWh)	31.68	30.26	31.78

the planner is introduced. The results of this variation is depicted in Figure 6.6. Due to the extra flexibility, the desired profile resembles the desired profile in Figure 6.4, where the production of electricity is steered towards the high price periods. Furthermore, due to the bigger buffer size and higher reserve capacity for the real-time controller, the planning is followed better.

To quantify the results discussed above, the average price for electricity and the Mean Absolute Percentage Error (MAPE) between the actual and desired profile are used.

In Table 6.1 the total amount of electricity produced and the average price for this production is given. In this table the values for the realization of the planning and a planning based on the forecasted heat demand are given. In the last column of this table values for a planning based on the actual heat demand are given to give

an indication of the potential of the planner.

The table shows, in line with Figure 6.4 up to Figure 6.6, that the total amount of electricity produced is less than was planned. However, the electricity that has been produced, was mostly produced in the high price periods, resulting in a good average electricity price for the produced electricity.

From Table 6.1 we can also see the potential of the planner by looking at the last column in this table, where the results of a planning using the actual heat demand is given. A bigger flexibility in the planner yields to a higher price for the produced electricity, as was expected. The higher flexibility, using a buffer of 20 kWh and bounds on 4-16 kWh clearly gives the best performance.

The MAPE of the realization and the planned profile are 0.12% for the 10 kWh buffer size and 1-9 kWh bounds. Reducing the bounds to 4-8 kWh improved the MAPE to a value of 0.08%. As expected, the extra reserve capacity for the real-time controller results into a realized profile closer to the desired profile. The MAPE of the 20 kWh case was 0.11%. Thus although the real-time controller has some reserve capacity for handling the forecast errors, the fleet has difficulties following a peaky profile. From the graphs shown in Figure 6.6 we can see that the real-time controller uses the reserve capacity to handle the forecast errors. However, around 9h in the morning, the capacity is not enough and deviations start to occur.

6.1.4 CONCLUSIONS

As shown in this case TRIANA is capable of reshaping the production profile of a fleet of micro-CHP appliances. Important in this approach is that the real-time controller, the last step of TRIANA, has enough reserve capacity to be able to cope with forecast errors. As shown in this use case by using a fixed buffer size, but different capacities that can be exploited by the planner, that a conservative planning yields better results. The use of a larger buffer introduces some extra scheduling freedom, leading to more extremely planned profiles. As result, due to forecasts errors, the real-time controller has difficulties handling the forecast errors caused by the restrictions set on the runtime by the micro-CHP. Therefore, a good tradeoff during planning between buffer size, forecast quality and allowed exploitation of the scheduling freedom must be made.

6.2 HEAT PUMP USE CASE

In this use case, the ability to reshape the electricity demand profile of a group of households equipped with a heat pump using TRIANA is analyzed. In older houses with less insulation quality, gas-fired micro-CHP appliances or conventional high efficiency boilers are often used due to their high production capacity of heat. However, the buildings and especially the insulation quality of newly build houses have improved. As a result, the heat demand of modern buildings is reduced significantly, allowing a converter with lower production capacity.

Heat pumps are increasingly regarded as an attractive option for domestic heating. Instead of burning natural gas, diverts a heat pump heat from the ground

or the surrounding air at a lower temperature to the house at a higher temperature using mechanical work. The process is similar to a refrigerator, but instead of cooling the thermodynamical process is used for heating. However, during summer the heat pump can be used for cooling as well.

Since no gas is required for heating, the gas demand of a house decreases significantly. Investments in the local gas distribution infrastructure in newly developed neighborhoods becomes unattractive due to lower penetration of gas applications [30] and higher insulation standards [70]. However, the lack of a gas distribution infrastructure means that all activities in a building that normally consume gas must be provided by another form of energy, mostly in the form of electricity. Common gas consuming activities, like heating and cooking, can be done using electricity, but they consume a lot of energy. Especially when the heat pumps are switched on during periods when people are also getting home, switching on their appliances and the electric cooking stove, a large amount of electricity has to be transported and provided to the houses. Since most people more or less have a similar living pattern, this leads to a high peak in the electricity demand. Since the electricity distribution network must be able to provide these high peak, extra investments have to be made in the distribution network, only supplying these peaks. These investments in the electricity distribution grid only to supply the peak lead to very high investment costs, and is due to the low utilization not very efficient financially.

Therefore, it is very beneficial to reduce the peaks in the electricity distribution network. By shifting the run cycles of the heat pump to low-demand periods, the peaks can be reduced. This spreading is possible by using a heat store in the house, which introduces the flexibility of when to use the heat pump.

In the rest of this section, first the modeled heat pump and its characteristics are described. Then the objective is given, and based on this objective an Integer Linear Programming (ILP) formulation and the corresponding results are given. Next, the obtained results by using the TRIANA approach are presented.

6.2.1 HEAT PUMP MODEL

The heat pump is classified as a converting device: it converts heat of one temperature to heat of another temperature, while consuming electricity. Often, heat pumps are used for heating, converting heat of a lower temperature to a higher temperature. Therefore, the device has four streams: the power supply (electricity in), source stream (heat in/out), sink stream (heat out/in) and loss (heat out).

The efficiency of a heat pump depends on the temperature difference between the source and the sink element (ΔT). The Coefficient Of Performance (COP) of the heat pump is defined as the ratio of the heat displacement (from the source to the sink stream) to the required work. A higher COP thus means a higher efficiency. At each instance in time, the device enforces this fixed conversion ratio between inputs and outputs. We assume that the efficiency of the heat pump is fixed.

The performance of a heat pump is bounded, which we model by limiting the electricity consumption. In our heat pump device model, the heat pump has

a number of different modulation levels. The modulation level determines the amount of electricity consumed by the pump and thus also the heat production of the pump. A heat pump with one modulation point corresponds to a type which does not support modulation.

Many heat pumps can in heating mode recover a large part of the drive energy and contribute this to the heat sink. Whether this is possible can be configured by connecting the loss either to the heat sink pool or to a loss consuming device. The loss is *not* modeled as part of the COP, because that does not properly represent the energy transfer between the heat source and the sink. This becomes particularly important when the heat source is fed from a finite or billed resource.

During simulation, 100 houses furnished with a heat pump and a heat store are modeled. Most houses are furnished with a 10 kWh heat store. The heat store to some extent decouples the production and consumption of heat, introducing flexibility regarding when and at which modulation level the heat pump operates. The start level of the heat store is chosen at 75% full and is the same for all houses.

The heat demand of these 100 houses is again from our heat demand database, where 100 days with a heat load between 38 kWh and 75 kWh are extracted (see Appendix A.2 for a detailed description). We consider this heat demand a representative heat demand for households with a heat pump installed.




Figure 6.7: Model of the house with a heat pump

The overall model of a house is depicted in Figure 6.7. Although there can be more devices present in the house, these are considered as non-steerable and have been abstracted away. Therefore, these are not included in the house model. The depicted heat exchange represents the heat extracted from the ground. For now, this source is assumed to be unlimited. The electricity exchange represents the grid connection of the house, which can provide all the electricity required by the heat pump. As described above, the heat production of the heat pump is determined by its COP and electricity consumption. Since we are using the heat pump for heating, both the heat output and the electrical loss energy-streams are connected to the heat store via a heat pool. The heat demand is subsequently supplied using the heat store.

In this use case, the heat pump has five modulation modes and a maximum electricity consumption of 2 kW. The modulation modes are divided evenly over the maximum electricity consumption, resulting in the following six heat production modes: 0 W, 400 W, 800 W, 1200 W, 1600 W and 2000 W. For COP, a value of 3.0 is

chosen. As we assume that the electrical loss can be fully recycled, an effective COP of 4.0 is attained, a value representative for current soil-water heat pump systems. Therefore, between 0 W and 8000 W of heat can be produced by the heat pump.

