

# Decentralized Energy Management with Profile Steering

Resource Allocation Problems in Energy Management

Thijs van der Klauw

Members of the graduation committee:

Prof. dr. J. L. Hurink	University of Twente (promotor)
Prof. dr. ir. G. J. M. Smit	University of Twente (promotor)
Prof. dr. M. J. Uetz	University of Twente
Prof. dr. ir. B. R. H. M. Haverkort	University of Twente
Prof. dr. C. Witteveen	Delft University of Technology
Dr. R. E. Hebner	University of Texas, Austin
Dr. B. J. Claessens	REstore, Antwerpen
Prof. dr. P. M. G. Apers	University of Twente (chairman and secretary)

## UNIVERSITY OF TWENTE.


Faculty of Electrical Engineering, Mathematics and Computer Science, Discrete Mathematics and Mathematical Programming (DMMP) group and Computer Architecture for Embedded Systems (CAES) group

# CTIT

CTIT Ph.D. thesis Series No. 17-424  
Centre for Telematics and Information Technology  
PO Box 217, 7500 AE Enschede, The Netherlands



This research is conducted within the Energy Autonomous Smart Micro-grids (EASI) project (project number 12700) supported by STW and Alliander. This research is supported by the Dutch Technology Foundation STW, which is part of the Netherlands Organisation for Scientific Research (NWO) and partly funded by the Ministry of Economic Affairs.



Copyright © 2017 Thijs van der Klauw, Enschede, the Netherlands. This work is licensed under the Creative Commons Attribution-NonCommercial 4.0 International License. To view a copy of this license, visit [http://creativecommons.org/licenses/by-nc/4.0/deed.en\\_US](http://creativecommons.org/licenses/by-nc/4.0/deed.en_US).

This thesis was typeset using  $\LaTeX$ , TikZ, and  $\TeX$ nicCenter. This thesis was printed by Gildeprint Drukkerijen, The Netherlands. The cover image is credited to NASA/Goddard Space Flight Center Scientific Visualization Studio.

ISBN 978-90-365-4301-9  
ISSN 1381-3617; CTIT Ph.D. Thesis Series No. 17-424  
DOI 10.3990/1.9789036543019

DECENTRALIZED ENERGY MANAGEMENT WITH PROFILE  
STEERING

RESOURCE ALLOCATION PROBLEMS IN ENERGY MANAGEMENT

PROEFSCHRIFT

ter verkrijging van  
de graad van doctor aan de Universiteit Twente,  
op gezag van de rector magnificus,  
prof. dr. T.T.M. Palstra,  
volgens besluit van het College voor Promoties  
in het openbaar te verdedigen  
op vrijdag 19 mei 2017 om 16.45 uur

door

Thijs van der Klauw

geboren op 31 juli 1989  
te Alphen aan den Rijn

Dit proefschrift is goedgekeurd door:

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

# ABSTRACT



Our energy supply chain is changing rapidly, driven by a societal push towards clean and renewable resources. However, these resources are often uncontrollable (e.g., wind and sun) and are increasingly being exploited on smaller scales (e.g., rooftop photovoltaic). This poses a reliability challenge for the operation of our energy supply chain, specifically for our electricity grid. In this grid, supply and demand must be matched at all times, since storage is virtually non-existent. Traditionally, the supply is controlled centrally and follows the load, where the latter is assumed to be uncontrollable. With the growing number of uncontrollable distributed renewable resources in the system, the centralized paradigm is quickly becoming infeasible.

To combat the decreasing flexibility due to loss of controllability on the generation side, often the exploitation of flexibility on the consumption side is considered. This flexibility comes from devices that can adapt their energy use, e.g., smart white goods or electric vehicles (*EVs*) with smart chargers. Such resources of flexibility on the consumer side are called distributed energy resources (*DERs*). With the expected growth of the number of *DERs* in future energy systems, their coordination offers potential to operate the grid more efficiently and allows the integration of more (uncontrollable) energy from renewables into the grid. Traditional steering approaches in the electricity grid do not scale well with the number of *DERs* and were not designed for the diversity (i.e., heterogeneity) of the envisioned *DERs*. Thus, new energy management approaches are required.

In this thesis we introduce and study such an energy management approach called profile steering. The profile steering approach decentralizes (part of) the computational effort to ensure scalability, making it a decentralized energy management approach. Profile steering relies on predictions and scheduling, meaning that it predicts the future system state and requirements and schedules the use of flexibility of the available *DERs* to best meet the system goals. We focus on the distribution grid, as a large part of the *DERs* are expected to be present in this part of the grid.

The profile steering approach influences the energy use of *DERs* using generic steering signals. We show that the approach can incorporate a broad class of such steering signals. This implies that the approach is flexible enough to be applied in many different situations. Furthermore, we exploit the hierarchical structure of the electricity grid to set up a corresponding hierarchical control structure. This structure allows us to incorporate local limitations into our approach, for instance maximum cable loading of the considered grid section.



As the developed approach is decentralized, we distribute (part of) the required computation to a local level, i.e., to a controller inside a home or embedded in a device. Such controllers often do not have large computational power. To ensure the computations can be feasibly executed on these local controllers with low computational power the resulting distributed scheduling problems have to be researched. We show that many of these problems fall into the class of resource allocation problems, which are well studied in literature. Several of these problems are extensions of known problems. Therefore, we apply some of the techniques found in literature and extend them to include common cases (current and futuristic) in residential energy settings.

In particular we consider buffering devices. Such buffering devices can utilize an internal buffer to decouple (part of) the time they require energy for their operation and the time this energy is taken from the grid. The first type of such a device we consider is the electric vehicle (*EV*). Scheduling the energy use of the *EV* is similar to a classical resource allocation problem if it can charge at any rate between zero and a given maximum. To solve this case we apply techniques from literature. However, if the *EV* can only be charged at a finite number of rates, the problem becomes  $\mathcal{NP}$ -hard, even if we are only interested in obtaining feasible solutions. To circumvent this issue we consider an adaptation of the problem for which we develop an efficient solution method giving results that are nearly identical to feasible solutions to the original problem.

In a follow up chapter we extend the results found for the *EV* to devices that also allow discharging, e.g., residential stationary batteries and *EVs* with vehicle-to-grid capabilities. Furthermore, we study heating, ventilation, and air conditioning systems as a special case. In these systems the energy losses depend on the energy present in the storage (in this case the house itself). Next to developing a method to control such a device, we also study the effect of prediction errors on our approach and show that we can effectively deal with them in the case of heating, ventilation, and air conditioning systems using an approach inspired by model predictive control.

We use simulations to show the validity of profile steering using several cases. We show that profile steering can also be used to achieve near optimal results when minimizing the degradation of a power transformer. This indicates that the benefits that can be expected from using our approach are not limited to energy markets, but also include increased lifetime of grid assets resulting in reduced investment costs in these assets.

Summarizing, the introduced profile steering decentralized energy management approach promises to be a valuable approach in the future (smart) electricity grid where it can unlock the potential of many residential *DERs* and assist in an effective and efficient energy transition.

# SAMENVATTING

vii



Ons energie distributie netwerk ondergaat snelle veranderingen, gedreven door een ideaal van een maatschappij gebaseerd op schone en herbruikbare bronnen. Deze bronnen zijn echter vaak onbeheersbaar (bijvoorbeeld de wind en de zon) en zij worden in steeds grotere mate op kleine schaal geëxploiteerd (bijvoorbeeld zonnepanelen op het dak). Deze veranderingen bedreigen de stabiliteit van onze energie distributie netwerken, met in het bijzonder ons elektriciteitsnetwerk. In het elektriciteitsnetwerk moeten vraag en aanbod altijd in balans zijn, omdat opslag praktisch niet voorkomt. Traditioneel gezien wordt het aanbod gestuurd zodat productie en vraag altijd overeenkomen, omdat men er vanuit gaat dat de vraag niet beheersbaar is. Maar, met het groeiend aantal onbeheersbare gedistribueerde bronnen in het systeem wordt deze traditionele centrale aanpak snel onbruikbaar.

Om de vermindering van flexibiliteit door het verlies van stuurbaarheid aan de aanbodkant tegen te gaan, wordt vaak flexibiliteit aan de vraagkant als optie gezien. Deze flexibiliteit komt van apparaten waarvan het gebruik van energie kan worden aangepast, bijvoorbeeld slim witgoed en elektrische auto's met slimme laders. Zulke apparaten, gezien als bronnen van flexibiliteit aan de kant van het verbruik, worden gedistribueerde energiebronnen genoemd. Met het verwachte aantal gedistribueerde energiebronnen in ons energiesysteem van de toekomst leveren zij de mogelijkheid om, met de juiste coördinatie, het net efficiënter te maken en een grotere mogelijkheid te bieden tot het integreren van energie van hernieuwbare bronnen. Traditionele aanpakken van de coördinatie van deze bronnen schalen niet naar een groot aantal gedistribueerde bronnen en zijn bovendien niet ontworpen voor een grote diversiteit aan bronnen. Daarom zijn nieuwe aanpakken nodig om het net te besturen.

In dit proefschrift introduceren en bestuderen wij een aanpak om het net te besturen, genaamd profile steering. Deze profile steering aanpak is een gedecentraliseerde aanpak omdat (een gedeelte van) de berekeningen decentraal uitgevoerd wordt om zo schaalbaarheid te bereiken. Profile steering maakt gebruik van voorstellingen en plannings, waarmee bedoeld wordt dat de aanpak toekomstige situaties in het systeem voorspelt en dat het gebruik van flexibiliteit door de beschikbare gedistribueerde bronnen wordt gepland, zodat dit gebruik overeen komt met de doelstellingen van het systeem. Wij leggen de nadruk op elektrische distributie netwerken, omdat daar een groot aantal van de gedistribueerde energiebronnen verwacht wordt in de toekomst.

Profile steering stuurt het gebruik van flexibiliteit van de gedistribueerde energie-



bronnen door middel van stuursignalen. Wij laten zien dat onze aanpak gebruik kan maken van een brede klasse van zulke stuursignalen. Dit impliceert dat de aanpak flexibel genoeg is om in vele verschillende situaties te worden toegepast. Verder gebruiken we de hiërarchische structuur van het elektriciteitsnet om een vergelijkbare structuur in de aansturing te bereiken. Met deze structuur in de aansturing wordt het mogelijk om lokale limieten, zoals maximale kabel belasting, mee te nemen in onze aanpak.

De ontwikkelde aanpak is decentraal; we distribueren (een gedeelte van) de vereiste berekeningen naar een lokaal niveau, bijvoorbeeld naar een systeem dat het verbruik binnen een huis aanstuurt of naar de apparaten zelf. De rekenkracht op dit lokale niveau is vaak beperkt. Om toch de uitvoerbaarheid van de aanpak te garanderen met deze beperkte rekenkracht, moeten de resulterende problemen, die op lokaal niveau moeten worden opgelost, worden bestudeerd. Wij tonen aan dat een groot gedeelte van deze problemen in de klasse van ‘resource allocation problems’ valt. Deze problemen zijn bekend en eerder bestudeerd in de literatuur. Een aantal van de problemen die wij tegenkomen zijn uitbreidingen van bekende problemen. Dit maakt het voor ons mogelijk om technieken uit de literatuur her te gebruiken en uit te breiden om veel voorkomende problemen in de decentrale aansturing van het net op wijkniveau aan te pakken.

In het bijzonder kijken wij naar apparaten die intern opslag bevatten om zo (een gedeelte van) hun energieverbruik van het net te kunnen afnemen voordat dit daadwerkelijk nodig is voor het gebruik van het apparaat. Een eerste apparaat dat hieraan voldoet, en dat wij bestuderen, is een elektrische auto. Het inplannen van het verbruik van een elektrische auto is vergelijkbaar met een klassiek resource allocation problem zolang we aannemen dat de auto op elk vermogen tussen nul en een gegeven maximum kan laden. In dit geval passen we technieken uit de literatuur toe. Als de auto, in tegenstelling tot eerder, beperkt is tot laden op een eindig aantal verschillende niveaus wordt het probleem  $\mathcal{NP}$ -hard, zelfs als we alleen geïnteresseerd zijn in een oplossing die aan de eisen voldoet. Om toch tot een werkbare oplossing te komen bekijken we een aanpassing van het probleem. Voor deze aanpassing ontwikkelen we een efficiënte oplossingsmethode die resultaten geeft die bijna identiek zijn aan oplossingen die aan de eisen van het originele probleem voldoen.

Hierna breiden we onze resultaten, gevonden voor de elektrische auto, uit naar apparaten die ook energie kunnen ontladen, bijvoorbeeld batterijen en elektrische auto's die ook kunnen ontladen naar het net. Verder bestuderen we ook luchtbehandelingssystemen als een speciaal geval, omdat in zulke systemen de energieverliezen afhangen van de hoeveelheid energie ‘opgeslagen’ in het systeem (in dit geval het gebouw zelf). Naast dat we een aanpak om zo'n systeem te besturen ontwikkelen, bestuderen we ook het effect van voorspellingsfouten op onze aanpak. We laten zien dat onze aanpak voor luchtbehandelingssystemen effectief met voorspellingsfouten kan omgaan, waarbij we gebruik maken van een aanpak geïnspireerd door ‘model predictive control’.



We tonen door middel van verschillende simulaties aan dat profile steering goed functioneert. In het geval dat we het verouderen van een transformator bestuderen laten we zien dat onze aanpak bijna optimale resultaten geeft. Dit geeft een indicatie dat de voordelen van onze aanpak niet beperkt blijven tot energiemarkten, maar ook verlenging van de levensduur van componenten in het net omvatten, wat resulteert in gereduceerde investeringen in het net.

Samenvattend belooft profile steering om van waarde te zijn in het toekomstige (slimme) elektriciteitsnet, waar de aanpak efficiënt gebruik kan maken van vele verschillende gedistribueerde energiebronnen en kan assisteren in een effectieve en efficiënte energietransitie.



X



# DANKWOORD



Voor je liggen ruim 200 pagina's die gezamenlijk het proefschrift vormen en waarop ik, even aangenomen dat de commissie het hier mee eens is, mag promoveren. Het is mijn verslaglegging van de wetenschappelijke tocht die ik de afgelopen vier jaar heb afgelegd aan de Universiteit Twente (en daar buiten). Zo'n proefschrift is echter alleen maar een verslaglegging van het wetenschappelijke aspect van het traject dat je als promovendus doorloopt, waarop uiteindelijk maar één naam komt. Dat terwijl er hordes mensen zijn die in meer of mindere mate hebben geholpen bij de totstandkoming van dit boekje, zowel op wetenschappelijk vlak als daarbuiten. Daarom wil ik hieronder graag een aantal mensen die voor mij belangrijk zijn geweest bedanken.

Allereerst wil ik graag mijn promotoren Johann en Gerard bedanken. Johann, jij hebt er voor gezorgd dat ik me direct op mijn plek voelde in Twente, onder andere door onze gedeelde passie voor schaatsen. Jouw interesse in de wetenschappelijke en niet-wetenschappelijke bezigheden van jouw collega's weet mij, en met mij vele anderen, altijd te motiveren. Hoe je alles wat je doet voor elkaar krijgt en weet te balanceren is voor mij soms een raadsel, waarvan (een gedeelte van) het antwoord misschien ligt in het verschil in uren slaap per nacht. Gerard, de manier waarop jij de CAES vakgroep weet te leiden, zodat er een sfeer ontstaat waarin zowel hard gewerkt als hard gelachen kan worden, vind ik zeer waardevol. Ook bij jou heb ik mij meerdere malen afgevraagd hoe je alles weet te balanceren!

Naast de begeleiding van Johann en Gerard is een belangrijk deel van dit proefschrift tot stand gekomen dankzij Marco. Alhoewel je nooit officieel aangewezen bent als mijn begeleider, heb ik wel enorm veel aan de discussies met jou gehad vanaf het moment dat je na het schrijven van jouw boekje de overstap maakte naar de 'energiegroep'. Vooral het feit dat ik je gevraagd en ongevraagd eigenlijk altijd kon komen storen met wetenschappelijke (en soms minder wetenschappelijke) vragen is voor mij zeer van waarde geweest. Ik hoop dat je net zo veel aan onze discussies hebt gehad als ik!

Verder wil ik graag de verschillende (oud-)collega's uit de CAES and DMMP groepen bedanken. Het is niet altijd even gemakkelijk om deel uit te maken van twee vakgroepen, want twee keer zoveel pauze's, vakgroepuitjes, kerstdiners, etc. is best vermoeiend ;). Al deze sociale bezigheden waren vaak een goede afleiding van het onderzoek, en soms een iets te goede afleiding. In het bijzonder heb ik genoten van de volgende dingen (in willekeurige volgorde): de interessante gesprekken over bijvoorbeeld Magic the Gathering (Jasper en Ruben), (computer)spelletjes (Ger-



win, Gijs, Tom en Christiaan), boeken (Marco) en sport (Rinse en Robert; dit jaar de halve van Enschede!); de discussies over bijvoorbeeld economie en politiek in binnen- en buitenland (Jordi, Jochem en Martijn); de puzzels (Guus; volgend jaar 100 punten!); en de soms gortdroge humor (Hermen), allemaal (voornamenlijk) tijdens de koffie/thee pauzes en de lunchwandelingen. Daarnaast wil ik graag Marjo, Marlous, Nicole en Thelma bedanken voor de ondersteuning bij verschillende (ze-ker niet onbelangrijke) randzaken als declaraties, vluchten, hotels, verzekeringen, etc. de afgelopen jaren.

Als AiO ben je vaak ‘veroordeeld’ tot een kantoor dat je met meerdere mensen moet delen, in mijn geval bijna de hele periode de vissenkomp, ook wel bekend als het energiehok. Ik wil hier graag mijn lotgenoten uit dit ‘hok’ bedanken (voor zover die hierboven nog niet genoemd zijn). Ik ga niet iedereen noemen want de lijst van bewoners van het energiehok is nogal lang (lees: ik wil niet de fout maken om iemand te vergeten!), maar weet dat jullie een belangrijke bijdrage hebben geleverd door middel van een goede werksfeer en (soms eindeloze) nuttige en minder nuttige discussies. Ik hoop dat we de whiteboards de komende periode nog maar flink mogen misbruiken (en dan niet om teksten van Johann te ontcijferen).

Ook vind ik het belangrijk te noemen dat mijn werk geen losstaand onderzoek is maar past in het geheel dat ondertussen door een vrij grote groep mensen in Twente wordt onderzocht. Hierbij wil ik allereerst mijn voorgangers Albert, Vincent, Maurice, Stefan en Hermen bedanken. Verder zou dit proefschrift er niet staan in de huidige vorm zonder dat ik gebruik heb kunnen maken van het template dat door vele mensen voor mij is gebruikt en bewerkt. Alhoewel het nog een open vraag is of het nou daadwerkelijk onder Windows kan draaien (huidig antwoord: misschien...). Verder wil ik in het bijzonder Gerwin als collega op hetzelfde project bedanken. Jouw input, soms vanuit een totaal andere hoek, was vaak van grote waarde voor mijn onderzoek. Ook heb ik zeer genoten van onze discussies, conferenties, uitstapjes, etc. met wat mij betreft als hoogtepunt twee keer de VS.

Naast mijn werkzaamheden als PhD student zijn mijn (vele) hobbies voor mij altijd van belang geweest. Gelukkig kon ik mijn passie voor zowel muziek als sport al snel kwijt in Twente bij respectievelijk SHOT en de Skeuvel. De muziek en het schaatsen hebben altijd als goede afleiding gefungeerd voor het werk en daarvoor wil ik graag alle (oud)-SHOT'ers en (oud)-skeuvels bedanken, met in het bijzonder mijn sectiegenoten (sexy saxy sectie!), trainingsgenoten (met of zonder gele muts/hesje) en trainers.

Via de muziek is de stap richting twee (ondertussen) mede-promovendi in Twente die het verdienen om apart genoemd te worden snel gezet: Ewert en Lisette. Jullie beiden zijn voor mij een belangrijk steunpunt geweest, ook in de periode dat het wat minder ging. Daarnaast zijn jullie (wellicht soms ongemerkt) vaak een bron van inspiratie geweest voor mij om alles er uit te halen wat er in zat zowel in de muziek als in de sport als in de wetenschap.

Ook buiten Twente hebben de vele mensen die ik vriend mag noemen als vrienden vaak voor de nodige afleiding gezorgd. Alhoewel ik de mensen die ik ken van de

studie in Utrecht (Anieke, Annemarie, Auke, Björn, Hannelore, Hilco, Max, Olaf en Sanne), van de middelbare school (Bartjan, Dennis, Marcel, Niels, Tjitse), van de basisschool (Gerben en Jens), of via andere wegen (Anna-vera, Daan, Fritsjan, Kasper, Vicent) tegenwoordig wat minder zie dan ik soms zou willen, zijn jullie (nog steeds) zeer belangrijk voor mij.



Dan rest mij nog om de belangrijkste personen te noemen, mijn ouders Cora en Kees en mijn zusje Dorien. Dank voor het feit dat jullie mij altijd een luisterend oor bieden en voor de mogelijkheid om toch altijd weer een beetje thuis te komen in Alphen, ook al zit ik soms wat verder (of een paar maanden héél ver) weg. Ik heb mij gedurende mijn jeugd en de tijd als (PhD) student altijd door jullie gesteund gevoeld. Zonder jullie zou ik niet de persoon zijn die ik nu ben en zou dit proefschrift er nooit gekomen zijn.

Tenslotte, omdat ik weet dat sommige lezers van puzzelen houden en omdat het hoog tijd is dat ik weer een taartweddenschap afsluit nog een puzzel. De eerste persoon die mij voor 1 februari 2018 het juiste antwoord op de volgende vraag (inclusief ontcijferde vraag!) weet te geven krijgt een zelfgemaakte taart. Een eerste hint om te beginnen: CK2B4PHVHP4B2KC. Verder staat alles wat je nodig hebt in dit proefschrift.

PQ QQDEFQ	O
QSBP PNCUSNIO	B
HS SBBEU	AK
UP UNHEU	AWT
LP YR	ACMGH
BGJL CVD	LJDXFYDST
QDG NAAITCI	FYMKUQRKPWAAN
FIS ANKAIWIGFRP	GJKILRAMXZWACBNANGOPG
CR MYEAGNKR	WWCCGVWROLISO
MHA BT	ULZNUJXHU
LP LJGCG	BZANB
ZKEL TLCTE	RWR
RM FRI	RO
KBIEWJPK IB	P
PQLQ FQWEFQZ	E

Thijs  
Enschede, januari 2017



# CONTENTS



<b>1</b>	<b>INTRODUCTION</b>	<b>1</b>
1.1	Introduction . . . . .	3
1.2	Illustrative Example . . . . .	4
1.3	Problem Statement . . . . .	7
1.4	Approach & Contributions . . . . .	9
1.5	Outline of the Thesis . . . . .	10
<b>2</b>	<b>BACKGROUND</b>	<b>15</b>
2.1	Introduction . . . . .	17
2.2	The Electricity System . . . . .	18
2.2.1	<i>Traditional Centralized System</i> . . . . .	18
2.2.2	<i>Current System</i> . . . . .	23
2.2.3	<i>Future System</i> . . . . .	24
2.3	Requirements for Energy Management Approaches . . . . .	26
2.4	Formulation of the Energy Management Problem . . . . .	28
2.5	Different Types of Residential Distributed Energy Resources . . . . .	30
2.5.1	<i>EF-Pi Devices</i> . . . . .	31
2.5.2	<i>Model for Buffering Devices</i> . . . . .	32
2.6	Related Work . . . . .	33
2.7	Conclusion . . . . .	38
<b>3</b>	<b>PROFILE STEERING</b>	<b>41</b>
3.1	Introduction . . . . .	43
3.2	Complexity of Scheduling Based EM . . . . .	44
3.2.1	<i>Complexity</i> . . . . .	45
3.3	Profile Steering Heuristic . . . . .	46
3.4	Steering Signals . . . . .	51
3.4.1	<i>Price Signals</i> . . . . .	51
3.4.2	<i>General Steering Signals</i> . . . . .	52



3.4.3	<i>Scheduling Devices under Steering Signals</i> . . . . .	53
3.5	Example Application of Profile Steering . . . . .	54
3.6	Hierarchical Control . . . . .	56
3.6.1	<i>Incorporating Bounds in Profile Steering</i> . . . . .	57
3.6.2	<i>Incorporating Local Limits in the Example</i> . . . . .	60
3.6.3	<i>Overloading of Cables</i> . . . . .	61
3.7	Conclusion . . . . .	63
4	<b>CHARGING OF ELECTRIC VEHICLES</b> . . . . .	<b>67</b>
4.1	Introduction . . . . .	69
4.2	Related Work . . . . .	70
4.2.1	<i>Optimization of a Single Vehicle</i> . . . . .	70
4.2.2	<i>Optimization of a Fleet of Vehicles</i> . . . . .	71
4.3	Electric Vehicle Scheduling Problem . . . . .	72
4.3.1	<i>Resource Allocation Problems</i> . . . . .	74
4.3.2	<i>Optimality Conditions</i> . . . . .	75
4.4	Solution Methods . . . . .	77
4.4.1	<i>Lagrangian Multiplier Approach</i> . . . . .	78
4.4.2	<i>Pegging Approach</i> . . . . .	81
4.5	Discrete Variant . . . . .	85
4.5.1	<i>Complexity of Charging over a Discrete Set</i> . . . . .	85
4.5.2	<i>Piecewise Linear Approximation</i> . . . . .	87
4.5.3	<i>Solution Methods</i> . . . . .	89
4.6	Simulation Study . . . . .	96
4.6.1	<i>Comparison with Pricing</i> . . . . .	96
4.6.2	<i>Comparison with State-of-the-Art</i> . . . . .	98
4.7	Conclusion . . . . .	99
5	<b>ENERGY STORAGE</b> . . . . .	<b>103</b>
5.1	Introduction . . . . .	105
5.2	Related Work . . . . .	106
5.3	Problem Formulation and Solution . . . . .	107
5.3.1	<i>A Model for Discharging</i> . . . . .	107
5.3.2	<i>Efficiently Solving the Battery Problem</i> . . . . .	110
5.3.3	<i>Discrete Buffer</i> . . . . .	116
5.4	Application of Buffer Devices . . . . .	121
5.5	Conclusion . . . . .	124



<b>6</b>	<b>HEATING, VENTILATION, AND AIR CONDITIONING SYSTEMS</b>	<b>129</b>
6.1	Introduction	131
6.2	Related Work	132
6.3	Thermal Model	133
6.3.1	<i>Model Determination</i>	134
6.3.2	<i>Model Verification</i>	135
6.4	HVAC Scheduling Problem	137
6.4.1	<i>HVACS Without Cumulative Bounds</i>	139
6.4.2	<i>Solution Approach for HVACS</i>	141
6.4.3	<i>Discrete Variant of HVACS</i>	144
6.5	Operational Control of HVACs	148
6.5.1	<i>Prediction Errors in Profile Steering</i>	149
6.5.2	<i>Base Case</i>	150
6.6	Simulation Study	151
6.7	Conclusion	152
<b>7</b>	<b>TRANSFORMER AGEING AND EM</b>	<b>159</b>
7.1	Introduction	161
7.2	Related Work	162
7.3	Transformer Ageing Model	163
7.3.1	<i>Thermal Ageing Model</i>	163
7.3.2	<i>Simplifying Assumption</i>	166
7.4	Using EM to Reduce Transformer Ageing	167
7.5	Simulation Study	170
7.5.1	<i>Considered Scenario</i>	170
7.5.2	<i>Optimizing Transformer Lifetime</i>	171
7.5.3	<i>Profile Steering</i>	172
7.5.4	<i>Uncontrolled Charging</i>	172
7.5.5	<i>Comparison of the Cases</i>	173
7.6	Conclusion	174
<b>8</b>	<b>SUMMARY AND CONCLUSION</b>	<b>179</b>
8.1	Summary of the Obtained Results	181
8.2	Conclusion & Discussion	182
8.3	Recommendations	185
8.3.1	<i>Dealing with uncertainty</i>	187





<b>A</b>	<b>MATHEMATICAL BACKGROUND</b>	<b>189</b>
A.1	Convex Optimization . . . . .	191
A.1.1	<i>Convex Sets</i> . . . . .	191
A.1.2	<i>Convex Functions</i> . . . . .	191
A.1.3	<i>Convex Optimization</i> . . . . .	193
A.2	Complexity . . . . .	194
A.2.1	<i>Notation</i> . . . . .	194
A.2.2	<i>Polynomial Time Complexity</i> . . . . .	195
	<b>ACRONYMS</b>	<b>199</b>
	<b>LIST OF SYMBOLS</b>	<b>201</b>
	<b>BIBLIOGRAPHY</b>	<b>205</b>
	<b>LIST OF PUBLICATIONS</b>	<b>219</b>
	<b>ABOUT THE AUTHOR</b>	<b>221</b>




XX



# 1

## INTRODUCTION

*ABSTRACT – This chapter serves as an introduction to the thesis. We briefly introduce the research area energy management for smart grids and illustrate some of the prominent problems in this area. Furthermore, we give a problem sketch by means of an illustrative example. This leads to the formulation of the main research question: How can we manage energy in the future smart grid? We briefly introduce our novel energy management approach called profile steering. Finally, we give an outline of the thesis.*






Figure 1.1: The share of electricity produced by renewables (data taken from [25] (NL) and [7] (Ger)).




Figure 1.2: Example energy profile of a Dutch house with rooftop photovoltaic panels with and without a battery (eight kWh).




Figure 1.3: The state of charge of the battery when used to greedily charge and discharge energy (corresponding to Figure 1.2).

## 1.1 INTRODUCTION

In December of 2015 world leaders gathered in Paris for the 21st United Nations conference on climate change (COP21). During this conference, representatives from countries all over the world signed an agreement to limit the emission of green houses gases and, thereby, holding the increase in global average temperature well below two degrees Celsius above pre-industrial levels [130]. Furthermore, the agreement states an intention of pursuing a target of one-and-a-half degrees of temperature rise. As of January 2017, the agreement has been signed by 194 countries, contributing to a total of 99% of the greenhouse gas emissions worldwide [131, 132]. The agreement has been ratified by enough parties for it to enter into force as of November 2016, with the ratifying parties accounting for over 80% of the global emissions as of January 2017.

Among experts there is a consensus of between 90% and 100% that we as humans are causing recent global warming [33]. Together with other important drivers such as health related issues and a desire to be independent of finite resources that often come from politically less stable regions, this is causing a need for change in our energy supply chain: the energy transition. This transition primarily includes a push towards energy savings and incorporating energy from clean and renewable sources such as wind and sun. In Figure 1.1 an example for the share of electricity production from renewables in the Netherlands and Germany is given. The transition is realized through agreements, plans, and policies worldwide, one of which is the Paris agreement mentioned above. Another example is the set of targets of the European Union in their 2020 and 2030 energy and climate packages [49, 50].

However, a society that no longer depends on carbon-based energy is not without difficulty. One of the challenges to be tackled is that our energy supply chain was designed using a paradigm completely based on the use of fossil fuels. The situation in our energy supply chain is quickly changing, in particular in the electricity grid (see, e.g., [103, 128]). Note that we often use the term electricity grid to mean the entire infrastructure used to supply us with electricity. This includes not only the physical cables and power electronics to transport and distribute power but also the generation and consumption assets and the hardware and software used for control.

One of the changes in the electricity grid is that households are increasingly installing rooftop photovoltaic (*PV*), which causes them to become net producers at some times (see Figure 1.2 for an example profile of a house with *PV*). This potentially causes problems in the grid. While novel devices, such as stationary batteries, in principle can assist with alleviating these problems, this is only possible if they are used properly. For example, consider a scenario where the battery is used to store local solar energy for the house as soon as the local production exceeds the demand, as illustrated in Figure 1.2. In this situation the battery is filled before the *PV* peak is over (Figure 1.3). The result is that power at the peak of the production provided by the *PV* still has to be transported away from the house.





One of the main challenges in the electricity grid is that production and consumption need to be balanced at all times. In the old system, there is a lot of flexibility available on the production side of the system in large-scale generators running on fossil fuels. This flexibility is traditionally used to match the production to the consumption as the generators can adapt their output to match fluctuations in demand. However, many of the renewable sources used to produce clean energy are of an uncontrollable nature, e.g., sun and wind. This means that the flexibility on the generation side of the system is decreasing.

Advancements in information and communication technology can offer potential solutions to deal with the loss of flexibility in production, by considering flexibility on the consumption side instead [120, 122]. These potential solutions arise as appliances on the consumer side are increasingly being equipped with hardware and software that allows them to act on input from the consumer. The same hardware and software can be used to let these appliances react on control signals from the electricity grid. As an example, one can consider running appliances when energy is cheap and/or available from local (renewable) production such as rooftop PV. In essence these smart appliances offer flexibility in their energy use, compensating for the loss of flexibility on the generation side. However, the traditional paradigm used to manage our energy supply chain (so-called energy management (*EM*)) was not designed with these (new) distributed sources of flexibility in mind and cannot exploit this flexibility. Hence, one of the major challenges in the future electricity grid is the design of an *EM* approach that is capable of exploiting the flexibility offered to the system by a large number of novel heterogeneous appliances on the consumption side. In this thesis we focus on the design of such an *EM* approach.

The remainder of this chapter serves as a general introduction to the research described in this thesis. It is outlined as follows. In the next section we discuss the changes occurring in the electricity grid and the complexity of the resulting problem using an example. This leads to the formulation of the central problem statement studied in this thesis in Section 1.3. Then, in Section 1.4 we discuss our approach to the problem and the contributions of this thesis to the field. Finally, in Section 1.5 we give an outline of the remainder of the thesis.

## 1.2 ILLUSTRATIVE EXAMPLE

In this section we explain changes and challenges we observe in our electricity grid by means of examples. In particular we focus on the distribution part of the electricity grid; the part of the grid used to connect smaller customers, such as households, to the main grid.

One of the many changes in the electricity grid is the introduction of local small-scale generation, particularly connected to the distribution grid. These local generators often exploit renewable sources such as wind and sun that are intermittent. As an example we consider the introduction of PV, specifically rooftop PV on houses. In several European countries, subsidies have been (and often still are) in place





to give financial support to people with rooftop PV panels [74]. This makes an investment in these panels rather attractive as the panels have a payback period of just a few years. To further stimulate home-owners to consider this option, many countries introduced beneficial energy metering policies for residential PV. As an example consider the Dutch ‘nul-op-de-meter’ (net metering) policy, which allows PV owners to directly subtract all production of their panels from their total energy use in a year [74].

Financial incentives such as the subsidies described above contributed to a growing amount of PV being installed in several European countries, with Germany being a front-runner. The PV panels are also often installed with the idea to produce (part of) the required energy within a neighbourhood locally. We note that the above mentioned policy of subtracting produced energy from consumed energy on a yearly basis effectively means the customer is allowed to use the grid as free and 100% efficient storage. This is because energy that is produced, e.g., in the middle of the day, when the PV production peak occurs, is generally of a larger volume than required in the residential areas at that time. The peak in consumption of such areas usually lies in the evening, however, at such times little or no energy from the sun is available. Thus, customers ‘store’ their PV energy in the grid during the day and ‘retrieve’ it back in the evening. Furthermore, due to this policy, the amount of installed PV panels on a house often yields a production equal to the total consumption of that house on a yearly basis. As an example see again Figure 1.2, which shows the energy profile of a house equipped with rooftop PV. The house is a net producer during the afternoon but needs to import energy at night. The above also implies that it makes sense from a financial point of view to place the panels facing south. On the other hand, facing the panels west would have caused their production peak to occur in the late afternoon or early evening, better matching the consumption peak of a typical household.

If only a small number of consumers use the grid as a large storage this is not harmful for the grid, as (local) overproduction in a home is easily consumed locally. This is currently the case in (nearly all parts of) the Netherlands, as the penetration of local generation is rather low compared to, e.g., Germany (see Figure 1.1). However, severe problems can occur at later stages during the transition if the penetration of PV panels in a neighbourhood reaches higher levels. In such a scenario peaks in local production can no longer be consumed locally. Instead the energy needs to be transported away from this neighbourhood, towards areas where the energy can be consumed. With significant levels of PV penetration the resulting production peaks are at times far larger than the typical consumption peaks occurring in the grid. As the grid has to be dimensioned to accommodate the largest peak, this potentially causes overloading of grid assets and thus causes (extra) investments in the distribution grid to be required. These problems have recently been encountered in several areas in Germany [106]. As an example see Figure 1.4 where we give the power flows for a rural area in Germany with a lot of PV installed at the beginning of May 2016. In this area the production from PV is far larger than the local consumption, for which the system was originally dimensioned.




Figure 1.4: An example of the overproduction of PV in a German neighbourhood. Data provided by Westnetz for the week of the first until the seventh of May 2016.

To combat the problems sketched, policies are changing. For example, to encourage local producers to consume as much of their own energy as possible, Germany now uses a feed-in tariff for rooftop PV that is lower than the price customers pay for energy [74]. One of the (newer) ways to increase self-consumption is through the use of *smart appliances*. These smart appliances can match their consumption with the local availability of energy and, thereby, increase the self-consumption of a household. Another option is using storage, such as a battery. At the moment battery prices are generally considered to be too high to be economically viable in this setting. However, battery prices are expected to continue to drop sharply, as we have seen over the last few years (see, e.g., Figure 5.2 on page 104). This means batteries might become economically feasible on a household level in the (near) future. In such a scenario it is likely, due to economical reasons, that the battery is sized to cover just the PV production on an average day rather than on all days. If we now consider the scenario where the battery is controlled to charge energy as soon as there is a surplus in production from PV and discharge energy whenever the demand exceeds production, the battery might be fully charged before the full peak production of PV kicks in on very sunny days. To illustrate the issue we equipped the house illustrated in Figure 1.2 with a battery of eight kWh, which covers the daily *average* excess of energy produced. However, if the battery is used in a greedy manner, i.e., charge when there is a surplus and discharge when there is demand, on a very sunny day the maximum production peak is not lowered as illustrated. While this does not affect the level of self-consumption, it still implies that the problem of high peaks in production on the neighbourhood level remains. Furthermore, as Figure 1.3 shows, it might be that the battery is only partly discharged during the evening and night as the total stored energy exceeds the total energy demand of the house for the given day. This implies the battery can store even less of the solar peak of the next day if it is not fully discharged first. This discharging of the excess energy in the battery at night can be used to provide neighbours with energy. However, to determine when and by how much the battery should be discharged the situation in the rest of the neighbourhood needs to be known or at least estimated.