6.2.2 OPTIMAL SOLUTION

To determine the quality of the steering, the optimal solution is determined using the following ILP formulation. First the simulated time period of one day is evenly divided into T_N time intervals. For each time interval $t \in \{1, \dots, T_N\}$ the heat demand $C_{h,t}$ of house $h \in H$ must be supplied. In our simulation, the heat pump supports six modes, i.e. $M = \{0, \dots, 5\}$. The variable $z_{h,t} \in M$ is introduced to describe the mode of each heat pump at house h and time interval t . Based on $z_{h,t}$ the heat production can be calculated via $P_z \cdot z_{h,t}$, where P_z is the heat production capacity of mode z in one time interval. Similarly, the electricity demand (in W) is determined using E_z via $E_z \cdot z_{h,t}$.

The goal of this use case is to decrease the peaks by flattening the electricity demand profile of the group of houses. In other words, the fluctuation of the electricity demand should be minimized. This results in the following objective function:

$$\min \sum_{i=2}^T \left| \sum_{j=1}^H P_z z_{j,i} - \sum_{j=1}^H P_z z_{j,i-1} \right| + \sum_{i=1}^T \left| \sum_{j=1}^H P_z z_{j,i} - \hat{C} \right|,$$

where \hat{C} is the average heat consumption determined via $\frac{1}{T_N} \sum_{i=1}^T \sum_{j=1}^H C_{i,j}$.

Since the heat is supplied from the heat buffer, the buffer level must always be maintained between a lower limit b_{\min} and an upper limit b_{\max} . The heat buffer is depleted as a result of supplying the heat demand and can be filled by generating heat using the heat pump. Therefore, the following constraint is added:

$$b_{\min} \leq b_{\text{start}} + \sum_{i=1}^t P_z \cdot z_{h,i} - \sum_{i=1}^t C_{h,i} \leq b_{\max} \quad \forall t \in T, h \in H,$$

where b_{start} is the begin level of the heat store (in Wh).

In the optimization as well as the simulation, a time interval length of six minutes is used. The maximum electricity consumption of the heat pump is 2000 W. Since an effective COP value of 4.0 is used, a maximum of 8000 W of heat can be produced, which is $8000/5 = 1600$ W per modulation level. Each time interval is six minutes, therefore $P_z = \frac{1600}{60/6} = 160$ Wh and $E_z = 400$ W.

The performance of our approach is quantified using multiple metrics. The first metric is the *diversity factor*, which is the ratio of the sum of the individual maximum demands to the maximum real demand of the system. In our case, this is $\frac{2000 \cdot 100}{E_{\max}}$, where E_{\max} is the highest peak in the demand. The second metric is 3σ , where σ is the standard deviation of the electricity consumption, expressing the variation of the load. A lower variation means less fluctuations, meaning that the demand can be supplied more efficiently. Furthermore, load duration curves are used to visualize the capacity utilization. In the load duration curve the demand

data is ordered in descending order of magnitude, rather than chronologically. A load duration curve is very suitable to visualize the requirements and utilization of the network capacity.




Figure 6.8: Results ILP solution

The load duration curve after solving the ILP is given in Figure 6.8(b). The start level of the heat store $b_{\text{start}} = 7500 \text{ Wh}$, as in the simulations. As can be observed, the start level of the heat store results in start up effects. For every time interval, there is a heat buffer level which provides the flexibility required for the given objective. However, it takes a while to reach this heat buffer level. The limited heat demand restricts the possible operation modes, resulting in a deviation from the desired profile. After this startup phase, it is possible to achieve a perfect flat electricity consumption profile with a maximum electricity demand of $5.36 \cdot 10^4 \text{ W}$. The corresponding 3σ value is $6.14 \cdot 10^4$ and the diversity factor is 3.73.

6.2.3 EFFECTIVENESS TRIANA APPROACH

The ILP formulation of the previous section gives the optimal solution, exploiting all the information available about future heat demand, the buffer etc. The ILP solver uses the real heat demand data, implying a perfect heat forecast and knowledge of each house including information about the heatpump, heatstore etc. In TRIANA, the whole planning process is separated into multiple steps, each executed on different locations.

The analyze how well the TRIANA approach can reach a flat profile, again a perfect heat demand forecast is assumed. Only now the simulator is used, using the TRIANA approach to steer the electricity consumption of the heat pumps. For the heat pump, an initial planning program is created and via an iterative approach the steering vectors are adjusted, as described in Chapter 4. The objective remained the same, flattening the electricity consumption of a group of houses.




Figure 6.9: Results perfect forecasts case

The results of this simulations are depicted in Figure 6.9. As can be seen in the figure, the high morning peak is reduced significantly. A more flattened demand is obtained, as visible in Figure 6.9(a). Looking at the load duration curve in 6.9(b), a far more flattened duration curve can be achieved. Most importantly, the highest demand is reduced significantly, requiring less (distribution) network capacity. Although the results still show some fluctuations, the changes are relatively low.

The 3σ value went from $6.24 \cdot 10^4$ (no planning) to $4.24 \cdot 10^4$ (with planning). The diversity factor went from 1.98 to 2.67. The consumption profile improved significantly: the highest peak decreased by 26% and the fluctuation (variation) with 32%, even when using a rather naive and straightforward planning method. Exploration of the optimal solution showed that there is even more potential to decrease peaks and fluctuations. Studying the simulation results in more detail showed that the differences between the optimal solution and the simulation using the TRIANA methodology are mainly caused by the planning methodology. Improving this planning methodology by adding optimization on the lowest level is expected to enhance the results significantly. This is left for future work.

6.2.4 USING FORECASTED HEAT DEMAND VALUES

The previous section showed that the TRIANA methodology can exploit the scheduling freedom introduced by the heat stores. Although improvements can be made with improved planners, the results are promising. To analyze the effect of forecasting errors, simulations are performed in which the forecasted heat demand values, as described in Chapter 3, are used. Using the simulated annealing approach, the best forecasting method is determined. The heat forecasts are made on an hourly time base and are evenly divided into ten time intervals, as determined by

the simulation time interval length.




Figure 6.10: Results forecasts case

In Figure 6.10 the results of using the forecasted heat demand data while planning is depicted. As can be seen in the figure, still a more flattened profile can be achieved. However, especially from 12h up to the end of the day more fluctuation can be seen. The corresponding 3σ value is $4.14 \cdot 10^4$ and the diversity factor 2.54.

Replanning

The developed planning program uses the complete scheduling freedom of the heat store, resulting in quick deviations in case a forecasting error occurs. Therefore, very likely a replanning is required. Normally, during replanning new improved forecasts are made. Since this is not (yet) implemented in the simulator, this is emulated by using forecasting function f . This function combines the 24 hour ahead forecasted heat demand and the real demand that is used for the simulation. Very short term forecasts often are of a higher quality than forecasts on the longer term, which should be reflected by f . Therefore, the closer the forecasted value is compared to the current time interval, the better the forecasts becomes. Based on this requirement, f is defined as:

$$f(u, t) = P_{t+u} \times \frac{u}{T_N - t} + R_{t+u} \times \left(1 - \frac{u}{T_N - t}\right),$$

where t is the current time interval, u the forecasted time interval compared to t , R_i the real heat demand for time interval i , P_i the forecasted heat demand for time interval i and T_N the total number of time intervals.

During simulation, dependent on the deviation and the deviation threshold (maximum allowed deviation in %), a replanning session is initiated. Multiple deviation thresholds have been simulated: 10%, 15%, 20% and 25%.

Table 6.2: Results all cases

Case	Replanning threshold	3σ	Diversity factor	Replanning sessions
ILP	-	$6.14 \cdot 10^4$	3.73	
Real	-	$6.24 \cdot 10^4$	1.98	-
Real + Planning	-	$4.24 \cdot 10^4$	2.67	-
Forecasted + planning	-	$4.14 \cdot 10^4$	2.54	-
Forecasted + planning	10%	$5.56 \cdot 10^4$	1.86	114
Forecasted + planning	15%	$5.88 \cdot 10^4$	1.80	81
Forecasted + planning	20%	$4.12 \cdot 10^4$	2.66	10
Forecasted + planning	25%	$4.09 \cdot 10^4$	2.60	1

The results of these simulations are depicted in Table 6.2. As mentioned above, the currently used planning program exploits all the scheduling freedom without any room to cope with forecasting errors. As a result, deviation thresholds of 15% and lower resulted in a high number of replanning sessions, up to almost half of the time intervals. This clearly is an unusable solution.

When large forecasting errors occur, strictly following the planning using the currently used planner causes a lot of deviations, and thus replanning. The replanning threshold should somehow meet the forecasting quality. Furthermore, the planning program should reserve some capacity to deal with forecasting errors locally, reducing the deviation from the planning.

In this case, proper replanning threshold values are around 20%. Using this threshold value, a low variation and still a high diversity factor can be achieved. Furthermore, ten replanning sessions is still acceptable.

6.2.5 CONCLUSIONS

The ILP solution shows the available potential of adding a heat store to the system to steer the consumption pattern of a group of houses with heat pumps. After a startup effect, as a result of solving only a single day, a completely flat electricity profile can be achieved. Although the results in Section 6.2.3 do not show a completely flat profile, still a big improvement in the electricity profile is obtained. The peaks are decreased by 26% and the fluctuation with 32%, even when using a rather naive and straightforward planning method.

When using forecasted heat demand data in the planning, more realistic results are obtained. Although the results differ from the perfect forecasting case, still a reduction of 25% in the peak demand can be obtained. Furthermore, a reduction of 34% on the fluctuations is achieved.

When using this planning method with forecasted heat demand data, the forecasting errors cause deviations from the planning. Dependent on the replanning threshold, these deviations will result in a replanning session. The replanning threshold should match with the forecasting quality and the flexibility of the plan-

ning determined by the planning program. It is advised that the planning program should reserve some capacity to deal with forecasting errors locally, reducing the deviation from the planning.

6.3 MULTIPLE DEVICES USE CASE

In the previous use cases the predicted heat demand data was used to steer the electricity profile of a house. In the micro-CHP case, the electricity production of the group of houses was steered, while in the heat pump use case the consumption pattern of the fleet was steered. The goal of this use case is to analyze the effects of two different types of devices, one generating electricity and the other one consuming electricity. Since the simulated micro-CHP can only produce 1 kW of electricity, while the heat pump consumes maximal 2 kW, the mix of the houses should balance the production capacity with the consumption capacity. Therefore, a neighborhood with 50 houses equipped with a micro-CHP appliance and 25 houses equipped with a heat pump is simulated. Each house is equipped with a 10 kWh heat store and the heat demand of each household is again extracted from our heat demand database. All the houses have a heat demand between the 30 kWh and 40 kWh. During simulation, a time interval length of six minutes is used and the simulation simulates 24 hours.

The objective of the global planning is to flatten the overall consumption profile, just as in the heat pump use case. During planning, a perfect prediction is assumed, since the goal of the use case is to analyze the effect of using a single price vector for different devices.

6.3.1 RESULTS




Figure 6.11: Electricity profiles multiple devices

The electricity profiles of the heat pump and micro-CHP appliance are depicted in Figure 6.11. Using this figure, the effects of the (shared) steering vector on each individual device can be seen. As can be seen on the figure, the planner tries to reduce the morning peak by lowering the peak caused by the heat pump and increasing the production of the micro-CHP appliances during that period. Furthermore, the micro-CHP profile is more flattened in the afternoon. The heat pump is less sensitive to the steering vector and does not change that much.




Figure 6.12: Results multiple devices case

The electricity profile and load duration curve of combined energy flow are depicted in Figure 6.12. The load duration curves shows a more flattened line, as desired by the overall objective. Furthermore, Figure 6.12(a) nicely shows the effect of the micro-CHP appliance smoothing the overall profile.

6.3.2 CONCLUSIONS

In this use case, two different appliances are steering with the same vector. Although the devices have different characteristics, they both respond to the shared steering vector. However, changes in the steering vector have bigger effects on the micro-CHP appliance than on the heat pump. This is caused by the definitions of the cost functions of both devices. The cost function of the heat pump is less dependent and therefore less sensitive to price changes. Therefore, the costs functions should be adjusted to make them evenly sensitive to the steering vector. Another possibility is to develop another planner, and adjust the costs functions to follow the determined planning as good as possible, while maintaining the same level of comfort for the residents and still properly controlling the devices. These adjustments are left for future work.

CONCLUSIONS

Traditionally, most western countries supply domestic electricity demand through generation in large central power stations, with subsequent transmission and distribution through networks. Although this scheme was designed decades ago, due to continuous monitoring and management of the physical electricity flow and maintenance of the network, the current energy grid has been working very stable and reliably. However, the increasing energy prices, environmental concerns and the continuously increasing demand for energy require a different and more sustainable electricity supply chain. Renewable distribution generation is a solution to reach a more sustainable electricity production. But, due to their nature, this generation is mostly uncontrollable. Fluctuations and imbalance caused by renewable generation must be compensated elsewhere in the grid, which currently is handled via even less efficient power plants. New technology like smart appliances and demand side load management introduces the possibility to shift from a demand driven supply chain into a more active, cooperative approach. In the smart grid producers, consumers and transmission networks are continuously cooperating to ensure a properly functioning grid. The emergence of smartening the grid, as described in Chapter 2, and updating the electricity supply chain is emphasized by the numerous initiatives worldwide, from the European Union, from governments, from industry as well as from the academic world. However, to reach a smarter grid, a number of technical, economical, legislative and ethical challenges have to be addressed. To tackle the technical challenges, Information and Communication Technology (ICT) is seen as one of the key enabling technologies.

A proposed control strategy for smart grids is TRIANA. The goal of this control strategy is to manage the energy profiles of individual devices in buildings to support the transition towards an energy supply chain which can provide all the required energy in a sustainable way. Since there are a lot of different (future) domestic technologies and building configurations, the control methodology should be able

to work generically and independently of these configurations. Furthermore, the methodology should be flexible such that new technologies can be added in the future. Consequently, the control methodology needs to be very flexible and generic. Based on the above considerations, TRIANA split up into a local and a global part: 1) local offline forecasting, 2) global offline planning and 3) local online control. Because of scalability reasons, the global planning has a hierarchical structure and can aggregate data and plannings on different levels.

In Chapter 3 the first step, the forecasting step, is described. Via forecasts, the flexibility of each device, called the scheduling freedom, is determined. This forecasting is performed up to one day ahead, to allow the exploitation of this scheduling freedom on the day ahead market. To determine the scheduling freedom of individual devices, forecasts are made for each individual device. In order to incorporate device specific information in a forecasting system, without the necessity of communicating all required information of each device to a central location, forecasting is performed by the local controller in each building. Using such an approach, the requirement of a scalable (forecasting) system can be met. The local controller can use locally harvested data and can be programmed to forecast device specific information, resulting in a flexible system. To not burden the residents with questions about their (expected) behavior, the forecasting system should be running completely autonomously. In the use case of individual heat demand forecasts to determine the scheduling freedom of a micro-Combined Heat and Power (CHP) appliance, neural network techniques are used. The possibility of autonomous learning of (non-linear) relations between the input- and outputdata makes neural network techniques a good candidate. By using different neural networks for each weekday, differences in behavior during the week can be incorporated. By adjusting the neural network structure, forecasts on different timescales can be generated. Furthermore, periodically evaluating the forecasting quality and, when necessary, adjusting the input for the neural network, results in a system adaptable to change. Via the Simulated Annealing searching algorithm, an automated search for a proper input set and parameters for the neural network can be determined. This yields to a good forecasting quality.