One way to combat the first difficulty mentioned above is by using a smarter battery management system that predicts when the peak in PV production will likely occur and then uses these predictions to ensure the battery is steered to reduce the peak. In the case of PV systems, the typical production peak lies around noon. This peak might in some cases coincide with a local consumption peak for cooking lunch in, e.g., Germany. Furthermore, it may be that other resources are available in the neighbourhood which can be ‘adapted’ to the production peak. For example, it might be possible to (also) charge an electric vehicle (*EV*) of a neighbour around noon to consume the produced PV energy. In such a case the system needs to make a trade-off between using the locally produced PV energy for the *EV* of the neighbour now or for storage in the battery to cover the own demand in the evening. The example illustrates that the local system benefits from knowledge of what happens in the rest of the neighbourhood. With a lot of potential sources of energy and (flexible) consumption in a neighbourhood, the design of a system that properly manages these sources becomes a challenging task.

### 1.3 PROBLEM STATEMENT

As mentioned in the previous section, the traditional EM approach can no longer be applied in the future grid due to the changes caused by the energy transition. To this end it is of paramount interest that new EM approaches are designed and studied, as we propose in this thesis. This leads to the main research question:

*How can we effectively and efficiently manage the flexibility provided by (future) energy resources in the electricity grid to facilitate the changes occurring due to the energy transition?*

With effective and efficient we mean that the approach should provide reasonable solutions to the problems that are currently observed (and are expected to occur) in the (future) electricity grid while being practically applicable. This means that we are not only searching for solutions that have a high theoretical performance, but we also require our solutions to be implementable in practice meaning that constraints on available resources, such as computational power, are respected. As an example, we consider the battery from the previous section and we assume multiple houses in a neighbourhood have such a battery installed. It is theoretically possible to compute the best way to coordinate the use of all batteries together to minimize the stress on the electricity grid at a central location where all relevant information is known. However, finding such a solution probably requires too much computational power and uses information that is privacy sensitive. As an outcome it is often probably better to trade performance in terms of finding the optimal solution for other desirable characteristics, such as lower computational complexity or independence of privacy sensitive information.

The above raises the question what these desirable characteristics are. For this we have to look at the requirements for EM approaches. We summarize these requirements as:



*An EM approach is required to employ scheduling in a scalable manner to use flexibility offered by a large set of heterogeneous resources in a feasible manner at the appropriate time while respecting user privacy.*

In light of the example above, scheduling means that the approach is capable of determining when flexibility is required the most, e.g., it determines when it is best to charge locally produced energy into the battery and when to discharge energy. By scalable we mean that the computational requirements of the approach should not exceed what is generally assumed to be available. This scalability requirement is of particular relevance in large systems, e.g., thousands of households each with a multitude of flexible appliances. The heterogeneity aspect is important because flexibility comes from many different appliances in the (future) grid, e.g., smart white goods, heat pumps, batteries, EVs, etc. This implies that the approach should be capable of exploiting the flexibility provided by many different devices. Furthermore, the approach needs to take all sorts of constraints into account to produce feasible solutions, both from the users perspective (e.g., an EV needs to be charged before departure) and from the grid perspective (e.g., currents flowing through a cable should not exceed the cable's limits). Some of these constraints as well as other relevant parameters entail privacy sensitive information, for example, the arrival and departure times of an EV disclose when the user is at home. It is often desirable to keep this information as local as possible. We discuss the requirements listed above in more detail in the next chapter.

In order to tackle the main problem in this thesis we study three different aspects of the main research question. We first focus on the coordination *between* different energy resources in the grid. This leads us to formulate a novel EM approach called profile steering. In several future EM approaches, including profile steering, some level of local decision making is required. By this we mean that on a device or household level decisions need to be made on when and how the available flexibility is used. This already happens to some extent in the current grid, through the use of day and night tariffs. With these tariffs, users have to make decisions whether they want to shift (part of) their electricity consumption to the cheaper night period or not. The study of these local or device level problems plays a central role in this thesis. Note that in many cases the available hardware to make the local decisions is very limited. Because of this reason it is important that local solutions can be found very efficiently with very low computational and hardware requirements (e.g., they cannot depend on external programs that require significant computational power such as commercial solvers that may be available for these problems).

Finally, the goals of different stakeholders in the electricity grid do not necessarily align. As an example we consider the discussion in the previous section about houses with rooftop PV. As long as the 'nul-op-de-meter' policy is in place, customers have no incentive to reduce their afternoon production peak. However, local peaks in production can potentially lead to more investments required in grid assets for the system operators. The above highlights that the different stakeholders in the system do not necessarily share the same goals. However, it is important that



an EM approach is capable of realizing different goals for different stakeholders by being able to consider trade-offs between these goals. We briefly touch upon this issue in this thesis. The three aspects mentioned above are captured by the three sub-questions we consider in this thesis, given below:

- » How can we effectively manage the coordination of flexibility of a set of heterogeneous devices?
- » What are the local decision problems in future EM approaches and how can we solve them on the local level?
- » Can an EM approach assist in realizing goals of different stakeholders in the (future) smart grid?

#### 1.4 APPROACH & CONTRIBUTIONS

In this thesis we propose a novel EM approach called *profile steering*. Profile steering is a decentralized approach, meaning that decisions on how and when to use flexibility are made locally. In the profile steering approach a top level controller uses steering signals to steer the use of flexibility provided by different devices. The flexible devices react to these steering signals by scheduling the use of their own flexibility. This is similar to current practice in most electricity grids, where consumers are offered a different tariff depending on their time of consumption (e.g., day and night tariff). However, our approach expands on current practice in two ways.

- » First, the steering signals are more general than energy prices, as we discuss in Chapter 3. This allows the central controller in our approach to communicate its goals towards the devices more accurately. The end result is that the obtained energy profiles better match the system goals.
- » Second, we implement two way communication, meaning that the flexible devices communicate their decisions, based on the received steering signals, back to the central controller. This is in contrast to current practice. We allow the central control to respond to the received decisions by updating the steering signals, causing the system to become iterative. Subsequent updating of steering signals and decisions of the devices is done until the result is satisfactory.

The profile steering approach we propose and study in this thesis fulfils all requirements we list for an EM approach. Furthermore, the general approach of steering signals we propose leads to energy profiles that generally better fit the system goals, e.g., reducing stress on grid assets. This leads us to believe that our profile steering approach can potentially be used in the (future) smart grid to facilitate the energy transition in a feasible and effective way.

Another problem tackled in this thesis is the aforementioned decision making of the devices. For profile steering (and other EM approaches) to be applicable,



devices need to be able to make decisions based on the steering signals they receive. For this we assume that devices locally are searching for their *optimal* decision with respect to the received signal and their local state. We show that, for a broad class of steering signals, the resulting problem for a lot of the devices is a convex optimization problem. Using techniques from convex optimization we formulate and study algorithms for the resulting device problems.

The results we obtain for the device level problem can be divided over three classes of devices. The first class of devices uses an internal buffer that only needs to be charged (e.g., an EV). The devices in this class can be modelled using a classical resource allocation problem. We apply results from literature to the continuous case of this model, i.e., when the energy consumption of the device is only limited by a lower and upper bound. We also consider the discrete case, i.e., the case that the energy consumption of the device is limited to a finite set of operational levels. While this model leads to an  $\mathcal{NP}$ -hard problem, we show that we can obtain good results with a minor modification combined with a greedy solution approach.

The second class of devices extends on the first in that the internal buffer can also be discharged (which is the case in, e.g., a stationary battery). We show that the corresponding optimization problem for the continuous case of this class can be solved using a divide and conquer approach, generalizing results found in literature for similar models. Furthermore, we extend the greedy approach used for the discrete case of the first class of devices to be applicable to the discrete case of the second class.

The third and final class of devices is an extension of the second class, which we obtain by adding losses that depend on the stored energy in the system. These losses play an important role in, e.g., heating and cooling systems. We extend the results obtained for both the discrete and continuous case of the second class to be applicable to the models we describe for the third class.

Finally, we show, through a simulation study, that using our approach to utilize locally produced energy as much as possible results in energy profiles that result also in minimal asset ageing. For this we use a model of transformer ageing and calculate how to best use the flexibility in a neighbourhood to minimize this ageing. The results show that the flattening of energy profiles through profile steering serves multiple system goals simultaneously, such as: minimization of transport losses, maximizing self-consumption, minimizing asset ageing, etc.

## 1.5 OUTLINE OF THE THESIS

The structure of this thesis is as follows. In Chapter 2 we provide some background for our approach and discuss related work. This chapter is used to give a more detailed overview of the current practice of energy management in the electricity grid resulting in a mathematical problem formulation.



In Chapter 3 we introduce the core concepts of the profile steering EM approach. Furthermore, we modify the basic approach such that it fits all the requirements on an EM approach listed in Chapter 2. In this chapter we also demonstrate the effectiveness of our approach through a simulation study.

Chapters 4 to 6 study models for various flexible devices. In Chapter 4 we study devices with an internal buffer that need to be charged (e.g., an EV). We extend this model to also include discharging in Chapter 5 (for, e.g., stationary batteries). A further extension is discussed in Chapter 6, where we include losses depending on the state of charge of the system. These losses play a crucial role in many heating and cooling systems. In each of these chapters we show the effectiveness of our approach by means of simulation studies, utilizing devices studied in the respective chapter.

Chapter 7 focusses on asset degradation and, in particular, on transformer ageing. This chapter studies a model of transformer ageing and discusses how an EM approach can minimize this. We compare the results of different EM approaches and study how they perform with respect to asset ageing.

The main part of this thesis concludes with a summary of the obtained results and conclusions in Chapter 8. This is followed by a discussion on these results, on the conclusions, and on future work. Finally, some mathematical background for the thesis is given in the Appendix.

Readers with a background in the field of energy management can consider skipping the background in Chapter 2. Furthermore, the concepts introduced in Chapter 3 are only required in Chapters 4 to 7 for a better understanding of the nature of the studied problems and the simulation results, hence these later chapters can be read without intimate knowledge of Chapter 3. The results in Chapter 5 build further on results obtained in Chapter 4. Furthermore, the results in Chapter 6 extend upon those in Chapter 5 and hence indirectly require knowledge from Chapter 4. In summary, Chapters 4 to 6 have a linear dependency. Chapter 7 can in theory be read as a standalone chapter, though some concepts from Chapters 3 and 4 are used in the simulation study presented there. The conclusion, given in Chapter 8, logically depends on the results from all other chapters. Finally, the concepts of convex optimization and complexity theory are heavily used throughout this thesis. We refer the reader to the Appendix for a brief introduction to these mathematical concepts. A visual representation of the dependency between the various chapters of this thesis is given in Figure 1.5

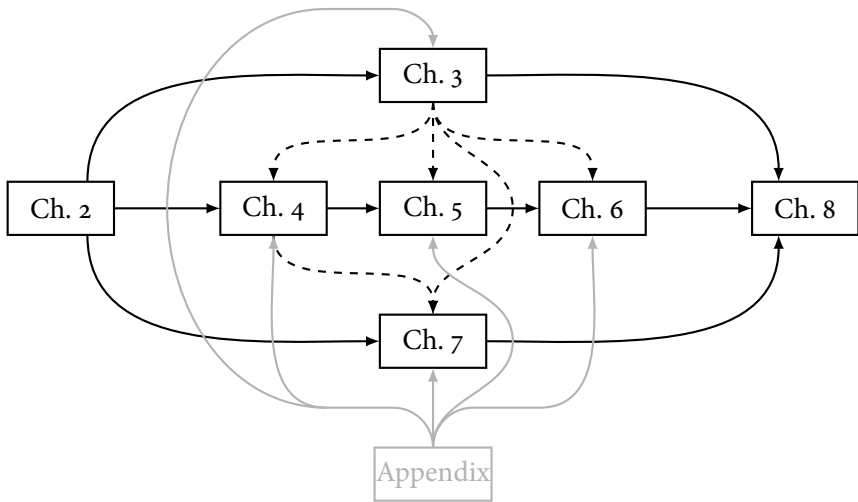


Figure 1.5: Flowchart depicting the dependencies between the various parts of this thesis.








# 2



## BACKGROUND

*ABSTRACT – In this chapter we discuss the changes we observe in our energy supply chain, with a particular focus on the management of the electricity grid. To do so, we begin by giving an overview of the original design of the grid and how it was controlled. The main changes happening in the system are the incorporation of energy from renewable sources in our system and the electrification of our energy use. Overall, these changes warrant a change in how we manage our energy supply chain, in particular the electricity grid, as the old centralized paradigm will no longer be applicable in the future. We discuss requirements on a future proof energy management approach of the electricity grid. We focus on the distribution grid and the distributed energy resources therein that are expected to play an important role in the future grid. Finally, we discuss some related work on energy management approaches that is relevant to the approach we introduce in this thesis.*






Figure 2.1: The share of electricity from renewable sources over the total electricity consumption in the Netherlands compared to the average in the EU in the period 2004 through 2014. Data taken from [51].




Figure 2.2: The SAIDI (System Average Interruption Duration Index) for the Dutch grid for the period of 1999 through 2013. Note that the data for planned interruptions is only available from 2006 onwards. Data taken from [35].




Figure 2.3: The unplanned SAIDI (System Average Interruption Duration Index) for the Dutch grid split per interruption caused on the different grid levels for the period of 1999 through 2013. Data taken from [35].

## 2.1 INTRODUCTION

Our energy supply chain is changing rapidly. The most important drive for the change is a desire to become independent of fossil fuels, which is an important goal as these fuels are linked to climate change, are finite, and often come from politically less stable regions. Thus, alternative sources of energy are considered more often, in particular energy from renewable sources such as wind and sun. The shift towards energy from renewable sources in the EU is depicted in Figure 2.1 though the Dutch share is well below the European average. Interestingly enough, many renewable sources were already exploited before the second industrial revolution driven by the electrification of our society, for example, the traditional Dutch windmills.

One of the major challenges of (most) renewable sources considered as alternatives to fossil fuels is that we cannot control them. This causes problems in the electricity supply chain where, in the traditional control paradigm: the generation follows the load. Such a paradigm is no longer applicable when a significant portion of the generation utilizes uncontrollable sources such as wind and sun. The traditional control paradigm has led to a very stable system. For example, consider the SAIDI index, an index used to indicate average interruption time of low voltage (*LV*) customers in a year, for the Netherlands, given in Figure 2.2. This index is in fact one of the lowest in Europe. i.e., the Netherlands has one of the more stable grids [35]. We further distinguish between interruption time caused by congestions on the various grid levels in Figure 2.3. This shows that a large share of the interruption time is caused by congestions in the medium voltage (*MV*) and *LV* distribution grid. Thus, in order to facilitate a smooth transition towards a society free from fossil fuels, a new control paradigm is required. A further change in the energy supply chain is a shift from large-scale central generation towards small-scale distributed generation (e.g., rooftop photovoltaic (*PV*) and residential scale combined heat and power (*CHP*) units).

This shift leads, as a side-effect, to a decrease in transportation losses due to a lower distance between generation and consumption. Also, local generation allows for easier use of by-products, e.g., heat, which would otherwise be wasted. This increases the overall efficiency of the system. Finally, access to locally generated energy increases the autonomy of local systems in many situations.

Next to the supply side of our energy supply chain also the demand side is changing. The increasing electrification of our energy use is the main influence, i.e., the share of our energy consumption through electricity is increasing. This change is largely motivated by the fact that most renewable sources only produce electricity. As an example, we consider the shift towards electric driving through electric vehicles (*EVs*) and plug-in hybrid electric vehicles (*PHEVs*). With sufficient electricity available from clean and renewable sources this shift significantly reduces the carbon footprint of our transportation sector.

The above sketched changes in our energy supply chain lead to complex challenges within these systems, in particular in the electricity supply chain, i.e., the electricity





grid. For this reason, we focus on the electricity system in this thesis. We tackle the aforementioned issue that the traditional centralized control paradigm can no longer be applied in the changing electricity system. As the systems of various other energy carriers (e.g., gas and heat networks) are intertwined with the electricity system on various levels, we do consider these other systems where applicable. As an example, generation of electricity and heat is often coupled, as in CHP units, which we study in Chapter 5.

This chapter serves as a background for the challenges considered in the remainder of this thesis. We first sketch the situation regarding the electricity grid as it was, currently is and is likely to become in the future. In particular we show that the current centralized control paradigm of the electricity grid is no longer valid. As a consequence we need a new approach and we outline the requirements on such an approach in Section 2.3. This leads to a formulation of the main problem considered in this thesis in Section 2.4. For this problem, many different appliances play a vital role, which we discuss in Section 2.5. In particular, we study the class of buffering devices, for which we introduce a general model. Then, in Section 2.6, we discuss the related work on the problem considered in this thesis. We wrap up with a conclusion in Section 2.7.

## 2.2 THE ELECTRICITY SYSTEM

Most electricity systems in the western world were designed decades ago using a central control paradigm. A large portion of these systems are now nearing the end of their predicted lifetime. Furthermore, due to the reasons outlined above, the control paradigm is changing. To this end we discuss the electricity system as it was designed, how it evolved to its current state, and what changes are envisioned for the future. The emphasis will be on the control part of the electricity grid. Only the details that are relevant for the control part are covered. For more information (on the Dutch grid) we refer the reader to [140].

### 2.2.1 TRADITIONAL CENTRALIZED SYSTEM

Our electricity supply chain was originally designed as a centralized system, where a small number of large-scale generation plants provide power to cover all demand. Such large plants benefit from the advantage of economy of scale. Because of the nature of the electricity system, i.e., storage of electricity is difficult and practically non-existent, supply and demand must be balanced at all times. To ensure this balance, the generators follow the demand using a central control paradigm. In this paradigm the demand is considered uncontrollable and the supply adapts to the demand. To fulfil the demand the generated electricity is transported from the plants via the electricity grid to the customer. The grid can be roughly subdivided into three levels based on the voltages used; high voltage (*HV*), medium voltage (*MV*), and low voltage (*LV*). A schematic overview of the grid and the different levels is given in Figure 2.4. While a higher voltage implies higher transportation

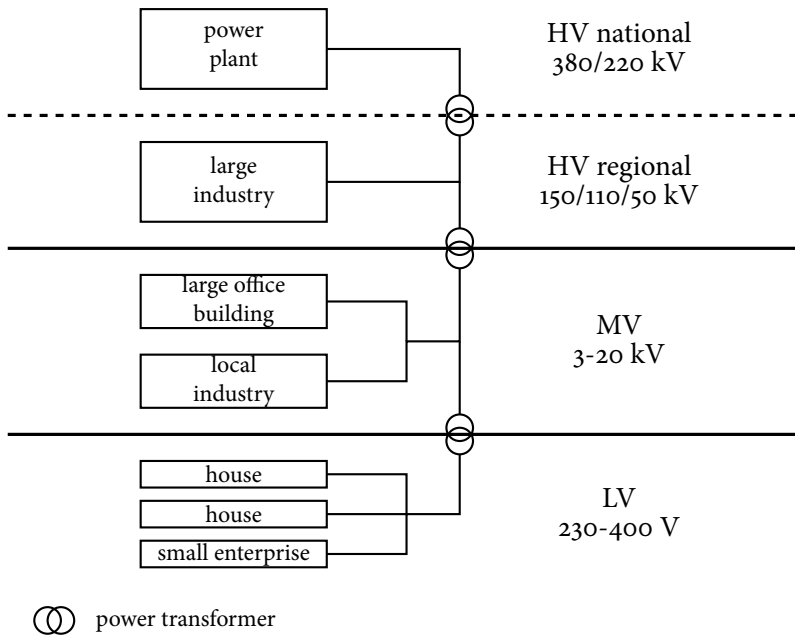


Figure 2.4: A schematic overview of the electricity grid in the Netherlands.

efficiency, i.e., reduced transportation losses relative to lower voltages, the downside is an increase in required hardware (in particular in their size and costs). For safety reasons lower voltages are used closer to the customers, particularly inside towns and cities.

The HV level is primarily used to transport electricity over longer distances, i.e., (trans)national and regional transportation. Conventional, large-scale generators are connected to this level. Furthermore, only a small number of very high load customers are connected directly to the HV level of the grid, for example aluminium smelters. Closer to the majority of the demand, the voltage is lowered using power transformers. The MV level is used to further distribute electricity within a demand area. Larger customers, such as large office buildings and local industry, are connected to this level. Finally, the LV level is used for the distribution to residential customers and small enterprises. The exact operating voltages of the grid vary between different countries. As an example we list the voltages used within the Dutch grid (see, e.g., [107]).

- » HV level: voltages of 380 and 220 kV are used for (inter)national transportation and voltages of 150, 110, and 50 kV are used for regional transportation.
- » MV level: voltages of 3-30 kV are used (local distribution and larger users).
- » LV level: voltages of 230-400 V are used (residential users and small enterprises).

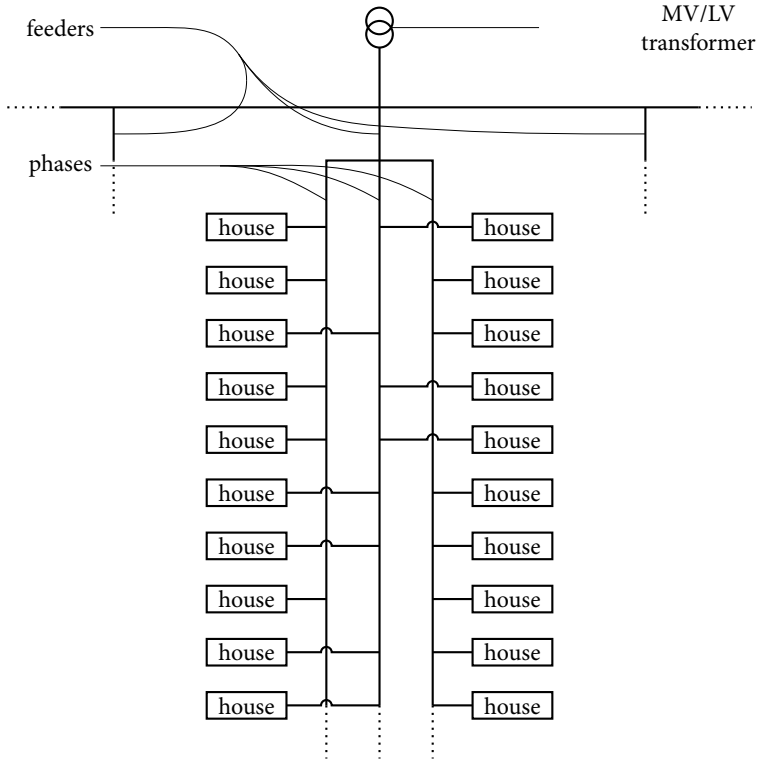


Figure 2.5: A schematic overview of a typical LV grid in the Netherlands.

In the previous paragraph we partitioned the electricity grid into three levels, based on the voltage used. Another often used partition of the grid is based on the distinction between the part that is used for transportation (HV) and the part that is used for distribution (MV and LV). This thesis primarily focusses on the management of the electricity produced and consumed by residential users. As nearly all residential customers are connected to the distribution grid, we focus on this part of the grid and in particular on the LV grid. The structure of LV grids in the Netherlands and many other countries generally follows the design we outline below. A transformer is used to change the voltage to the low voltage level. From this transformer several feeders run to the various areas, usually streets, supplied by the transformer. These feeders consist of four conductors; the three phases providing power and the neutral conductor. In the Netherlands, most existing residential connections are connected to a single phase of the feeder, with three phase connections primarily used for small enterprises with high power appliances and new residential connections. We note that the choice of the phase to which a house is connected is often random and not well documented. A schematic overview of a typical Dutch LV grid is given in Figure 2.5





Until the mid 1990's, the electricity supply chain was vertically integrated in most of the western world, i.e., in most countries a single company (state owned and/or regulated) owned the entire supply chain from the generators down to the customer connection. These companies were responsible for the entire supply chain, from generating the power to the delivery to and billing of the customers. This ended with legislation passed in both the US and Europe to split the ownership of the electricity supply chain. Since then, the ownership of generation, transmission and distribution assets is split in many countries, particularly in Europe. Retailers are now responsible for selling power to customers, which they buy from generators, owned by different companies. The transmission and distribution parts of the grid are owned and operated by transmission system operators (*TSOs*) and distribution system operators (*DSOs*) respectively. These system operators facilitate the interconnection between generation and customer such that the energy sold by a retailer to a customer can be delivered. Because asset ownership of the electricity grid causes a natural monopoly [102], *TSOs* and *DSOs* are generally state regulated. We note that, while the distinction between retailers and grid operators is often clear in Europe, this is not the case in the US.

To facilitate the energy trading between energy producers and retailers, several markets exist. On these markets electricity production and consumption can be traded on various time scales, ranging from long term contracts (i.e., months in advance) to short term (i.e., day ahead or intra-day). The retailers forecast their energy demand and purchase electricity using these markets. The exact schedules for the energy generation plants are made by the energy producers (or the system operator in an integrated system), typically a day ahead, solving so called unit commitment problems (*UCPs*) and related problems [22, 145]. The system operators are responsible for the balance between supply and demand. They ensure this balance by means of spinning reserve of online generators, i.e., the capability of a running generator to quickly adjust its production in response to fluctuations in demand. This spinning reserve is traded through so called capacity markets.

The safe operation of the transportation grid requires monitoring and management of the grid by the operators (the *TSOs*). To this end, the transportation grid generally has measurement equipment in place to monitor the state of the grid and ensure it is operated within safe margins. On the other hand, the distribution grid, and in particular the LV grid, is typically managed using a fit and forget strategy by the *DSOs*. Cables and other assets are dimensioned using a forecast on future expected required capacities upon installation and are assumed to operate within the boundaries without active monitoring and management. The main reason for this paradigm is the large number of customers typically connected to distribution grid assets, causing load diversification due to the law of large numbers. This means that demand profiles seen by the grid assets are usually smooth and predictable because of the large number of customers connected below an asset. As an example, we compare the load profile, measured every 5 minutes, of a house and a neighbourhood transformer in Figure 2.6. In the figure we give, for both the house and the transformer, data from two days with a week in between. While the two profiles




Figure 2.6: Comparison of the load profiles of two days for a house (a) and a transformer (b). The data was provided by Alliander and consists of two days with a week in between.

given for the house are quite different, the neighbourhood profiles are roughly the same. Furthermore, because generators were traditionally only connected to the top level of the grid, power was flowing unidirectionally downstream, i.e., from the power plants through the HV, MV and eventually the LV parts of the grid to the customer. The above no longer applies in several cases due to the new emerging trends in the electricity grid. For example, the production peaks of PV installed on houses in the same neighbourhood coincides. Also, some heavy loads such as heat pumps can become synchronized, causing a load that is too large to handle for the (local) grid.

Due to the mostly passive role of customers, specifically of residential customers, the interaction for most customers with the players in the electricity system is limited to a single supplier. While originally this was (part of) the, typically state owned, electricity company responsible for the entire infrastructure, this changed to independent suppliers after the market liberalization in the 1990's. Information about the customers electricity use is measured ('metered') locally and collected sparingly for billing purposes, e.g., once a year. The electricity bill paid by cus-



tomers for their consumed electricity is based on the collected measurement data. While small, residential customers pay a fixed price per unit of consumed energy, larger customers also get billed for, e.g., their maximum consumption peak in the considered interval. In other words, the energy tariff structure generally increases in complexity for larger customers. We note that more complex tariff structures increase the requirements on the measurement on when and how much energy is consumed by the customer.

### 2.2.2 CURRENT SYSTEM

Caused by a drive towards a de-carbonization of our energy supply chain, the number of clean, renewable sources has rapidly increased in many parts of the western world. Furthermore, these generators of renewable energy are for a large part installed on a smaller scale, closer to the customer (see, e.g., the data available for the US [134]). The installation of these units is far from uniform across the grid, due to several factors, including social-economic, legislative, and geographical reasons. This non-uniform spread of the renewable generators causes excessive stress on certain parts of the distribution grid, e.g., parts of the LV grids with large amounts of residential PV and low consumption are already being overloaded in Germany [103] (see also Figure 1.4 on page 6). In such cases the local generation peak is far larger than the demand in the area at the same time. This causes a large inverted power flow upstream, i.e., in the opposite direction than originally intended during system design. This peak may overload the transformer. The current practice is to upgrade grids where and when this occurs, allowing the surplus of energy to be transported through the MV and HV parts of the grid to areas with a net demand. However, this practice is costly and becomes infeasible when the overall penetration of these (renewable) resources at the customer level increases [106].

Another problem with the adoption of renewable generation is that a large share of these generators are based on uncontrollable sources such as wind and sun. Without the proper mix of different available technologies and suitable control, a gap between supply and demand is likely to exist for a small portion of the year (e.g., a few weeks) [103]. To be able to supply energy in these rare periods, backup plants need to be kept operational, just for these periods, which is unfavourable for their profitability. Without a change in the control paradigm, this can cause a large increase in operational costs of the system as well as hamper the transition to a system based purely on clean, renewable sources [22, 145].

The energy use of consumers follows the aforementioned trend of electrification; an increasing share of our total energy consumption is in the form of electricity. In particular large and novel loads such as EVs and heat pumps (*HPs*) are introduced. The electricity grid was not designed with these new, large loads in mind. Therefore, the introduction of these loads causes the system to be operated closer to or at the limits, implying that safe and secure operation of the grid is becoming increasingly difficult. To ensure the safe and secure operation, investments in additional infras-



structure or other approaches that can ensure a safe introduction of the increasing demand in the grid are required [30, 67, 112].

For the residential customer, a recent change in the system is the introduction of smart metering. The traditional electro-mechanical measurement equipment used for metering and billing is being replaced by more advanced sensory equipment (see, e.g., [52]). These smart meters introduce two different types of opportunities. First, these meters collect data with finer granularity, allowing the customers a higher level of insight in their temporal energy use. This, consequently, opens up possibilities for saving costs through reduction of (unnecessary) energy use. Furthermore, because these smart meters are also equipped with more advanced communication infrastructure, the data they collect can give the retailer, and potentially other third parties, a deeper insight in the behaviour of its customers (after customer consent is given). The more advanced measuring and communication capabilities also allow more intricate tariff structures to be implemented. As an example, the meter can be used to communicate time-varying prices to a customer a day ahead of their use. From the beginning of 2017 limited use of these time-varying prices is possible in the Netherlands [79].

### 2.2.3 FUTURE SYSTEM

The trends mentioned above are expected to continue in the future, even at an increasing rate as technologies in these areas are adopted and become cheaper and more widely available [103]. To ensure a stable and secure energy supply, two trends are emerging. The first is expansion of transportation capacity and interconnection between various regions. This change is especially visible in Europe where many countries strive for a higher level of interconnectivity with their neighbours [62]. This higher level of interconnectivity ensures that local surpluses or shortages of energy can be more easily solved by exporting or importing energy to/from another area.

The second notable trend is a drive towards local balancing of energy. With an increase of local production, such as rooftop PV, it makes sense to balance this with local consumption to reduce losses and stress on the grid. This leads to the emergence of so-called micro-grids; parts of the electricity grid that balance consumption and production locally as much as possible and, in some cases, can even disconnect (for some period) from the main grid, particularly in case of contingencies upstream [63].

Both trends have to do with reduced flexibility on the generation side of the electricity supply chain, which is expected to continue decreasing in the future. Thus, flexibility needs to be found elsewhere if we want to ensure a safe and economic operation of our (future) grid. There is a general consensus that to ensure safe and economic operation, the grid needs to be updated to become a *smart grid*.

Recently, also flexibility on the consumption side of the supply chain is considered as an alternative to lost flexibility on the generation side. Such flexibility can come



in the form of shifting consumption or even shedding of loads. The latter generally causes a large loss of user comfort and hence is only used in emergency situations, e.g., during temporary grid congestion. This flexibility can be used to, e.g., better match the demand to the available supply or alleviate grid congestions in cases of stress (situations where grid assets are operated near or at their physical limits). The management of loads, and potentially local production assets, to shift or shed their electricity consumption/production is commonly referred to as demand side management (*DSM*). We note that some authors use the term demand response for the management of loads instead. *DSM* was historically introduced for larger electricity consumers [129]. Such consumers are financially compensated to decrease their load in case of grid congestions such as large imbalances between generation and consumption.

To facilitate the introduction of *DSM* on a residential scale, (external) control over the loads on the distribution grid level is required. Traditionally this control is not present for most of the devices used by residential customers, beyond the user manually turning the device on or off. However, new devices that have an external control option are being introduced into the market under the term smart appliances. With the increasing penetration of these smart appliances in the residential sector, as well as advancements in information and communication technologies (*ICT*), *DSM* on a residential scale is becoming increasingly feasible. Due to the relatively small scale of the loads on this level, a large number of these loads need to be jointly managed to have a significant impact on the overall load profile of a distribution grid. The assets considered for residential scale *DSM* are part of the class of distributed energy resources (*DERs*). This class of devices is not limited to loads and includes all devices that can offer some form of flexibility, e.g., also batteries and PV installations with controllable inverters.

To effectively manage the (future) grid with a large number of *DERs* present, the old, centralized control paradigm is no longer applicable for two reasons. The first reason is that the centralized paradigm does not scale to a large number of assets. For example, the problems solved to find generation schedules (*UCPs*) become intractable for a large number of assets [21]. Furthermore, the paradigm is not designed to handle the large diversity of assets envisioned in the future (smart) grid. Part of the reason for this is that, in the old paradigm, assets at the customer side, specifically those in the distribution grid, are considered uncontrollable. However, many devices are currently being equipped with some form of local control, e.g., smart washing machine (*WM*) and dish washer (*DW*).

To this end we study, in this thesis, a new approach to manage all types of *DERs* that emerge in the (smart) grid. We call an approach that ensures the safe and effective management of all (future) *DERs* in the grid an energy management (*EM*) approach. We note that, while the approach is not limited to this case, we primarily focus on residential grids and the management of *DERs* that already are present or are envisioned to become available in these grids. To study such *EM* approaches, we first study what the requirements for these approaches are.

## 2.3 REQUIREMENTS FOR ENERGY MANAGEMENT APPROACHES

An EM approach needs to manage the flexibility offered by the various DERs present in the grid effectively, from both a control and economic perspective. The type of problems that can potentially be tackled by EM approaches form a very large and diverse landscape. As an example, when considering the balance between consumption and generation, which needs to be managed on many different time scales from seasonal to sub-second balance. Ensuring a balance between supply and demand on an inter-day scale is closely related to system planning, while sub-second balance is generally handled via droop control of spinning reserves or control of batteries.

We note that EM shows similarities with the recent topic called transactive energy (TE) [60]. Within TE a framework is being proposed that facilitates an effective integration of various (future) DERs. The main difference between EM and TE, as we understand it, is that the framework considered for TE focusses mainly on markets. Within these markets the proposed TE framework is designed to ensure value discovery of the services offered by DERs. In other words, the framework largely focusses on the financial benefits of these services and their inclusion in traditional energy markets. However, the total value of certain services offered by DERs for the entire system is often hard to gauge. Furthermore, different stakeholders in the energy supply chain are sometimes excluded from working together due to current regulatory structures [102, 106]. Therefore, we focus on the technical aspects of an EM approach. Thus, the requirements we formulate below are of a technical nature, while the required economical and regulatory means to effectively implement such a system are outside the scope of this thesis.

In this thesis we focus on an EM approach that works on time scales of minutes to days, the time-scale that is also traditionally used in energy and ancillary markets. Nevertheless, sub-second control and system expansion planning are also of importance. We envision that problems in both sub-second control and system expansion planning can be tackled by systems that run concurrently to our approach, potentially in an integrated manner.

The requirements for an EM approach are: the system can

- » Scale to the many DERs present in the (future) smart grid (*scalability*).
- » Use flexibility offered by DERs at the right time (*scheduling*).
- » Handle the many different types of DERs (*heterogeneity*).
- » Ensure the flexibility of the DERs is used within the physical limitations of the system (*feasibility*).
- » Keep privacy sensitive information as local as possible (*privacy*).

Note that these requirements largely correspond with the technical requirements set for TE architectures [60]. In the following we elaborate on the individual issues.





### Scalability

As mentioned before, the classical centralized approaches used to manage the electricity grid, such as the operational planning of generation, do not scale to a large number of assets. This conflicts with the expected large number of DERs in the future (smart) grid, particularly in residential grids. Thus, to ensure feasible operation, new EM approaches are required that scale to thousands, or, in some cases even, millions of DERs. Since the underlying mathematical problems are often hard and intractable to solve to optimality for such a large number of devices, heuristics can offer a solution.

### Scheduling

An important aspect of EM approaches is that they use the flexibility offered by the various DERs at the most opportune time. As an example, it is beneficial for the system to run a smart WM such that its energy consumption coincides with the (local) PV production peak. To be able to achieve this, the approach needs to look ahead and *schedule* the flexibility use of the DERs. Hence scheduling is an integral part of a proper EM approach (see also Section 1.2).

To properly schedule the use of flexibility, *predictions* of the system parameters are required. For example, to properly schedule that the use of the WM coincides with the PV production peak, a prediction of when this peak occurs is required. Furthermore, a prediction of when the WM is available is needed together with an estimate of the consumption of other devices to ensure feasible and optimal operation. Since perfect predictions are generally impossible to make in the EM setting, especially on a local level, proper control is required to adjust and implement the schedules to minimize errors and avoid infeasibilities [6]. We call such control *operational control*. The above results in an approach that consists of three steps or phases; predictions, planning and (operational) control.

### Heterogeneity

In more traditional EM approaches, the energy resources considered are typically of a similar nature. For example, in UCPs the schedules for the generators typically have to satisfy similar constraints. This no longer applies in the future smart grid where a large variety of DERs is present; an EM approach needs to be able to deal with a large diversity (in size, constraints, reaction time, granularity, etc.) in the available resources. It is important to note that, especially on the residential level, many DER assets have to serve the user's comfort as primary objective. The DER asset in question is not allowed to sacrifice this comfort, or only up to a level specified by the user, to achieve the system goals. This increases the complexity of the problem significantly [28]. We discuss the diversity of the DERs expected in the future (smart) grid in more detail in Section 2.5.



### *Feasibility*

The primary function of many DERs present in the distribution grid is generally something other than supporting the electricity grid, e.g., an EV's primary purpose is transportation. The constraints and preferences of the user can deviate from what is optimal from a grid perspective. For example, the owner of the EV might want the vehicle to charge as fast as possible upon arrival, which may increase the evening consumption peak, putting extra stress on the grid. To ensure the safe and feasible operation of the grid and all the DERs connected to it, two types of constraints need to be taken into account. The first type of constraints are those set by the user, e.g., the vehicle has to be charged before hour  $t$ , together with the technical constraints, e.g., the vehicle cannot charge at a higher or lower rate than a given value. As these constraints only affect the usage of the flexibility of the DER for which they are specified we call these the *local constraints*.

On the other hand, the constraints put on the overall use of (groups of) DERs by the grid, e.g., maximum aggregated charging of a group of EVs to prevent overloading of cables, are called *grid constraints*. Note that we focus on constraints on the aggregated energy consumption/production of groups of devices in this thesis, as these can be used to model capacity constraints. These capacity constraints are currently already being violated in several grids, e.g., in Germany (see [106]), and are expected to be violated more often in the future [67] with the higher penetration of (large) domestic electric loads. A feasible EM approach thus needs to consider the local constraints when scheduling the flexibility of each individual DER together with grid constraints on groups of schedules of the DERs.

As noted in Section 2.1, currently the electricity grid is very reliable. To ensure the system remains as reliable during the transition towards energy from renewable sources it is of particular importance that new EM approaches respect all constraints.

### *Privacy*

Many parameters detailing the flexibility of DERs include privacy sensitive information. For example, the arrival and departure times of an EV provide information on the occupancy of a home. Because of this privacy issue, the approach should keep this information as close as possible to the user.

## 2.4 FORMULATION OF THE ENERGY MANAGEMENT PROBLEM

The core task of an EM approach is to solve a scheduling problem that specifies how to use the flexibility offered by the various DERs present in the grid. In order to create such schedules, we mentioned that predictions and operational control are required. As this thesis focusses on the problem of deriving schedules for the use of flexibility of the DERs, we assume that predictions are available that are sufficiently accurate. Furthermore, we assume that a suitable operational control approach





is in place to realize the generated schedules. In the following we consider the scheduling problem from a mathematical point of view.

The schedules we consider are for a given time horizon (e.g., a day) consisting of equal length time intervals (e.g., fifteen minutes). Although the length of the time intervals can be chosen arbitrarily, we focus on the scale of minutes or longer. The main reason is that, with a granularity below that of minutes, certain system dynamics start to play crucial roles in the application of an EM approach. For example, the reliability and speed of the communication infrastructure becomes important when schedules are made and/or readjusted within seconds. Because these dynamics are outside the scope of this thesis, we focus on a time granularity where such issues do not play an important role. Note that a time scale of minutes to hours also coincides with most market structures used to operate the electricity grid [102].

To introduce the scheduling problem more formally, we consider a given time horizon  $\mathcal{T}$  consisting of time intervals  $1, 2, \dots, T$ . An EM approach schedules the flexibility use of a DER  $m$ , resulting in an energy profile  $\mathbf{x}^m$  of  $T$  elements, for each of the available DERs in a set  $\mathcal{M} = \{1, 2, \dots, M\}$ . The goal of an EM approach is to determine schedules that are optimal with respect to the given system goals while the resulting energy profiles are feasible. To this end we define  $\mathbf{x} = (\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^M)$  as the vector of the energy profiles of all devices. The system goals are expressed using an arbitrary objective function  $f(\mathbf{x})$ . This function indicates how well the energy profiles resulting from a particular set of schedules fit the system goal. As an example, using  $f(\mathbf{x}) = \|\sum_{m=1}^M \mathbf{x}^m\|_2$  ensures schedules are made for the devices that flatten the aggregated energy profile, and as such reduce stress on the network.

When scheduling the flexibility provided by DERs, especially those in a residential setting, we mentioned that two types of constraints have to be considered: local and grid constraints. To ensure that the local constraints are satisfied, we use a local constraint set  $X^m$  for each device  $m$ . Thus we require that  $\mathbf{x}^m \in X^m$  for every device  $m \in \mathcal{M}$ . Grid feasibility is ensured through the use of a constraint set  $X$  that applies to the aggregated energy profiles resulting from the schedules together. In other words, we require that  $\mathbf{x} \in X$ .

From the above we obtain the following EM scheduling (EMS) problem:

**Problem 2.1 (EMS).**

$$\min_{\mathbf{x}} f(\mathbf{x}), \quad (2.1)$$

$$s.t. \mathbf{x}^m \in X^m \quad \forall m \in \mathcal{M}, \quad (2.2)$$

$$\mathbf{x} \in X. \quad (2.3)$$

This problem is central in scheduling based EM approaches. As an example, the classical UCP solved to obtain generation schedules for the central generators is



obtained from Problem *EMS* in the following manner. Take  $\mathcal{M}$  as the set of generators and  $f$  the sum of the cost of generation according to profile  $\mathbf{x}^m$  over each  $m$ . Furthermore, take  $X^m$  as the feasible set for generator  $m$  and  $X$  as the feasible set for the network, i.e., the set of schedules that satisfy network constraints while matching the load.

When considering schedules in the EM domain, we note a particular subtlety that is often overlooked. A schedule results in a profile given in terms of either power or energy. While technically the power drawn by a device can differ greatly over a considered time interval, a common assumption is that this does not occur. In other words we assume that the power consumption/production is (near) constant per time interval and equal to the average power consumption/production. Using this assumption power and energy values can be obtained from one another through multiplication with an appropriate factor. In this thesis we prefer the use of energy values and assume values given in terms of power have been transformed using the appropriate factor. For example, we assume a constraint on the maximum power a device can draw from the grid in a time interval is transformed into the maximum amount of energy the device can draw from the grid in this time interval.

As mentioned before, the traditional methods for solving the UCP and other variants of Problem *EMS* do not scale to a large number of devices/DERs. Thus, a new approach is required. A popular approach that has been much discussed recently is decentralized energy management (*DEM*), where (part of) the computation is decentralized to ensure scalability. For example, the devices (or their users) make their own schedules based on steering signals received from a central controller. These steering signals give information on how desirable each possible scheduling decision is. Note that in practice this approach is already applied in many countries, by means of a day and night tariff. This price signal incentivizes customers to shift their consumption to the night when the cheaper tariff applies and when energy is generally cheaper for the retailer too. We come back to the use of a *DEM* approach in the next chapter, where we introduce our own scheduling based *DEM* approach called profile steering.

## 2.5 DIFFERENT TYPES OF RESIDENTIAL DISTRIBUTED ENERGY RESOURCES

We already mentioned that one of the challenges for an EM approach is the heterogeneity of the many DERs envisioned for the smart grid. A non-exhaustive list of potential DERs on the residential scale is: EVs, white goods, (electric) heat generation combined with storage, rooftop PV, residential-scale generators, heating, ventilation, and air conditioning (*HVAC*) systems, batteries, and pool pumps. A proper EM approach exploits the specific flexibility offered by each of these devices. To be able to facilitate the proper interaction between such an EM approach and the actual devices a platform is required that facilitates this communication. An example of such a platform is the Energy Flexibility Platform and interface (*EF-Pi*)

[55] (which was previously known as FPAI).

### 2.5.1 EF-Pi DEVICES

Within the EF-Pi platform DERs are divided in four different *appliance classes*:

- » uncontrollable,
- » time-shiftable,
- » buffer,
- » unconstrained.

These classes are based on the type of flexibility in the energy use of these devices. Below we elaborate on each class.

#### *Uncontrollable*

Uncontrollable devices are those devices of which the energy consumption and/or production cannot be altered. This class encompasses nearly all devices currently found in residential areas. Examples are televisions, lighting, and (rooftop) PV. While these devices do not offer flexibility to be used by an EM approach they still play an important role via their energy profile (this profile is sometimes referred to as the base load). As the energy profile of these devices cannot be altered, the other devices that do offer flexibility need to schedule their flexible energy use around the energy use of the uncontrollable devices.

As a subclass of the uncontrollable devices we have the *curtailable* devices. These devices can be steered to consume/produce less energy or in some cases even switch off at certain times. An example is PV with a controllable inverter. These inverters can reduce the output of the PV system in case of, e.g., over voltages or severe frequency deviations in the (distribution) grid. Such inverters are already required in Germany and California, due to significant levels of PV penetration [88]. However, curtailing such devices is often undesirable from either an environmental and/or a user comfort point of view. Therefore, such curtailment should often only be used as a last resort. Because of this, we do not consider curtailable devices in this thesis. However, we believe that many curtailable devices can be incorporated in the approach discussed in this thesis with minor adoptions to the presented approach.

#### *Time-shiftable*

Time-shiftable devices are devices with a fixed energy profile that can be shifted in time by changing their start time. Examples are pool pumps and smart white goods, such as WMs and DWs. These devices can shift the start of their program in time to a more opportune time, e.g., when more solar energy is available. The flexibility of these devices is generally characterized by a time window in which the devices' program has to be started and an energy profile of the program to be run. Some of these devices also offer the opportunity to temporarily interrupt their program at



specific points, thereby offering a higher level of flexibility. The scheduling decision for these devices consists of a time to start the program and potentially a set of decisions on where and how long to interrupt the program.

### *Buffer*

This class consists of two types of devices: 1) pure storage devices, e.g., batteries and heat vessels and 2) devices that utilize internal storage for their primary functions, e.g., EVs and laptops. Note that a device of the latter type can be obtained when combining a pure storage with another device, e.g., a heat pump or CHP unit combined with a heat vessel. The internal storage of the latter type of devices allows the energy consumption/production to be decoupled from the use of the device for its primary function, e.g., using a heat vessel allows for the decoupling of the production of heat and electricity by a CHP unit from the actual heat use. Typically, the devices considered here consume a relatively large amount of energy and can therefore have a large impact on the distribution grid.

The constraints on such devices is that enough energy has to be present in the (internal) buffer for the desired primary operations, allowing flexibility usually in both when and how much energy is consumed/produced. Thus, the scheduling decision for these devices is often how much energy is to be consumed/produced by the device for every time interval. This implies that a large amount of flexibility is offered. Combining this with the typical large load of such devices implies that most devices in this class should be an integral part of any effective future EM approach. To this end we discuss a model for such devices in the next section. On the other hand, the scheduling problem for many of these devices, as we discuss in Chapters 4, 5, and 6, are typically more difficult to solve than those considered for time-shiftable devices.

### *Unconstrained*

Unconstrained devices are those devices without user specified constraints, i.e., the primary purpose of such devices is to support the grid within their own technical constraints. The typical example for this class of devices is a (residential-scale) generator. Thus the primary goals of such devices is to provide the best energy profile for the grid. As such devices are both rare on a residential scale and the control of most of these devices on a larger scale is already covered by traditional EM approaches, we do not consider this class of devices in this thesis.

#### 2.5.2 MODEL FOR BUFFERING DEVICES

In the previous section we mentioned that the class of buffering devices plays an important role in future EM approaches. In the following we introduce a model for the buffer inside such devices. The buffer itself is used to store energy, where the *internal state* of the buffer gives the total amount of energy currently stored. This amount of energy is called the state of charge (*SoC*) and is usually given as the total




Figure 2.7: A schematic overview the buffer model used in this thesis.

value stored (e.g., 4 kWh) or the fraction of the maximum storage capacity used without a unit (e.g., 1.0 or 100% if the device is full).

The SoC of the buffer depends on both an *inflow process* and an *outflow process*. Via the inflow process additional energy is stored inside the buffer while energy is taken from the buffer through the outflow process. Depending on the exact device in question one or both of these processes are (partially) controllable. Note that, if neither of the two is controllable, the device belongs to the class of uncontrollable devices instead. A schematic overview of the buffer model used in this thesis is given in Figure 2.7.

As an example we consider an EV. The outflow process is assumed to be uncontrollable as the required energy for driving depends on the required trips, and road and traffic conditions which we assume we cannot influence. The best we can do is to predict the outflow. On the other hand, the inflow process in an EV is the charging process. We assume that this process is controllable by controlling the amount of energy that is charged into the EV in every time interval considered.

## 2.6 RELATED WORK

It appears that the need for new approaches in EM, and thus studies in this area, has only arisen in the last ten to fifteen years due to an increasing interest in clean, renewable energy and the emergence of enabling technologies in, e.g., communication technology. The number of publications in this area increased drastically in the last five years [142, Fig 2]. However, the first studies and discussions on this subject date back earlier than this, with the first being published in 1980 to the best of our knowledge [118]. In this work, Schweppe *et al.* argue that economic principles can be used to balance supply and demand between individual customers on a minute to hour scale. Furthermore, they propose to ensure the shorter term balance by responding to frequency deviations with eligible devices, potentially (partially) displacing the need for conventional spinning reserve. The marketing interface to customers, introduced to facilitate the billing of customers under the proposed paradigm, is very similar to smart meters currently deployed in the system.



Many concepts similar to the one introduced in 1980 have since been covered, as evidenced by the many surveys (recently) published on EM (see, e.g., [6, 15, 39, 120, 142]). Similar to the energy marketplace proposed by Schweppe *et al.*, several authors study the effect of implementing more intricate tariff structures for (smaller) customers. A popular structure is time-of-use (*ToU*) pricing [125, 147]. In *ToU* pricing the retailer broadcasts prices for different time intervals during the day, e.g., fifteen minute or hourly intervals, at predetermined times (e.g., a day or an hour in advance) to the customer. This communication can be facilitated by, e.g., the smart meter. Generally, these prices are assumed to reflect the prices seen by the retailers on the electricity market, thus allowing the retailer to better reflect its actual costs to the consumer. The differences in prices between time intervals incentivizes customers to shift their consumption from intervals with high prices to intervals with low prices, e.g., to better match supply and demand. Another possibility is that a new market player, often called an aggregator, exploits price differences on the market through the use of the flexibility provided by its customers. The exact role and potential of such a player is, however, often unclear and could be hampered by regulations [77, 137].

As the communication between supplier, or potentially the grid operator, and the customer is one-way in such pricing schemes, the response to the broadcasted prices can only be estimated on before hand. This can be done through price elasticity models, which determine the (expected) responsiveness to price fluctuations by customers [11]. It has been noted however that, when sufficient flexibility is available, *ToU* pricing (or other linear pricing mechanisms) potentially only shift the peak instead of shaving it [15, 91]. This means that the highest consumption is shifted in time to the interval with the lowest price without a significant reduction of the actual peak. Thus, the actual stress put on the grid is potentially hardly reduced by these means. Added to this, *ToU* pricing can lead to power quality issues (due to synchronous behaviour) [66]. As an alternative, differentiated pricing schemes are considered, which send different prices to different customers [20, 95]. While such methods can be effective, there are potential ethical and legislative reasons against such systems.

The above mentioned potential issues of linear pricing, i.e., pricing per consumed kWh, can be alleviated by using tariff structures that better reflect the actual costs of generation and transportation. As an example, Mohsenian-Rad *et al.* [94] consider a setting in which the system is optimized towards minimizing actual generation costs. They formulate the resulting problem as a non-cooperative game and show that a Nash equilibrium exists that is equal to the system optimum. Furthermore, they show that, with the proper incentives for the individual users, the system wide optimum is obtained in this game theoretic setting, i.e., all users are individually incentivized to take actions that benefit the overall system. Similar conclusions are also reached in [10]. Next to the mentioned structures, also other tariff structures are considered in literature (see, e.g., the discussion on price-based programs in [39]).



Another popular approach is to directly control devices on the demand side. In such an approach the use of flexibility of the available devices is generally controlled through the use of optimization, i.e., a controller decides how to best use the available flexibility for a given set of objectives. Traditionally, such control is reserved for large customers and used in emergency situations, i.e., during grid congestion [129]. However, with recent advancements in ICT, such as the introduction of smart meters, direct load control on a residential scale becomes increasingly viable [46, 119]. In these traditional approaches, larger non-critical loads can be deferred during times of congestion.

However, with current technologies, such approaches do not need to be limited to load deferral during grid congestion, but can be implemented in a broader EM approach designed to operate the grid more efficiently and economically as a whole. However, as mentioned before, dealing with this problem in a central way leads to a problem that is too large and complex to be tackled by the traditional approaches when going beyond the scale of a house or a few large generators. In other words, it is not feasible from a computational perspective to centrally gather knowledge of the flexibility provided by every available DER and compute a schedule of the flexibility use in this central controller followed by a dispatch of the computed schedules to the DERs. Furthermore, the data required by such a central approach is often privacy sensitive. To tackle these issues, heuristics and approximations are required to ensure that a solution that is good enough can be obtained within reasonable time.

One approach to the problem is to decompose it into several sub-problems that can be solved individually. For example, Toersche *et al.* [127] decompose the problem using a nested Dantzig-Wolfe decomposition resulting in tractable sub-problems that can in practice be solved concurrently. The authors note a performance degradation when using smaller group sizes, i.e., when the problem is decomposed into a larger number of sub-problems, due to overhead of the remaining central problem. This work is part of a larger set of publications on a DEM approach called Triana [14, 20, 95, 96].

The Triana approach consists of three steps; prediction, planning, and operational control, mimicking the structure we outlined for a scheduling-based approach in Section 2.3. The prediction step is done on a household level for each of the relevant parameters, e.g., uncontrollable load, available flexibility, and generation from renewables. These parameters serve as an input for the neighbourhood level planning. The scalability of the planning step is ensured through the aforementioned decomposition. The decomposition is set up in such a way that the problem and subproblems form a tree [126]. The scheduling problem, present at the root of the tree as a master node, is solved by iteratively cascading price signals downwards through the tree to the leaves, which represent the various flexible devices. These devices schedule the use of their flexibility based on the received price signal and communicate the scheduled profile upwards. On each intermediate level, a selection is made between the different schedules each device has communicated, such



that the overall profile best fits the received price signals on this level.

One of the downsides of such an approach is that, in order for the approach to work in practice, commonly a 'copper plate' is assumed for the distribution grid present in the neighbourhood. This assumption means that, regardless of individual consumption patterns of the houses, as long as the neighbourhood profile stays between the operational bounds of the transformer, no problems occur elsewhere in the network. This assumption turns out to be false in practice in many scenarios [66, 67] and is sometimes even amplified by the fact that the approach uses price signals. Such price signals cause the devices to choose between either consuming as much energy as possible, if the price is low enough, or not consuming any energy, if the price is too high. As a consequence the devices never operate using a state that consumes a moderate amount of energy (e.g., an EV will never charge with less than maximal power). While grid constraints can be taken into account during the planning phase [66], the scheduled profiles still show significant fluctuations on a local level due to the price signals.

Furthermore, such an approach relies heavily on reasonably accurate predictions as input for the planning. Toersche shows [126] that it is possible to use an auction structure, similar to the PowerMatcher we discuss below, to determine requests of when and how devices should use their flexibility. To ensure that the requests match the available flexibility, even in case of predictions errors, the system uses simulations to predict possible outcomes and uses these to determine its control policy. This leads to an effective, though somewhat computationally heavy, approach. The approach can be combined with the aforementioned (decomposed) planning approach. Toersche [126] also outlined the profile steering approach in his thesis, which relies on the same fundamental ideas as the EM approach we study in this thesis. However, we extend the approach to include more general steering signals than just desired profiles (see Chapter 3).

The PowerMatcher [83] approach avoids the extreme on/off reactions when sending (linear) prices, by using an auction structure where devices bid onto the market by giving a consumption or production level for each price. This approach is very similar to the third step in the Triana approach, i.e., the operational control step. For example, an EV provides a curve that relates its power consumption, i.e., its charging level, for the current time interval to the price of electricity. Similarly, producing devices provide an energy production level for each possible price. The bids of all consumers and producers are then aggregated separately into two curves giving the total consumption and production for each possible price respectively. Placing mild restrictions on the individual bids, the two aggregated curves always intersect. The price at this intersection is called the market clearing price. This market clearing price is sent to the devices, which can then determine their corresponding consumption/production for the coming interval by evaluating their bid at the received market price.

A downside of the PowerMatcher approach is that it only considers the current time interval. This leads to the unwanted effect that the flexibility of devices is





sometimes already used up before the real problems appears, e.g., a large peak of the uncontrollable consumption occurs in the network at a later time. When combining a scheduling-based approach with the PowerMatcher, it is possible to mitigate these effects [101].

The pricing signal used by Triana as well as the market clearing price used by the PowerMatcher could in theory be directly linked to financial incentives for the users, i.e., the prices sent correspond directly to those used to bill the customer. However, in practice, these prices mostly serve solely as a steering signal, i.e., a signal directing the individual users/devices to schedule the use of their flexibility to best match the system goals. The participation in the system is often envisioned using different incentives, where a user can sometimes decide on a case-by-case basis to participate, i.e., if they feel the reward outweighs the downsides [118]. Such a system is thus effectively a mix between direct control and participation purely incentivized through the price of energy.

Instead of prices, different signals can be used to steer a group of appliances. An example for this is the priority signal used by the Intelligator [28, 29, 65]. This approach reduces the need for local intelligence by deducing certain important parameters on a central level using self-learning. The Intelligator approach consists of three steps; aggregation, optimization and control. Devices use local bidding functions, similar to those used in the PowerMatcher approach, to communicate the desire to consume energy right now compared to delaying the consumption. Furthermore, the relative energy, a ratio between already consumed energy and (expected) total energy consumption, is also communicated for each device. Both the bidding functions and the given values of relative energy are aggregated. The amount of available flexibility is then deduced from the aggregated bid function and relative energy for the coming period using self-learning. This leads to bounds that are used to schedule the use of flexibility over a predetermined scheduling horizon consisting of time intervals. Finally, the use of energy for the current time interval is communicated back to the devices using the aforementioned priority signal, which the devices can use to determine their energy consumption. The above steps are repeated in each time interval, similar to a classical model predictive control (MPC) approach [5].

The approach has been shown to work for a large cluster of EVs [141] but assumes EVs only arrive and depart at the start of time slots. This restriction has been lifted in [36], where the problem includes a real-time coordination aspect to adopt the priority signal for EVs that arrive or leave during a time interval. A similar approach has been applied for domestic water heaters in [72]. In this approach for the water heaters, the aggregated model of flexibility is based on tracer devices, which give a representative model for the entire population of water heaters. The authors also include the dispatch mechanism into the optimization step, which further improves the results.

The advantage of the Intelligator approach is that no privacy sensitive information needs to be communicated. Furthermore, local constraints on, e.g., user comfort



can be guaranteed since devices can always consume when required regardless of the priority signal sent by the central controller. However, the approach depends on an appropriate self-learning technique to determine the relevant parameters used in the optimization. For example, Claessens *et al.* [28] show that under a strong repetitiveness assumption on the parameters, an intuitive approach gives adequate results. However, in case the parameters are hard to learn the system operates sub-optimally [27].

## 2.7 CONCLUSION

In this chapter we reviewed the energy system and in particular the electricity system. Our review sketched the system in the past, how it evolved to its current form and how it is expected to evolve in the future. From the (expected) changes sketched it is evident that, in order to ensure safe and stable operation of our electricity supply chain, new approaches are needed. In particular, with the increasing penetration of uncontrollable renewable energy sources and measurement and control capabilities on all levels of the system and a greater need for flexibility in consumption, EM approaches will play an important role in the future (smart) grid.

To facilitate the integration and proper use of DERs in the future (smart) grid, a new EM approach is required. The main requirements for such an approach are: scalability, scheduling, heterogeneity, feasibility, and privacy. We discussed and motivated each of these requirements. Furthermore, we gave a mathematical problem formulation of the (scheduling) problem that an EM approach needs to tackle. In this problem many different types of DERs will play an important role.

We classified the different types of DERs that we assume will play a role in the future smart grid in four classes: uncontrollable, time-shiftable, buffer, and unconstrained. We argued that the buffer class will play an important role in the future grid, particularly in residential areas. Hence this thesis focusses on the use of devices falling into the buffer class. For this, we described a generic model for storage devices that will aid us in later chapters.

Finally, we reviewed some of the related work on EM approaches. In this review we distinguished between approaches that incentivize the use of the available flexibility through adapting different tariff structures in the market versus those directly controlling loads. Furthermore, we discussed several approaches (i.e., Triana, PowerMatcher, and Intelligator) that effectively combine the two using different incentives for users to participate without assuming that the approach has full control of the appliances of the customers. However, none of these approaches fulfils all the requirements we listed for an EM approach.

Summarizing, the requirements listed in this chapter lead us to formulate our own EM approach in the next chapter. This is a heuristic approach that is inspired by the Triana approach as reviewed in this chapter.








# 3

## PROFILE STEERING

**ABSTRACT** – *In the previous chapter we sketched a changing energy supply chain. As the old paradigm of central control is no longer applicable in the future grid, a new energy management approach is required. In this chapter we discuss our profile steering decentralized energy management approach. This approach ensures that the flexibility of all distributed energy resources can be used efficiently. The approach uses a hierarchical control structure to ensure that it scales to the large amount of distributed energy resources expected in the future grid. Furthermore, we show that local limitations of the grid, i.e., maximum power flows through assets, are important to consider in future grids with large loads and much local production. We adapt our approach to incorporate these limitations ensuring that the approach uses the provided flexibility in a feasible way.*



This chapter is based on [TvdK:7] and [TvdK:3].




Figure 3.1: The cumulative share of different sources of renewable energy in the total energy consumption of the Netherlands over the period 1990 to 2015. Data taken from [51].




Figure 3.2: The cumulative share of different sources of renewable energy in the electricity consumption of the Netherlands over the period 1990 to 2015. Data taken from [51].

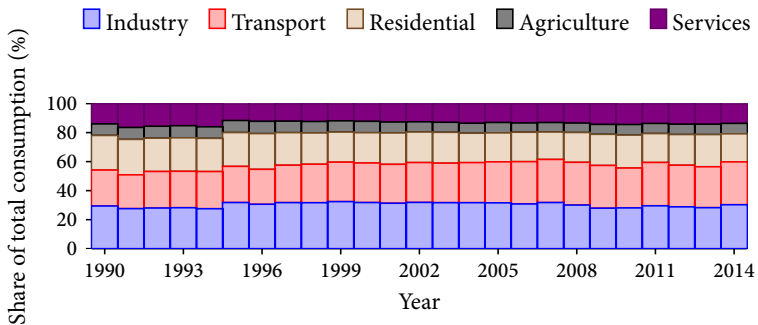


Figure 3.3: Share of the different sectors in the total energy consumption in the Netherlands over the period 1990 to 2014. Data taken from [51].

### 3.1 INTRODUCTION

In the previous chapter we sketched a rapidly changing energy supply chain, in particular the electricity grid. The changes are mostly caused by a shift towards energy from carbon free sources and the electrification of our energy supply chain. The introduction of energy from renewable sources in our energy supply chain is depicted in Figure 3.1. This change is particularly notable in the electricity grid (see Figure 3.2). Most of electricity produced from new sources comes from smaller units that are often situated closer to the demand, i.e., in the distribution grid. For example, consider rooftop photovoltaic (*PV*) that are very popular in Germany. Furthermore, many sources of flexibility (e.g., electric vehicles (*EVs*), smart white goods, and heat pumps) have a relatively small impact on their own but can offer large benefits in groups of large sizes.

We argued that traditional approaches to control our energy supply chain, and in particular our electricity grid, are no longer applicable in this changing landscape. Thus, a new energy management (*EM*) approach is needed. We focus on the residential sector in this thesis. Figure 3.3 lists the share of the different sectors in the total energy consumption of the Netherlands over the period of 1990 to 2014. The residential sector takes up a significant share of the total energy consumption. Furthermore, the approaches and ideas we present in this thesis also apply to other customers (e.g., offices, schools, etc.) located in the distribution grid. In other words, a large share of the total energy consumption is covered by our approach.

In this thesis we focus on the electricity grid. The traditional approach to manage the electricity grid is through a centralized approach. If all (future) distributed energy resources (*DERs*) are to be considered by such an approach, all the information on the flexibility provided by the *DERs* needs to be known in a central location. This causes a very large communication requirement and dependency as well as an issue concerning privacy. Furthermore, the problem to be solved in such a central approach is  $\mathcal{NP}$ -hard, as we show in this chapter. This means that we expect these problems to be computationally difficult to solve in reasonable time as the computational times of methods solving (sub)problems to optimality grow exponentially in the size of the instances.

In the previous chapter we listed several requirements on a new *EM* approach. Recapping on these; we need a scalable scheduling based approach capable of handling many heterogeneous devices in a feasible manner for all levels of the grid while respecting privacy sensitive information. To tackle the scalability, communication, and privacy issues, we consider a heuristic approach, namely decentralized energy management (*DEM*). In a decentralized approach the system is controlled through a hierarchical control structure. In such a control structure, the *DERs* make their own schedule of the use of flexibility. How and when this flexibility is used by the *DERs* is coordinated through the use of steering signals send to the (controllers of the) *DERs* by higher level controllers. Note that this is already common practice in many countries through the use of different energy tariffs throughout the day,



indicating when energy is expected to be abundant (cheap periods) compared to when it is constrained (expensive periods), e.g., day and night tariffs.

The scalability to a large amount of DERs is ensured through the decentralization. Furthermore, less communication is required than in a central approach, as only the steering signals and schedules are communicated instead of all relevant information required to obtain them. This also ensures that privacy sensitive information is kept locally, (partially) solving the privacy issue. Note that it is possible to make small groups of DERs and schedule them together rather than individually. This reduces the computing power requirement for each of the individual DERs.

As mentioned in the previous chapter, purely steering on prices, as is currently common practice and often considered for the smart grid, can have adverse effects. Furthermore, if the communication in a decentralized scheduling system is one way, i.e., only the steering signals are sent and the DERs do not communicate their schedule back, it is hard to predict how the system will react and regularly might overshoot or not reach the desired effect. To this end we study a system with communication in both directions; a centralized unit sends steering signals to the DERs. Then, these DERs schedule their flexibility use based on the signals and communicate these schedules back to the central units. The approach may react to the schedules made by the DERs by updating the steering signals, resulting in an iterative process.

The remainder of this chapter is outlined as follows. In the next section we further study the EM scheduling (*EMS*) problem as introduced in the previous chapter. In particular we study the complexity of this problem and show it to be already  $\mathcal{NP}$ -hard in a simple setting. To avoid the computational difficulty associated with such problems we introduce a basic variant of our decentralized, heuristic approach in Section 3.3. The steering signals used by this approach are elaborated in Section 3.4. In Section 3.5 we detail the results of a simulation study showing the effectiveness of the proposed approach. The simulation results reveal that local limitations of the physical grid are not always respected by the approach. As this violates the feasibility requirement we stated, we adapt our approach to rectify this problem in Section 3.6. Finally, some conclusions and a discussion are presented at the end of the chapter, in Section 3.7

## 3.2 COMPLEXITY OF SCHEDULING BASED EM

In this section we study the complexity of the *EMS* problem as introduced in Chapter 2. For this, we first consider time-shiftable DERs and how they are scheduled in an EM approach. As discussed in the previous chapter, a time-shiftable appliance offers flexibility to the system by allowing its program to be shifted in time. In other words, the device has to run a fixed program, which is synonymous with a fixed energy profile from the perspective of the grid. However, the time at which the program starts can be shifted to better fit the system goals.







More formally, a time-shiftable appliance is characterized by a time window in which the device has to run. This time window is generally characterized by two parameters, the interval of arrival  $t^a$  and the deadline for starting  $t^d$ . The interval  $t^a$  specifies the time interval in the scheduling horizon at which the device becomes available, i.e., it is the earliest time interval in which the program of the device can be started. On the other hand, the interval  $t^d$  gives the latest possible starting time of the program. Generally the intervals  $t^a$  and  $t^d$  result from user defined constraints. Furthermore, the energy profile of the program of the device is given by a vector  $\mathbf{p} := (p_0, p_1, \dots, p_{\ell-1})$ , i.e.,  $\mathbf{p}$  defines the energy use of the device for each of the  $\ell$  time intervals that the program runs. Recall that  $T$  is the scheduling horizon used in Problem *EMS*. We assume that  $\ell \leq T - t^d$ , i.e., that the entire program can be completed within the considered time horizon for every possible starting time. The scheduling decision to be made by the device is to determine the start time  $t^s$  of the program, where  $t^s$  has to be chosen such that:

$$t^a \leq t^s \leq t^d. \quad (3.1)$$

For a chosen start time, the energy use  $x_t$  of the device for time interval  $t$  is given by:

$$x_t = p_{t-t^s}, \quad (3.2)$$

where we assume that  $p_t = 0$  for  $t < 0$  or  $t \geq \ell$ . Using the notation introduced in Chapter 2, we obtain that the feasible set  $X^m$  for the energy profiles of a time-shiftable device  $m$  is given by:

$$X^m = \{\mathbf{x}^m \mid \exists t^s : t^a \leq t^s \leq t^d \wedge x_t^m = p_{t-t^s} \forall t\}. \quad (3.3)$$

We note that this feasible set can readily be extended to include multiple programs that have to be executed sequentially.

### 3.2.1 COMPLEXITY

Above we defined the scheduling decision for a time-shiftable appliance in a scheduling based EM approach. In this subsection we study the complexity of the general *EMS* problem tackled by a scheduling based EM approach (see Chapter 2 for scheduling based EM approaches, for an introduction to the complexity of problems see the Appendix). We show that Problem *EMS* is already  $\mathcal{NP}$ -hard when scheduling a set of time-shiftable devices. For this we consider a set  $\mathcal{M}$  of time-shiftable devices. Note that, to indicate the parameters of a device  $m$  we add the superscript  $m$  to the parameters of the notation introduced above. As constraints of the *EMS* problem, we add only a single constraint on the energy profile resulting from the joint schedule of the time-shiftable devices, given by  $\mathbf{x}$ , limiting the total energy consumed by the time-shiftable devices by a vector  $\mathbf{x}^{max}$ . This limit can for instance be imposed by maximum loads on a cable or transformer. In other words, we assume that the set  $X$  is given by:

$$X = \{\mathbf{x} \mid \sum_{m \in \mathcal{M}} \mathbf{x}^m \leq \mathbf{x}^{max}\}. \quad (3.4)$$



Using these assumptions we obtain the following result.

**Theorem 3.1.** *Determining if a feasible solution to the EMS problem exists is  $\mathcal{NP}$ -complete, even for instances where we only consider two time intervals ( $T = 2$ ), time-shiftable devices that run for a single time interval ( $\ell = 1$ ) and can be scheduled in both time intervals (i.e.,  $t^{a,m} = 1, t^{d,m} = 2 \forall m$ ), and a uniform upper bound on the consumed energy ( $x_t^{max} = x^{max} \forall t$ ).*

*Proof.* The problem belongs to  $\mathcal{NP}$  because it is straightforward, given starting times  $t^{s,m}$  for each  $m$ , to check the feasibility of these start times. Also it is straightforward to calculate the total power profile  $\mathbf{x}$  and check if  $\sum_{m \in \mathcal{M}} x_t^m \leq x_t^{max}$  for every  $t$ . To prove  $\mathcal{NP}$ -completeness we reduce the ( $\mathcal{NP}$ -complete) partition problem to Problem EMS. Let the multiset  $S = \{s_1, s_2, \dots, s_K\}$  specify an instance of the partition problem. We construct a corresponding instance of Problem EMS as follows. We define  $M := K, T := 2, x_t^{max} := \frac{1}{2} \sum_{k=1}^K s_k$ , and  $\ell^m := 1, p_1^m := s_m$  for all  $m \in \mathcal{M}$ . A feasible solution to Problem 2.1 now translates into a feasible partition  $(S_1, S_2)$  by taking:

$$S_1 = \{s_m \in S | t^{s,m} = 1\}, \quad (3.5a)$$

$$S_2 = \{s_m \in S | t^{s,m} = 2\}, \quad (3.5b)$$

and vice versa. Since the reduction given above uses a polynomial number of steps, the theorem follows.  $\square$

We note that the result is independent of the objective function  $f$  for the EMS problem. Furthermore, we note that it is also possible to prove a similar result for a broad class of objective functions  $f$  without imposing constraints on the joint schedule  $\mathbf{x}$ , i.e., without using the (uniform) limit  $\mathbf{x}^{max}$ .

The complexity result above implies that centrally computing optimal schedules for all smart appliances is not feasible in many situations, especially when the set of devices becomes large and the allowed solution times and computational power are limited. To tackle this issue we propose a decentralized, heuristic approach where the computation is split over the devices by letting them make their own schedule. We describe this heuristic approach in more detail in the next section.

### 3.3 PROFILE STEERING HEURISTIC

To solve the (hard) problem of scheduling smart appliances, we propose a heuristic approach that splits the computation by letting the smart devices, or their controllers, make their own schedules. The heuristic coordinates the use of flexibility between the different devices by first requesting an initial schedule of each device, followed by an iterative phase where each device is asked to send a (new) candidate schedule. In each iteration a single device that is selected greedily is asked to update its schedule to its proposed candidate schedule. We call this approach



*profile steering* as we steer the total profile formed by the individual schedules of the devices based on the objectives of the system. For example we steer the profile of the devices directly based on the desired profile of the system. Profile steering works by using a centralized controller that steers the scheduling of the individual appliances. This centralized controller knows the system objectives and can thus coordinate between the schedules made by the devices. To ensure the resulting schedules properly align with the system goals the controller works in two phases: an *initial phase* and an *iterative phase*.

In the *initial phase*, the controller asks each individual device to create an initial schedule and to send back the corresponding energy profile. This initial schedule can be constructed in several manners, e.g., by scheduling the device as soon as possible, by maximizing user comfort, or by scheduling towards the system objectives. In the last case, these system objectives, captured by the objective function  $f$  in the EMS problem, need to be known by the device or device controller. For example, if the system goal is to flatten the energy profile, then each device is asked to create a schedule resulting in a profile that is as flat as possible in the initial phase. We often use the Euclidean norm as the objective function to model the desired flattening of the profiles. In this case each device is asked to create a schedule such that  $\|\mathbf{x}^m\|_2$  is minimized. The advantages of scheduling towards the system objectives in the initial iteration is that in many situations the resulting overall profile already fits the objectives reasonably well. By this we mean that the combination of all initial schedules leads to a result already quite close to the optimal solution. After construction of these (initial) schedules all devices send the resulting profiles back to the central controller, which aggregates these profiles into an initial total (scheduled) profile.