The forecasts are used in the second step, the planning. Based on these forecasts, and the desired objective, a target profile for the group of buildings is determined. In Chapter 4, an iterative approach is given to achieve this desired profile. The desired profile is subsequently divided along the hierarchical structure of grid controllers. Each grid controller is then responsible for reaching their desired profile. Due to this subsequent division of the large optimization problems into subproblems via a hierarchical structure, a fast scalable system is achieved. By choosing a proper structure, communication requirements can be kept low.

Based on the forecasts determined at the first step, and the steering signals received from a grid controller (generated in the second step), the local controllers adjust the runtimes and/or operation modes of the devices controlled in the third step. Since both the forecasts and planning are performed at each building, all the required information for planning a (group of) devices is locally available. The local controller can use this information, and use device specific constraints to generate a

planning for the device. In this approach, the available distributed computational power available in the grid is exploited. Since less information about the device and its environment needs to be sent outside the building, less privacy sensitive data is exposed. Regarding the determination of the steering signals, addressing each building individually by using different steering signals gives the best results.

Since the planned schedules are based on forecasts, forecasting errors can lead to deviations of the planning. By detecting these deviations and performing a replanning, based on the actual situation and improved short term forecasts, deviations can be minimized. Replanning can be performed on different levels within the grid, where a deviation of one building might be compensated elsewhere within the neighborhood. If the deviation cannot be handled locally, the controller can signal other controllers higher in the hierarchical structure.

Via various use cases, the effectiveness of TRIANA has been tested. The freezer use case of Chapter 4 and the heat pump use case in Chapter 6.2 show that it is possible to flatten the consumption profile of a large fleet of devices. In the freezer case, the fleet of devices consisted of consuming devices, while in the heat pump case converting devices are used. Although the constraints of the devices were different, TRIANA was able to achieve a more flattened demand profile for both use cases. The constraints set by the devices limited the effectiveness of the steering.

The multiple devices use case showed that when steering different kinds of devices, the sensitivity to the steering signals of both local planner should be in balance. Although the devices have different characteristics, they both respond to the shared steering signal. However, changes in the steering signal have bigger effects on the micro-CHP appliance than on the heat pump. This is caused by the definitions of the cost functions of both devices. The cost function of the heat pump is less dependent and therefore less sensitive to price changes. Therefore, the costs functions should be developed such that they are evenly sensitive to the steering signal. Another possibility is to create a planner which generates a planning, respecting the characteristics and limitations of the device. The costs functions can then be extended by adding additional costs for deviating from the generated planning.

The Virtual Power Plant (VPP) use case shows that TRIANA is flexible enough to be used for different kind of objectives. Where in the above mentioned use case a flattened profile was desired, the VPP required a different, varying profile. The desired profile could not be reached perfectly, but this was caused by forecast errors and therefore limitations on feasibility of the desired profile. This can be solved by generating a more conservative planning and introducing some reserve capacity during planning to cope with forecast errors. Furthermore, after detecting the forecasting errors, TRIANA offers opportunities via replanning to improved short term forecasts to adapt to the actual situation.

To analyze the possibilities of TRIANA, a generic energy model and corresponding simulation software has been developed. The energy model and simulator provide useful tools to analyze current and future smart grid technology and their impact on the grid. The model's flexibility enables the possibility to simulate the whole energy supply chain and even future (smart) devices. The simulator is able to

simulate a realistic mix of devices by providing a wide variety of available devices, with possibilities to add stochastic variation to the characteristics of the devices and loads. By distributing the simulation over multiple machines, a large, complex simulation can still be performed within reasonable time.

7.1 CONCLUSIONS

Based on the results of the previous chapters, the research questions introduced in the first chapter can be answered:

- *What is the optimization potential of devices located in buildings/houses?*
Via the presented use cases it is shown that there are quite some devices present in the building that have flexibility on how to use them. Devices with little direct interaction with the residents, like freezers, fridges, boilers, heat pumps can easily adjust their consumption profile. Simulation shows that the consumption profile of a large group of devices can be reshaped significantly, without any loss of comfort for the residents.
- *How can this optimization potential be exploited?*
To exploit the potential of domestic devices, the current (inflexible) consuming devices should be replaced by smart devices, capable of cooperating with other devices present in the grid. Using different controllers in the grid, each responsible for a part of the whole system, the devices can be controlled to reach different kind of objectives. Using TRIANA, objectives like peak shaving or achieving an predefined profile can be reached.
- *How can a control system autonomously determine the optimization potential of devices?*
The optimization potential of devices is dependent on the class of devices, and the amount of interaction with the devices. Based on these criteria, different influence factors determine the optimization potential of a device. Via forecasts of the most important influence factors of a device, the flexibility and thus optimization potential of a device can be determined. Controllers present in the buildings have knowledge about the devices in the building and harvest information about the residents. This information can be used by controllers to generate the required forecasts. Chapter 3 of this thesis describes the forecasting problem, and via a simulated annealing process a proper forecasting scheme and the resulting optimization potential for each device can be determined autonomously.
- *What is a proper control system and methodology to utilize the optimization potential, taking the size and timing constraints of the system?*
In Chapter 4 the hierarchical approach of TRIANA is described. By performing the forecasting, planning and control steps in a hierarchical, distributed way, a scalable solution is achieved. This also allows optimization and coordination on different levels in the grids, each with different objectives. Due to the

used mathematical analysis and optimization techniques, the behavior of the devices in the grid after steering becomes more reliable and predictable. The combination of high quality forecasts, exploiting the optimization potential via planning while reserving some capacity to cope with forecast errors and real-time control is a proper control strategy for smart grids.

Using TRIANA, the behavior of distributed generation, storage and consumption technology can be adjusted to reshape the overall energy profile. In this thesis, TRIANA has been analyzed and used to achieve global objectives and planning strategies. Molderink [62] has analyzed the real-time control system and possibilities to reach more locally focussed objectives in more depth. For both cases, it is shown that TRIANA is capable of working towards the set objectives.

Summarizing, optimizing the behavior of distributed generation, storage and consumption technologies has the potential to increase the efficiency of conventional power plants and to facilitate the introduction of large scale renewable generation. A large scale introduction of new technologies for production, consumption and storage allows maintaining grid stability and ensures a reliable and affordable supply. TRIANA is able to optimize the behavior of domestic devices to work towards local and global objectives in a predictable way. We believe that TRIANA, with the hierarchical tree structure of control nodes, is a good solution for a scalable, generic and efficient Smart Grid control strategy.

7.2 RECOMMENDATIONS FOR FUTURE WORK

In this work, initial results and the proof of concept of TRIANA are presented. Although the results are promising, still improvements are possible.

In the first step, only the long term forecasts (one day in advance) are researched. The possibilities to perform short term forecasts, based in recent events, needs to be researched. Using these short term forecasts, problems in reaching the desired objective can be detected and dealt with earlier. Therefore, the results are expected to improve. Furthermore, other forecasts than heat demand forecasts should be investigated as well.

In the second step, on multiple level improvements can be achieved. Algorithms deciding when and on which level to perform a planning and replanning are still desired. Furthermore, algorithms to determine how many iterations of a planning session are required can decrease the time required to execute a complete (re)planning session and reduce the amount of communication required. A suitable communication protocol can decrease the communication requirements even further. On the lowest level in the hierarchical structure, more efficient and improved device planners can result in improved results in reaching the desired objective and decrease the amount of time required for planning.

After improving each individual step of TRIANA, an overall system analysis with an overview of which optimization technique, hierarchical structure, planning algorithm etc. are suitable for which objective further mature the control strategy.