The second phase is an *iterative phase*, in which schedules of individual appliances are updated to achieve a better overall profile regarding the system goals. To this end, each appliance is asked to construct a *candidate schedule* that best fits the overall system objective given the current total profile. More precisely, let  $\mathbf{x}$  be the current total profile with objective value  $f(\mathbf{x})$ . Now, device  $m$ , with current profile  $\mathbf{x}^m$ , is asked to construct a *candidate schedule*, resulting in a profile  $\hat{\mathbf{x}}^m$ . This candidate schedule should be constructed such that the system goals, captured in the objective function  $f$ , are maximally improved (i.e., decreased) by switching from the original schedule of the device to the candidate schedule. In other words, the candidate schedule is constructed such that the resulting energy profile  $\hat{\mathbf{x}}^m$  minimizes  $f(\mathbf{x} - \mathbf{x}^m + \hat{\mathbf{x}}^m)$ . This schedule needs to be feasible, i.e.,  $\hat{\mathbf{x}}^m \in X^m$ . After constructing a candidate schedule, the device sends the corresponding (new) energy profile back to the central controller.

After receiving all profiles corresponding to candidate schedules, the centralized controller decides which device gives the best improvement when updating this device's current schedule to the proposed candidate schedule. If the improvement is large enough, the centralized controller then asks *only* this device to update its schedule to the candidate schedule. The reason that only the schedule of a single



device is updated is to prevent oscillation and overshoot effects. The process is then restarted by again asking each device to propose a candidate schedule, using the updated total profile. This iterative process is repeated until the central controller decides that no proposed candidate schedule significantly improves the system objectives. The approach is summarized in Algorithm 3.1, where we denote schedules of devices using their resulting energy profile.

---

**Algorithm 3.1** Profile steering algorithm *PS*


---

- 1:  $x = \text{Function } PS(f, \mathcal{M}, \mathcal{T}, X, X^m, \epsilon)$
  - 2: Request an initial schedule  $\mathbf{x}^m$  of each device
  - 3:  $\mathbf{x} := \sum_m \mathbf{x}^m$  {Aggregate the initial schedules}
  - 4: **repeat**
  - 5:   **for**  $m \in \mathcal{M}$  **do**
  - 6:      $\hat{\mathbf{x}}^m = \arg \min_{\bar{\mathbf{x}}^m \in X^m} f(\mathbf{x} - \mathbf{x}^m + \bar{\mathbf{x}}^m)$  {Construct candidate schedule for  $m$ }
  - 7:      $\delta^m = f(\mathbf{x}) - f(\mathbf{x} - \mathbf{x}^m + \hat{\mathbf{x}}^m)$  {Improvement made by  $m$ }
  - 8:   **end for**
  - 9:    $\hat{m} = \arg \max_m \delta^m$  {Find device with best improvement}
  - 10:    $\mathbf{x}^{\hat{m}} = \hat{\mathbf{x}}^{\hat{m}}$  {Update schedule of  $\hat{m}$ }
  - 11: **until**  $\hat{\delta}^m < \epsilon$  {Repeat as long as sufficient progress is made}
  - 12: Return  $\mathbf{x}$
- 

By distributing the computation of the (device specific) schedules to the (controllers of the) actual devices, the problem is distributed and parallelized over these devices (see Figure 3.4). This implies that the approach is scalable. The added benefit of using the proposed profile steering approach is that we immediately solve (part of) the privacy issue. This is because the devices no longer have to communicate their exact preferences and limitations, only schedules and candidate schedules. However, some information can still be derived from these schedules, especially when many requests are made, i.e., in case of many iterations in the iterative phase of the algorithm.

Central to the profile steering approach is the bidirectional communication between the centralized controller and the devices. The central controller sends steering signals to the devices. In turn, the devices respond with the best schedule with respect to these steering signals. How we achieve the latter is discussed in more detail in the next section. Note that it is possible to introduce intermediate levels in the control hierarchy, of which we give an example in Figure 3.4. It is possible to group devices and/or controllers, where the controllers themselves represent a group of devices, and introduce an intermediate controller for this group. As the central control is agnostic to how the schedules it receives are obtained, the intermediate controller can also use the profile steering approach to transfer the received steering signals to its children and ask them for their schedules. Then, the intermediate controller can provide a schedule, as requested by its parent controller, through aggregation. Using this principle the devices can be split over a tree struc-

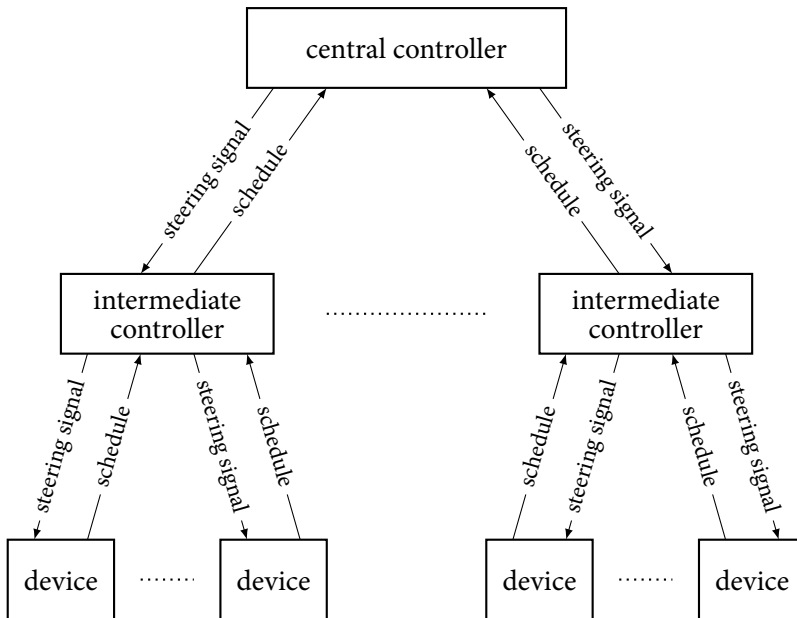


Figure 3.4: The control hierarchy of the profile steering approach, as given in Algorithm 3.1, extended with one intermediate level.

tured hierarchy, where the intermediate controllers send steering signals based on the signals they receive from their parent. The devices remain the only entities that have to schedule the use of flexibility, as schedules on intermediate levels are obtained only through aggregation of schedules received from below.

To design a hierarchy of controllers, we may use the structure of the underlying grid. The low voltage (*LV*) part of the electricity grid is usually designed radially, i.e., it has a tree structure. This implies that there is a natural mapping from the electricity grid to a control hierarchy. We gave an example of a typical Dutch *LV* network in Figure 2.5 (Chapter 2). For this example a natural control hierarchy in the *LV* grid is given in Figure 3.5. In this hierarchy, first devices are grouped together per house. Next, the houses are grouped together per phase. Note that the phase level can be removed in case a three phase connection is common. The three phases of a single feeder are grouped together at the feeder level. Finally, the feeders behind the *LV* transformer are combined and governed by the central controller for this *LV* grid. In the figure we depicted a structure with three feeders, however, it is possible to add more. Note that it is also possible to add higher levels to the control hierarchy, e.g., part of the medium voltage (*MV*) grid.

Note that additional bookkeeping is required on the candidate schedules/profiles if multiple layers are used with profile steering. To illustrate this we consider the following example.

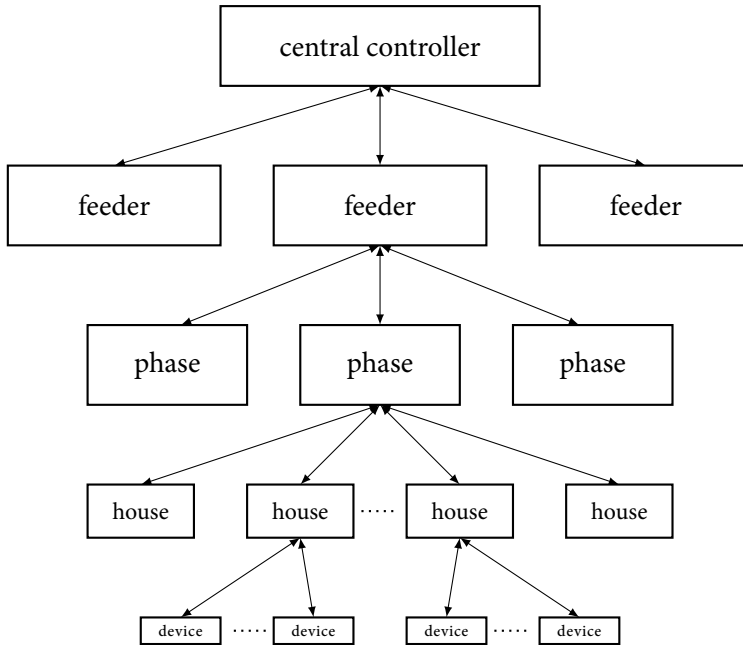


Figure 3.5: The natural control hierarchy that follows from the structure of the Dutch LV grid as depicted in Figure 2.5. Note that we do not show all children of the entities on each level.

**Example 3.1.** We consider profile steering applied to an arbitrary set of devices using a hierarchical control approach with two levels of controllers. These levels consist of a level with two intermediate controllers to which the devices are connected and a level with the central controller to which both intermediate controllers are connected. In the initial phase of profile steering the central controller asks the intermediate controllers for their initial schedules, which these intermediate controllers obtain from the devices. Then, in the iterative phase the central controller asks both intermediate controllers to propose candidate schedules. To obtain these candidate schedules both intermediate controllers apply the profile steering approach. Now, the intermediate controllers already know the initial schedule of the devices and hence only use the iterative phase of profile steering. This results in candidate schedules being requested from the devices by the intermediate controllers with these candidate schedules being committed if they best fit the signals the intermediate controllers received from the top level controller. After both intermediate controllers finish their application of profile steering they send the resulting schedule to the central controller. Note that the devices have already updated their initial schedules based on the requests of the intermediate controllers. After receiving both candidate schedules of the intermediate controllers, the central controller chooses exactly one to be used. Thus, only the schedules of the devices below the intermediate controller that is chosen can be used and the schedules

of the devices below the other intermediate controller need to be reverted back to their initial schedules.

The above example illustrates that, when using a hierarchical control structure with profile steering, the updated schedules of devices, which were obtained through requests of intermediate level controllers, sometimes have to be reverted back due to choices made by higher level controllers. This can be done by letting the device save the accepted schedule for each level in the hierarchical structure, i.e., by letting a device save a schedule for each controller above it in the hierarchy. In the initial phase, the initial schedule made by the device is the same for all hierarchical levels. Afterwards, in the iterative phase, the schedules are updated only if they are accepted by the appropriate controller. Furthermore, if a higher level controller rejects the schedule of an intermediate controller, the schedules of the devices below this intermediate controller are reverted back to the one saved for this higher level controller. In the example above this means the devices have two schedules; the first for the intermediate controller and the second for the top level controller. While the intermediate controller is applying profile steering only the first of these two schedules is updated. Then, if the top level controller accepts the candidate schedule of the intermediate controller to which the device is connected, the device updates its second schedule to be equal to the first. On the other hand, if the top level controller does not select this schedule, the device reverts its first schedule back to the original schedule (given by the second schedule).

### 3.4 STEERING SIGNALS

The profile steering approach, introduced in Section 3.3, depends on bidirectional communication between the central controller and the devices. The central controller sends the current total profile  $\mathbf{x}$  together with the system's objective  $f$  to the devices. We call this information sent by the global controller the *steering signal*, as it steers the profile of the local devices in a certain direction. The devices respond with a candidate profile  $\hat{\mathbf{x}}^m$ . Note that it is possible to forego the bidirectional communication, by only making an initial schedule for each device, based on the received steering signal, and using this schedule for each device.

The results attained by the approach depend heavily on the choice of the steering signal, as we discuss below. Steering signals are often applied in decentralized control structures (see, e.g., [29, 126]). Below we first discuss the traditional pricing signals and then extend on these.

#### 3.4.1 PRICE SIGNALS

Steering signals are currently already in use in electricity grids in many countries. For example, many energy retailers differentiate between the tariff for energy consumed during on-peak hours (when energy is in high demand and expensive) and off-peak hours (when energy is cheaper). Furthermore, many larger customers



(e.g., large industry) pay a fee based on their maximum consumption next to the fee based on their total energy consumption [39]. These approaches can be used as a starting point and adapted to better reflect the requirements of the future grid.

A popular approach in the literature is dynamic pricing (also often referred to as time-of-use (*ToU*) pricing) or other (linear) pricing approaches. However, such pricing structures, and specifically linear pricing, can result in outcomes far from the desired effect and in some cases can even have adverse effects, as we discussed in Chapter 2. Furthermore, (linear) pricing signals are ambiguous in communicating the goals of the system, as the following example shows.

**Example 3.2.** *Consider the situation where an EV has to be charged in two hours upon arrival. Furthermore, we assume the user does not care how the charging is done exactly in these two hours, as long as the battery is full afterwards. We also assume that the ideal scenario for the grid is that the EV equally spreads its charging over the two hours. When sending a pricing signal with the same price for both hours, any division of the required charging over both hours is optimal when minimizing costs. Furthermore, when sending a slightly lower price for one of the two hours, the user is financially incentivized to shift all its charging to the hour with the lower prices.*

This example illustrates that (linear) dynamic pricing leads to extreme behaviour (shifting all or nothing), which can be far from ideal for the overall system. Note that, the more systems that react automatically on linear pricing signals the more extreme the behaviour will be.

### 3.4.2 GENERAL STEERING SIGNALS

The profile steering approach is designed to allow the use of a very broad class of steering signals, as any objective function that can be feasibly transferred to signals for the controllers can be used. This implies that the approach is capable of asking the devices directly to optimize towards the system's objectives rather than in an indirect and ambiguous way through dynamic prices.

As an example we consider the case that the (central) controller has a *desired profile* in mind for the group of devices it controls. This desired profile  $\mathbf{d}$  is simply an energy profile for the considered scheduling horizon, i.e.,  $\mathbf{d} := (d_1, d_2, \dots, d_T)$ , it could for example be a constant, e.g., flat, profile to minimize stress on grid assets, or a profile that best fits renewable energy generation. As a measure of fitness to the system goals we use, e.g., a norm on the difference between the current profile ( $\mathbf{x}$ ) and desired profile ( $\mathbf{d}$ ), i.e.,  $f = \|\mathbf{x} - \mathbf{d}\|$ . The norm can be chosen depending on a specific aspect of the difference between the desired profile and scheduled profile that is to be minimized. For example, the Euclidean norm (2-norm) minimizes squared deviations, thus equally spreading differences between the scheduled and desired profile over the time intervals minimizes this norm. On the other hand the Manhattan norm (1-norm) minimizes only the deviations, hence it does not care where the differences between scheduled and desired profile are, as long as the total (absolute) difference is minimal.





The freedom of choosing the steering signal ensures that the profile steering approach is broadly applicable. The choice of the steering signal may significantly complicate the scheduling problem that has to be solved by each individual appliance. We come back to this issue in the next section.

Furthermore, in the examples and applications throughout this thesis, we frequently combine a desired profile  $\mathbf{d}$  with the Euclidean norm, i.e.  $f_t = \|\mathbf{x} - \mathbf{d}\|_2$ . Where in many cases we take the (original) desired profile to be the zero profile, i.e.,  $\mathbf{d} = \mathbf{0}$ . This is mathematically the same as steering towards any other flat desired profile when the total consumed energy over the planning horizon does not change between different schedules. We made the choice for this objective because it steers at flattening the energy profile, which reduces stress on the grid assets in a grid. Furthermore, this objective also ensures that locally produced energy is used locally as much as possible. Finally, minimizing the squared energy consumption is also closely related to minimizing grid losses.

### 3.4.3 SCHEDULING DEVICES UNDER STEERING SIGNALS

The introduced profile steering approach distributes the computation of schedules to the devices using steering signals as discussed above. It remains to determine how the devices obtain their schedules based on these steering signals. In Chapter 2 we separated DERs into four classes; uncontrollable, time-shiftable, buffer, and unconstrained devices. We argued that the most important classes to be considered for residential EM are the time-shiftable and buffer classes.

We first consider a time-shiftable, as discussed in Section 3.2. Time-shiftables offer flexibility by allowing to shift the program that they need to run in time. This leads to a feasible set of schedules for the devices where each schedule corresponds with a feasible start time. In general we assume that these feasible start times correspond with time intervals used by our approach. As the number of time intervals considered is generally small, this implies that the number of feasible start times, and hence the number of feasible schedules, is small for these devices. This results in a small discrete set of possible solutions, which can be explored exhaustively to obtain the best possible schedule for the received steering signal. While straightforward implementations are already rather efficient, more intricate methods can provide a speed-up such as the convolution approach used by Toersche [126, Section 5.4.2.6].

Besides the time-shiftables, the buffer class also plays an important role in residential EM. DERs in this class generally offer flexibility in both *when* and *how much* they consume and/or produce. This makes their scheduling problem more difficult, as the set of potential schedules is no longer finite and hence an exhaustive search is no longer possible. In Chapters 4 to 6 we consider various devices in this class and formulate efficient algorithms for the scheduling problems they need to solve in the profile steering approach.

### 3.5 EXAMPLE APPLICATION OF PROFILE STEERING

In the previous section we described the profile steering DEM approach. This approach iteratively steers the energy profile of the flexible appliances in a given grid towards a desired profile. To demonstrate the potential of this approach, we consider an example with 121 houses in the Dutch town of Lochem. We have access to a detailed model available of the LV grid used to distribute energy to these houses (provided by the DSO Alliander). This allows us to not only study the overall supply and demand matching but also how energy flows through the grid, including several important aspects of power quality. We consider a futuristic scenario regarding the penetration of renewables and flexible appliances. However, note that this neighbourhood is already actively working towards a large penetration of renewables and flexible devices, mostly in the form of rooftop PV panels and EVs [67]. For the scenario we assume 50 houses are mounted with rooftop PV panels, where the total surface area of the panels of each house is randomly selected to be between 8 and 16 square metres. Furthermore, we assume the following flexible appliances are present;

- » 54 EVs with an internal battery of either 12 or 24 kWh (both equiprobable) and a maximum charging rate of 7.7 kW.
- » 121 smart washing machines (*WMs*), each with built-in tumble dryer, with a time-shiftable load profile as found in [121].
- » 121 smart dish washers (*DWs*), with a time-shiftable load profile as found in [121].
- » 30 stand-alone batteries with a capacity of 10 kWh and a maximum power of 3.7 kW, both for charging and discharging.

Each house has exactly one smart WM and one smart DW. Furthermore, the batteries are only present in houses with rooftop PV panels. We generated the availability of the smart appliances as well as the load of the uncontrollable devices using the profile generator described in [68]. For this example, we specifically assumed EVs can sometimes be charged during the day (using energy produced by the PV panels), e.g., due to the owner having a day off or working from home. This is in contrast to the usual operation of charging in the evening and/or at night.

We applied the profile steering approach to the described scenario in a simulation study. For the simulation we chose the minimization of the Euclidean norm of the total profile as the objective. In the simulation we steer the energy profile of the entire neighbourhood for a week (using data for the 31st week of a year). As a comparison we also simulated a no control case, where devices consume energy as soon as possible, i.e., the smart white goods run as soon as they become available and the EVs charge as fast as possible. The batteries are not used in this case. The results of the simulation for the entire neighbourhood are given in Figures 3.6 (energy profile at the transformer) and 3.7 (minimum and maximum voltages in the network), together with some grid statistics in Table 3.1 (note that the last column



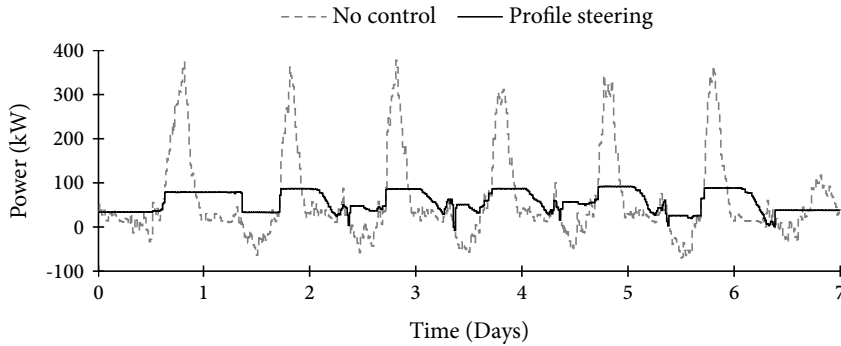


Figure 3.6: Energy profile for the transformer of the neighbourhood of 121 houses with either no control (devices run as soon as they become available) or using profile steering towards a flat profile.

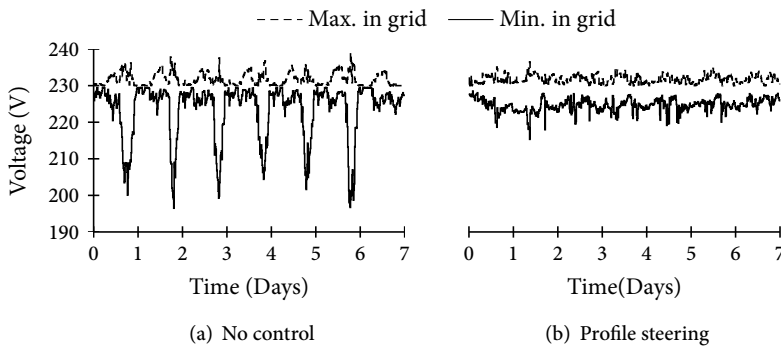


Figure 3.7: Lowest and highest voltages observed in the (simulated) network corresponding to the energy profile given in Figure 3.6 for profile steering and the no control case.

in this table gives statistics for a case we consider later on). For the profile steering approach we steered houses to flatten their own profile as much as possible in the initial phase of the algorithm. This ensured that the resulting profile was already quite flat, resulting in a small amount of iterations required in the iterative phase.

The results show that a significant peak occurs during the evening hours if no control is applied, i.e., when devices run as soon as they are available. This peak occurs because the majority of the EVs charge on top of the usual evening consumption peak. Due to this peak the voltages drop significantly in the network. When applying profile steering the charging of the EVs is shifted to the night or even to the afternoon, when there is energy available from the PV panels, if the EV is available long enough. Furthermore, the batteries are used to store energy produced by the PV panels in the morning/afternoon. The stored energy is used in the evening to



Table 3.1: Comparison of several grid statistics between our profile steering approach and no control.

	No control	Profile steering	Including limits
Total losses (%)	3.2	1.1	1.1
Lowest voltage (V)	196.5	215.3	215.3
Highest voltage (V)	239.1	236.5	236.4
Max. peak (kW)	380.6	93.7	93.9
Cable load (%)	151.4	84.1	67.5

compensate for the peak in demand. Overall the resulting neighbourhood profile is significantly flattened, causing a significant reduction in stress on the network assets. Also, the surplus of energy is minimal, causing a minimal flow of electricity through the transformer towards the MV grid. This also ensures that the voltage is much more stable. When comparing the grid statistics (the first and second column in Table 3.1), the grid losses, given as a percentage of the total energy supplied by the transformer, are reduced by about a factor three. Significant improvements in voltages, peak load, and cable load are also observed.

### 3.6 HIERARCHICAL CONTROL

In the previous sections we introduced the profile steering DEM approach to manage the energy profiles of flexible devices and showed that it can significantly flatten the energy profile of a neighbourhood. We argued before that such an approach should also take the local limitations of the grid into consideration. In principle, steering towards a flat profile should already ensure that asset stress is minimized. However, there are cases where significant overloading of network assets is still possible. As an example, we go back to the simulation study performed in Section 3.5. In particular, we consider a single house equipped with PV that offers a lot of flexibility during the day due to the availability of a lot of smart appliances (DW, WM, EV, and battery). The schedules of the appliances and the total profile of the house are depicted in Figure 3.8. The results for the considered house show that all smart appliances consume as much energy as possible during the afternoon. This coincides with the peak availability of solar energy in the neighbourhood. The total power drawn by the house peaks over 9.2 kW, which is the standard fuse limit for Dutch households with a single phase connection ( $40A \times 230V$ ). Thus, the resulting schedule for the household appliances of this house is not feasible. However, the overall profile at the transformer is acceptable for this time period.

In the depicted scenario we have the situation that a single house is overloaded in the scheduled energy profile. However, the same can occur for other network assets. An example of this can be observed if the appliances connected to a single feeder together try to compensate for a large production peak, e.g., of PV, elsewhere in the network. In general, in such situations the safe operation of distribution networks can no longer be guaranteed by the traditional fit and forget strategy. This is mainly

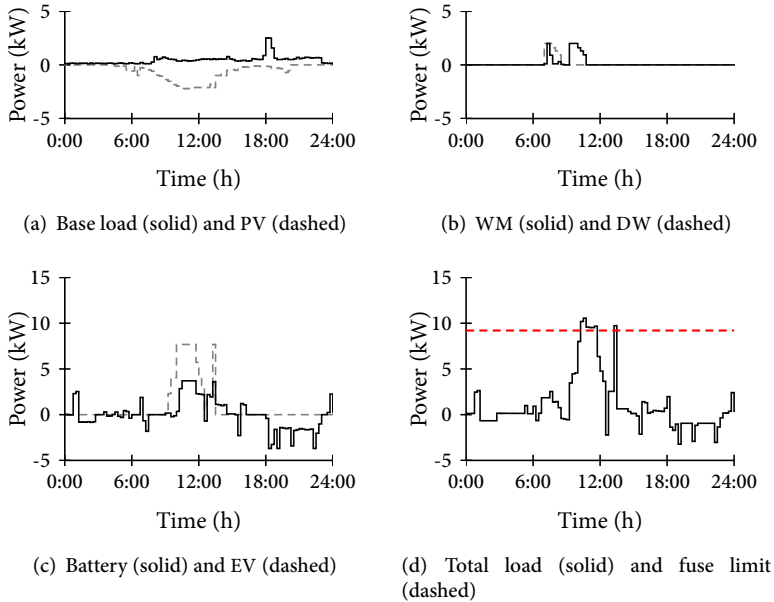


Figure 3.8: Profiles for a single house of the simulated neighbourhood with a lot of flexibility available during the afternoon PV peak.

due to an expected increase in large loads and an increasing simultaneity factor, e.g., in the case of a significant penetration of EVs that need to be charged in the evening. Further evidence of this is, for example, provided by Hoogsteen *et al.* [67].

The given examples illustrate the importance of a DEM approach also considering the local limitations of the grid, ensuring that the constructed schedules are also feasible with respect to the physical grid. To ensure this we can exploit the natural mapping between the physical (LV) grid structure and the control hierarchy we discussed in Section 3.3 (see Figure 3.5). The controllers that govern intermediate levels (e.g., a feeder or a phase) can consider the aggregate schedule for their children and adapt the steering signal they send whenever this aggregated schedule exceeds the constraints on the asset they govern. It remains to adapt the basic profile steering approach such that these constraints are taken into account.

### 3.6.1 INCORPORATING BOUNDS IN PROFILE STEERING

By design each controller within the hierarchy corresponds to a physical part of the grid. We assume that the controller knows the physical limitations of this part. To ensure these limitations are taken into account when schedules are made, we incorporate them in the steering signals. Let the vectors  $x^{min}$  and  $x^{max}$  represent the known limitations on the aggregated profile for a given controller. Note that,



in addition to the upper bound given in Problem 2.1, we also incorporate a lower bound on the profile. This lower limit can be used to ensure that the energy delivered back through the assets governed by the controller does not exceed certain limits. Thus, this allows us, e.g., to limit peaks in local production. In terms of the scheduling problem for the controller, the methodology needs to be adapted to ensure that the sum of the profiles resulting from the schedules made by the children of a controller lies between the given bounds. This means that the set  $X$ , used to denote the set of feasible aggregated schedules, includes the constraint,  $\mathbf{x}^{min} \leq \sum_{m \in \mathcal{M}} \mathbf{x}^m \leq \mathbf{x}^{max}$ . Note that we assume that the energy flow through an asset can be calculated by summing the energy profiles of the devices connected below this asset. This implies we neglect losses in the grid.

As discussed in Section 3.3, the profile steering approach works in two phases. In the initial phase of the controller, no information is known about the schedules of the children of the controller; they still have to be made. To ensure feasibility we extend the approach to include an additional iterative phase after the initial phase. Thus, the approach now consists of an initial phase and two iterative phases. In the first iterative phase we are purely looking for a feasible set of schedules, while we optimize these schedules in the second iterative phase. We note that this is similar to the two phase simplex method, commonly used to solve linear programs [144].


In the initial phase no schedules are known, hence we let (device) controllers make an initial schedule without considering the limitations, as these depend on the schedules of the other controllers. This is the same as in the basic profile steering approach. Then we come to the first of the two iterative phases, the one responsible for ensuring feasibility. In this phase we first check if the set of initial schedules is feasible, i.e., if the combined schedules do not violate any local limitations. If this set is feasible we can immediately continue with the second iterative phase, as this phase requires a feasible set of schedules as starting point. If not, we need to update the set of schedules to a feasible one. To do this, we can use the basic profile steering approach with an objective aiming to minimize violation of the given bounds. As an example, we can use:

$$f(\mathbf{x}) = \sum_{t \in \mathcal{T}} \max \left\{ \sum_{m \in \mathcal{M}} x_t^m - x_t^{max}, x_t^{min} - \sum_{m \in \mathcal{M}} x_t^m, 0 \right\}. \quad (3.6)$$


Next, we consider the second iterative phase of the profile steering approach, used to optimize the set of schedules. In this phase the controller knows the current schedule  $\mathbf{x}^m$  of child  $m$ . For child  $m$ , we can deduce bounds on the candidate schedule to ensure feasibility of the sum of its schedule added to the schedules of the other children. We calculate these bounds  $\mathbf{x}^{max,m}$  and  $\mathbf{x}^{min,m}$  by taking the difference between the given bound and the current schedule of the *other* devices:

$$\mathbf{x}^{min,m} = \mathbf{x}^{min} - \sum_{m' \neq m} \mathbf{x}^{m'}, \quad (3.7a)$$

$$\mathbf{x}^{max,m} = \mathbf{x}^{max} - \sum_{m' \neq m} \mathbf{x}^{m'}. \quad (3.7b)$$



(a) Scheduled total load profile minus the EV



(b) Limits imposed on the EV schedule

Figure 3.9: The total scheduled load profile for the house considered in the example in Section 3.5 minus the scheduled EV profile given in (a) can be used to calculate the upper bound on a candidate EV schedule given in (b).

**Example 3.3.** An example of the limits on a candidate schedule is given in Figure 3.9. For this example we consider the same house as in the example in Section 3.5. For this house we specifically consider the EV and deduce the bounds on the profile resulting from its schedule. These bounds are given by the limit of the house (9.2 kW for the fuse) minus the aggregated schedule of the other devices, which are both given in Figure 3.9 (a). We only calculate the upper bound on the schedule of the EV, which we give in Figure 3.9 (b). Note that the limits resulting from this are stricter than the maximum charging allowed by the EV around six o'clock in the morning and around noon. In the schedule constructed without limits imposed, the EV charged at maximal power around noon to compensate for the PV production peak at that time in the neighbourhood. This led to a total schedule for the house giving a consumption profile that was too high around noon.

In general, to ensure that a child of a controller using profile steering can calculate the bounds on a new candidate schedule, it needs to know the current total schedule  $\mathbf{x}$  (of the controller) and the limits  $\mathbf{x}^{min}$  and  $\mathbf{x}^{max}$ . Thus, we include this information into the steering signal sent by the controller. Next we need to consider how a controller that receives such bounds can incorporate them in their approach to derive a schedule.



We first consider the case that the controller receiving a steering signal including bounds is a controller using profile steering. This implies that this controller corresponds with a physical part of the grid, for which local limits need to be taken into account next to the limits received in the steering signal. We obtain the local limits used by this controller by taking, for both lower and upper bound, the stricter between the bound received in the steering signal and the local limits applicable to this controller. In other words, this controller sends the stricter of the two limits in the steering signals to its children. For example, consider the case that a controller for a house with a total maximum allowed load of ten kWh for a time interval receives a limit of twenty kWh in the received steering signal. This house controller then incorporates a limit of ten kWh in the steering signal it sends to the devices in the house.

Next we consider the case that the controller receiving a steering signal including bounds is a device controller. In this case the controller needs to ensure that the schedule produced for its device is feasible with respect to both the limitations put on the device and the limits received in the steering signal. This asks for an adaptation of the corresponding scheduling algorithms and the details on how this can be incorporated into the scheduling approaches of various devices is discussed in later chapters.

We summarize the above discussed adaption of the basic profile steering approach to include local grid limitations in Algorithm 3.2. First the initial schedules of the children are requested (line 2). These schedules are made without considering the limits on the sum of the schedules. Next we check if the sum of the initial schedules is feasible (lines 3-4) and, if this is not the case, we arrive at the first iterative phase of the adapted approach, used to obtain a feasible set of schedules. In this phase we use the basic profile steering approach (Algorithm 3.1) with the objective as given in (3.6) (lines 5-7). Next, we check if the resulting sum of the schedules is feasible. If it is not, i.e., we still have that  $\hat{f}(\mathbf{x}) > 0$ , then we conclude that we cannot obtain a set of feasible schedules and we return what we have (lines 8-10). If, on the other hand  $\hat{f}(\mathbf{x}) \leq 0$  we arrive at the second iterative phase of the adapted approach. In this phase we iteratively request candidate schedules of the children while ensuring that the sum of the schedules stays within the limits, until no further improvement can be found (lines 11-20).

### 3.6.2 INCORPORATING LOCAL LIMITS IN THE EXAMPLE

To show the effectiveness of the approach, we take the example described in Section 3.5 and incorporate local limits in this example. More precisely, we only consider a limit on the aggregated profile of the houses, namely a limit of 9.2 kW imposed by the fuse of a single phase connection. The achieved results for the house already considered in Figure 3.8 are given in Figure 3.10. The energy profiles of the battery and EV are slightly adjusted to ensure the total profile of the household does not exceed the limit imposed by the fuse.



---

**Algorithm 3.2** Profile steering with grid limitations algorithm *PSL*


---

```

1:  $\mathbf{x} = \text{Function } PSL(f, \mathcal{M}, \mathcal{T}, \mathbf{x}^{min}, \mathbf{x}^{max}, \epsilon)$ 
2: Request an initial schedule  $\mathbf{x}^m$  of each device
3:  $\mathbf{x} := \sum_m \mathbf{x}^m$  {Aggregate the initial schedules}
4:  $\hat{f} = \sum_t \max\{\mathbf{x} - \mathbf{x}^{max}, \mathbf{x}^{min} - \mathbf{x}, 0\}$ 
5: if  $\hat{f}(\mathbf{x}) > 0$  then
6:    $\mathbf{x} = PS(\hat{f}, \mathcal{M}, \mathcal{T}, 0)$  {Call Algorithm 3.1 with the objective function given in
   (3.6) to make  $\mathbf{x}$  feasible}
7: end if
8: if  $\hat{f}(\mathbf{x}) > 0$  then
9:   return  $\mathbf{x}$  {If the schedule remains infeasible return the best obtained}
10: end if
11: repeat
12:   for  $m \in \mathcal{M}$  do
13:      $\mathbf{x}^{min,m} = \mathbf{x}^{min} - (\mathbf{x} - \mathbf{x}^m)$  {Determine lower bound}
14:      $\mathbf{x}^{max,m} = \mathbf{x}^{max} - (\mathbf{x} - \mathbf{x}^m)$  {Determine upper bound}
15:      $\hat{\mathbf{x}}^m = \arg \min_{\bar{\mathbf{x}}^m} \{f(\mathbf{x} - \mathbf{x}^m + \bar{\mathbf{x}}^m) \mid \mathbf{x}^{min,m} \leq \bar{\mathbf{x}}^m \leq \mathbf{x}^{max,m}\}$  {Construct
     candidate schedule for device  $m$  within bounds}
16:      $\delta^m = f(\mathbf{x}) - f(\mathbf{x} - \mathbf{x}^m + \hat{\mathbf{x}}^m)$  {Improvement made by device  $m$ }
17:   end for
18:    $\hat{m} = \arg \max_m \{\delta^m\}$  {Find device with best improvement}
19:    $\mathbf{x}^{\hat{m}} = \hat{\mathbf{x}}^{\hat{m}}$  {Update schedule of  $\hat{m}$ }
20: until  $\delta^{\hat{m}} < \epsilon$  {Repeat as long as sufficient progress is made}
21: Return  $\mathbf{x}$ 

```

---

The differences between the scheduled energy profile of the entire neighbourhood with and without limits on the total household profile are only very minimal. The small difference can be explained by the observation that the overloading of a fuse in a single house occurs rarely in the considered scenario. Furthermore, simultaneous overloading of the fuses of multiple houses nearly never happens. If we would consider more extreme scenarios the probability that the limits imposed by the physical grid are violated will increase. As a consequence the resulting total profile may differ more from the profile obtained if these limits are not taken into account. The grid statistics (see the last column of Table 3.1) change slightly. The most notable change is the reduction in the maximum load put on the cables. This is caused by the fact that the highest loaded cable was previously a cable running directly from a house to the feeder, on which the maximum load has now been reduced.

### 3.6.3 OVERLOADING OF CABLES

Above we discussed how we can incorporate the physical grid into our profile steering EM approach. We showed, through an example, how this can effectively ensure that limitations on energy flows through grid assets can be taken into account. To




Figure 3.10: The results for the same house as in Figure 3.8 when we incorporate local limits into the example of Section 3.5.

achieve this we mapped the physical grid structure, e.g., transformer, feeder, phase, to controllers in our control hierarchy. However, energy flows in the grid can cause overloading at different points than modelled in this way, as we show below.

**Example 3.4.** Consider the grid given in Figure 3.11. In this figure a feeder is depicted where the first twelve houses (given in the ‘does not contribute to overloading’ box) are equipped with renewable generation, e.g., rooftop PV, such that they together are net energy producers during the afternoon. To minimize energy flows through the cable our profile steering approach can now ask the other houses (given in the ‘contribute to overloading’ box) connected to this feeder to consume as much energy as possible locally. If these houses together generate a schedule that results in a profile with high energy consumption, then the net energy flow at the beginning of the feeder might be low while the energy flowing through the section of the cable between the producing and consuming houses causes overloading.

Example 3.4 shows that it is important to be able to consider groups of houses that are supplied by a particular section of a feeder that can become congested in our approach. Note that it is often not possible to pinpoint where such problems will occur in a grid. However, if they are detected the approach can be adapted by adding an additional layer in the control hierarchy, as depicted in Figure 3.12, splitting the control of the houses on the feeder in question over two groups. The

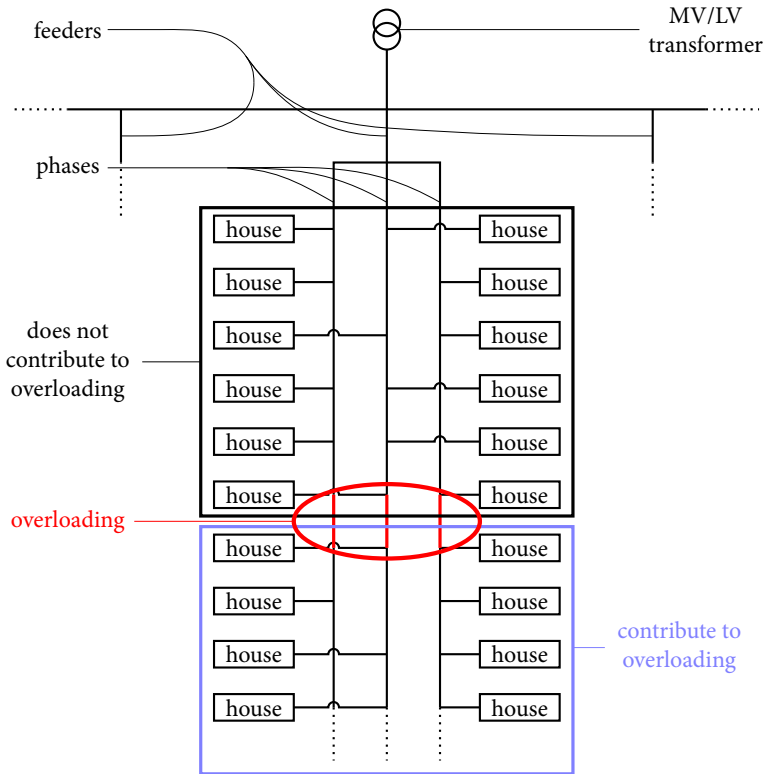


Figure 3.11: An example of an overloading situation in the LV grid from Figure 2.5. The assets in the ‘contribute to overloading’ box (potentially) contribute to the overloading of the cable given in the red ellipsoid, while the assets in the ‘does not contribute to overloading’ box don’t.

first group governs the houses that are not responsible for the congestion, while the second group governs those houses that are. The controller governing the houses that caused the overloading take the maximum cable load into consideration when making schedules, ensuring that no congestions occur. Note that, in order to apply the approach described above, measurement equipment on various points in the LV grid is required.

### 3.7 CONCLUSION

In this chapter we presented the profile steering DEM approach and discussed how the use of flexibility provided by smart devices can be scheduled by it. The approach overcomes the scalability issue of traditional scheduling approaches (e.g., solving unit commitment problems (*UCPs*)) in the energy domain. Furthermore, it also par-

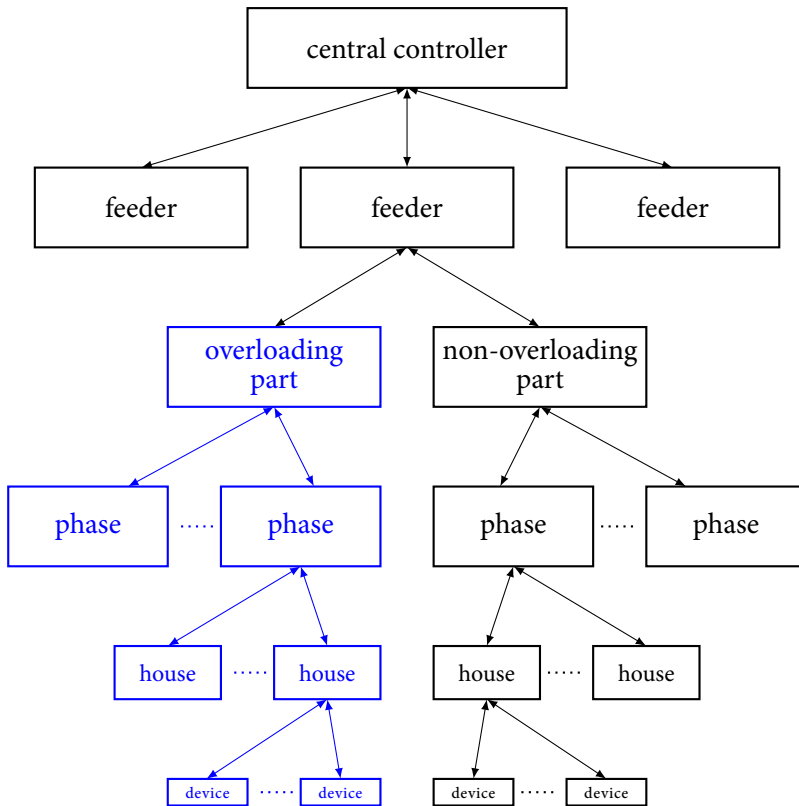


Figure 3.12: Adopting the control hierarchy given in Figure 3.5 to ensure the profile steering approach can be adopted to alleviate the problem.

tially solves the privacy issue and can incorporate local grid limitations through the proposed hierarchical control structure. The heterogeneity of the devices, expected in the future smart grid, is taken into account by the use of non-device specific steering signals. Devices can handle their specific constraints locally by making a schedule based on the received steering signal. The devices only communicate back their scheduled profile to the next controller upstream, thereby keeping privacy-sensitive device-specific information rather local. For the class of time-shiftable devices we argued that schedules can be obtained efficiently through exhaustive searches independent of the used objective. In the next chapters we discuss the aspect of creating schedules and profiles for the class of buffer devices. In summary, the profile steering approach meets all the requirements of a DEM approach, sketched in Chapter 2.

The scheduling approach presented in this chapter requires predictions of the available flexibility in the system to function properly. Furthermore, if these predictions are unreliable, the produced schedules can be suboptimal or even infeasible. We



believe the system can be adapted to become event-based. In such a system, schedules are only made for devices when their relevant parameters are known, i.e., at the time the devices become available. Note that in this case predictions of the expected flexibility available later on in the system are still required to make a schedule that best fits the system goals. For a further discussion on dealing with uncertainty we refer the reader to Section 8.3.1.

The presented basic profile steering approach extends the current billing structure used for nearly all energy consumers, as well as the much discussed extension of ToU pricing in several ways. The steering signals are more general than a pricing structure. We argued that simple pricing structures can cause extreme behaviour and that by using more complex structures, e.g., non-linear pricing (prices that depend on the consumption rate), as steering signals some of these drawbacks can be overcome. As a general example we used a desired profile as the steering signal. We combined this desired profile with a norm that measures the deviation between the desired profile and the current schedule. This allows the approach to steer the total schedule towards the desired goals.

We showed that our basic profile steering approach can produce schedules that result in profiles violating local limitations in the grid. To overcome this issue we proposed to use a hierarchical control structure that mimics the physical grid. We showed that we can incorporate local limitations of grid assets in such a control structure in our extended profile steering approach. This extended approach is capable of scheduling the use of flexibility while respecting grid limitations on various levels in the grid (i.e., house, feeder, transformer).

Our presented approach is based on two-way communication; devices receive steering signals and send the schedule they make based on these signals back to controllers upstream. As the approach works iteratively, we can update the steering signals to take the (updated) schedules of others into account. This results in the capability of devices to react on what is happening elsewhere, effectively incorporating a feedback loop into the approach. However, this feedback is not necessary for the approach to work, though it significantly improves the results. Without the feedback loop, steering signals are sent only once and the devices schedule their use of flexibility once without communicating this schedule upwards. For this to effectively work in practice, it is important to have proper knowledge on how the various connected devices will react to the steering signal used by the central controller.

The scalability issue present in EM is solved by letting the devices produce their own schedule based on the steering signals they receive. An open question that remains is if and how these schedules can be made efficiently. It is important to note that the computing power present in smart devices is (very) limited and is expected to remain so in the future [126]. Thus, efficient methods for scheduling the use of flexibility provided by these devices are required. In the next part of this thesis we discuss these scheduling problems and several efficient solution methods in more detail.

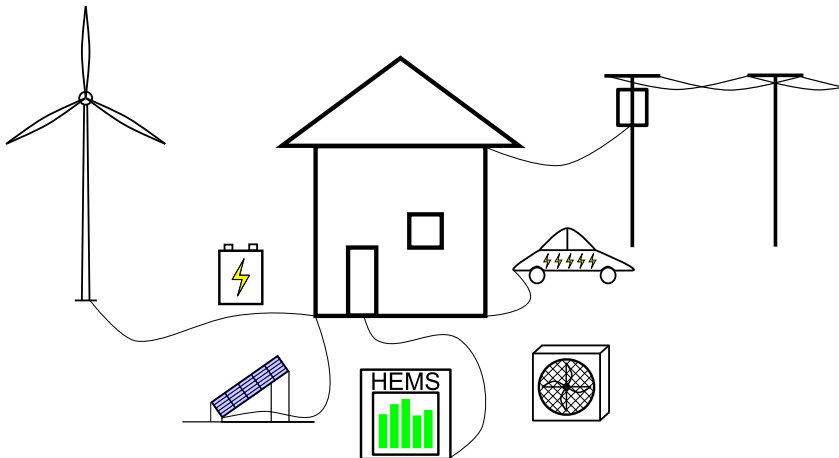


# 4



## CHARGING OF ELECTRIC VEHICLES

**ABSTRACT** – *In this chapter we discuss electric vehicles (EVs), which are a prime example of the electrification of our energy supply chain. EVs offer a large potential to make our transportation sector sustainable. In this chapter we model the scheduling problem an EV faces when it receives steering signals from a controller using the profile steering approach. We show that this problem falls into the class of resource allocation problems. This allows us to use solution approaches for classical resource allocation problems. Furthermore, we extend the model to limit the charging of the vehicle to discrete levels, resulting in an  $\mathcal{NP}$ -hard problem. We approximate this problem using convex combinations. For this approximation we develop an efficient solution approach. We use the obtained solution approaches to simulate a case study of a neighbourhood with a large number of EVs showing the effectiveness of our approach.*



This chapter is based on [TvdK:4] and [TvdK:1]




Figure 4.1: The number of EVs sold per year and the market share of EVs in the EU. Data taken from [48].




Figure 4.2: The number of EVs sold per year and the market share of EVs in the Netherlands. Data taken from [48].

Country	2020 EV stock target (millions)	EV market share of new vehicles 2016-2020 (%)	EV share of total stock in 2020 (%)
Austria	0.2	13	4
China	4.6	6	3
Denmark	0.2	23	9
France	2.0	20	6
Germany	1.0	6	2
India	0.3	2	1
Ireland	0.1	8	3
Japan	1.0	4	2
Netherlands	0.3	10	4
Portugal	0.2	22	5
South Korea	0.2	4	1
Spain	0.2	3	1
United Kingdom	1.6	14	5

Table 4.1: Table denoting EVs penetration targets for several countries with ambitious goals. Data taken from [76].



## 4.1 INTRODUCTION

Our energy supply chain, and specifically the electricity grid are undergoing rapid changes, as sketched in Chapters 1 and 2. One of the major changes is the electrification of our energy supply chain; devices that traditionally use other sources of energy switch to electricity instead. One of the drivers for this switch to electricity for these devices is the abundance of clean and renewable sources available for electricity generation, e.g., wind and sun. An example of this electrification is the introduction and rise in popularity of EVs. In this thesis, we use the term EV to cover both full electric vehicles and plug-in hybrid electric vehicles (*PHEVs*), vehicles that combine an electric motor and a battery system with a conventional combustion engine, where the battery can be charged by plugging into the grid.

To operate the electricity grid of the future in a safe and efficient manner, as mentioned in Chapters 1 and 2, a completely new approach is required. One such approach, presented in Chapter 3, is the decentralized energy management (*DEM*) approach called profile steering. The coordination mechanism of this approach distributes the scheduling problems to the individual devices, using steering signals to communicate the system goals to these devices. It remains to study methods that determine the schedules of the individual devices, based on the received steering signals. In this chapter we focus on the EV as a flexible device. We formulate its (local) scheduling problem and study efficient solutions to this problem.

The number of EVs is rapidly growing in many countries and their uptake is expected to increase worldwide in the coming decades. As an illustration we give the total number of EVs sold as well as the market share of EVs on the total automotive market in the EU in Figure 4.1. While a steep increase in the total number sold as well as the market share is visible, the total share remains low (around 1.1%). The story for the Netherlands is somewhat different, as depicted in Figure 4.2. For the Netherlands extensive subsidiary programs have run to promote electric driving, resulting in a far larger market share in 2015 (nearly 10%). The decline in market share between 2013 and 2014 can be explained by the termination of some subsidiary programs. Furthermore, many countries have set ambitious targets, which we give in Table 4.1.

An EV is a perfect example of a flexible device. Vehicles are commonly used to commute between home and work, a relatively short distance trip with predictable characteristics for many residential consumers. Furthermore, most vehicles only spend little time on the road compared to the time they spend parked and are thus potentially available for charging [147]. On the other hand, the energy that needs to be charged into a depleted EV battery is typically several times larger than the daily energy consumption of a Dutch household [117]. Compare, e.g., a Tesla with a 85 kWh battery and the average Dutch daily electricity consumption of a household of 10 kWh. Furthermore, if EVs are charged at maximum power (i.e., as fast as possible) directly upon arrival at home in the evening this can have potential negative effects on the distribution grid [30]. We observed these negative effects in





the form of large peak loads in the simulation study performed in Chapter 3 (see Figure 3.6). However, as the vehicles are expected to remain stationary the whole night, a perfect opportunity for a DEM approach to manage the charging arises.

The profile steering DEM approach can be used for the management of this charging of an EV to, e.g., minimize stress on the network by flattening the total load profile. For this, a schedule for the EV has to be made by a controller in, e.g., the charging station to which the vehicle is connected. This controller determines the schedule for the charging of the EV based on the received steering signals and (predicted) parameters of the vehicle, e.g., arrival time, departure time, required charging and maximal charging. As such a controller is generally expected to have low computational power, we require efficient solution methods to this scheduling problem. In this chapter, we study two variants of the EV charging problem, both falling into the category of resource allocation problems. For both problems we present efficient solution methods.

The remainder of this chapter is outlined as follows. In the next section we briefly discuss related work on EV charging. Then, in Section 4.3, we introduce the EV charging problem considered in this chapter. For this problem we formulate necessary and sufficient optimality conditions. Next, in Section 4.4, we discuss two solution methods for resource allocation problems commonly applied in literature and apply them to the EV problem. In Section 4.5 we discuss an extension of the EV charging problem that ensures the EV is only charged at specific charging levels. As this extension turns the problem into a hard problem mathematically speaking, we adapt the problem slightly to be able to derive efficient solution approaches. Then, in Section 4.6 we simulate the profile steering approach coordinating the charging of multiple EVs in a neighbourhood and compare this approach to state-of-the-art methodologies. Finally, we draw some conclusions and discuss the obtained results.

## 4.2 RELATED WORK

In this section we briefly sketch related work on EV charging. We note that we focus primarily on problems that are closely related to the problem we study in this chapter (see Section 4.3 for our problem definition). In the literature on the scheduling of EV charging there is a clear distinction between work discussing the optimal charging of a single vehicle and work aimed at optimizing the charging of a fleet of vehicles. The latter case is often implemented using steering signals, similar to our profile steering approach, which are used to decompose the overall problem into optimization problems for each vehicle separately.

### 4.2.1 OPTIMIZATION OF A SINGLE VEHICLE

In [115] a detailed model of the battery use in a PHEV is described. The model is used to determine the optimal balance between using the battery and conventional fuel to supply the energy required for a detailed trip of the vehicle. Simultaneously, an optimal charging strategy for the battery is determined when the vehicle is



plugged into the grid. The charging part is optimized with respect to costs of charging under a time-varying pricing scheme of electricity. The option of offering regulation services to the grid is also considered which leads to a higher net profit in the simulation results. The authors solve the optimization problem using dynamic programming, for which the range of possible charge and discharge values is a discrete set, similarly to the second variant studied in this chapter. We note that the steering signals used in our profile steering approach can be interpreted as time-varying prices to ensure economic scheduling of the vehicle for the owner. Furthermore, regulation services to the grid can also be adopted in the steering signal in several cases. This implies that our solution methods are more general and can be used to solve the problems considered in [115].

The model for EVs in [147] is used to compare different strategies for EV charging based on electricity prices and driving data obtained in Denmark in 2003. A comparison is made between charging during the night and charging at any time during the day. In both cases the option of gaining additional profit by selling energy through vehicle-to-grid (V2G) utilization is also considered. The problem is solved as a mixed integer program, where the optimization is done with respect to the time-varying prices. While the models presented in this chapter do not consider V2G, we make extensions to the model in Chapter 5 to allow for this. This implies that our solution methods can be used to solve the problem considered in [147] in a much more efficient manner than through a mixed integer problem.

#### 4.2.2 OPTIMIZATION OF A FLEET OF VEHICLES

The charging of a fleet of EVs is often optimized from a grid perspective. Examples of objectives are the minimization of the generation costs of the required energy for charging or the minimization of peaks in the load profile.

In [56] the charging of a fleet of EVs is considered on top of a given base load. The optimization goal is assumed to be a general convex function of the sum of the base load and total EV load. It is shown that an overall optimal solution can be found by an iterative algorithm. This iterative algorithm optimizes the charging of each EV separately with respect to a control signal of time-varying prices and the squared distance to the charging profile of the previous step. The authors give no indication on how this local charging problem is solved. The problem considered in [56] fits exactly into the continuous problem discussed in this chapter, implying that the local charging problems can be solved efficiently by our solution approaches.

The methodology proposed in [141] attempts to minimize the cost of the generation of electricity required to charge a fleet of EVs. This methodology aggregates the scheduling possibilities of all EVs in the fleet together into a single master problem. This master problem is then solved using dynamic programming. A control signal, based upon the solution of the master problem, is then sent to the individual EVs. The actual charging done by the EVs is based upon the control signal and the proximity to their deadline. The master problem is very similar in nature to the problems considered in this chapter. However, we allow for more general, con-



vex objective functions. We note that the approach used in [141] is similar to the approach of the Intelligator (see, e.g., [29]).

The approaches found in literature listed here focus primarily on EVs. Our profile steering approach is more general in that it incorporates all types of distributed energy resources (*DERs*). Furthermore, the solution methods we develop and discuss in this chapter can be applied to other appliances that have an internal battery that needs to be charged without allowing discharging to the grid.

### 4.3 ELECTRIC VEHICLE SCHEDULING PROBLEM

When considering the EV we note that it fits into the buffer class of *DERs* as introduced in Chapter 2. The internal battery of the EV corresponds with the buffer in the model, whereby the state of charge indicates how much energy is left inside the battery. The inflow process of the battery is the charging done while the vehicle is plugged in. We assume that this process is controllable. The outflow process is the energy used for driving the vehicle. Because the in- and outflow processes are decoupled in time, we only need to make sure that the inflow process, i.e., the charging of the vehicle, is done in such a way that the energy need of the next outflow process is satisfied. Also, note that some EVs recover energy internally, e.g., from the braking system. As we cannot control such a system, we assume the total energy requirement for the drives already takes such recoveries into account.

To integrate EV charging into the profile steering DEM approach, we have to consider the problem of scheduling the EV charging with respect to received steering signals. As outlined in Chapter 3, the steering signals consist of an objective function  $f(\mathbf{x})$  potentially together with a lower bound  $\mathbf{x}^{min}$  and an upper bound  $\mathbf{x}^{max}$  on the charging schedule. The steering signals thus define the objective of the optimization problem together with bounds on the schedule. This schedule consists of decisions for a finite set of equal length time intervals, i.e., we need to find a decision vector  $\mathbf{x}$  that describes the amount of energy charged into the EV for every time interval. Furthermore, we need to ensure that the charging volume of the EV is equal to the specified level between arrival and departure, where in general this level is fully charged. This leads to the following optimization problem.

**Problem 4.1** (*EVC*).

$$\min_{\mathbf{x}} f(\mathbf{x}), \quad (4.1a)$$

$$s.t. \quad \sum_{t=t^a}^{t^d} x_t = C, \quad (4.1b)$$

$$\mathbf{x}^{min} \leq \mathbf{x} \leq \mathbf{x}^{max}. \quad (4.1c)$$

We call this problem the EV charging (*EVC*) problem. Within Problem *EVC*, constraint (4.1b) ensures that the charging volume between arrival at  $t^a$  and departure



at  $t^d$  is equal to the specified amount  $C$ . In other words the EV is charged with the specified amount  $C$  between arrival and departure. Furthermore, constraint (4.1c) ensures the schedules stays between specific bounds, which result from the bounds given by the steering signal and device specific bounds, e.g., the maximum rate of energy transfer allowed by the charger. In the notation introduced in Chapter 3 the feasible set of schedules  $X^m$  is given as:

$$X^m = \{\mathbf{x} \mid \mathbf{x} \text{ satisfies (4.1b) and (4.1c)}\}. \quad (4.2)$$

Since the EV cannot charge energy while not connected, we have that  $x_t = 0$  for  $t < t^a$  and for  $t > t^d$ . In particular, we can restrict the scheduling horizon for the EV to be exactly the intervals between  $t^a$  and  $t^d$ , i.e., we assume w.l.o.g. that  $t^a = 1$  and  $t^d = T$ .

We note that Problem *EVC* in the given form cannot be used to cover discharging of the vehicle, i.e., V2G. This is because constraint (4.1b) only restricts the state of charge (SoC) of the battery at the end of the optimization horizon. This means that, with V2G enabled, no constraint is given that forbids the SoC of the battery from dropping below zero or rising above the capacity of the battery during the scheduling horizon. This implies that we have to assume that  $\mathbf{x}^{min} \geq 0$ . In Chapter 5 we extend the model and problem to include the case that the vehicle can discharge energy to the grid when beneficial and present also corresponding efficient solution approaches.

Note, that we furthermore may assume w.l.o.g. that the lower bound on the charging done on each time interval is zero, i.e., that  $\mathbf{x}^{min} = 0$ . This follows from the fact that we can transform any instance of Problem *EVC* to an instance with  $\mathbf{x}^{min} = 0$  using the variable substitution  $\mathbf{y} = \mathbf{x} - \mathbf{x}^{min}$ .

To ensure that Problem *EVC* and also other device scheduling problems we encounter later on are tractable we make the following assumption.

**Assumption 4.1.** *The objective function  $f(\mathbf{x})$ , communicated to the devices in the profile steering approach, are convex and separable, i.e.,  $f(\mathbf{x}) = \sum_{t=1}^T f_t(x_t)$ , with each  $f_t$  convex.*

In essence, the objective function  $f(\mathbf{x})$  determines how well the schedule  $\mathbf{x}$  fits the system goals. When the function is separable, the fitness of the schedule for a specific time interval is independent of the fitness on other time intervals. This means that the fitness of the scheduled value for interval  $t$ , given by  $x_t$ , can be expressed by an explicit function  $f_t(x_t)$ , which only depends on  $x_t$ . Note, that this, e.g., implies that the change of the charging level between two consecutive intervals cannot be part of the objective. The convexity assumption implies that each  $f_t$  is convex, i.e., it is "changing at an increasing rate". For more background on convexity we refer the reader to the Appendix.

We note that the Euclidean norm, which we often use to steer the load profile towards a flat profile within profile steering, in its basic form is not a separable



function. However, if we instead consider the square of the Euclidean norm, we get a separable (and convex) function. This modification does not alter the optimal load profiles found by the devices, since a schedule with a smaller (non-negative) objective remains smaller after squaring the objective.

#### 4.3.1 RESOURCE ALLOCATION PROBLEMS

With Assumption 4.1, Problem *EVC* becomes a problem in the class of nonlinear continuous resource allocation problems. Resource allocation problems consider a bounded resource or resources that has/have to be optimally divided over various activities or tasks. In Problem *EVC* the resource is the offered flexibility in the charging, which has to be divided over the time intervals. The problem is continuous since the decision variable, given by  $\mathbf{x}$ , has a continuous feasible set. Furthermore, because we in general do not assume the objective function  $f(\mathbf{x})$  to be linear, the problem is in its general form nonlinear.

In resource allocation problems, often the resource constraint, given by (4.1b), is a less than or equal constraint  $\sum_{t=1}^T x_t \leq C$ , i.e., at most a given amount of the resource can be used. We note that this variant can be reduced to our formulation in the following way. We first solve the relaxation of Problem *EVC* obtained by dropping the resource constraint (4.1b). We note that this turns the problem into  $T$  disjoint one-dimensional convex optimization problems by Assumption 4.1, which are generally easy to solve. Then we calculate the total resource usage  $\sum_{t=1}^T x_t$  for these  $T$  disjoint problems. If  $\sum_{t=1}^T x_t \leq C$  we have found an optimal solution. On the other hand, if  $\sum_{t=1}^T x_t > C$ , then, by the convexity of the objective, it follows that this constraint must be tight in an optimal solution. Hence, we can solve Problem *EVC* to obtain an optimal solution in this case. Note that, when considering EVs, a constraint of the form  $\sum_{t=1}^T x_t \geq C$  (i.e., the vehicle has to charge at least  $C$  units of energy) can be treated similarly to the above.

Continuous resource allocation problems have various application fields. Interestingly enough one of those fields is the field of traditional energy management, where the problem needs to be solved to, e.g., let running generators respond to fluctuations in demand [82, 135, 136]. Another such application field is green computing, where the energy consumption of a processor is minimized, while respecting constraints on the tasks to be processed. Since the relation between processor speed and energy consumption is superlinear, it can be beneficial to slow down the processor to conserve energy if the restrictions on the tasks allow this. The technique used in this field is often referred to as either dynamic voltage and frequency scaling (*DVFS*) or speed scaling. For more details we refer the reader to [58, 87, 146]. The problems solved in this field are similar to those we encounter in DEM, as already noted by Tang *et al.* [124].

Another application area is the area of shipping logistics (see, e.g., [71, 100]). The problem studied in this field is that of finding optimal vessel routes and speed schedules on the obtained routes. Herein a resource allocation problem is used to find the optimal speed of vessels for a fixed route. This problem is solved as



a subroutine of the main solution method to find the combination of routes and speed schedules.

The application of resource allocation problems is not limited to the fields mentioned above. Another field of application lies in economics, in portfolio optimization [42, 90]. Furthermore, the problem is also of importance in strategic planning, e.g., in the field of weapon allocation games [37]. An elaborate history of the class of nonlinear continuous resource allocation problems can be found in the excellent surveys recently published by Patriksson [108] and Patriksson and Strömberg [109].

#### 4.3.2 OPTIMALITY CONDITIONS

Resource allocation problems fall in the class of convex optimization problems (for a background on the concepts used here, see the Appendix). For these problems general solution methods exist based on the well known Karush-Kuhn-Tucker (KKT) conditions [80, 85] or on generalizations of these conditions. Furthermore, these conditions (or a generalization of them) often lead to tailored solution methods for specific problems in this class. In our context, the generalization to non-smooth objective functions using subderivatives is of interest. For our convex objective function  $f(\mathbf{x})$  in particular the left and right derivatives, given by  $f^-(\mathbf{x})$  and  $f^+(\mathbf{x})$ , play an important role. This generalized version of the KKT optimality conditions for Problem EVC can be formulated as follows.

**Lemma 4.1.** *A feasible solution  $\mathbf{x}$  to EVC is optimal if and only if  $\sum_{t \in \mathcal{T}} x_t = C$  and there exists a multiplier  $\lambda$  such that, for all  $t$ :*

$$x_t^{\min} < x_t < x_t^{\max} \quad \Rightarrow \quad f_t^-(x_t) \leq \lambda \leq f_t^+(x_t), \quad (4.3a)$$

$$x_t = x_t^{\min} \quad \Rightarrow \quad f_t^+(x_t) \geq \lambda, \quad (4.3b)$$

$$x_t = x_t^{\max} \quad \Rightarrow \quad f_t^-(x_t) \leq \lambda, \quad (4.3c)$$

where  $f_t^-$  and  $f_t^+$  denote the left and right derivatives of  $f_t$  respectively.

For completeness sake we give a direct proof of the above conditions, i.e., a proof without use of the generalized KKT conditions.

*Proof.* To prove that the conditions are necessary, we consider a solution  $\mathbf{x}$  for which the conditions do not hold and show that we can construct a (feasible) solution with lower objective value. Assume, that for  $\mathbf{x}$ , no  $\lambda$  can be found for which the conditions of the lemma hold. This implies that there must be two indices  $t'$  and  $t$  with  $x_{t'} < x_{t'}^{\max}$  and  $x_t > x_t^{\min}$  such that  $f_{t'}^+(x_{t'}) < f_t^-(x_t)$ . By the convexity of each  $f_t$ ,  $f_t^-$  and  $f_t^+$  exist. Furthermore, they are left-continuous and right-continuous respectively and it follows that

$$f_{t'}^+(x_{t'} + \epsilon) < f_t^-(x_t - \epsilon), \quad (4.4)$$



for  $\epsilon$  small enough. This implies that the solution where  $x_{t'}$  is replaced by  $x_{t'} + \epsilon$  and  $x_t$  by  $x_t - \epsilon$  is a better feasible solution. This in turn implies that  $\mathbf{x}$  is not an optimal solution.

Next we prove the sufficiency of the conditions. To do so, we consider a feasible solution  $\mathbf{x}$  that satisfies the conditions and an arbitrary feasible solution  $\mathbf{y} \neq \mathbf{x}$ . We consider two sets of indices:

$$I^+ := \{t \mid y_t > x_t\}, \quad (4.5a)$$

$$I^- := \{t \mid y_t < x_t\}. \quad (4.5b)$$

Note that, since both  $\mathbf{x}$  and  $\mathbf{y}$  are feasible it follows from (4.3) that  $f_t^+(x_t) \geq \lambda$  for  $t \in I^+$  and  $f_t^-(x_t) \leq \lambda$  for  $t \in I^-$ . Furthermore, it follows that

$$\sum_{t \in I^+} y_t - x_t = \sum_{t \in I^-} x_t - y_t. \quad (4.6)$$

We now obtain

$$\begin{aligned} \sum_{t \in I^+} f_t(y_t) - f_t(x_t) &\geq \sum_{t \in I^+} (y_t - x_t) f_t^+(x_t), \\ &\geq \lambda \sum_{t \in I^+} y_t - x_t, \\ &\geq \sum_{t \in I^-} (x_t - y_t) f_t^-(x_t), \\ &\geq \sum_{t \in I^-} f_t(x_t) - f_t(y_t), \end{aligned} \quad (4.7)$$

where the first and last inequality follow from the convexity of each  $f_t$  and the other two inequalities follow from the fact that  $\mathbf{x}$  satisfies (4.3) and (4.6). Summarizing, we obtain that  $f(\mathbf{y}) - f(\mathbf{x}) = \sum_{t \in I^+} f_t(y_t) - f_t(x_t) + \sum_{t \in I^-} f_t(y_t) - f_t(x_t) \geq 0$  for any feasible  $\mathbf{y}$ . This implies that  $\mathbf{x}$  is indeed optimal.  $\square$

It is worthwhile to note that, in case the objective functions  $f_t$  are differentiable for every  $t$ , the conditions (4.3) simplify slightly. This is because, for  $f_t$  differentiable, we have that  $f_t^-(x_t) = f_t'(x_t) = f_t^+(x_t)$  for all  $x_t$ . Thus, in (4.3) the subderivatives can be replaced with the derivate and (4.3a) simplifies to  $x_t^{\min} \leq x_t \leq x_t^{\max} \Rightarrow f_t'(x_t) = \lambda$ . Below we give an example of an optimal solution to an instance of Problem EVC demonstrating the balance between the derivatives given by Lemma 4.1.

**Example 4.1.** Consider an instance of Problem EVC with  $T = 4$ ,  $C = 5$ ,  $\mathbf{x}^{\min} = (0, 0, 0, 0)$ ,  $\mathbf{x}^{\max} = (2, 2, 2, 2)$ , and objective functions:

- »  $f_1(x_1) = x^2 - 4x + 5$ ,
- »  $f_2(x_2) = 3x^2 - 6x + 4$ ,
- »  $f_3(x_3) = \frac{1}{5}x^2 + \frac{8}{5}x + \frac{11}{5}$ ,
- »  $f_4(x_4) = \frac{3}{2}x^2 - \frac{9}{2}x + \frac{35}{8}$ .



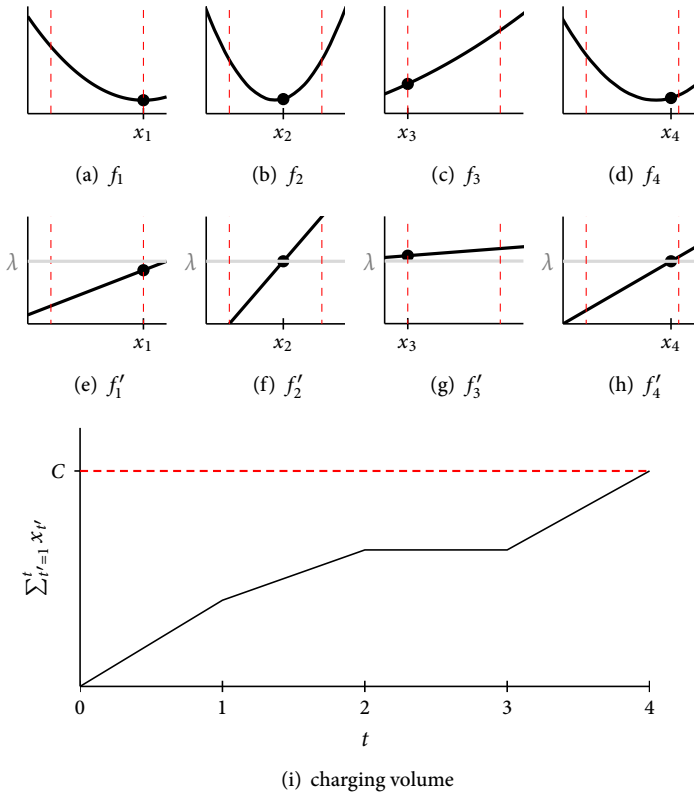


Figure 4.3: Example of the balance between the derivatives as given by Lemma 4.1. Where (a)-(d) depict the objective functions with their respective derivatives given in (e)-(h) and the charging volume given in (i).

For the optimal solution  $\mathbf{x} = (2, 1\frac{1}{6}, 0, 1\frac{5}{6})$  we have  $\lambda = 1$ . This solution is depicted in Figure 4.3. Note that for the optimal solution it holds that  $f'_2(x_2) = f'_4(x_4) = \lambda$ ,  $f'_1(x_1) < \lambda$  and  $f'_3(x_3) > \lambda$ .

#### 4.4 SOLUTION METHODS

In Section 4.3 we formulated Problem *EVC* together with necessary and sufficient optimality conditions. The necessary and sufficient optimality conditions we formulated in Lemma 4.1 hold for convex separable  $f_t$  and simplify if we also assume that each  $f_t$  is differentiable. A distinction can be made between two types of solution approaches to Problem *EVC* with continuous  $f$ , both based on Lemma 4.1. The first type of solution approach attempts to find the value of the Lagrangian multiplier  $\lambda$  as given in Lemma 4.1. Such an approach is often referred to as a La-



grangian multiplier approach. The other approach builds a solution to Problem *EVC* from solutions to relaxations of *EVC*, in which constraints (4.1c) are dropped. This approach is an iterative one where in each iteration one or more variables are fixed at their optimal value (either  $x_t^{min}$  or  $x_t^{max}$ ). This second approach is often called pegging.

#### 4.4.1 LAGRANGIAN MULTIPLIER APPROACH

For the Lagrangian multiplier approach for problem *EVC* with differentiable objective we note the following. We first consider the case that  $f_t$  is *strictly* convex. In this case we know that  $f'_t$  is a strictly increasing function. Furthermore, since  $f_t$  is convex,  $f'_t$  is continuous. Hence, if we guess the value of the multiplier to be  $\hat{\lambda}$ , then, using Lemma 4.1, we can uniquely determine the associated value of the decision variable  $\hat{x}_t$  using:

$$\hat{x}_t(\hat{\lambda}) = \begin{cases} x_t^{min} & \text{if } \hat{\lambda} \leq f'_t(x_t^{min}), \\ g_t(\hat{\lambda}) & \text{if } f'_t(x_t^{min}) < \hat{\lambda} < f'_t(x_t^{max}), \\ x_t^{max} & \text{if } \hat{\lambda} \geq f'_t(x_t^{max}), \end{cases} \quad (4.8)$$

where  $g_t$  is the inverse of  $f'_t$  (which is well-defined because  $f'_t$  is continuous and strictly increasing).

To find an optimal solution we use (4.8) to calculate the value  $\hat{C}(\hat{\lambda}) = \sum_t \hat{x}_t(\hat{\lambda})$ , which depends on the guess  $\hat{\lambda}$ . If  $\hat{C}(\hat{\lambda}) = C$ , we have found an optimal solution. This follows directly from the fact that the solution is feasible and satisfies (4.3). Furthermore, since each  $f'_t$  is a strictly increasing function, each  $g_t$  is also strictly increasing. From this it follows that  $\hat{C}(\hat{\lambda})$  is also continuous and strictly increasing, allowing us to use, e.g., bisection search methods to determine the optimal  $\lambda$ .