Besides the technical improvements of TRIANA, other important issues like Smart Grid economics and legislative restrictions are also very important field. Although a Smart Grid may be technical feasible, economical feasibility is critical to ensure society will adopt the Smart Grid. Our energy flow model can be extended to include financial flows, like investment costs, energy pricing, maintenance etc. Adding such an extension to the model and the simulator will allow the simultaneous analysis of technical and economical feasibility of Smart Grid solutions. If we can show that Smart Grids are technically and economically feasible, with all the benefits for the environment, a more sustainable energy supply chain and society can be within hand's reach.

ACRONYMS

AMI	Advanced Metering Infrastructure.
CCS	Carbon Capture and Storage.
CHP	Combined Heat and Power.
COP	Coefficient Of Performance.
DC	Direct Current.
DG	Distributed Generation.
DS	Distributed Storage.
DSM	Demand Side (Load) Management.
DSO	Distribution System Operator.
GUI	Graphical User Interface.
HVAC	Heating, Ventilating, and Air Conditioning.
HVDC	High Voltage Direct Current.
IC	Integrated Communications.
ICT	Information and Communication Technology.
IEA	International Energy Agency.
IEEE	Institute of Electrical and Electronics Engineers.
ILP	Integer Linear Programming.
MAPE	Mean Absolute Percentage Error.
METAR	Meteorological Aerodrome Reports.
MIP	Mixed Integer Problem.
MPC	Model Predictive Control.
MPE	Mean Percentage Error.
NIST	National Institute for Standardization and Technology.
PV	Photovoltaics.
RAM	Random Access Memory.

SA	Simulated Annealing.
SCADA	Supervisory Control And Data Acquisition.
SoC	State of Charge.
TOU	Time Of Use.
TSO	Transmission System Operator.
UPS	Uninterruptible Power Supply.
VPP	Virtual Power Plant.

DETAILS USE CASES

A.1 FREEZER USE CASE CHAPTER 4

In this use case, a large fleet of buildings with a smart freezer are used to analyze the possibilities of the distributed planning algorithm. Different group sizes and planning topologies are used, but in all instances, the houses contains the same (pseudo)random freezer.

Each freezer present in the house must maintain its internal temperature within a certain bound, which in this case is between -18°C and -28°C . All freezers have the three different mode:

Off In the off state, the cooling element is switched off. Due to losses, the temperature increases slowly each time interval. The power consumption during the off state is 2 Watt.

Cooling mode In this state, the cooling element is switched on, causing a decrease of the internal temperature. Using this mode, a minimal internal temperature of -23°C can be reach. The power consumption during this state is 100 Watt.

Extra cooling mode Once the internal temperature is -23°C or lower, the freezer required more power to decrease the internal temperature even further. Using this mode, the freezer can decrease the temperature to a minimum of -28°C at the expense of consuming more power. This mode can be useful in case the electricity price is expected to be higher in the future. In this state, the power consumption is 140 Watt.

In this use case, 50 or 200 freezers are simulated, each with different characteristics. First, the internal temperature at the first time interval is pseudo-randomly

determined using the π distribution. In this distribution, the decimals of π are used. The start temperatures of the freezers are then determined via

$$T_0 = -23 + 0.5 \times \pi_i,$$

where π_i is the i^{th} decimal of π and i is increased each time a value is drawn from the π distribution. Note that $\pi_i \in [0, 9]$ and that the corresponding start temperatures thus are between $[-22.5, 18.5]$.

The cold loss of the freezer is determined by the insulation of the freezer, which is different for each freezer. This loss per time interval, defined in $^{\circ}\text{C}$, is determined via

$$0.1 + \frac{\pi_i \bmod 2}{10}.$$

The cooling capacity determines the efficiency of the cooling element and described the temperature decrease per time interval in case the freezer is cooling. The cooling capacity, again in $^{\circ}\text{C}$ per time interval, is determined via

$$0.3 + \frac{\pi_i \bmod 5}{10}.$$

A.1.1 PSEUDORANDOM FORECASTING ERRORS

In the replanning use case, pseudo random timestamps when an interaction with the freezer occurs is simulation. The result of this interaction is a increase of the internal temperature of the freezer, and this increase is also pseudorandomly picked.

A pseudorandom timestamp is determined by selecting a time interval in which this interaction occurs using

$$\frac{30 + (\pi_i \bmod 7) \times 10}{100} \times T_N,$$

where T_N is the total number of time intervals. The interaction thus occurs somewhere in in time interval 72 and 216, which is between 7:12h and 21:36h. The temperate increase (in $^{\circ}\text{C}$) during this interval is determined via $2 \times \frac{\pi_i}{10}$, resulting in a temperate increase between $[0.0, 1.8]^{\circ}\text{C}$.

A.2 HEAT PUMP USE CASE

In this use case, 100 houses equipped with a heat pump and 10 kWh heat store are simulated. Each heat store is initially 75% filled.

The heat pumps have five modulation modes and a maximum electricity consumption of 2 kW. The modulation modes are divided evenly over the maximum electricity consumption, resulting in the following six heat production modes: 0 W, 400 W, 800 W, 1200 W, 1600 W and 2000 W. For Coefficient Of Performance (COP), a value of 3.0 is chosen. As we assume that the electrical loss can be fully recycled, an effective COP of 4.0 is attained, a value representative for current soil-water heat

pump systems. Therefore, between 0 W and 8000 W of heat can be produced by the heat pump.

The real heat demand (D_r) and predicted heat demand (D_p) data is given in the table below. For each house number, the total real and predicted heat demand (in kWh) is given. To give some insight how this heat demand is distributed over time and how well the predictions are, small demand profiles are given. The black lines are the real heat demand profiles, the grey ones the forecasted profiles. To save space, no ticks and labels have been drawn in these graphs. In the graphs, for each time interval (horizontal axis) the heat demand for that time interval (vertical axis) is depicted. The minimum value on the vertical axis is 0, the maximum value is 16370. The unit of the heat demand is $kW\tau$, where τ is the interval length, which in this simulation is six minutes.

H	P	D_r	D_p	H	P	D_r	D_p
1		70.2	73.1	51		43.7	48.1
2		42.6	42.6	52		45.8	60.8
3		52.9	53.4	53		62.4	65.7
4		71.6	72.0	54		48.6	46.4
5		44.8	46.8	55		48.6	62.1
6		70.6	70.4	56		53.0	61.8
7		47.8	52.2	57		64.0	74.6
8		35.9	36.5	58		56.6	64.6
9		42.1	48.6	59		35.4	52.2
10		39.5	41.3	60		60.8	62.1
11		52.8	54.1	61		37.4	37.2
12		63.2	62.6	62		41.5	49.1
13		48.6	46.4	63		40.7	58.6
14		64.4	70.7	64		38.4	52.7
15		37.0	39.3	65		56.1	61.7
16		58.8	54.2	66		41.0	54.0
17		36.0	39.2	67		48.4	71.7
18		58.6	66.3	68		36.9	45.4
19		59.5	68.4	69		35.1	39.6
20		53.3	59.0	70		69.5	73.8
21		64.8	70.2	71		40.7	62.5