Next we consider the case that one or more of the  $f_t$  are convex but not strictly convex. In this case the inverse  $g_t$  is not defined as a real-valued function, but rather as a set function. However, because  $f_t$  is convex, it follows that  $g_t$  always returns an interval. This allows us to specify the functions  $g_t^{min}$  and  $g_t^{max}$  that return the endpoints of this interval. With this we can reformulate (4.8) to:

$$\hat{x}_t(\hat{\lambda}) \in \begin{cases} \{x_t^{min}\} & \text{if } \hat{\lambda} \leq f'_t(x_t^{min}), \\ [g_t^{min}(\hat{\lambda}), g_t^{max}(\hat{\lambda})] & \text{if } f'_t(x_t^{min}) < \hat{\lambda} < f'_t(x_t^{max}), \\ \{x_t^{max}\} & \text{if } \hat{\lambda} \geq f'_t(x_t^{max}). \end{cases} \quad (4.9)$$

For a guess  $\hat{\lambda}$  it may be that some values of  $\hat{x}$  can be freely chosen from an interval. Hence the value of  $\hat{C}$  may be given by an interval  $[\hat{C}^{min}(\hat{\lambda}), \hat{C}^{max}(\hat{\lambda})]$ . If  $C \in [\hat{C}^{min}(\hat{\lambda}), \hat{C}^{max}(\hat{\lambda})]$ , an optimal solution exists for this value of the multiplier  $\hat{\lambda}$ . Because  $f'_t$  is non-decreasing, both  $g_t^{min}(\hat{\lambda})$  and  $g_t^{max}(\hat{\lambda})$  (and hence also both  $\hat{C}^{min}(\hat{\lambda})$  and  $\hat{C}^{max}(\hat{\lambda})$ ) are non-decreasing functions. Furthermore  $\hat{C}^{min}(\bar{\lambda}) > \hat{C}^{max}(\hat{\lambda})$  for  $\bar{\lambda} > \hat{\lambda}$ . Thus again, a bisection search method can be applied to find



the value of  $\lambda$ . After such a  $\lambda$  is found, exact values for each  $x_t$  still have to be determined. Note that, because not all  $f_t$ 's are strictly convex, the existence of a unique optimal solution is no longer guaranteed. In fact, any combination of  $x_t$ 's that satisfies both (4.8) and  $\sum_t x_t = C$  is optimal in this case. Thus a greedy approach can be used to determine values of the  $x_t$ 's such that  $\sum_t x_t = C$ .

The above process can be computationally difficult, depending on the complexity of finding the inverses  $g_t$ . However, in DEM many objectives can be captured using quadratic objective functions, e.g., linear functions for pricing or the sum of squared differences for the Euclidean norm. To this end we study a Lagrangian multiplier approach in the case of quadratic objectives, i.e.,  $f_t(x_t) := \frac{1}{2}a_t x_t^2 + b_t x_t + c_t$  and  $f_t'(x_t) = a_t x_t + b_t$ . In the following we assume that  $a_t > 0$  to simplify the description. However, we note that the approach can be adapted to include the case that  $a_t$  is zero for one or more values of  $t$ .

For the given objective (4.9) reduces to:

$$\hat{x}_t(\hat{\lambda}) = \begin{cases} 0 & \text{if } \hat{\lambda} \leq b_t, \\ \frac{\hat{\lambda} - b_t}{a_t} & \text{if } b_t < \hat{\lambda} < a_t x_t^{\max} + b_t, \\ x_t^{\max} & \text{if } \hat{\lambda} \geq a_t x_t^{\max} + b_t, \end{cases} \quad (4.10)$$

Where we used that w.l.o.g. we assumed that  $x_t^{\min} = 0$ . In the following, we call a time interval  $t$  used if  $x_t > 0$  and maximally used if  $x_t = x_t^{\max}$ .

Let us assume we have a guess  $\hat{\lambda}$  for the value of  $\lambda$ . Using (4.10) we can compute the schedule  $\mathbf{x}$  for this  $\hat{\lambda}$  leading to a total charging of  $\hat{C}$ . The goal is to update  $\hat{\lambda}$  until  $\hat{C} = C$  since then  $\mathbf{x}$  is feasible and hence optimal by Lemma 4.1. To this end, let us first sort the time intervals such that  $b_1 \leq b_2 \leq \dots \leq b_T$ . Note that for any  $\hat{\lambda} \leq b_1$ , it follows that  $\mathbf{x} = 0$  and hence  $\hat{C} = 0$ . Thus we only have to consider  $\hat{\lambda} \geq b_1$  and we can start our approach using  $b_1$  as the initial value of  $\hat{\lambda}$ .

Next we assume that we have a solution  $\mathbf{x}$  for some value  $\hat{\lambda}$  with  $I$  the index set of intervals that are used but not maximally used, i.e.,  $I := \{t \mid 0 < x_t < x_t^{\max}\}$ . When we increase  $\hat{\lambda}$  by  $\Delta$ , where  $\Delta$  is sufficiently small such that no indices are added to or deleted from  $I$ , then  $x_t$  increases by  $\frac{1}{a_t}\Delta$  for each  $t \in I$ . Thus,  $\hat{C}$  increases by  $\eta\Delta$ , with  $\eta = \sum_{t \in I} \frac{1}{a_t}$ . We note that an interval  $t$  becomes used whenever  $\hat{\lambda} > b_t$  and it becomes maximally used when  $\hat{\lambda} \geq a_t x_t^{\max} + b_t$ . Thus,  $I$  only changes when  $\hat{\lambda}$  becomes larger than a value in either of the following two sets:

$$A := \{b_1, b_2, \dots, b_T\}, \quad (4.11a)$$

$$B := \{a_1 x_1^{\max} + b_1, a_2 x_2^{\max} + b_2, \dots, a_T x_T^{\max} + b_T\}. \quad (4.11b)$$

Note that we sorted the time intervals such that  $A$  is non-decreasing.

In our approach, we iteratively increase the value of  $\hat{\lambda}$  to the next value in the combined set  $A \cup B$ . While doing so, we calculate the increase in  $\hat{C}$ . During the



process we ensure that  $\hat{C} \leq C$ , by ensuring that we never increase  $\lambda$  by more than  $\frac{C-\hat{C}}{\eta}$ . Furthermore, as soon as  $\hat{C} = C$  we have found an optimal solution, which we can calculate using (4.10).

The approach sketched above is summarized in Algorithm 4.1. For the complexity of the algorithm, note that each of the steps in the while loop can be done in constant time when using two counters to track the number of deleted elements from the front of  $A$  and  $B$  respectively. Since at least one element is deleted from either  $A$  or  $B$  in each but the final iteration of the while loop, the while loop runs for at most  $O(T)$  iterations. Thus, the sorting step on both  $A$  and  $B$  dominates the complexity, giving a total complexity of  $O(T \log T)$ . We note that Hochbaum and Hong solve the same problem in time  $O(T)$  by exploiting the fact that median find has linear complexity [64]. However, the linear time algorithm for median find is inefficient for small to medium sized  $T$  [18]. Thus, since many scheduling problems encountered in DEM use a relatively small number of time intervals, we expect our approach to be more efficient in practice. Furthermore, we expect our approach to be easier to implement.

---

**Algorithm 4.1** Water filling approach  $WF$  for  $EVC$  with  $f_t = \frac{1}{2}a_t x_t^2 + b_t x_t + c_t$

---

```

1: x = Function  $WF(f, \mathbf{x}^{max}, C)$ 
2: Set  $A, B$  according to (4.11) and order them both non-decreasingly
3: Take  $\alpha$  as the first value from  $A$  with associated interval  $t$  and delete it from  $A$ 
4: Set  $\hat{\lambda} := \alpha$ ,  $\hat{C} := 0$  and  $\eta := \frac{1}{a_t}$ 
5: while  $\hat{C} < C$  do
6:   Take  $\alpha$  the first element from  $A$  with associated time interval  $t$ 
7:   Take  $\beta$  the first element from  $B$  with associated time interval  $t'$ 
8:    $\gamma := \min\{\alpha, \beta\}$ ,  $\Delta := \min\{\gamma - \hat{\lambda}, \frac{C-\hat{C}}{\eta}\}$  {Calculate increase of  $\hat{\lambda}$ }
9:    $\hat{\lambda} = \hat{\lambda} + \Delta$ ,  $\hat{C} = \hat{C} + \Delta\eta$  {Increase  $\hat{\lambda}$  and  $\hat{C}$ }
10:  if  $\hat{C} < C$  and  $\alpha = \gamma$  {Interval  $t$  becomes used} then
11:     $A = A \setminus \{\alpha\}$ 
12:     $\eta = \eta + \frac{1}{a_t}$ 
13:  else if  $\hat{C} < C$  and  $\beta = \gamma$  {Interval  $t'$  is now maximally used} then
14:     $B = B \setminus \{\beta\}$ 
15:     $\eta = \eta - \frac{1}{a_{t'}}$ 
16:  end if
17: end while
18:  $\mathbf{x} = \hat{\mathbf{x}}(\hat{\lambda})$  {Using (4.10)}
19: Return  $\mathbf{x}$ 

```

---

**Example 4.2.** As an example we applied Algorithm 4.1 to the instance of Problem  $EVC$  given in Example 4.1. The results are depicted in Figure 4.4. In Figure 4.4(a) a graphical representation of the optimal solution is given. In Figure 4.4(b) the values of the various variables used in the algorithm are given for the iterations of the while loop.



We note that the current value of  $\gamma$  is the marked red value in either A or B. After iteration 4, it holds that  $\hat{C} = 5 = C$ , thus the while loop terminates.

Algorithm 4.1 is known in the literature as the water filling approach, where it has been applied in various fields [97, 148]. The reason for the name water filling is that an optimal solution can be found by pouring an amount of water equal to  $C$  on the graph given in Figure 4.4(a). In this graph, a floor has been added for every time interval equal to  $b_t$  and a ceiling equal to  $a_t x_t^{max} + b_t$ . Furthermore, the width of the interval is scaled such that  $a_t x_t + b_t$  increases by the appropriate amount whenever additional water is poured on the graph. The value of each  $x_t$  is given by the amount of water in bar  $t$ . Since the water evenly distributes itself over the intervals the derivatives automatically balance in this approach. This implies that the result is an optimal solution. We note that a minor modification from the original approach, where upper bounds were not present, is needed. To prevent blocks from becoming isolated, like the block for  $t = 4$  in Figure 4.4, the infeasible region is assumed to be porous without holding water.

#### 4.4.2 PEGGING APPROACH

The pegging approach for solving continuous nonlinear resource allocation problems is based on two observations. The first is that for a lot of these problems the relaxation when dropping the constraints on the individual usage of the resource, i.e., constraints (4.1c) in Problem EVC, is easy to solve. In particular, if each  $f_t$  is differentiable, it follows from Lemma 4.1 that, for an optimal solution to the relaxation the derivatives must be in balance, i.e.  $f'_t(x_t) = \lambda$  for all  $t$ . This means that the objective cannot be improved by increasing a variable  $x_t$  while decreasing another variable by the same amount. After such a solution  $\mathbf{x}$  of the relaxation is found two sets  $I^+$  and  $I^-$  are considered, which contain the intervals that violate  $x_t \leq x_t^{max}$  and  $x_t \geq x_t^{min}$  respectively. More formally:

$$I^+ := \{t \mid x_t > x_t^{max}\}, \quad (4.12a)$$

$$I^- := \{t \mid x_t < x_t^{min}\}. \quad (4.12b)$$

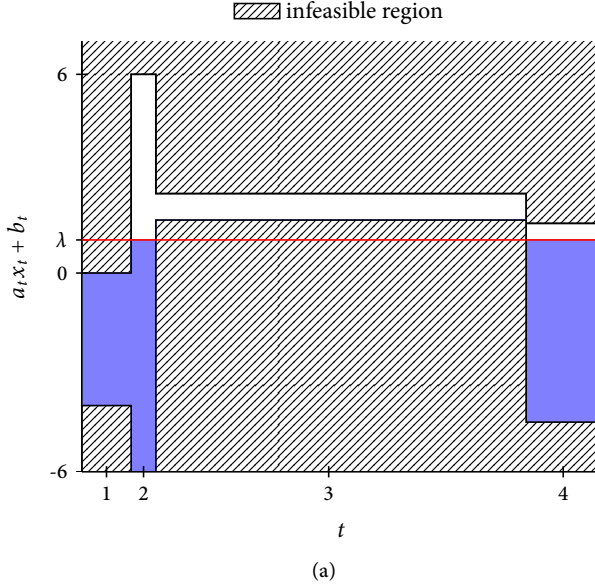
Using these sets, we can calculate the total violations of the upper and lower bounds on the  $x_t$ 's, given by  $\Delta^+$  and  $\Delta^-$  respectively:

$$\Delta^+ := \sum_{t \in I^+} x_t - x_t^{max}, \quad (4.13a)$$

$$\Delta^- := \sum_{t \in I^-} x_t^{min} - x_t. \quad (4.13b)$$

The second important observation is formulated in the following lemma.

**Lemma 4.2.** *Consider Problem EVC and an optimal solution  $\mathbf{x}$  to the relaxation obtained by dropping constraints (4.1c) together with the corresponding sets  $I^-$  and  $I^+$  and values  $\Delta^-$  and  $\Delta^+$ . Furthermore, assume that  $\mathbf{x}$  is not feasible to the original problem. Then, there exists an optimal solution  $\bar{\mathbf{x}}$  to the original problem such that:*



Step	A	B	$\hat{\lambda}$	$\hat{C}$	$\eta$
Initial	$\{-6, -4\frac{1}{2}, -4, 1\frac{3}{5}\}$	$\{0, 1\frac{1}{2}, 2\frac{2}{3}, 6\}$	-6	0	$\frac{1}{6}$
1	$\{-4\frac{1}{2}, -4, 1\frac{3}{5}\}$	$\{0, 1\frac{1}{2}, 2\frac{2}{3}, 6\}$	-4.5	$\frac{1}{4}$	$\frac{1}{2}$
2	$\{-4, 1\frac{3}{5}\}$	$\{0, 1\frac{1}{2}, 2\frac{2}{3}, 6\}$	-4	$\frac{1}{2}$	1
3	$\{1\frac{3}{5}\}$	$\{0, 1\frac{1}{2}, 2\frac{2}{3}, 6\}$	0	$4\frac{1}{2}$	$\frac{1}{2}$
4	$\{1\frac{3}{5}\}$	$\{1\frac{1}{2}, 2\frac{2}{3}, 6\}$	1	5	$\frac{1}{2}$

(b)

Figure 4.4: Results obtained when applying Algorithm 4.1 to the instance of EVC given in Example 4.1. (a) gives a graphical representation of the optimal solution and (b) gives the values of the variables used in the algorithm for the iterations of the while loop.

- » If  $\Delta^+ \geq \Delta^-$  then  $\bar{x}_t = x_t^{max}$  for all  $t \in I^+$ .
- » If  $\Delta^+ \leq \Delta^-$  then  $\bar{x}_t = x_t^{min}$  for all  $t \in I^-$ .

*Proof.* We first consider the case where  $\Delta^+ \geq \Delta^-$ . Assume that for an optimal solution  $\bar{x}$  there is an index  $t^* \in I^+$  such that  $\bar{x}_{t^*} < x_{t^*}^{max} < x_{t^*}$ . Since  $\sum_t \bar{x}_t = C = \sum_t x_t$  it follows that there must be an index  $t'$  for which  $\bar{x}_{t'} > x_{t'}$ . Since  $\Delta^+ \geq \Delta^-$  it follows that we can assume that also  $\bar{x}_{t'} > x_{t'}^{min}$ . Furthermore, from Lemma 4.1, it follows that  $f_{t^*}^-(x_{t^*}) \leq f_{t^*}^+(x_{t'})$  and  $f_{t'}^-(\bar{x}_{t'}) \leq f_{t^*}^+(\bar{x}_{t^*})$ . Summarizing we obtain:

$$f_{t'}^+(x_{t'}) \leq f_{t'}^-(\bar{x}_{t'}) \leq f_{t^*}^+(\bar{x}_{t^*}) \leq f_{t^*}^-(x_{t^*}) \leq f_{t'}^+(x_{t'}). \quad (4.14)$$



Therefore, all the inequalities in (4.14) must be equalities. This implies that  $f'_{t^*}$  and  $f'_{t'}$ , exist, are constant, and are equal on the intervals  $[\bar{x}_{t^*}, x_{t^*})$  and  $(x_{t'}, \bar{x}_{t'}]$  respectively. This in term implies that we can increase  $\bar{x}_{t^*}$  and decrease  $\bar{x}_{t'}$  without changing the objective value. We can repeat the above argument until  $\bar{x}_{t^*} = x_{t'}^{max}$ .

The proof for the case  $\Delta^+ \leq \Delta^-$  is symmetrical.  $\square$

The observation given in Lemma 4.2 ensures that whenever a relaxation of Problem *EVC* is solved either we obtain an optimal solution or we get a set of indices for which the optimal values are at a specific bound. This implies that the problem can be solved iteratively, whereby in every iteration the relaxation of Problem *EVC* is solved. This solution is either optimal or it gives us several variables that are fixed at a bound, the latter is often referred to as pegging these variables. These pegged variables can consecutively be removed from the problem, resulting in a smaller remaining problem to be solved. This process terminates in at most  $T$  steps with an optimal solution.

In the following we study a pegging approach for the specific case that each  $f_t$  is strictly quadratic, i.e.,  $f_t(x_t) = a_t x_t^2 + b_t x_t + c_t$  with  $a_t > 0$ . For this pegging approach, we assume some variables might already be pegged and denote the set of variables that are currently not pegged by  $I$ . We denote by  $\hat{C}$  the amount of charging that still has to be done by the unpegged variables. When solving the relaxation in this case, we know that in the corresponding optimal solution  $\mathbf{x}$  the derivatives of each  $f_t$  for  $t \in I$  are equal. Thus, for any two time intervals  $t', t \in I$  it holds that  $a_{t'} x_{t'} + b_{t'} = a_t x_t + b_t$ , or, equivalently;

$$x_t = \frac{a_{t'} x_{t'} + b_{t'} - b_t}{a_t}. \quad (4.15)$$

Furthermore, since  $\sum_t x_t = \hat{C}$ , we obtain

$$\sum_{t \in I} \frac{a_{t'} x_{t'} + b_{t'} - b_t}{a_t} = \hat{C}. \quad (4.16)$$

From this we obtain

$$x_{t'} = \frac{\hat{C} + \sum_{t \in I} \frac{b_t}{a_t} - b_{t'}}{\sum_{t \in I} \frac{1}{a_t}}, \quad (4.17)$$

which we can use to calculate the value of  $x_t$  for every  $t \in I$ . The resulting pegging approach for *EVC* with quadratic objective is summarized in Algorithm 4.2. The complexity of the approach is  $O(T^2)$ . This follows from the observation that all steps before a potential recursive call can be calculated in linear time. Furthermore, in the worst case only a single variable is pegged in each call, resulting in a total complexity of  $O(T + T - 1 + \dots + 1) = O(T^2)$ .

**Example 4.3.** *As an example we applied Algorithm 4.2 to the instance of EVC given in Example 4.1. The values of the various variables used in the algorithm are given*

---

**Algorithm 4.2** Pegging approach *PEG* for *EVC* with  $f_t = \frac{1}{2}a_t x_t^2 + b_t x_t + c_t$ 


---

```

1:  $\mathbf{x} = \text{Function } PEG(f, \mathbf{x}^{max}, I, \hat{C})$ 
2:  $\Delta^+, \Delta^- := 0$ 
3:  $I^+, I^- := \emptyset$ 
4: for  $t \in I$  do
5:   Calculate  $x_t$  using (4.17)
6:   if  $x_t < x_t^{min}$  then
7:      $\Delta^- := \Delta^- + x_t^{min} - x_t$ 
8:      $I^- = I^- \cup \{t\}$ 
9:   else if  $x_t > x_t^{max}$  then
10:     $\Delta^+ := \Delta^+ + x_t - x_t^{max}$ 
11:     $I^+ = I^+ \cup \{t\}$ 
12:   end if
13: end for
14: if  $\Delta^+ \geq \Delta^-$  and  $\Delta^+ > 0$  then
15:   Set  $x_t = x_t^{max}$  for all  $t \in I^+$ 
16:    $\hat{C} = \hat{C} - \sum_{t \in I^+} x_t^{max}$ 
17:    $I = I \setminus I^+$ 
18:   Calculate  $x_t$  for  $t \notin I^+$  using  $PEG(f, \mathbf{x}^{max}, I, \hat{C})$ 
19: else if  $\Delta^- > \Delta^+$  then
20:   Set  $x_t = x_t^{min}$  for  $t \in I^-$ 
21:    $\hat{C} = \hat{C} - \sum_{t \in I^-} x_t^{min}$ 
22:    $I = I \setminus I^-$ 
23:   Calculate  $x_t$  for  $t \notin I^-$  using  $PEG(f, \mathbf{x}^{max}, I, \hat{C})$ 
24: end if
25: Return  $\mathbf{x}$ 

```

---

Step	$I$	$C$	$(x_1, x_2, x_3, x_4)$	$\Delta^+$	$I^+$	$\Delta^-$	$I^-$
0	{1, 2, 3, 4}	5	$(2\frac{9}{14}, 1\frac{3}{14}, -\frac{11}{14}, 1\frac{13}{14})$	$\frac{9}{14}$	{1}	$\frac{11}{14}$	{3}
1	{1, 2, 4}	5	$(2\frac{1}{4}, 1\frac{1}{12}, 0, 1\frac{4}{6})$	$\frac{1}{4}$	{1}	0	$\emptyset$
2	{2, 4}	3	$(2, 1\frac{1}{6}, 0, 1\frac{5}{6})$	0	$\emptyset$	0	$\emptyset$

Table 4.2: Values of the various variables used in Algorithm 4.2 on the instance of Problem *EVC* given in Example 4.1.

in Table 4.2. In the first call of the algorithm (recursive depth zero)  $\Delta^- > \Delta^+$  and hence  $x_3$  is pegged to be equal to zero. Then, in the first recursive call (depth of one),  $\Delta^+ > \Delta^-$  and  $x_1$  is pegged to be equal to  $x_1^{max} = 2$ . Finally, the optimal solution is found in the second recursive call since  $\Delta^- = \Delta^+ = 0$ .

The pegging approach can also be used to solve extensions of Problem *EVC* to optimality. For example, Bretthauer and Shetty studied Problem *EVC* with a quadratic







objective [23], where they replaced the resource constraint (4.1b) with

$$\sum_t a_t x_t \leq C, \quad (4.18)$$

and  $a_t \geq 0$  for all  $t$ . They also added generalized upper bound constraints, modelled by

$$\sum_{t \in I_k} b_t x_t \leq C_k, \quad k = 1, 2, \dots, K, \quad (4.19)$$

where  $I_1, I_2, \dots, I_K$  are disjoint subsets of the index set  $\mathcal{T}$ . Their solution approach divides the problem into several subproblems, which they solve using the pegging approach.

The same authors also showed in [24] that the pegging approach applies to a generalization of Problem *EVC* where the resource constraint (4.1b) is replaced with the more general constraint

$$\sum_t g_t(x_t) \leq C, \quad (4.20)$$

with each  $g_t$  convex and differentiable. They study both the continuous and discrete variant, where the latter is obtained by the further addition of the constraint that each  $x_t \in \mathbb{Z}$ . Below we formulate a discrete version of the problem that only allows a finite set of values for each  $x_t$ , which makes the problem difficult. The standard approach for discrete resource allocation problems no longer applies. Furthermore, in the next chapters we extend the problem to include both a lower and an upper bound on the total charged volume up to each time interval.

## 4.5 DISCRETE VARIANT

In Section 4.3 we introduced the EV charging problem *EVC* and classified it as a continuous resource allocation problem. Solutions to this problem, e.g., obtained through the methods discussed in Section 4.4, are used in the profile steering DEM approach to obtain charging schedules for EVs and other battery powered devices that only charge. However, in practice not always each of the achieved schedules may be applicable. This is because of the modelling assumption that the allowed charging levels have a continuous range, i.e.,  $\mathbf{x}$  is considered as continuous, but this may not hold in all cases. For example, several EV charging stations currently available only allow charging at specific amperages [81]. In the following we extend Problem *EVC* to incorporate such constraints.

### 4.5.1 COMPLEXITY OF CHARGING OVER A DISCRETE SET

To model the mentioned restriction we replace constraint (4.1c) by the restriction that  $x_t$  must lie in a finite set. To this end we assume that a finite set  $Z_t := \{z_t^0, z_t^1, \dots, z_t^{m_t}\}$  with  $z_t^j < z_t^{j+1}$  given for every index  $t$ . If we replace constraint (4.1c) by the constraint that  $x_t$  must lie in this set  $Z_t$  we obtain the following optimization problem:

**Problem 4.2** (*dEVC*).

$$\min_{\mathbf{x}} f(\mathbf{x}), \quad (4.21a)$$

$$s.t. \sum_{t=1}^T x_t = C, \quad (4.21b)$$

$$x_t \in Z_t \quad t = 1, 2, \dots, T. \quad (4.21c)$$

We call this problem the discrete EV charging problem (*dEVC*). Similar to the assumption that  $\mathbf{x}^{min} = 0$  for Problem *EVC*, we can assume that  $z_t^0 = 0$  for every  $t$ , since we can always apply the transformation  $y_t = x_t - z_t^0$  for every  $t$ . Furthermore, we note that we explicitly do not assume that the sets  $Z_t$  are equal for every  $t$ . This allows us to model, e.g., grid restrictions at certain times in the feasible set of operational points.

Discrete (or integer) resource allocation problems have been studied in the literature. However, these integer problems differ from the problem we consider here. The distinction lies in the fact that the common assumption in the literature is that  $x_t \in \mathbb{Z}$  for every  $t$  is added as a constraint to arrive at a discrete resource allocation problem. For such problems a proximity results holds (see [64] and the references therein). This proximity result states that for any optimal solution to the integer problem of any resource allocation problem, there is an optimal solution to the continuous variant that lies relatively close (in fact, this result holds for the more general class of convex optimization problems with submodular constraints [64]. For an introduction to optimization problems with submodular constraints see, e.g., [12]). This proximity result implies that an optimal solution to the integer problem can efficiently be obtained from an optimal solution to the continuous problem. Thus, an efficient solution method to the continuous problem also (indirectly) results in an efficient method to solve the integer problem.

The common assumption that  $x_t \in \mathbb{Z}$ , which is equivalent to assuming that the distance between two consecutive points in each  $Z_t$  is the same, does not hold for our problem. This impacts the complexity of Problem *dEVC*.

**Lemma 4.3.** *The decision problem of determining if a feasible solution to Problem dEVC exists is  $\mathcal{NP}$ -complete, even if all sets  $Z_t$  are the same.*

*Proof.* The problem of existence of a feasible solution is in  $\mathcal{NP}$ , since verification of the feasibility of a given solution can be done in polynomial time. The  $\mathcal{NP}$ -completeness in case the sets  $Z_t$  may differ follows from a reduction from the partition problem. For this we define  $Z_t := \{0, p_t\}$  and  $C = \frac{1}{2} \sum_t p_t$ , where  $p_1, p_2, \dots, p_T$  are the integers from the set to be partitioned.

For the case that all  $Z_t$  are equal we assume that  $Z_t = \{z_0, z_1, \dots, z_m\}$  for every  $t$  and use a reduction from even/odd partition. In the even/odd partition problem a





set  $P = \{p_1, p_2, \dots, p_{2l}\}$  of  $2l$  non-negative integers is given with total sum  $2B$ . The problem is to determine if there is a subset  $A \subset \{1, 2, \dots, 2l\}$  such that  $\sum_{i \in A} p_i = B$  and for  $i = 1, 2, \dots, l$  we have:  $2i - 1 \in A$  iff  $2i \notin A$ .

Consider an instance  $I$  of even/odd partition and let  $k := \lfloor \log_2(2B) \rfloor$ , i.e.  $k$  is the unique integer such that  $2^k \leq 2B < 2^{k+1}$ . To transform this instance of even/odd partition to an instance  $I'$  of  $dEVC$  we take  $T = m = 2l$ . Furthermore, we choose  $z_{2t-1} = p_{2t-1} + 2^{k+t}$  and  $z_{2t} = p_{2t} + 2^{k+t}$  for  $t = 1, 2, \dots, l$  and  $C = B + \sum_{i=1}^l 2^{k+i}$ . In the following we show that  $I$  is a yes-instance iff  $I'$  is a yes-instance.

First assume that  $I$  is a yes-instance. Thus there exists a subset  $A$  with  $\sum_{i \in A} p_i = B$  and  $2i \in A$  iff  $2i - 1 \notin A$ . By defining  $x_{2t} = p_{2t} + 2^{k+t}$  if  $2t \in A$  and  $x_{2t} = 0$  otherwise, and  $x_{2t-1} = p_{2t-1} + 2^{k+t}$  if  $2t - 1 \in A$  and  $x_{2t-1} = 0$  otherwise, it now follows that  $\sum_{t=1}^T x_t = \sum_{i \in A} p_i + \sum_{i=1}^l 2^{k+i} = C$ .

On the other hand, assume that  $I'$  is a yes-instance for the  $dEVC$  problem and  $\mathbf{x}$  a corresponding solution. Note that the  $l$  most significant bits of  $C$  in binary representation are 1 and correspond to  $2^{k+1}, 2^{k+2}, \dots, 2^{k+l}$ . Since  $\sum_{t=1}^T x_t = C$ , it follows that for  $i = 1, 2, \dots, l$  there is exactly one of the values  $p_{2i} + 2^{k+i}$  or  $p_{2i-1} + 2^{k+i}$  contained in  $\sum_{t=1}^T x_t$ . If we now define  $A$  as the subset of  $\{1, 2, \dots, T\}$  of the indices of these values that are contained in  $\sum_{t=1}^T x_t$ , it follows that  $2i \in A$  iff  $2i - 1 \notin A$ . Furthermore, by construction we have  $\sum_{i \in A} p_i = \sum_{t=1}^T x_t - \sum_{i=1}^l 2^{k+i} = B$ .

This shows that checking whether a feasible solution to Problem  $dEVC$  exists is  $\mathcal{NP}$ -complete, even if we assume that the  $Z_t$  are the same for every index  $t$ .  $\square$

Lemma 4.3 shows that determining if a feasible solution to Problem  $dEVC$  exists is already difficult. In the context of the profile steering DEM approach we need efficient algorithms to solve the scheduling problems for the individual devices. This implies that we should not aim for an optimal approach but for a heuristic approach instead to incorporate a finite set of potential charging levels for the EV into our scheduling problems.

#### 4.5.2 PIECEWISE LINEAR APPROXIMATION

Before discussing a possible approach for Problem  $dEVC$ , we first consider the interpretation of a schedule  $\mathbf{x}$  for an EV. Such a schedule gives, for every time interval  $t$ , the amount of energy  $x_t$  to be charged. As discussed in Chapter 2, we assume that this energy is charged at a (near) constant power throughout the time interval. For instance, if we schedule the charging of the EV in fifteen minute time intervals, then a scheduled value of one kWh charged in a time interval translates into four kW of power drawn from the grid for these fifteen minutes. The limitation to a finite set mentioned previously is exactly on this amount of *power* drawn at any given point in time. However, limiting the EV to only charge an amount of *energy* equal to charging at any of the values in the finite set for a full time interval does not represent the options available in practice. Different amounts of energy can be charged into the EV by changing to a different (allowed) power level *during* the



time interval. For instance, charging for 7.5 minutes at four kW and for the other 7.5 minutes at two kW results in a total of 0.75 kWh charged. In fact, any value between 0.5 and 1 kWh can be charged into the EV by a combination of charging at two and four kW in a fifteen minute time interval. This strategy is practically applicable, since tests have shown that switching between different charging levels can be done in seconds or even faster [81].

To incorporate the above into our scheduling model, we allow the decision variable  $x_t$  to take on any convex combination of the points in the set  $Z_t$ . To do this, we modify problem *dEVC* into:

**Problem 4.3** (*rdEVC*).

$$\min_y \sum_{t=1}^T F_t(x_t), \quad (4.22a)$$

$$\text{s.t. } \sum_{t=1}^T x_t = C, \quad (4.22b)$$

$$x_t = \sum_{j=0}^{m_t} y_t^j z_t^j \quad i = 1, 2, \dots, n, \quad (4.22c)$$

$$\sum_{j=0}^{m_t} y_t^j = 1 \quad i = 1, 2, \dots, n, \quad (4.22d)$$

$$y_t^j \geq 0 \quad t = 1, 2, \dots, T; j = 0, 1, \dots, m_t. \quad (4.22e)$$

We denote the objective function for time interval  $t$  here with  $F_t$  instead of  $f_t$ , as the same amount of energy charged can be obtained using different charging levels. Furthermore, we call this the relaxed discrete EV charging problem (*rdEVC*). Within Problem *rdEVC*, the variable  $y_t^j$  is a multiplier that indicates the fraction of time interval  $t$  the vehicle charges with a power of  $z_t^j$ . We note that the use of convex combinations is also common practice in the field of DVFS [86, 87], where similar problems are studied.

As mentioned, the objective function  $F_t$  expresses the objective for a specific charging volume instead of a charging level. Below we discuss how we can obtain  $F_t$  from the objective functions  $f_t$  of the original problem. The most straightforward way to do so is to use the original objective function at the convex combination of the charging levels, i.e.,  $F_t(x_t) = f_t(\sum_j y_t^j z_t^j)$ . However, as discussed before, this does not adequately reflect factors such as the stress put onto the network. This is because charging at both a high and a low power in an interval puts considerably more stress on the network than charging at the average of the two. To take this into account we instead consider the convex combination of the values of the original objective function at the various charging levels, i.e.,  $F_t(x_t) = \sum_j y_t^j f_t(z_t^j)$ .

For this case we note that, since we are minimizing, any optimal solution will always pick a combination of charging levels that gives the best objective. In other



words, to obtain a charging volume  $x_t$ , the multipliers  $y_t^j$  are always chosen such that  $F_t(x_t) = \sum_{j=0}^{m_t} y_t^j f_t(z_t^j)$  is minimal. Since  $f_t$  is convex, this implies that at most two consecutive multipliers  $y_t^{j-1}$  and  $y_t^j$  are non-zero. This in turn implies that the value of  $F_t$  at  $x_t$  is given by  $y_t^{j-1} f_t(z_t^{j-1}) + y_t^j f_t(z_t^j)$  with  $z_t^{j-1} \leq x_t < z_t^j$  and  $y_t^{j-1} z_t^{j-1} + y_t^j z_t^j = x_t$ . In other words, we can replace  $F_t$  by the piecewise linear convex function obtained by linearizing  $f_t$  on each of the intervals  $(z_t^{j-1}, z_t^j)$  for  $j = 1, 2, \dots, m_j$ . In summary, solving *rdEVC* with  $F_t = \sum_{j=0}^{m_t} y_t^j f_t(z_t^j)$  reduces to solving *EVC* with a convex piecewise linear objective.

#### 4.5.3 SOLUTION METHODS

To solve Problem *EVC* with a piecewise linear objective function, we can use any solution approach (see Section 4.4) that is applicable with non-differentiable objectives. However, as mentioned before, we are interested in very fast and efficient solution methods, since the considered problems generally have to be solved often and on embedded platforms. To this end, we study Problem *EVC* with a piecewise linear objective in more detail. The objective function  $f_t$  is piecewise linear with the pieces on the intervals  $(z_t^0, z_t^1), (z_t^1, z_t^2), \dots, (z_t^{m_t-1}, z_t^{m_t})$ , where  $z_t^j$  is called a breakpoint of  $f_t$ . We use  $s_t^j$  to denote the slope of the piece with left breakpoint  $z_t^{j-1}$  and right breakpoint  $z_t^j$ , which is given by  $s_t^j := \frac{f_t(z_t^j) - f_t(z_t^{j-1})}{z_t^j - z_t^{j-1}}$ . Note that by the convexity of  $f_t$ , it follows that  $s_t^1 \leq s_t^2 \leq \dots \leq s_t^{m_t}$  and we can assume, w.l.o.g., that these inequalities are strict. Furthermore, we call a piece *used* by a solution  $\mathbf{x}$  if  $x_t$  is at least as large as the left breakpoint of this piece and *completely used* by  $\mathbf{x}$  if  $x_t$  is at least as large as the right breakpoint. Furthermore, we call a piece *active* for solution  $\mathbf{x}$  if it is used but not completely used, i.e., a piece with breakpoints  $z_t^{j-1}$  and  $z_t^j$  is active if  $z_t^{j-1} \leq x_t < z_t^j$ . Note that only one piece can be active per time interval. As an example see Figure 4.5, where a function is depicted for which the first two pieces are used, the first piece is completely used and the second piece is active.

When we increase the charging on interval  $t$  for a solution  $\mathbf{x}$  by  $\delta$ , with  $\delta$  sufficiently small, the objective increases by  $\delta$  multiplied by the slope of the active interval for  $t$ . To construct a greedy solution to Problem *EVC* with piecewise linear objective, we can iteratively increase the charging on the time interval that gives the lowest increase in the objective value, i.e., the time interval with active piece with the smallest slope. In each iteration we increase the charging until either a new piece becomes active or until the total charging is equal to  $C$ . Below we show that this greedy approach is indeed optimal. To do so, we show that, in an optimal solution, any piece with a slope smaller than the slope of a used piece must be completely used.

**Lemma 4.4.** *Consider an instance of Problem *EVC* where each  $f_t$  is a piecewise linear function. Furthermore, assume that a piece with slope  $s$  is used in an optimal solution (i.e., the corresponding  $x_t$  is at least as large as the left breakpoint of this piece). Then,*




Figure 4.5: Example of a piecewise linear objective function  $f_t$  with value  $x_t$  where the first two pieces are used and the second piece is active.

any piece with a slope smaller than  $s$  is completely used by an optimal solution (i.e., the corresponding  $x_t$  is at least as large as the right breakpoint of this piece).

*Proof.* Consider an optimal solution  $\mathbf{x}$  and assume there is a piece with slope  $s' < s$ , which is not completely used, with corresponding objective function  $f_{t'}$ . Let  $s^*$  be the slope of the active piece for interval  $t'$ . By the convexity of  $f_{t'}$  it follows that  $f_{t'}^+(x_{t'}) \leq s^* \leq s' < s \leq f_{t'}^-(x_t)$ . This is a contradiction with Lemma 4.1.  $\square$

From Lemma 4.4 it follows that the greedy approach we described above is optimal. To implement this greedy approach we need to sort the pieces such that their slopes are non-decreasing and then iteratively increase the charging on the time interval with the smallest slope, deleting this slope from the sorted list, until  $\sum_t x_t = C$ . Note that the smallest slope considered in each iteration of this greedy approach does indeed correspond to the slope of an active piece, by the convexity of the functions  $f_t$ . A straightforward implementation considers all pieces simultaneously and sorts them based on their slopes, resulting in a complexity of  $O(M \log M)$ , where  $M$  is the sum of the number of pieces  $m_t$  of all  $T$  objective functions. However, we can reduce the complexity to  $O(M \log T)$  by only considering a single piece for each objective function at a time, as only the currently active piece needs to be considered. The resulting greedy approach is summarized below in Algorithm 4.3. We use  $Z$  to denote the set containing all the breakpoints of all functions  $f_t$ .

**Lemma 4.5.** *Algorithm 4.3 solves Problem EVC with piecewise linear convex objective functions to optimality in  $O(M \log T)$  time, where  $M$  is the sum of the number of pieces of the objective functions.*

*Proof.* The feasibility of the algorithm follows from the convexity of the  $f_t$ 's and the optimality follows directly from Lemma 4.4.




---

**Algorithm 4.3** Greedy approach *pwLEVC* for *EVC* with piecewise linear objective
 

---

```

1: x = Function pwLEVC(f, Z, C)
2: Take  $S = \{s_1^1, s_2^1, \dots, s_T^1\}$  and order  $S$  non-decreasingly
3: Set  $x_t = 0$  for all  $t$ 
4: while  $C > 0$  do
5:   Take the first slope  $s_t^j$  from  $S$  with associated  $t$  and  $j$ 
6:    $\delta := \min\{C, z_t^j - z_t^{j-1}\}$ 
7:    $x_t = x_t + \delta$ 
8:    $C = C - \delta$ 
9:    $S = S \setminus \{s_t^j\}$ 
10:  if  $j < m_t$  then
11:    Insert  $s_t^{j+1}$  into the ordered set  $S$ 
12:  end if
13: end while
14: Return x

```

---

To prove the stated time complexity, note that the first sorting in Line 2 can be done in time  $O(T \log T)$ . Furthermore, the heaviest operation in the while loop is the insertion of the slope of a piece in the ordered vector of at most  $T - 1$  other slopes. This can be done in time  $O(\log T)$  when using an appropriate data structure. Since the while loop runs for at most  $M$  iterations and  $M \geq T$ , it follows that the total complexity is  $O(M \log T)$ .  $\square$

**Example 4.4.** We apply Algorithm 4.3 to an instance of *EVC* with piecewise linear objective. We take the instance of *EVC* described in Example 4.1 and replace the objective functions by their piecewise linear approximations using:

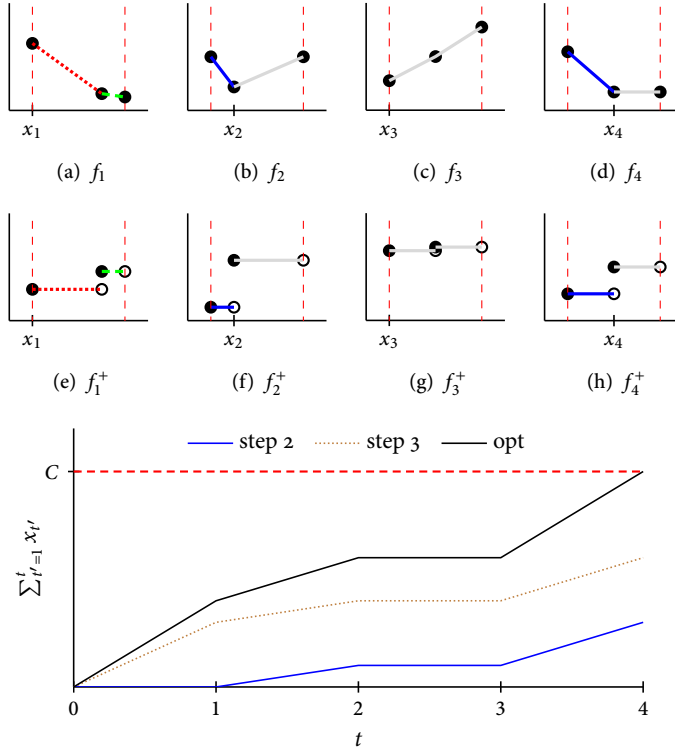
$$\gg Z_1 = \{0, 1\frac{1}{2}, 2\},$$

$$\gg Z_2 = \{0, \frac{1}{2}, 2\},$$

$$\gg Z_3 = \{0, 1, 2\},$$

$$\gg Z_4 = \{0, 1, 2\},$$

as sets of breakpoints. The example is depicted in Figure 4.6. In the figure the 3rd iteration of the while loop is described, where the currently completely used pieces (the first piece of both  $f_2$  and  $f_4$ ) are in blue. The next piece in the ordered set  $S$  is the first piece of  $f_1$ , which is in dashed red. After this piece is completely used, the second piece of  $f_1$  is inserted into the ordered set  $S$ . For all the iterations of the algorithm, the ordered set  $S$  is given in Figure 4.6(j). The red crossed out slope is the first and hence smallest slope in the set for the iteration, the corresponding piece is used in the iteration. The slope marked in green for each iteration belongs to the piece added to  $S$  after the piece with the red slope is completely used. This is the piece that is now active for its corresponding time interval.



(i) cumulative sum

Step	S	$(x_1, x_2, x_3, x_4)$	C
Initial	$\{s_2^1, s_4^1, s_1^1, s_3^1\}$	$(0, 0, 0, 0)$	5
1	$\{\overset{1}{s_2^1}, s_4^1, s_1^1, \overset{2}{s_2^1}, s_3^1\}$	$(0, \frac{1}{2}, 0, 0)$	$4\frac{1}{2}$
2	$\{\overset{1}{s_4^1}, s_1^1, \overset{2}{s_4^1}, \overset{2}{s_2^1}, s_3^1\}$	$(0, \frac{1}{2}, 0, 1)$	$3\frac{1}{2}$
3	$\{\overset{1}{s_1^1}, \overset{2}{s_1^1}, s_4^1, \overset{2}{s_2^1}, \overset{2}{s_2^1}, s_3^1\}$	$(1\frac{1}{2}, \frac{1}{2}, 0, 1)$	2
4	$\{\overset{2}{s_1^1}, \overset{2}{s_4^1}, \overset{2}{s_2^1}, \overset{2}{s_2^1}, s_3^1\}$	$(2, \frac{1}{2}, 0, 1)$	$1\frac{1}{2}$
5	$\{\overset{2}{s_4^1}, \overset{2}{s_2^1}, s_3^1\}$	$(2, \frac{1}{2}, 0, 2)$	$\frac{1}{2}$
6	$\{\overset{2}{s_2^1}, s_3^1\}$	$(2, 1, 0, 2)$	0

(j) Set S during the algorithm

Figure 4.6: Application of Algorithm 4.3 to the instance of *EVC* described in Example 4.4. Figures (a)-(d) give the objective functions and (e)-(h) the derivatives. Figure (i) gives the cumulative sum at the start of iterations two and three of the while loop. Here, the blue pieces were used at the start of iteration two and the red piece was used in this iteration. The table given in (j) gives the values of various variables throughout the iterations of the while loop in the algorithm.





In Section 4.4 we mentioned the linear time complexity algorithm of Hochbaum and Hong [64] for Problem *EVC* with quadratic objective. A similar approach can be used to solve *EVC* in time linear in  $M$ , the total number of breakpoints. Instead of sorting the set of slopes of the pieces between the breakpoints, we iteratively guess the highest slope among pieces used by an optimal solution. For this guess we use the median slope  $m$  in the set of all slopes, which we can find in linear time [18]. In case of an even number of slopes, we pick either of the two middle values. After this we divide the pieces over two sets:

$$S_1 := \{s_t^j \mid s_t^j \leq m\}, \quad (4.23a)$$

$$S_2 := \{s_t^j \mid s_t^j > m\}. \quad (4.23b)$$

We then construct a schedule  $\mathbf{x}$  using all the pieces in  $S_1$  and calculate  $\hat{C} := \sum_{t=1}^T x_t$ . Next we consider three cases.


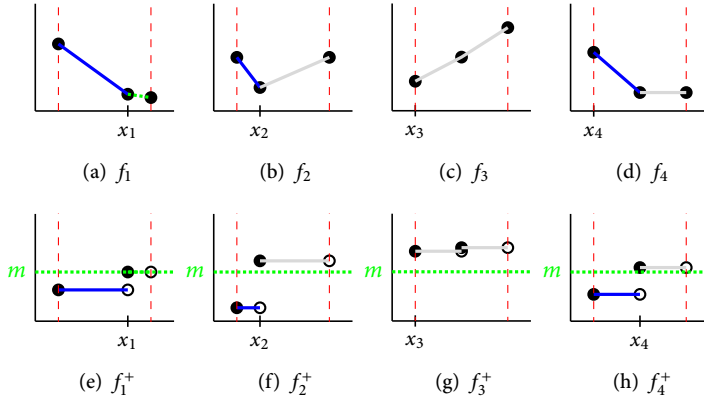
- » If  $\hat{C} = C$  we have found an optimal solution.
- » If  $\hat{C} > C$ , we check if we can obtain a feasible, and hence optimal, solution, by using the piece with slope  $m$  only partially. If this is not the case we know that an optimal solution cannot use pieces from  $S_2$  nor the piece with slope  $m$ . Thus, we recursively call the algorithm using only the pieces in  $S_1 \setminus \{m\}$
- » If  $\hat{C} < C$ , we know that an optimal solution must use all the pieces in  $S_1$ . Thus, we mark these pieces as used and recursively call the algorithm on the pieces in  $S_2$  with  $C := C - \hat{C}$ , since the remaining pieces must be used for this amount.

This approach is summarized in Algorithm 4.4, where we use similar notations as in Algorithm 4.3.

**Example 4.5.** *We apply Algorithm 4.4 to the instance described in Example 4.4. In the initial application of the algorithm (recursive depth of zero), the 8 pieces are split equally over two sets with the median piece being the second piece of  $f_1$ . The pieces in the first set  $S_1$ , have a total cumulative sum  $\hat{C} = 3\frac{1}{2}$ , hence all the pieces in  $S_1$  must be fully used in an optimal solution. Thus, the algorithm is recursively called on the pieces of in  $S_2$  with  $C = 1\frac{1}{2}$ . In this second application (recursive depth of one) the pieces in  $S_1$  have a total cumulative sum of  $\hat{C} = 2\frac{1}{2}$ . Furthermore, since the piece with median slope, the second piece of  $f_2$ , has a length of  $1\frac{1}{2}$ , the optimal solution is found by partially using this piece.*

**Corollary 4.1.** *Algorithm 4.4 solves Problem *EVC* with piecewise linear objective to optimality in time  $O(M)$ , with  $M = \sum_i m_i$ , i.e., the total number of pieces.*

*Proof.* The optimality of the algorithm follows immediately from the fact that pieces are used in order, hence the algorithm finds the same solution as Algorithm 4.3, which is optimal by Lemma 4.5. The complexity of  $O(M)$  follows from the fact that



(i) cumulative sum

depth	$S_1$	$S_2$	$m$	$C$	$\hat{C}$
0	$\{s_1^1, s_1^2, s_2^1, s_4^1\}$	$\{s_2^2, s_3^1, s_3^2, s_4^2\}$	$s_1^2$	5	$3\frac{1}{2}$
1	$\{s_2^2, s_4^1\}$	$\{s_3^1, s_3^2\}$	$s_4^1$	$1\frac{1}{2}$	$2\frac{1}{2}$

(j) Set S during the algorithm

Figure 4.7: Application of Algorithm 4.4 to the instance of EVC described in Example 4.4. Figures (a)-(d) give the objective functions and (e)-(h) the right derivatives. Figure (i) gives the cumulative sum for the pieces in  $S_1$  in the initial application of the algorithm. The cumulative sum of all the pieces minus the piece with median slope is also plotted together with that of the optimal solution. The table given in (j) gives the value of various variables in the different recursive calls of the algorithm.




---

**Algorithm 4.4** Median find approach *mpwEVC* for EVC with piecewise linear objective

---

```

1: x = Function mpwEVC(f, Z, C)
2: Take S as the set of slopes of all the pieces,  $S_1, S_2 := \emptyset$  and  $\hat{C} := 0$ 
3: Take  $l_t^j := z_t^j - z_t^{j-1}$  for each piece
4: Set  $x_t = 0$  for all t
5: Take m the median of S with associated  $t_m$  and  $j_m$ 
6: for t, j do
7:   if  $s_t^{j_m} > m$  then
8:      $S_2 = S_2 \cup \{s_t^{j_m}\}$ 
9:   else
10:     $S_1 = S_1 \cup \{s_t^{j_m}\}$ ,  $\hat{C} = \hat{C} + l_t^{j_m}$ , and  $x_t = x_t + l_t^{j_m}$ 
11:   end if
12:   if  $\hat{C} \geq C$  then
13:     if  $\hat{C} - l_{t_m}^{j_m} \leq C$  then
14:        $x_{t_m} = x_{t_m} - (\hat{C} - C)$ 
15:     else
16:       Take  $\hat{Z}$  the set of breakpoints for pieces in  $S_1 \setminus \{m\}$ 
17:        $\mathbf{x} = \text{mpwLSRA}(f, \hat{Z}, C)$ 
18:     end if
19:   else
20:     Take  $\hat{Z}$  the set of breakpoints for pieces in  $S_2$ 
21:      $\hat{\mathbf{x}} = \text{mpwLSRA}(f, \hat{Z}, C - \hat{C})$ 
22:      $\mathbf{x} = \mathbf{x} + \hat{\mathbf{x}}$ 
23:   end if
24: end for
25: Return x

```

---

all the steps before a potential recursive call in the algorithm can be solved in time linear in the number of pieces formed by the breakpoints in *Z*. For the first call this number of pieces is exactly *M* and for subsequent, recursive calls the number is halved each time. Hence the complexity is  $O(M + M/2 + M/4 + \dots) = O(M)$ .  $\square$

We note that the actual running time of the median find algorithm with linear asymptotic complexity is high for small to medium sized sets [18]. Thus, while Algorithm 4.4 has a lower asymptotic complexity than Algorithm 4.3, the latter might be more efficient in practice, particularly for decentralized energy management applications, where the number of intervals is typically low.

As mentioned before, for Problem *EVC* with piecewise linear objective, which can be solved efficiently using either Algorithm 4.3 or 4.4, we assumed that devices can switch between different operational levels somewhere during a time interval. However, excessive switching may cause undesirable wearing of the device,



e.g., the lifetime reduction that may occur for heat pumps and compressors in air-conditioning systems due to frequent on/off cycling or increased degradation of the internal battery within EVs from excessive switching between different charging and discharging currents [81]. Thus, solutions that avoid excessive switching are desirable. Furthermore, for a solution that switches between different operational levels it still needs to be determined when these different operational levels are used during the time interval. The design of an approach that considers this dynamic is outside the scope of this thesis. However, we expect that for solutions which switch little during time intervals, such an approach has lower computational complexity and requires less communication. If we now look at the way how solutions are produced by Algorithms 4.3 and 4.4, we see that they have  $x_t \notin Z_t$  for at most one index  $t$ . This leads to the following corollary.

**Corollary 4.2.** *For an instance of Problem EVC with piecewise linear objective, there exists an optimal solution  $\mathbf{x}$  such that  $x_t \notin Z_t$  for at most one  $t$ .*

As a consequence, the optimal solutions we obtained for Problem EVC with piecewise linear objective do not cause much extra wearing of the devices over any feasible solution to Problem dEVC.

## 4.6 SIMULATION STUDY

With the profile steering approach developed in Chapter 3 and the algorithms described in this chapter specifically for the EV, we can assess the potential benefits of our approach. We already gave an example application of the profile steering approach in Chapter 3, in which we demonstrated the effectiveness of the approach compared to the case where no control is applied to DERs. The simulation performed in this chapter is meant to compare our approach with other approaches. To this end we make two comparisons. First, we compare our approach with a time-of-use (ToU) pricing scheme. The prices we use are designed such that the highest peak in consumption of the flexible devices is shifted to the time intervals for which the consumption by the other devices is lowest. Second, in Section 4.6.2, we compare our approach with state-of-the-art research.

### 4.6.1 COMPARISON WITH PRICING

For the comparison with ToU pricing we consider a neighbourhood of 121 houses in the Dutch town of Lochem for which measurement data is available to us. To be able to determine the ToU prices efficiently we consider a case where only EVs are present. We aim to determine the optimal prices that can be sent to the consumers with respect to peak shaving, i.e., the prices that reduce the peak consumption the most. To compute these prices we consider the base load of the neighbourhood, i.e., the load without any EVs present. We only consider uniform prices among the customers, i.e., we assume every customer receives the same prices. Furthermore, we assume all EVs arrive at 18:00 and depart at 07:00 the next morning. Therefore,




Figure 4.8: Load on the transformer in a neighbourhood with many EVs for a single day when applying either no control, profile steering, or the best possible uniform prices.

intervals with a higher base load need to receive a higher price. To determine the prices we sort the intervals between 18:00 and 07:00 non-decreasingly and assign increasing prices with respect to this sorting. In this manner intervals with a high base load receive a high price.

To simulate a highly loaded scenario we assume every household has a steerable EV available. Furthermore, we assume that each EV has the same characteristics for the purpose of this simulation, namely a charge requirement of 12 kWh that needs to be satisfied with a charging power between 0 kW and 3.8 kW (i.e., we disregard V2G). In the prices case we assume that each EV is steered to minimize costs. We compare the ToU prices approach to our profile steering approach, which we set to steer towards a flat profile. Furthermore, to achieve better convergence and performance with respect to power quality we steer towards locally flat profiles in the initial phase of the approach, i.e., in the initial phase each EV is asked to flatten its load profile as much as possible. For completeness sake we also include the case when no control is applied to the EVs, i.e., when they charge as soon as possible, assuming that all vehicles arrive at 18:00.

The resulting load profiles on the transformer are given in Figure 4.8. A clear peak is visible when the EVs arrive at home in the no control scenario. While the pricing strategy shifts this peak away from the period with a high base load in the evening, the peak is only slightly reduced. This is because the combined load of the EVs charging at maximum power is far larger than the highest peak observed in the base load. The profile steering approach is capable of flattening the load profile.

Besides the power flow through the transformer, we also simulated the power flows through the low voltage (*LV*) network as we have details on the network structure in the considered area. We give the maximum and minimum observed voltages in the simulated network for every time interval in Figure 4.9. In both the no control and




Figure 4.9: Lowest and highest voltages observed in the (simulated) network corresponding to the load given in Figure 4.8 for the different approaches.

the pricing approach the voltages drop significantly during the charging of the EVs. They are barely within the limits specified by power quality requirements in the Netherlands (NEN-EN 50160:2010 [99], 207-253V). Our profile steering approach gives much better results with respect to voltages. This is partially because we steer towards locally flat profiles in the initial phase of our profile steering approach.

#### 4.6.2 COMPARISON WITH STATE-OF-THE-ART

Next we compare profile steering with the state-of-the-art Triana approach presented by Hoogsteen *et al.* in [66], which uses pricing signals as steering signals. We note that in this approach the pricing signals can differ between houses (and even between devices) and are not meant to reflect actual tariffs. We extend the scenario given above by adding a 3 kWh battery (max power 3.7 kW), 12 solar panels (1.65 m<sup>2</sup>, 18% efficiency, facing south), and a smart washing machine (WM) (to be scheduled between 08:00 and 17:00) to each house. In both approaches we assume that the start time of the smart WM is controllable as well as the charging and discharging of the battery. To be able to apply our profile steering approach to the WM and battery we use approaches described in Chapters 2 and 5 respectively. We also include the no control scenario, for completeness.

The results of the simulation are given in Figures 4.10 and 4.11. The no control case gives a small peak when all the WMs are turned on in the morning and a very large peak in the evening when all the EVs arrive at home. The state-of-the-art approach gives a significant improvement in the load profile by flattening it. The resulting voltage profile is also smoothed out significantly. The profile steering approach gives further improvements on top of the state-of-the-art approach in the form of a smoother load profile during the evening and night. The load of the EVs is spread more evenly than by the state-of-the-art approach. Furthermore, the voltage profile is also improved by our profile steering approach as the minimum voltage observed in the grid is higher and more stable than for the state-of-the-art approach.

To get some more insight, we record some statistics on the considered grid in Ta-




Figure 4.10: We compare the load on the transformer given by no control, state-of-the-art research and our profile steering approach.




Figure 4.11: Lowest and highest voltages observed in the (simulated) network corresponding to the load given in Figure 4.10 for the different approaches.

ble 4.3. The table shows several advantages of profile steering: the losses are reduced, the voltage is more stable, the highest peak on the transformer is reduced, and the maximum cable load is reduced. While such advantages are already achieved by the state-of-the-art approach over the no control case, further improvements are possible using our profile steering approach.

## 4.7 CONCLUSION

In this chapter we studied the problem of scheduling the charging of an EV, based on received steering signals. We made the assumption that this steering signal translates into a convex and separable objective function for the considered scheduling problem. With this assumption the problem transforms into a resource allocation problem, which we formulated as Problem *EVC*.

Resource allocation problems are well studied in the literature. We listed the two main approaches, Lagrangian multiplier methods and pegging approaches, for such



Table 4.3: Comparison between our profile steering approach, a state-of-the art approaches, and no control.

	No control	State-of-the-art	Profile steering
Total losses (kWh)	26.6	3.0	1.5
Lowest voltage (V)	201.8	219.6	223.2
Highest voltage (V)	236.9	241.5	235.7
Max. peak (kW)	563.0	174.8	131.2
Cable load (%)	143.2	53.0	32.1

problems and gave example instances of Problem *EVC* on which we applied these approaches with detailed algorithms. In particular, we studied the case where the objective function is convex quadratic, as this is the case for many practical scenarios. The presented Lagrangian multiplier method has a higher asymptotic complexity than the best known approach in the literature ( $O(T \log T)$  compared to  $O(T)$ ). However, we expect the former to outperform the latter in most DEM applications due to the fact that the number of time intervals, and hence variables, is often small.

As an extension to the classical continuous resource allocation problem, we considered the discrete EV charging problem, given as *dEVC*. This problem arises when the scheduled amount of energy  $x_t$  that can be charged into the battery of the vehicle in interval  $t$  is no longer assumed to be continuous, but must come from a discrete feasible set  $Z_t$ . This set  $Z_t$  gives the amount of energy charged into the vehicle when a constant power value is used for the entire interval. This discrete problem differs from the traditional discrete allocation problem, which is found often in literature, where the assumption is that  $x_t \in \mathbb{Z}$  (potentially with some bounds on  $x_t$  too). We showed that determining if Problem *dEVC* has a feasible solution is already an  $\mathcal{NP}$ -complete problem, even when all the feasible sets are the same, i.e.,  $Z_t = Z$  for all  $t$  and some  $Z$ .

To tackle the hard problem of discrete EV scheduling, we assumed that the vehicle is allowed to use two or more different possible power levels within an interval. This results in the possibility for the vehicle to charge any amount of energy between the lowest and highest value in  $Z_t$  in interval  $t$  through the use of several different levels throughout one interval. To adapt this change into our formulation we modified the objective value, resulting in a problem that is equivalent to *EVC* with piecewise linear objective. We showed that a greedy approach, preferring pieces with smaller slopes, is optimal for these instances of Problem *EVC* with piecewise linear objective. Furthermore, we formulated a linear-time approach for these types of problems. However, this approach might be outperformed in practice by our greedy approach for instances where the number of time intervals is small.

Finally, we compared the solutions found to the relaxed version of our discrete problem, where we allow modulation between different power levels. We showed that our algorithm obtains solutions to the relaxation of the discrete problem that is infeasible for at most one time interval for the original (discrete) problem. This





implies that potential device wearing due to frequent modulation between charging options is minimal for our given solutions.

The problems presented in this chapter do not take charging efficiency of the EV into account. If the losses are a (fixed) fraction of the amount charged, these losses can be accounted for by increasing the required charging by the appropriate amount. However, more complicated relations between the amount of charging and charging losses result in problems that fall outside the specific scope of this chapter. However, if the losses are a convex function of the amount charged in an interval, given by  $g_t(x_t)$ , the problem remains a convex resource allocation problem. In this case the resource constraint, given, e.g., by (4.1b), is replaced by:

$$\sum_t g_t(x_t) = C, \quad (4.24)$$

which is similar to (4.1b) and can be solved by the pegging approach described in Section 4.4.

Another practical limitation is that the maximum charging rate of the vehicle depends on the SoC. This behaviour, seen in vehicles currently available on the market, is typically modelled using the constant-current constant-voltage (*CC-CV*) charging model [97]. Such charging means that before a certain threshold value of the SoC is reached the current is kept constant and afterwards the voltage has to be kept constant. In particular, this charging method means that the charging rate when the battery is nearly full is both hard to control and low. However, the methods discussed in this chapter still apply to the first part of the charging, which contains the largest part of the energy to be charged. Note, that for instance, Mou *et al.* [97] list the threshold value to be around 85%.

Finally, we mentioned that the presented models only apply when the EV is assumed to do no discharging to the grid, i.e., we disregard V2G. V2G applications can have a positive impact on the grid, e.g., during a temporary high load on the grid. To unlock the positive impact of such applications, we need to adapt the model of the EV to ensure that the SoC is always feasible. We discuss this extension, together with several solution methods, in the next chapter.




# 5



## ENERGY STORAGE

**ABSTRACT** – *In this chapter we study the use of energy storage technologies in future smart grids, specifically focussing on electric energy storage. Such technologies will play an important role in the future smart grid, as they offer a lot of flexibility that can be used for various goals. Our starting point is the model of the previous chapter which we extend with the option to discharge energy. This extended model covers different types of storage. We show that the resulting device level problem can be solved efficiently using a divide and conquer strategy. Furthermore, when the allowed operational levels of the device are limited to a finite set, we obtain a reformulation for which we find an optimal greedy approach. We demonstrate the effectiveness of the derived approaches in a simulation study in which we study so-called soft-islanding of a group of sixteen houses within the Dutch grid.*



This Chapter is based on [TvdK:1] and [TvdK:9].




Figure 5.1: An example application of storage to shift overproduction from PV in the afternoon to the demand peak in the evening.




Figure 5.2: Price trend of Li-ion storage where we display the (expected) average costs of cells, automotive systems and stationary system. Data taken from [92].




Figure 5.3: Deployment of grid storage systems in the US. Data taken from [61].

## 5.1 INTRODUCTION

Due to the changes occurring in our electricity grid, as introduced in Chapter 2, energy storage can play an important role in the future grid. Batteries are a prime example of energy storage. Their potential applications in the (future) smart grid are numerous. For example, they can support the introduction of renewables by bridging the time gap between production and consumption or act on reserve markets to ensure frequency stability in the grid. An example of the first is given in Figure 5.1. As batteries are quite expensive, the economic viability of a stand-alone battery to support the grid currently only occurs in a few rare situations [106]. However, this is likely to change in the (near) future with a continued decrease of the prices of storage systems. As an example we illustrate the (expected) price trend of Li-ion systems in Figure 5.2. Bloomberg [17] reported a 65 percent decrease in battery pack costs with this trend expected to continue in the future. Furthermore, we depict the current and expected deployment of storage systems in the US in Figure 5.3.

Currently batteries are mainly used in appliances that require the storage of electrical energy for their operation (e.g., laptops and electric vehicles (EVs)). Several of these devices, in particular EVs, are expected to offer a large amount of flexibility in the future grid. Although their primary flexibility is in the time and amount of charging, their full potential is only unlocked when the internal battery can also be used to discharge energy when this is beneficial for the grid, the so-called vehicle-to-grid (V2G) applications in the case of EVs.

Besides the storage of electrical energy, other forms of energy storage are also envisioned to play an important role in the future smart grid. One example is the storage of heat in heat buffers, which is an interesting flexibility option for modern households. Because heating is increasingly done using electricity, which is one part of the increased electrification of our energy supply chain, heat storage facilities will influence the electricity grid. With electricity based heating systems the heat buffer allows the storage of the heat (hot water) produced, e.g., by a heat pump, until it is required. This enables the system to shift the use of electricity in time, without sacrificing user comfort. Such heat storage systems are typically less expensive and less influenced by ageing than batteries, causing such systems to be more widely applied. Furthermore, for several applications the house itself can also be considered as a heat storage, using the thermal inertia of the building to, e.g., delay the use of electricity for heating.

Energy storage can work on a large variety of time scales. Fast reacting storage devices, such as super capacitors and some types of batteries, can be used to ensure balance between supply and demand in smaller systems on a short time scale. Such systems can play an important role in future grids where this balance can no longer be ensured through traditional means. On the other hand, long-term storage, such as the conversion of electricity into hydrogen or synthetic gas, or the use of well-insulated underground thermal storage tanks allows for seasonal storage of energy





[38]. Long-term storage is particularly useful to combat the variance in seasonal availability of renewable sources such as sun and wind, and seasonal demand variations, observed in, e.g., heating. In this thesis we focus on storage that operates on the scale of minutes to hours, i.e., storage that can store energy within or between days but not on longer time scales of weeks or months. While the model we develop is generally applicable to any energy carrier, we focus on batteries, both stand-alone and inside devices such as EVs and heat buffers on a residential scale.

The remainder of this chapter is outlined as follows. In the next section we discuss related work, where we focus on the applications of battery systems. Then, in Section 5.3 we adapt the model introduced in the previous chapter to include state of charge (SoC) bounds, extending the model to cover discharging. We study the resulting optimization problem and give an efficient solution approach. Next, in Section 5.4 we consider an application of the developed models and algorithms. We finish with some conclusions in Section 5.5.

## 5.2 RELATED WORK

The list of technologies that (potentially) offer feasible energy storage is very long. However, not all technologies fit every application. For an overview of (most of) the available technologies we refer the reader to one of the surveys in the area of energy storage (e.g., [26, 41, 53, 54]). Typically, the various technologies are characterized using several indicators that play an important role in their applicability. A non-exhaustive list of these indicators is; maturity, power and energy rating, efficiency, operational lifetime, and costs. We note that in the area of batteries (i.e., chemical energy storage) no clear dominant technology is available. While lead-acid and lithium-ion are by far the most popular technologies, others have promising advantages over the most established ones and are being explored for various applications.

As mentioned before, the applications for storage in the (future) smart grid are near endless. The aforementioned indicators help to determine which technology is most suitable for a particular application or set of applications. We list several studies on the potential applications of storage systems below to give an overview.

One of the most studied applications of storage devices, specifically of batteries, in the electricity system is their use to reduce peaks (i.e., peak shaving) in consumption and/or production in the grid [84, 113, 116]. This is generally assumed to lead to deferrals in grid investments [105, 110]. Furthermore, such applications usually have the additional effect of increasing the level of self-consumption of locally produced energy.

Another application is energy arbitrage, i.e., acting with the storage asset on the energy market by exploiting differences in prices in time [59, 125]. Such applications can increase the overall system efficiency and reduce overall costs due to a reduction in the need for inefficient and costly peak power plants. However, national energy



prices might not reflect the situation in the local grid (e.g., congestion), causing potential conflicts of interest between the trader and the local grid operator [102].

Besides trading energy on the day-ahead or intra-day market, storage assets can also be used for frequency regulation, i.e., by charging or discharging more or less power based on the frequency changes in the system [98, 123]. Only those storage assets that can react very quickly are applicable for this case. Another application of storage assets is the smoothing of generation profiles of renewables [43, 47, 89]. The latter application benefits the owner of the generation unit, as they usually have to pay a penalty for deviations from their promised (i.e., predicted) profile. Both frequency regulations as well as smoothing of generation from renewables reduce the need for conventional spinning reserve. In turn, this potentially aids the transition towards a system based (purely) on energy from renewable sources.

For many applications of storage technologies, the asset only has to operate for a limited amount of time in a year. Due to this the system is free to be used for other applications for a large portion of time, as long as the goals do not conflict, potentially increasing the economic viability. Note that in several cases this combination is not (yet) allowed by law. For example, in Europe most distribution system operators (*DSOs*) are not allowed to trade on the energy market, preventing an increase in the number of viable business cases for storage devices [106].

### 5.3 PROBLEM FORMULATION AND SOLUTION

In the previous chapter we considered the problem of scheduling an EV under steering signals received from the profile steering decentralized energy management (*DEM*) approach. We noted that we only required a single constraint on the SoC of the internal battery because we assumed the vehicle could not discharge to the grid. In this section we extend the problem formulation given for the EV (Problem 4.1) to include the case when the car is allowed to discharge energy while connected, i.e., we allow V2G applications. We first introduce the resulting modifications of the model. Next we discuss how this model can be applied to other storage assets: a stand-alone battery and a heat producer combined with a heat buffer. After the model introduction, we study the resulting optimization problem and give an efficient solution approach.

#### 5.3.1 A MODEL FOR DISCHARGING

The model for buffering devices we introduced in Chapter 2 has to be adapted in the following manner if we allow the discharging of energy from the vehicle's battery. Allowing discharging implies that now the outflow is also partially controllable, namely for the time intervals that the vehicle is plugged in, whereas before the outflow process was uncontrollable and only occurred during driving. Note that in the case where we only consider charging, the SoC of the vehicle only had to be considered for the last time interval. This changes when discharging is allowed, because we can no longer assume the SoC remains feasible for all time intervals if we



just check if enough energy is charged into the battery at the end of the scheduling horizon.

In the following we describe the scheduling problem that arises in this situation in more detail. As before we assume the vehicle receives a steering signal that consists of an objective function  $f(\mathbf{x})$  potentially together with an lower bound  $\mathbf{x}^{min}$  and an upper bound  $\mathbf{x}^{max}$  on the energy use. Recall that we assumed the objective function to be convex and separable. Within the planning we have to ensure that the SoC of the battery inside the vehicle remains feasible. For this, we have to related the SoC for time interval  $t$  to the decisions made in our scheduling problem. These decisions are given by the amounts  $x_{t'}$  of energy which are charged to or discharged from the battery in interval  $t' \leq t$ . More precisely, to calculate the SoC for time interval  $t$  we have to take the initial content of the storage devices together with the cumulative sum of the values of  $x_{t'}$  with  $t' \leq t$ . In other words, the SoC for interval  $t$  is given by  $\sum_{t'=1}^t x_{t'}$  added to the initial SoC of the battery. To ensure that the SoC of the battery is feasible for time interval  $t$ , we introduce two bounds  $B_t$  and  $C_t$ , which we call the cumulative lower and upper bounds respectively. These bounds restrict how much total (dis)charging can be done up to and including time interval  $t$ , i.e.,  $B_t$  is a lower bound and  $C_t$  is an upper bound on  $\sum_{t'=1}^t x_{t'}$ . For an EV,  $B_t$  is the initial SoC of the battery and  $C_t$  is the difference between the maximum capacity of the battery and the initial SoC. Furthermore, the lower bound  $B_t$  can be used to, e.g., ensure that enough energy is present in case an unplanned trip to, e.g., the hospital has to be made, next to ensuring that the SoC stays within the limits of the buffer. The bounds  $B_t$  and  $C_t$  on the cumulative sum replace the single SoC Constraint (4.1b) in Problem *EVC*. Again we assume that the time horizon coincides with the time window where the car is plugged in. This leads us to the following optimization problem.

**Problem 5.1 (BC).**

$$\min_{\mathbf{x}} f(\mathbf{x}), \quad (5.1a)$$

$$\text{s.t. } B_t \leq \sum_{t'=1}^t x_{t'} \leq C_t \quad t = 1, 2, \dots, T, \quad (5.1b)$$

$$\mathbf{x}^{min} \leq \mathbf{x} \leq \mathbf{x}^{max}. \quad (5.1c)$$

We call this problem the battery charging (*BC*) problem. As before, we can assume w.l.o.g. that  $\mathbf{x}^{min} = 0$ , as we can apply a translation to obtain this case. While this might seem to imply that we assume that the battery can only charge, discharging of energy (given by a negative value of  $x_t$ ) is still possible in the solution to the original problem (after we translate the obtained solution back to original problem). This is because we ‘push up’ the values of the solution after the translation. To illustrate this we consider the following example instance of Problem *BC*.





**Example 5.1.** We consider a battery that can charge or discharge at most one kWh of energy per time interval, starts with an initial charge of one kWh, has a capacity of two kWh, and has to be scheduled for three time intervals (we disregard the objective function in this example and focus on the constraints). This results in an instance of Problem BC with bounds on the charging in the interval given by  $\mathbf{x}^{min} = -\mathbf{1}$  and  $\mathbf{x}^{max} = \mathbf{1}$ . Furthermore  $B_t = -1$  and  $C_t = 1$  for all  $t$  (at most one unit can be charged or discharged before the battery is full/empty). To obtain an instance where  $\mathbf{x}^{min} = 0$ , we use the translation  $\mathbf{x}' = \mathbf{x} + 1$  (in this new instance we use an apostrophe to denote the parameters). In this case the lower bound on the energy use, given by  $\mathbf{x}'^{min}$ , indeed becomes zero and the upper bound  $\mathbf{x}'^{max}$  becomes two. The values of  $B_t$  now differ per interval. For interval  $t = 1$  Constraint (5.1b) now reads  $B_1 \leq x'_1 - 1 \leq C_1$ . From this it follows that  $B'_1 = B_1 + 1 = 0$  and  $C'_1 = C_1 + 1 = 2$ . For interval  $t = 2$  we obtain  $B_2 \leq x'_1 - 1 + x'_2 - 1 \leq C_2$  from which we obtain  $B'_2 = 1$  and  $C'_2 = 3$ . Finally, for interval  $t = 3$  we obtain that  $B'_3 = 2$  and  $C'_3 = 4$ . An example of a feasible solution  $\mathbf{x}'$  is now given by  $(0, 2, 1)$ . This solution translates back into a schedule for the battery given by  $\mathbf{x} = (-1, 1, 0)$ .

Note that in the above example the bounds  $B_t$  and  $C_t$  both form an increasing sequence in  $t$ . We can assume, w.l.o.g., that this is always the case, after applying the translation that ensures that  $\mathbf{x}^{min} = 0$ . In summary, the bound  $B_t$  gives a lower bound on the total charging and discharging of the storage device, to ensure no more energy than is available is discharged and energy requirements for the (uncontrollable) process is available. This implies that the difference between the cumulative sum  $\sum_{t'=1}^t x_{t'}$  and  $B_t$  gives the SoC of the storage device for time interval  $t$ . On the other hand, the bound  $C_t$  gives an upper bound on the total charging and discharging of the storage device to ensure the device is never charged with more energy than it can hold.