Continued on next page

H	P	D _r	D _p	H	P	D _r	D _p
22		39.0	39.2	72		49.5	59.7
23		43.2	44.4	73		45.1	49.8
24		61.6	62.4	74		63.8	73.9
25		56.6	61.2	75		41.0	54.3
26		67.3	62.8	76		37.0	51.0
27		51.6	53.9	77		55.5	65.5
28		47.7	46.9	78		43.5	56.3
29		58.7	56.7	79		58.1	70.1
30		69.1	71.3	80		71.2	65.9
31		65.7	67.1	81		38.4	47.3
32		43.7	41.3	82		51.8	61.2
33		38.8	45.5	83		41.3	55.5
34		45.9	46.7	84		52.1	58.4
35		59.2	72.1	85		48.7	64.3
36		64.3	61.5	86		42.6	71.1
37		51.4	55.0	87		73.9	62.3
38		72.2	67.0	88		47.8	48.4
39		45.8	63.4	89		63.8	71.5
40		37.6	43.2	90		52.7	57.3
41		62.3	65.4	91		53.2	49.0
42		44.7	44.3	92		59.2	71.9
43		48.6	57.4	93		55.3	69.0
44		40.7	44.3	94		37.5	36.6
45		44.8	41.6	95		38.9	56.4
46		62.7	67.7	96		61.0	50.5
47		65.0	68.8	97		38.8	64.5
48		62.2	74.6	98		53.8	67.8
49		51.2	50.6	99		45.9	52.4
50		36.2	42.2	100		65.3	61.7

LIST OF PUBLICATIONS

REFEREED

- [VB:1] V. Bakker, A. Molderink, J. L. Hurink, and G. J. M. Smit. Domestic heat demand prediction using neural networks. In *Proceedings of Nineteenth International Conference on Systems Engineering, Las Vegas, Nevada, USA*, pages 189–194, Los Alamitos, August 2008. IEEE Computer Society Press.
- [VB: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*, Oxford, March 2010. Elsevier Ltd.
- [VB:3] 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.
- [VB:4] 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.
- [VB:5] 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.
- [VB:6] 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 1985-9406.
- [VB:7] 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.
- [VB:8] 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.
- [VB:9] 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.
- [VB:10] 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, July 2011. ISSN 0377-2217.
- [VB:11] A. Molderink, V. Bakker, J. L. Hurink, and G. J. M. Smit. Algorithms for balancing demand-side load and micro-generation in islanded operation. In *Proceedings of the nineteenth international conference on systems engineering, Las Vegas*, pages 115–120, Los Alamitos, August 2008. IEEE Computer Society Press.
- [VB:12] 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.
- [VB:13] 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.
- [VB:14] 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 Press.
- [VB:15] 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, USA*, December 2010. IEEE Power & Energy Society.
- [VB:16] 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.
- [VB:17] 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.
- [VB:18] 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 transaction on Smart Grid*, 1(2):109–119, September 2010. ISSN 1949-3053.

- [VB:19] H. A. Toersche, V. Bakker, A. Molderink, S. Nykamp, J. L. Hurink, and G. J. M. Smit. Controlling the heating mode of heat pumps with the TRIANA three step methodology. In *Proceedings of the third IEEE PES Conference on Innovative Smart Grid Technologies*, USA, January 2012 (accepted). IEEE Power & Energy Society.
- [VB:20] W. A. Wiggers, V. Bakker, A. B. J. Kokkeler, and G. J. M. Smit. Implementing the conjugate gradient algorithm on multi-core systems. In J. Nurmi, J. Takala, and O. Vainio, editors, *Proceedings of the International Symposium on System-on-Chip (SoC 2007)*, Tampere, number 07ex1846, pages 11–14, Piscataway, NJ, November 2007. IEEE. ISBN 1-4244-1367-2.

NON-REFEREED

- [VB:21] V. Bakker, A. Molderink, J. L. Hurink, and G. J. M. Smit. Using heat demand prediction to optimise virtual power plant production capacity. In *Proceedings of the Nineteenth Annual Workshop on Circuits, Systems and Signal Processing (ProRISC)*, Veldhoven, pages 11–15, Utrecht, November 2008. Technology Foundation STW.
- [VB:22] 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.
- [VB:23] 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.
- [VB:24] 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.
- [VB:25] 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.
- [VB:26] A. Molderink, V. Bakker, J. L. Hurink, and G. J. M. Smit. Islanded house operation using a micro chp. In *18th Annual Workshop on Circuits*, Veldhoven, The Netherlands, Proceedings of the STW programs 2007, pages 324–330, Utrecht, November 2007. Stichting PATO.
- [VB:27] 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.

- [VB:28] A. Molderink, V. Bakker, J. L. Hurink, G. J. M. Smit, and A. B. J. Kokkeler. Domestic electricity usage regulation using μ CHP appliances. *Power Systems Design Europe*, 5 (1):41–44, January 2008. ISSN 1613-6365.
- [VB:29] 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.
- [VB:30] 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. STW Technology Foundation.
- [VB:31] 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.

BIBLIOGRAPHY

- [1] S. Abu-sharkh, R.J. Arnold, J. Kohler, R. Li, T. Markvart, J.N. Ross, K. Steemers, P. Wilson, and R. Yao. Can microgrids make a major contribution to UK energy supply? *Renewable and Sustainable Energy Reviews*, 10(2):78–127, Sept 2004.
- [2] European Environment Agency. Household energy consumption, jan 2001. URL http://www.eea.europa.eu/data-and-maps/indicators/ds_resolveuid/0afd9fec59ea51fdabf2c531d497bbda.
- [3] International Energy Agency. Smart grids roadmap. Technical report, International Energy Agency, 2010.
- [4] Hesham K. Alfares and Mohammad Nazeeruddin. Electric load forecasting: Literature survey and classification of methods. *International Journal of Systems Science*, 33(1):23–34, 2002. doi: 10.1080/00207720110067421.
- [5] Godfried Augenbroe and Jan Hensen. Simulation for better building design. *Building and Environment*, 39(8):875 – 877, 2004. ISSN 0360-1323. doi: DOI: 10.1016/j.buildenv.2004.04.001. URL <http://www.sciencedirect.com/science/article/B6V23-4C9YYD9-1/2/d19011905a2587a4c178b259fe842610>. Building Simulation for Better Building Design.
- [6] V. Bakker, A. Molderink, J. L. Hurink, and G. J. M. Smit. Domestic heat demand prediction using neural networks. In *Proceedings of Nineteenth International Conference on Systems Engineering, Las Vegas, Nevada, USA*, pages 189–194, Los Alamitos, August 2008. IEEE Computer Society Press.
- [7] 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*, Oxford, March 2010. Elsevier Ltd.
- [8] A. Battaglini, J. Lilliestam, C. Bals, and A. Haas. The supersmart grid. Technical report, European Climate Forum, 2008.
- [9] Alberto Bemporad. Model predictive control design: New trends and tools. In *Proceedings of the 45th IEEE Conference on Decision & Control*, 2006.

- [10] Abdullatif E. Ben-Nakhi and Mohamed A. Mahmoud. Cooling load prediction for buildings using general regression neural networks. *Energy Conversion and Management*, 45(13-14):2127 – 2141, 2004. ISSN 0196-8904. doi: DOI: 10.1016/j.enconman.2003.10.009. URL <http://www.sciencedirect.com/science/article/B6V2P-4B5B98S-1/2/404650e5f75519184759b9ec47e5992b>.
- [11] S. Bertolini, M. Giacomini, S. Grillo, S. Massucco, and F. Silvestro. Coordinated micro-generation and load management for energy saving policies. In *Proceedings of the first IEEE Innovative Smart Grid Technologies Europe Conference*, 2010.
- [12] F. Bliiek, A. van den Noort, B. Roossien, R. Kamphuis, J. de Wit, J. van der Velde, and M. Eijgelaar. Powermatching city, a living lab smart grid demonstration. In *Innovative Smart Grid Technologies Conference Europe (ISGT Europe), 2010 IEEE PES*, 2010.
- [13] C. Block, D. Neumann, and C. Weinhardt. A market mechanism for energy allocation in micro-CHP grids. In *41st Hawaii International Conference on System Sciences*, pages 172–180, Jan 2008.
- [14] 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.
- [15] 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 CHPappliances. In *19th Annual Workshop on Circuits, Veldhoven, The Netherlands*, Proceedings of the STW programs 2008, Utrecht, November 2008.
- [16] 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.
- [17] 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*, Australia, February 2010. PCO.
- [18] 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.
- [19] D.W. Bunn and E.D. Farmer. *Comparative models for electrical load forecasting*. J. Wiley & Sons, Inc., 1985.

- [20] R. Caldon, A.R. Patria, and R. Turri. Optimisation algorithm for a virtual power plant operation. In *Universities Power Engineering Conference, 2004. UPEC 2004. 39th International*, volume 3, pages 1058–1062 vol. 2, Sept. 2004.
- [21] A. Chadwick, S. Holloway, and N. Riley. Deep subsurface CO₂ sequestration - a viable greenhouse mitigation strategy. *Geoscientist*, 11:4–5, 2001.
- [22] D.P. Chassin and L. Kiesling. Decentralized coordination through digital technology, dynamic pricing, and customer-driven control: The gridwise testbed demonstration project. *The Electricity Journal*, 21:51–59, 2008.
- [23] European Commission. Towards smart power networks. Technical report, European Commission, 2005.
- [24] 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. doi: 10.1109/PCT.2007.4538359.
- [25] K. H. van Dam, M. Houwing, Z. Lukszo, and I. Bouwmans. Agent-based control of distributed electricity agent-based control of distributed electricity generation with micro combined heat and power - cross-sectoral learning for process and infrastructure engineers. *Computer Aided Chemical Engineering*, 2007.
- [26] Arjen de Jong, Ernst Jan Bakker, Jan Dam, and Hans van Wolveren. Technisch energie- en CO₂-besparingspotentieel in Nederland (2010-2030) CO₂-besparingspotentieel van micro-WKK in Nederland (2010-2030). Technical report, Werkgroep Decentraal, July 2006.
- [27] A.L. Dimeas and N.D. Hatziargyriou. Agent based control of virtual power plants. In *Intelligent Systems Applications to Power Systems, 2007. ISAP 2007. International Conference on*, pages 1–6, Nov. 2007. doi: 10.1109/ISAP.2007.4441671.
- [28] K. Dodrill. Understanding the benefits of the smart grid. Technical report, U.S. Department of Energy, 2010.
- [29] E. Dotzauer. Simple model for prediction of loads in district-heating systems. *Applied Energy*, 73(3-4):277–284, NOV-DEC 2002. doi: PII:0306-2619(02)00078-8.
- [30] AG Energiebilanzen e.V. Energieverbrauch in Deutschland, Daten für das 1. Quartal 2011, 2009. URL <http://www.ag-energiebilanzen.de/viewpage.php?idpage=64>.
- [31] Excel Energy. Smartgridcity: Design plan for boulder, colo. Technical report, Excel Energy, 2008.
- [32] M. Erol-Kantarci and H. T. Mouftah. Tou-aware energy management and wireless sensor networks for reducing peak load in smart grids. In *Proceedings of the IEEE Vehicular Technology Conference Fall*, 2010.

- [33] Eurostat. Energy mix fact sheet 2004. Technical report, European Committee, 2007. URL http://ec.europa.eu/energy/energy_policy/doc/factsheets/mix/mix_nl_en.pdf.
- [34] Desertec Foundation. Clean power from deserts, whitebook 4th edition. Technical report, Desertec Foundation, 2008.
- [35] P. Fraser. Distributed generation in liberalised electricity markets. Technical report, International Energy Agency, 2002.
- [36] E.R. Furling, M. Piemontesi, P. Prasad, and D. Sukumar. Advances in energy storage techniques for critical power systems. In *The Battcon 2002 proceedings*, 2002.
- [37] M. Geidl and G. Andersson. A modeling and optimization approach for multiple energy carrier power flow. In *Power Tech, 2005 IEEE Russia*, pages 1–7, june 2005. doi: 10.1109/PTC.2005.4524640.
- [38] Martin Geidl, Gaudenz Koeppel, Patrick Favre-Perrod, Bernd Klockl, Goran Andersson, and Klaus Frohlich. Energy hubs for the future. *Power and Energy Magazine, IEEE*, 5(1):24–30, 2007.
- [39] D.J. Hammerstrom, R. Ambrosio, T.A. Carlon, J.G. DeSteese, G.R. Horst, and R. Kajfasz. Pacific Northwest GridWise testbed demonstration projects, part I and II. Pacific Northwest National Laboratory, July 2007.
- [40] E. Handschin and F. Uphaus. Simulation system for the coordination of decentralized energy conversion plants on basis of a distributed data base system. In *Power Tech, 2005 IEEE Russia*, pages 1–6, June 2005. doi: 10.1109/PTC.2005.4524428.
- [41] A. J. Heller. Heat-load modelling for large systems. *Applied Energy*, 72(1): 371 – 387, 2002. ISSN 0306-2619. doi: DOI:10.1016/S0306-2619(02)00020-X. URL <http://www.sciencedirect.com/science/article/B6V1T-45HWR3W-1/2/59380b8b267b16c422c686c5f64a7315>.
- [42] P. Hines, S. Hamilton, R. Yinger, C. Varanian, A. Feliachi, and K. Schoder. Integrated, agent-based, real-time control systems for transmission and distribution networks. Technical report, Grid-Interop Forum 2007, 2007.
- [43] M.P.F. Hommelberg, B. Roossien, C.J. Warmer, J.K. Kok, F.J. Kuijper, and J.W. Turkstra. Aggregatie van micro-wkk's in een virtuele centrale, July 2007.
- [44] M.P.F. Hommelberg, C.J. Warmer, I.G. Kamphuis, J.K. Kok, and G.J. Schaeffer. Distributed control concepts using multi-agent technology and automatic markets: An indispensable feature of smart power grids. In *IEEE Power Engineering Society General Meeting, 2007*, 2007.

- [45] M.P.F. Hommelberg, B.J. van der Velde, C.J. Warmer, I.G. Kamphuis, and J.K. Kok. A novel architecture for real-time operation of multi-agent based coordination of demand and supply. In *Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century, 2008 IEEE*, pages 1–5, July 2008. doi: 10.1109/PES.2008.4596531.
- [46] R. Jank. Advanced local energy planning (ALEP), a guidebook. Technical report, International Energy Agency, 2000.
- [47] W. Jansen and W. van Gemers. Nieuwe energie-onderzoeksfaciliteit in Groningen. press release, September 2008.
- [48] J.C.P. Kester and P.J.M. et al. Heskes. A smart MV/LV-station that improves power quality, reliability and substation load profile. In *20th international conference on Electricity Distribution, 2009*.
- [49] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by Simulated Annealing. *Science*, 220(4598):671–680, 1983. doi: 10.1126/science.220.4598.671. URL <http://www.sciencemag.org/content/220/4598/671.abstract>.
- [50] H.P.A. Knops, L.J. de Vries, and A.F. Correljé. Beleidskeuzes voor de inrichting van de elektriciteits- en de gasector in Nederland. Technical report, Wetenschappelijk Instituut voor het CDA, Nov 2004.
- [51] S. Koch, M. Zima, and G. Andersson. Active coordination of thermal household appliances for load management purposes. In *IFAC Symposium on Power Plants and Power System Control, 2009*.
- [52] J.K. Kok, C.J. Warmer, and I.G. Kamphuis. Powermatcher: Multiagent control in the electricity infrastructure. In *4th international joint conference on Autonomous agents and multiagent systems*, pages 75–82. ACM, Jul 2005.
- [53] B. Krose and P. van der Smagt. *An introduction to neural networks*. The University of Amsterdam, 1993.
- [54] National Energy Technology Laboratory. The transmission smart grid imperative. Technical report, U.S. Department of Energy, 2009.
- [55] National Energy Technology Laboratory. A vision for the smart grid. Technical report, U.S. Department of Energy, 2009.
- [56] National Renewable Energy Laboratory. Getting started guide for homer version 2.1. Technical report, National Renewable Energy Laboratory, 2005.
- [57] I. Lampropoulos, G.M.A. Vanalme, and W.L. Kling. A methodology for modeling the behavior of electricity prosumers within the smart grid. In *Innovative Smart Grid Technologies Conference Europe (ISGT Europe), 2010 IEEE PES*, pages 1–8, oct. 2010. doi: 10.1109/ISGTEUROPE.2010.5638967.

- [58] J.F. Manwell and J.G. McGowan. Lead acid battery storage model for hybrid energy systems. *Solar Energy*, 50(5):399–405, 1993.
- [59] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953. doi: 10.1063/1.1699114. URL <http://link.aip.org/link/?JCP/21/1087/1>.
- [60] 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.
- [61] A. Molderink, M.G.C. Bosman, V. Bakker, J. L. Hurink, and G. J. M. Smit. On the effects of mpc on a domestic energy efficiency optimization methodology. In *Proceedings of the IEEE International Energy Conference & Exhibition*, 2010.
- [62] Albert Molderink. *On the three-step control methodology for Smart Grids*. PhD thesis, University of Twente, May 2011.
- [63] H. A. Nielsen and H. Madsen. Predicting the heat consumption in district heating systems using meteorological forecasts. Technical report, Informatics and Mathematical Modelling, Technical University of Denmark, DTU, Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, 2000. URL <http://www.imm.dtu.dk/~han/pub/efp98.pdf>.
- [64] Steffen Nissen. Neural networks made simple. URL <http://leenissen.dk/fann/>.
- [65] National Institute of Standards and Technology. NIST framework and roadmap for smart grid interoperability standards, release 1.0. Technical report, National Institute of Standards and Technology, 2010.
- [66] J. Oyarzabal, J. Jimeno, J. Ruela, A. Englar, and C. Hardt. Agent based micro grid management systems. In *International conference on Future Power Systems 2005*, pages 6–11. IEEE, Nov 2005.
- [67] Ping-Feng Pai and Wei-Chiang Hong. Support vector machines with simulated annealing algorithms in electricity load forecasting. *Energy Conversion and Management*, 46(17):2669 – 2688, 2005. ISSN 0196-8904. doi: DOI:10.1016/j.enconman.2005.02.004. URL <http://www.sciencedirect.com/science/article/B6V2P-4FWV286-1/2/63348ba0d432bb12a18711fd3533f817>.
- [68] Yiqun Pan, Mingming Zuo, and Gang Wu. Whole building energy simulation and energy saving potential analysis of a large public building. In *Eleventh International IBSPA conference*, pages 129–136, July 2009.
- [69] J.M. Pearce, B.A.T. Al Zahawi, and R. Shuttleworth. Electricity generation in the home: modelling of single-house domestic combined heat and power. *Science, Measurement and Technology, IEE Proceedings -*, 148(5):197–203, September 2001. ISSN 1350-2344. doi: 10.1049/ip-smt:20010584.

- [70] C. Petersdorff, T. Boermans, and J. Harnisch. Mitigation of CO₂ emissions from the EU-15 building stock. *Environmental Science and Pollution Research*, 13(5): 350–358, 2006.
- [71] European SmartGrids Technology Platform. Vision and strategy for Europe's electricity networks of the future. Technical report, European SmartGrids Technology Platform, 2006.
- [72] Martin Riedmiller and Heinrich Braun. A direct adaptive method for faster back-propagation learning: The RPROP algorithm. In *IEEE International Conference on Neural Networks*, pages 586–591, 1993.
- [73] J. Scott, P. Vaessen, and F. Verheij. Reflections on smart grids for the future. Dutch Ministry of Economic Affairs, Apr 2008.
- [74] Elena Serban and Daniela Popescu. Prediction of domestic warm-water consumption. *W. Trans. on Comp.*, 7(12):2032–2041, 2008. ISSN 1109-2750.
- [75] Malte C. Thoma, Erge Thomas, Rainer Becker, Anselm Kröger-Vodde, and Christof Wittwer. Active management of electrical networks with a high share of distributed generation. In *2nd International Conference on Integration of Renewable and Distributed Energy Resources*, 2006.
- [76] United States Department of Energy. The micro-CHP technologies roadmap. *Results of the Micro-CHP Technologies Roadmap Workshop*, December 2003.
- [77] P.C. van der Laag and G.J. Ruijg. Micro-warmtekrachtsystemen voor de energievoorziening van nederlandse huishoudens. Technical report, ECN Schoon Fossiel, 2002.
- [78] Centraal Bureau voor de Statistiek. Statline - energiebalans. <http://statline.cbs.nl/StatWeb/publication/?DM=SLNL&PA=70846ned-&D1=31,59&D2=14,29-31&D3=5&D4=l&HDR=T&STB=G1,G2,G3&VW=T>, 2011. URL <http://statline.cbs.nl/StatWeb/publication/?DM=SLNL&PA=70846ned&D1=31,59&D2=14,29-31&D3=5&D4=l&HDR=T&STB=G1,G2,G3&VW=T>.
- [79] C.J. Warmer, M.P.F. Hommelberg, B. Roossien, J.K. Kok, and J.W. Turkstra. A field test using agents for coordination of residential micro-chp. In *Intelligent Systems Applications to Power Systems, 2007. ISAP 2007. International Conference on*, pages 1–4, Nov. 2007. doi: 10.1109/ISAP.2007.4441634.
- [80] R.M. Weatherly and E.H. Page. Efficient process interaction simulation in Java: implementing co-routines within a single Java thread. In *Proceedings of the 2004 Winter Simulation Conference*, volume 2, pages 1437–1443, Dec 2004.
- [81] Christoph Weber and Adriaan Perrels. Modelling lifestyle effects on energy demand and related emissions. *Energy Policy*, 28(8):549 – 566, 2000. ISSN 0301-4215. doi: DOI:10.1016/S0301-4215(00)00040-9.

URL <http://www.sciencedirect.com/science/article/B6V2W-40962RW-6/2/86b34f4d622df6397d6f0b4e561554fd>.

- [82] A. Weidlich, S. Karnouskos, J. Ringelstein, P. Selzam, A. Dimeas, V. Lioliou, C. Warmer, K. Kok, S. Drenkard, D. Ilic, M. Griesemer, and P. Goncalves da Silva. Technology trends for smarthouse/smartgrid. Technical report, SmartHouse/SmartGrid project, 2010.
- [83] G. Wood and M. Newborough. Dynamic energy-consumption indicators for domestic appliances: environment, behaviour and design. *Energy and Buildings*, 35(8):821 – 841, 2003. ISSN 0378-7788. doi: DOI:10.1016/S0378-7788(02)00241-4. URL <http://www.sciencedirect.com/science/article/B6V2V-47TNTC3-1/2/23dbae79d919d8e74544cdb0adb5093e>.
- [84] Andrew Wright and Steven 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, April 2007.
- [85] R.M. Yao and K. Steemers. A method of formulating energy load profile for domestic buildings in the UK. *Energy and Buildings*, 37(6):663–671, June 2005. doi: DOI10.1016/j.enbuild.2004.09.007.
- [86] Ye Yao, Zhiwei Lian, Shiqing Liu, and Zhijian Hou. Hourly cooling load prediction by a combined forecasting model based on analytic hierarchy process. *International Journal of Thermal Sciences*, 43(11):1107 – 1118, 2004. ISSN 1290-0729. doi: DOI:10.1016/j.ijthermalsci.2004.02.009. URL <http://www.sciencedirect.com/science/article/B6VT1-4CB627D-2/2/61d22c330dc7aa92645d6a686e4b0b8c>.