To ease the discussion, we can assume that the SoC of the battery at the end of the scheduling horizon has to be equal to a fixed level, given by  $C_T$ , i.e., we assume  $B_T = C_T$ . This is because of the following transformation. If  $B_T < C_T$  we add an additional time interval  $T+1$  to the problem with  $f_{T+1}(x_{T+1}) = 0$ ,  $x_{T+1}^{min} = 0$ ,  $x_{T+1}^{max} = C_T - B_T$  and  $B_{T+1} = C_{T+1} = C_T$ . A solution to the transformed instance can be translated to a solution of the original problem with the same objective by discarding the last time interval and vice versa. In essence this transformation adds an additional artificial time interval  $T+1$  at the end of the scheduling horizon. For  $T+1$  we add an additional demand on the amount of energy that needs to be charged, given by  $B_{T+1}$  (which we take equal to  $C_T$ ). By doing this we ensure that all feasible solutions to the (transformed) instance do the same total amount of charging and discharging. This makes it easier to find an optimal solution later on. For clarity we depict the cumulative bounds  $B_t$  and  $C_t$  together with a feasible solution for an instance which we transformed twice in Figure 5.4 and for which we give an interpretation.

We note that Problem BC can also be used to schedule a stand-alone battery, if we neglect losses. For such a battery both the inflow and outflow process are completely




Figure 5.4: Example of how the solution to an instance of Problem  $BC$  that is transformed twice (through a translation and the addition of an extra interval) can be interpreted.

controllable. Constraints (5.1b) ensure the battery operates within feasible SoC regions. In other words, the lower bound  $B_t$  ensures that the total amount of charging and discharging done by the battery does not exceed the energy initially available in the battery (the initial SoC). The upper bound  $C_t$  ensures the battery is never charged with more energy than it can stored (i.e., with more energy than the difference between the initial SoC and the capacity of the battery). Furthermore, these bounds can be used to ensure that certain SoC targets are met throughout the scheduling horizon. Constraints (5.1c) ensure the (dis)charging done does not exceed the rated power of the device.

Furthermore, Problem  $BC$  can also be used to model a heat producer, such as a combined heat and power ( $CHP$ ) unit or a heat pump ( $HP$ ), together with a heat buffer. For such a combination of devices the inflow process is controllable through the (controllable) operation of the heater, while the outflow is uncontrollable (heat demand must be satisfied whenever present). This heat demand is often not given as a single value at the end of the scheduling horizon, but by cumulative values over time. In other words, the heating device needs to ensure enough energy is present in the buffer to ensure the heat demand over time can be satisfied. This can also be modelled using the cumulative bounds of Problem  $BC$  as given in Constraints (5.1b). For this case, the lower bound  $B_t$  gives the total heat demand up to time interval  $t$ , which has to be produced by the schedule. On the other hand, the upper bound  $C_t$  gives the total heat demand up to time interval  $t$  plus the buffer capacity, to ensure that the buffer is never charged with more energy than it can hold.

### 5.3.2 EFFICIENTLY SOLVING THE BATTERY PROBLEM

In this section we discuss an efficient solution approach for Problem  $BC$ . The base of our solution approach is that we first generate a candidate solution  $\mathbf{y}$  by solving Problem  $EVC$  (from Chapter 4) with  $C = C_T$  and the bounds on the resource usage given by Constraints (5.1c). This solution will generally be infeasible for




Figure 5.5: The cumulative sum for the candidate solution  $\mathbf{y}$ , obtained by ignoring the cumulative bounds, and the optimal solution  $\mathbf{x}$ , for an instance of Problem BC.

several of the cumulative resource constraints, i.e.  $\sum_{t'=1}^t y_{t'}$  might be smaller than  $B_t$  or larger than  $C_t$  for several  $t$ . However, this solution contains some valuable information. Since the objective functions are convex, intuitively it seems to be true that the constraint that is violated most in the (infeasible) candidate solution should be tight (i.e., met with equality) in an optimal solution to Problem BC. The reason that this statement is only made for the maximal violation is that making the solution feasible for the index with maximal violation may already ensure that the bounds for other, previously infeasible indices, are met. As an example, consider the cumulative sum given in Figure 5.5. Because the maximum violation of the cumulative bound occurs at time interval two for  $\mathbf{x}$ , the charging on time interval one is already reduced enough to ensure feasibility of the upper cumulative bound on time interval one

A similar property is in fact proven and used in both the applications of DVFS [69] and vessel speed optimization [71] to construct efficient algorithms. However, in those settings the objective function is the same for every time interval, whereas this is not always the case in our profile steering DEM approach. In the following, we prove that this property still holds when we assume each  $f_t$  to be an arbitrary continuous and convex function.

**Lemma 5.1.** *Consider an instance of Problem BC with  $C_T = B_T$  and let  $\mathbf{y}$  be an optimal solution to the corresponding instance of Problem EVC obtained by ignoring the cumulative bounds for all indices except the last. Assume that  $\mathbf{y}$  is not feasible for the considered instance of Problem BC and let  $k$  be the index at which the cumulative bound is maximally violated, i.e.  $k = \arg \max_t \{ \sum_{t'=1}^t y_{t'} - C_t, B_t - \sum_{t'=1}^t y_{t'} \}$ . Then, there exists an optimal solution  $\mathbf{x}$  to the considered instance of Problem BC such that, if  $\sum_{t=1}^k y_t > C_k$  then  $\sum_{t=1}^k x_t = C_k$  and, on the other hand, if  $\sum_{t=1}^k y_t < B_k$  then  $\sum_{t=1}^k x_t = B_k$ .*

*Proof.* Let  $\mathbf{x}$  be an optimal solution to the considered instance of Problem BC and assume that  $\sum_{t=1}^k y_t > C_k$  and  $\sum_{t=1}^k x_t \neq C_k$ . Since  $\mathbf{x}$  is feasible, it follows that




Figure 5.6: An example of the cumulative sum for candidate solution  $y$  and optimal solution  $x$  with the indices  $l < n \leq k < n^* \leq m$ .

$\sum_{t=1}^k x_t < C_k$ . Let  $l$  be the last index before  $k$  for which the upper cumulative bound is met by  $x$  with equality, i.e.  $l := \max\{t < k \mid \sum_{t'=1}^t x_{t'} = C_t\}$ , and define  $l := 0$  if this set is empty. Furthermore, let  $m$  be the first index after  $k$  for which the upper cumulative bound is met by  $x$  with equality, i.e.  $m := \min\{t > k \mid \sum_{t'=1}^t x_{t'} = C_t\}$ . Note that since  $\sum_{t=1}^T x_t = C_T$  by assumption,  $m$  is well defined and  $m \leq T$ .

In the following we first show that we can find indices  $n$  and  $n^*$  with  $l < n \leq k < n^* \leq m$  for which  $y_n > x_n$  and  $y_{n^*} < x_{n^*}$ . Figure 5.6 gives such indices for an example. In this example the maximum violation occurs for  $k = 7$ , while the optimal solution has tight upper bounds for  $l = 5$  and  $m = 9$ . Furthermore,  $y_7 > x_7$  and  $y_9 < x_9$ , so  $n = 7$  and  $n^* = 9$ . Note that, for the cumulative sum, as plotted in Figure 5.6, the value of  $x_t$  is given by the increase of the cumulative sum between intervals  $t-1$  and  $t$ , i.e., by the *difference* between the plotted value for time intervals  $t$  and  $t-1$ .

Since  $\sum_{t=1}^l x_t = C_l$  and  $\sum_{t=1}^k x_t < C_k$  it follows that  $\sum_{t=l+1}^k x_t < C_k - C_l$ . Furthermore, by definition of  $k$  we have  $\sum_{t=1}^l y_t - C_l \leq \sum_{t=1}^k y_t - C_k$ , which implies that  $\sum_{t=l+1}^k y_t \geq C_k - C_l$ . Combining both inequalities we obtain that  $\sum_{t=l+1}^k y_t > \sum_{t=l+1}^k x_t$  and hence there exists an index  $n$  with  $l < n \leq k$  such that  $y_n > x_n$ . Similarly, since  $\sum_{t=1}^m x_t = C_m$  and  $\sum_{t=1}^k x_t < C_k$  it follows that  $\sum_{t=k+1}^m x_t > C_m - C_k$ . Also, since  $\sum_{t=1}^n y_t - C_m \leq \sum_{t=1}^k y_t - C_k$  it follows that  $\sum_{t=k+1}^m y_t \leq C_m - C_k$ . Combining this we obtain that  $\sum_{t=k+1}^m y_t < \sum_{t=k+1}^m x_t$  and hence there exists an index  $n^*$  with  $k < n^* \leq m$  such that  $y_{n^*} < x_{n^*}$ .

From the optimality conditions stated in the previous chapter (Lemma 4.1) we obtain that  $f_n^-(y_n) \leq f_{n^*}^+(y_{n^*})$ . Furthermore, by the convexity of  $f_n$  and  $f_{n^*}$  respectively it follows that  $f_n^+(x_n) \leq f_n^-(y_n)$  and  $f_{n^*}^+(y_{n^*}) \leq f_{n^*}^-(x_{n^*})$ . Thus we obtain:



$$f_n^+(x_n) \leq f_n^-(y_n) \leq f_{n^*}^+(y_{n^*}) \leq f_{n^*}^-(x_{n^*}). \quad (5.2)$$

Note that, for each  $t$  with  $l < t < m$ , we have  $\sum_{t'=1}^t x_{t'} < C_t$ . This implies that taking  $x_n = x_n + \epsilon$  and  $x_{n^*} = x_{n^*} - \epsilon$  does not violate feasibility, since  $l < n \leq k < n^* \leq m$ . Furthermore, (5.2) implies that this change does not increase the objective value for sufficiently small  $\epsilon$ . Summarizing, this change of the solution  $\mathbf{x}$  increases  $\sum_{t=1}^k x_t$  but does not increase the objective value. As we can repeat this process with the new solution, we finally get an optimal solution with  $\sum_{t=1}^k x_t = C_k$ , proving the lemma.

The proof for the case that  $\sum_{t=1}^k y_t < B_k$  and  $\sum_{t=1}^k x_t > B_k$  is symmetric.  $\square$

Lemma 5.1 forms the base of our solution approach for Problem *BC*. In this approach, we first ignore the intermediate cumulative bounds of *BC* and afterwards we iteratively satisfy the ignored constraints using a divide and conquer approach. We start by calculating an optimal solution for the instance of *EVC* in which we set  $C = C_T$ . For this optimal solution, we determine the index  $k$  where this solution maximally violates the cumulative bounds. By Lemma 5.1 we know that there exists an optimal solution for which the corresponding bound is tight at index  $k$ , meaning that we can set both  $B_k$  and  $C_k$  to the value of the violated bound (i.e., either to  $B_k$  or to  $C_k$ ). Note that this splits the original instance of *BC* into two independent instances of *BC*; one for the indices up to and including  $k$  and one for the indices after  $k$ . These problems can be solved separately following the same procedure. Hence we can recursively solve problems of the form *EVC*, until we no longer have violations of the intermediate cumulative bounds. Combining the individual solutions of these instances then gives a solution to the original instance of Problem *BC*, which is optimal by Lemma 5.1 and the fact that the individual solutions are optimal for their respective time intervals.

This sketched procedure is summarised in Algorithm 5.1. In this algorithm we use  $\mathbf{f}_{t \rightarrow t'}$  to denote the vector  $(f_t, f_{t+1}, \dots, f_{t'})$  of objective functions and similar notations for the parameters  $\mathbf{x}^{max}$ ,  $B_t$ , and  $C_t$  and decision variables  $x_t$ . Furthermore,  $optEVC(\mathbf{f}_{1 \rightarrow T}, \mathbf{x}_{1 \rightarrow T}^{max}, C)$  denotes a call to an algorithm that solves an instance of *EVC* with objective functions  $\mathbf{f}_{1 \rightarrow T}$ , and parameters  $\mathbf{x}_{1 \rightarrow T}^{max}$  and  $C$ . To solve such problem instances one can use, for example, the Lagrangian approach or the pegging approach (see Section 4.4). This algorithm outputs a solution vector  $\mathbf{x}_{1 \rightarrow T}$  that is optimal for this instance. Example 5.2 below gives some insight in the structure of the algorithm.

**Example 5.2.** *Fig. 5.7 depicts an application of Algorithm 5.1. The problem instance we use is the one from the previous chapter given in Example 4.1 with added lower and upper cumulative bounds. These bounds are depicted in the middle plot showing the cumulative sums. Above the cumulative sum the objective functions and derivatives are plotted together with the original solution  $\mathbf{x}$  to Problem *EVC*. This solution serves as a candidate solution to Problem *BC*, however, it does not consider the cumulative*

---

**Algorithm 5.1** Recursive algorithm *optBC* for problem *BC*


---

```

1:  $\mathbf{x}_{1 \rightarrow T} = \text{Function } \text{optBC}(\mathbf{f}_{1 \rightarrow T}, \mathbf{x}_{1 \rightarrow T}^{max}, \mathbf{B}_{1 \rightarrow T}, \mathbf{C}_{1 \rightarrow T})$ 
2:  $\mathbf{y}_{1 \rightarrow T} = \text{optEVC}(\mathbf{f}_{1 \rightarrow T}, \mathbf{u}_{1 \rightarrow T}, C_T)$ 
3: if  $\mathbf{y}_{1 \rightarrow T}$  is feasible then
4:    $\mathbf{x}_{1 \rightarrow T} = \mathbf{y}_{1 \rightarrow T}$ 
5: else
6:    $k = \arg \max \{ \sum_{t=1}^k y_t - C_k, B_k - \sum_{t=1}^k y_t \}$ 
7:   if  $\sum_{t=1}^k y_t > C_k$  then
8:      $B_k = C_k$ 
9:      $B_t = B_t - C_k$  and  $C_t = C_t - C_k$  for  $t = k + 1, k + 2, \dots, T$ 
10:  else
11:     $C_k = B_k$ 
12:     $B_t = B_t - B_k$  and  $C_t = C_t - B_k$  for  $t = k + 1, k + 2, \dots, T$ 
13:  end if
14:   $\mathbf{x}_{1 \rightarrow k} = \text{optBC}(\mathbf{f}_{1 \rightarrow k}, \mathbf{x}_{1 \rightarrow k}^{max}, \mathbf{B}_{1 \rightarrow k}, \mathbf{C}_{1 \rightarrow k})$ 
15:   $\mathbf{x}_{k+1 \rightarrow T} = \text{optBC}(\mathbf{f}_{k+1 \rightarrow T}, \mathbf{x}_{k+1 \rightarrow T}^{max}, \mathbf{B}_{k+1 \rightarrow T}, \mathbf{C}_{k+1 \rightarrow T})$ 
16: end if
17: Return  $\mathbf{x}_{1 \rightarrow T}$ 

```

---



bounds. This candidate solution is not feasible, since  $\sum_{t=1}^2 x_t > C_2$ . Based on this, we split the problem in two subproblems. In the first subproblem, we have to decrease  $x_1$  and/or  $x_2$ . Note that by doing this we obtain  $x'_1 < u_1$  and thus we find  $\lambda_1$  with  $f'_1(x'_1) = f'_2(x'_2) = \lambda_1$ . In the second subproblem we increase  $x_3$  and  $x_4$ . Doing this gives us  $x'_4 = u_4$ , thus we find  $\lambda_2$  such that  $f'_3(x'_3) = \lambda_2 > f'_4(x'_4)$ . Combining the solutions to the subproblems, we obtain the optimal solution  $\mathbf{x}'$  to Problem *BC*. This solution is given in the bottom eight plots.  $\square$

It remains to determine the complexity of Algorithm 5.1. Let  $F_{EVC}(T)$  denote the complexity of the algorithm *optEVC*( $f_{1 \rightarrow T}, u_{1 \rightarrow T}, C_T$ ), called by *optBC* to solve an instance of *EVC* with  $T$  intervals. Furthermore, let  $F_{BC}(T)$  be the complexity of Algorithm 5.1 for instances with  $T$  intervals. We then obtain the following recursive relation for  $F_{BC}(T)$ :

$$F_{BC}(T) = O(T) + F_{EVC}(T) + F_{BC}(k) + F_{BC}(T - k) \quad (5.3)$$

Assuming that the asymptotic complexity of  $F_{EVC}(T)$  is at least linear in  $T$  (i.e.,  $F_{EVC}(T)$  has a complexity of  $\Omega(T)$ , see the Appendix for a background), it follows that  $F_{EVC}(k) + F_{EVC}(T - k) \leq F_{EVC}(T)$ . This implies that  $F_{BC}(T) = O(T^2 + TF_{EVC}(T))$ . We note that, for the case that the objective functions are quadratic, Hochbaum and Hong [64] provide an  $O(T)$  algorithm to solve Problem *EVC*. Combining this with our approach yields a complexity of  $O(T^2)$  for problem *BC* with quadratic objective.

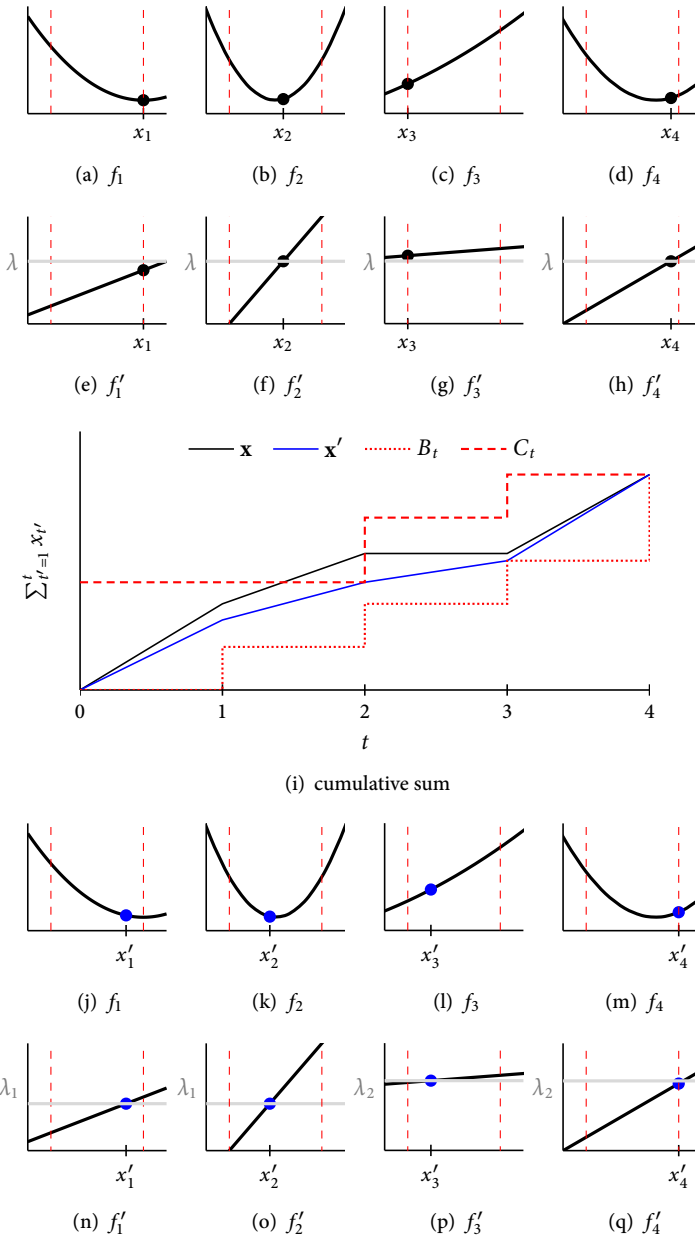


Figure 5.7: Example of the application of Algorithm 5.1. Figures (a)-(d) depict the objective functions with their respective derivatives given in (e)-(h) for the original solution to the obtained instance of *EVC*, with the cumulative sum given in (i). After splitting the instance at  $t = 2$  (the obtained solution is infeasible for interval  $t = 2$ ) we obtain the solution depicted in (j)-(q).

### 5.3.3 DISCRETE BUFFER

In the previous chapter, we considered a variant of problem *EVC*, where we restricted the feasible set of the decision variable  $x_t$  to a discrete set  $Z_t$  for every  $t$ . The same restriction can be applied to problem *BC*. As Problem *BC* is more general than *EVC*, it readily follows from Lemma 4.3 that this problem is also  $\mathcal{NP}$ -hard. As for Problem *EVC*, we again consider the case where we allow convex combinations of the points in  $Z_t$ . This leads to the following problem:

**Problem 5.2** (*dBC*).

$$\min_y \sum_{t=1}^T F_t(x_t) = \sum_{t=1}^T \sum_{j=0}^{m_t} y_t^j f_t(z_t^j), \quad (5.4a)$$

$$\text{s.t. } B_t \leq \sum_{i'=1}^t x_{i'} \leq C_t \quad t = 1, 2, \dots, T, \quad (5.4b)$$

$$x_t = \sum_{j=0}^{m_t} y_t^j z_t^j \quad t = 1, 2, \dots, T, \quad (5.4c)$$

$$\sum_{j=0}^{m_t} y_t^j = 1 \quad t = 1, 2, \dots, T, \quad (5.4d)$$

$$y_t^j \geq 0 \quad t = 1, 2, \dots, T; j = 0, 1, \dots, m_t, \quad (5.4e)$$

We call this problem the discrete battery charging (*dBC*) problem. By similar reasoning to that in the previous section we may again assume that  $B_T = C_T$ . Furthermore, similarly to the equivalence between *dEVC* and *EVC* in the previous chapter, this problem is equivalent to an instance of *BC* with piecewise linear objectives. Thus, combining Algorithm 5.1 from above and Algorithm 4.3 from the previous chapter gives us an algorithm that runs in time  $O(T^2 + TM \log T)$ , which reduces to  $O(TM \log T)$  since  $T \leq M$  (recall that  $M$  denotes the total number of pieces of the objective functions, i.e.,  $M = \sum_{t=1}^T m_t$ ). Furthermore, combining Algorithms 5.1 and 4.4 gives an algorithm that runs in time  $O(T^2 + TM) = O(TM)$ .

Similarly to the greedy approach for Problem *EVC* with piecewise linear objective, presented in the previous chapter, we now construct a greedy approach to solve *BC* with piecewise linear objective that prefers pieces with a smaller slope. Again we use an ordered set  $S$  for the slopes of the next pieces to be used for every time interval, from which we greedily pick the first (i.e., the one with smallest slope). However, note that we also need to satisfy the cumulative bounds  $B_t$  and  $C_t$  for every  $t$ . Furthermore, note that to satisfy the lower bound  $B_t$ , only pieces of the functions  $f_{t'}$  with  $t' \leq t$  are of relevance. Therefore, we iteratively build up a solution that satisfies the lower bounds  $B_1, B_2, \dots, B_k$  in iteration  $k$ . To do this, during iteration  $k$ , we only consider pieces of the functions  $f_1, f_2, \dots, f_k$  and increase the use of







these pieces until  $\sum_{t=1}^k x_t = B_k$ . This means that, in iteration  $k$ , we only add pieces of  $f_t$  with  $t \leq k$  to the set  $S$ .

It remains to incorporate the upper cumulative bound  $C_t$ . We note that, for a given solution  $\mathbf{x}$ , we cannot increase the use of the pieces of function  $f_t$  by more than  $C_{t^*} - \sum_{t'=1}^{t^*} x_{t'}$  for every  $t^* \geq t$ , since otherwise we violate the upper bound  $C_{t^*}$ . Thus, to ensure feasibility, we introduce a variable  $V_t := \min_{t^* \geq t} \{C_{t^*} - \sum_{t'=1}^{t^*} x_{t'}\}$  to track how much we can increase  $x_t$ , using the pieces of  $f_t$ , without violating the upper bounds  $C_t, C_{t+1}, \dots, C_T$ . Note that  $V_t$  depends on  $\sum_{t'=1}^{t^*} x_{t'}$  for all  $t^* \geq t$ , which in turn depends on  $x_{t'}$  for all  $t' \leq t^*$ . Thus, after we increased the use of the pieces of some  $f_{t^*}$  by  $\delta$ , we need to update  $V_t$  for every  $t$ . Furthermore, note that whenever we increase  $x_{t^*}$  by  $\delta$ ,  $\sum_{t'=1}^t x_{t'}$  is also increased by  $\delta$  for every  $t \geq t^*$ . This implies that  $V_t$  decreases by exactly  $\delta$  for every  $t \geq t^*$ . On the other hand, for  $t < t^*$ , we note that  $V_{t-1} = \min\{V_t, C_t - \sum_{t'=1}^{t-1} x_{t'}\}$ . This can be used to iteratively update  $V_{t^*-1}, V_{t^*-2}, \dots, V_1$ .

Finally, to obtain a better complexity, similar to the approach we took in Algorithm 4.3, we only consider at each moment in the algorithm at most one piece per function  $f_t$ . The resulting procedure is summarized in Algorithm 5.2, where we use the same notation as in the algorithms previously presented in this thesis. Example 5.3 demonstrates an application of the algorithm.

**Example 5.3.** *Figure 5.8 depicts an application of Algorithm 5.2. For this instance we added the cumulative bounds used in Example 5.2 to the instance used in Examples 4.4 and 4.5 in the previous chapter. A colour coding is used to denote (parts of) the pieces that are used in each iteration of the main for loop of the algorithm. The same colour coding is used in the plot that depicts the cumulative sum up to interval  $k$  for each of the partial solutions constructed. Furthermore, the solution marked in the top plots is the optimal solution the algorithm produces after completion. In the table the values of various variables used inside the algorithm are given for each iteration of the while loop in the algorithm. Furthermore,  $S_{start}$  denotes the set  $S$  at the start of the considered iteration of the while loop, and  $S_{end}$  the set  $S$  at the end of the iteration. New additions to  $S$  are depicted in green. In particular we note that  $s_2^2$  is deleted from  $S$  at the end of iteration 4, since  $V_2 = 0$ . Thus the algorithm must use  $s_3^1$  instead to ensure that the bound  $B_3$  is met.*

**Lemma 5.2.** *The greedy approach for BC with piecewise linear convex objective, given in Algorithm 5.2, gives an optimal solution to BC with piecewise linear convex objective functions. The algorithm runs in time  $O(TM)$ .*

*Proof.* The feasibility of the algorithm follows from the fact that the pieces are added in increasing order and that any new piece is either the first piece of an objective function (Line 5) or a piece for which the previous piece is already fully used (Line 26). Also, note that the lower cumulative bounds are satisfied since the constructed solution enforces this for every index in the main for loop of the algorithm (Lines 4 - 32). Finally, note that in every step of the main for loop,  $V_t \leq$




---

**Algorithm 5.2** Greedy approach *pwlBC* for *BC* with piecewise linear convex objective

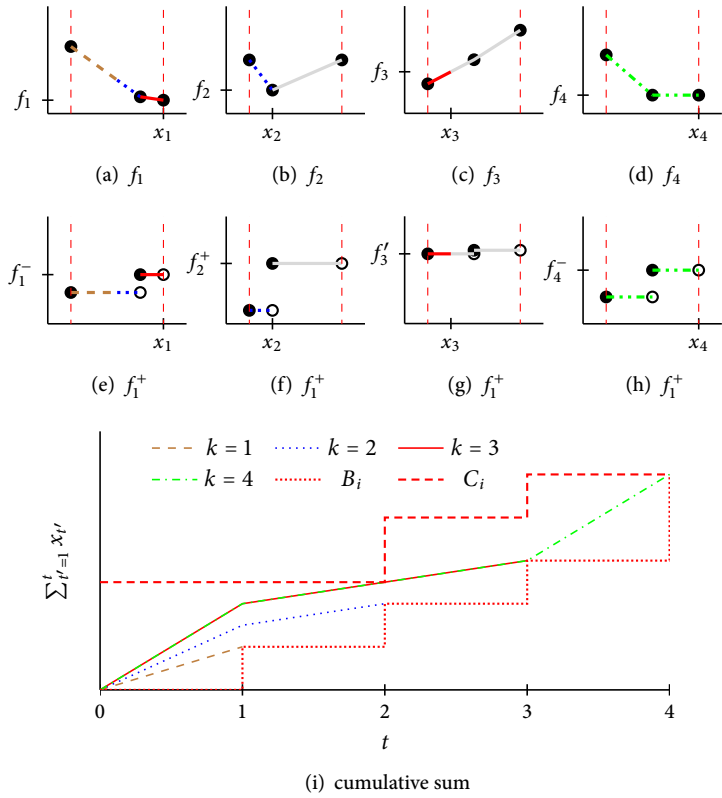
---

```

1:  $\mathbf{x}_{1 \rightarrow T} = \text{Function } pwlBC(f_{1 \rightarrow T}, Z, \mathbf{x}_{1 \rightarrow T}^{max}, B_{1 \rightarrow T}, C_{1 \rightarrow T})$ 
2:  $S = \emptyset$ 
3:  $V_t = C_t$  and  $x_t = 0$  for all  $t$ 
4: for  $k=1,2,\dots,T$  do
5:   Insert  $s_k^1$  into  $S$  such that  $S$  remains ordered non-decreasingly
6:    $R := B_k - \sum_{t=1}^k x_k$ 
7:   while  $R > 0$  do
8:     Take  $t$  and  $j$  such that  $s_t^j$  is the first piece from  $S$ 
9:      $\delta := \min\{R, V_t, z_t^j - z_t^{j-1}\}$ 
10:     $x_t = x_t + \delta$ 
11:     $R = R - \delta$ 
12:    for  $l=t,t+1,\dots,T$  do
13:       $V_l = V_l - \delta$ 
14:    end for
15:     $W := V_t.$ 
16:    for  $l=t-1,t-2,\dots,1$  do
17:       $W = \min\{W, V_l\}$ 
18:       $V_l = W$ 
19:    end for
20:    for  $s_{t'}^j \in S$  with  $V_{t'} = 0$  do
21:       $S = S \setminus \{s_{t'}^j\}$ 
22:    end for
23:    if  $\delta = z_t^j - z_t^{j-1}$  then
24:       $S = S \setminus \{s_t^j\}$ 
25:      if  $j < m_t$  and  $V_t > 0$  then
26:        Insert  $s_t^{j+1}$  into the ordered vector  $S$ 
27:      end if
28:    else if  $V_t > 0$  then
29:       $z_t^{j-1} = z_t^{j-1} + \delta$ 
30:    end if
31:  end while
32: end for
33: Return  $\mathbf{x}_{1 \rightarrow T}$ 

```

---



Iteration	$k$	$S_{start}$	$S_{end}$	$(x_1, x_2, x_3, x_4)$	$\Sigma$	$B_k$	$(V_1, V_2, V_3, V_4)$
Initial	0	$\emptyset$	$\emptyset$	$(0, 0, 0, 0)$	0	0	$(\frac{5}{2}, \frac{5}{2}, 4, 5)$
1	1	$\{s_1^1\}$	$\{s_1^1\}$	$(1, 0, 0, 0)$	1	1	$(\frac{3}{2}, \frac{3}{2}, 3, 4)$
2	2	$\{s_2^1, s_1^1\}$	$\{s_1^1, s_2^2\}$	$(1, \frac{1}{2}, 0, 0)$	$\frac{3}{2}$	2	$(1, 1, \frac{5}{2}, \frac{7}{2})$
3	2	$\{s_1^1, s_2^2\}$	$\{s_1^2, s_2^2\}$	$(\frac{3}{2}, \frac{1}{2}, 0, 0)$	2	2	$(\frac{1}{2}, \frac{1}{2}, 2, 3)$
4	3	$\{s_2^2, s_2^2, s_3^1\}$	$\{s_3^1\}$	$(2, \frac{1}{2}, 0, 0)$	$\frac{5}{2}$	3	$(0, 0, \frac{3}{2}, \frac{5}{2})$
5	3	$\{s_3^1\}$	$\{s_3^1\}$	$(2, \frac{1}{2}, \frac{1}{2}, 0)$	3	3	$(0, 0, 1, 2)$
6	4	$\{s_4^1, s_3^1\}$	$\{s_4^2, s_3^1\}$	$(2, \frac{1}{2}, \frac{1}{2}, 1)$	4	5	$(0, 0, 1, 1)$
7	4	$\{s_4^2, s_3^1\}$	$\{s_3^1\}$	$(2, \frac{1}{2}, \frac{1}{2}, 2)$	5	5	$(0, 0, 0, 0)$

(j) Values of variables during the algorithm

Figure 5.8: Example of an application of Algorithm 5.2 to the instance of BC described in Example 5.3. Figures (a)-(d) give the objective functions and (e)-(h) the right derivatives. Figure (i) gives the cumulative sum for the different iterations of the main for loop in the algorithm. The table given in (j) gives the value of various variables in the different iterations of the for loop of the algorithm.



$C_{t^*} - \sum_{t'=1}^{t^*} x_{t'}$  for  $t^* \geq t$ . Hence, in Line 9, we ensure that the upper cumulative bounds are satisfied.

To prove optimality of the algorithm, consider an instance of *BC* with piecewise linear convex objective functions. Let  $\mathbf{y}$  be an optimal solution, and let  $\mathbf{x}$  be the solution produced by Algorithm 5.2. Furthermore, let  $t$  be the smallest index for which  $x_t \neq y_t$ . We consider two cases.

*Case 1*  $x_t > y_t$ : Let  $t' > t$  be the smallest index for which  $x_{t'} < y_{t'}$ . This index must exist since  $\sum_{t=1}^T x_t = B_T = \sum_{t=1}^T y_t$  and  $x_t > y_t$ . Furthermore, let  $s_t$  and  $s_{t'}$  be the slopes of the pieces on which  $x_t$  and  $x_{t'}$  lie respectively. In case  $x_t \in Z_t$  (i.e.,  $x_t$  is an endpoint to two consecutive pieces), we pick the piece with  $x_t$  as endpoint and in case  $x_{t'} \in Z_{t'}$  we pick the piece with  $x_{t'}$  as begin point. By the fact that the right and left derivatives of the convex piecewise linear functions  $f_t$  and  $f_{t'}$  are non-decreasing it follows that  $s_t \geq f_t^+(y_t)$  and  $f_{t'}^-(y_{t'}) \geq s_{t'}$ . Also, by the optimality of  $\mathbf{y}$  it follows that  $f_t^+(y_t) \geq f_{t'}^-(y_{t'})$ . Note that, since both  $\mathbf{x}$  and  $\mathbf{y}$  are feasible, we can decrease  $x_t$  and increase  $x_{t'}$ , until either  $x_t$  is equal to  $y_t$  or  $x_{t'}$  is equal to  $y_{t'}$ , without violating feasibility. Finally, note that doing so does not increase the objective value.

*Case 2*  $x_t < y_t$ : This case can be treated completely symmetrical to the previous case.

Note, that the above process can be repeated until  $\mathbf{x} = \mathbf{y}$  without increasing the objective value. This shows that  $\mathbf{x}$  is indeed optimal.

Finally, we show that the time complexity of the algorithm is  $O(TM)$  with  $M$  the sum of the number of pieces of each  $f_t$ , i.e.  $M = \sum_{t=1}^T m_t$ . Note that a slope  $s_t^j$  can only be added to and subsequently removed from  $S$  once. Furthermore, if a slope is picked to be used in the while loop and not removed it follows that  $R = 0$  after this iteration and hence the while loop is finished. Finally, note that the steps inside the while loop can all be executed in time  $O(T)$ , since the size of  $S$  is never more than  $T$ . Hence, the total complexity of the for loop in the algorithm is  $O(TM)$ , which clearly dominates the complexity of the other steps.  $\square$

We note that the asymptotic complexity of Algorithm 5.2 is slightly lower than the complexity of an algorithm that recursively applies Algorithm 5.1 using Algorithm 4.3 from the previous chapter to obtain solutions for the resulting *EVC* problems. However, the latter method might be more favourable in practice, specifically if the number of intervals for which the cumulative bounds are tight is low, as can be expected in many applications [143, 146]. Furthermore, the asymptotic complexity of Algorithm 5.2 is the same as the complexity when we recursively apply Algorithm 5.1 using Algorithm 4.4 from the previous chapter. As noted before however, the latter method is probably only efficient for large instances, i.e., instances where the number of breakpoints is very large.

From a practical point of view, it is again important to consider for how many intervals  $t$  solutions to *dBC* use different values in the feasible set  $Z_t$  (as we also



did in the previous chapter). We obtain the following result, which is similar to Corollary 4.2 for Problem *rdEVC*.

**Corollary 5.1.** *There exists an optimal solution to *dBC* such that for any two indices  $t < t^*$  with  $x_t \notin Z_t$ ,  $x_{t^*} \notin Z_{t^*}$ , there exists an index  $k$  in  $\{t, t + 1, \dots, t^* - 1\}$  with  $\sum_{t'=1}^k x_{t'} \in \{B_k, C_k\}$ . In other words, between two intervals with ‘fractional’ solutions,  $x$  meets either the lower or upper cumulative bound tightly.*

This result follows immediately from the fact that we can recursively apply Algorithm 4.3 to Problem *dBC* by Lemma 5.1. In practice we do not expect these cumulative bounds to be met often. This expectation together with Corollary 5.1 implies that, in most practical cases, the expected number of time intervals for which two operational values are chosen instead of one is low.

## 5.4 APPLICATION OF BUFFER DEVICES

In the previous section we have shown that we are capable of scheduling several different types of devices in the buffer class within our profile steering DEM approach. As a proof of concept we now present the results of a simulation study where several of these devices are involved. The simulations were originally performed to study the potential of soft-islanding (i.e., balancing local production and consumption as much as possible) for a future neighbourhood of houses in the Netherlands. For this case sixteen houses are considered where each house has its own local production of energy in the form of rooftop photovoltaic (*PV*) panels and where a central CHP unit providing heat for the entire set of houses is present. The idea behind this setup is that the demand for heat is higher at times of the year when the *PV* production is lower, causing this gap to be filled by the electricity production of the CHP unit. The central CHP unit is connected to a large central heat buffer. This heat buffer decouples the production of heat and electricity by the CHP unit from the heat demand, providing the unit with flexibility in when and how much it produces. Furthermore, smart devices in the form of a washing machine (*WM*) (combination of washer and dryer) and dish washer (*DW*) are present in each house. Finally, in a second simulation, local storage in the form of a battery is added to each house to determine the effect this has on the results. A schematic overview of the situation is given in Figure 5.9.

When setting up this simulation, first reasonable sizes of both production (the CHP unit and *PV*) and storage (heat buffer and batteries) were determined. The size of the CHP unit together with the heat buffer was determined first, such that it could always supply the thermal demand of the houses. For the analysis we constructed a combined heat demand pattern of the spacial and hot water demands of the houses. These heat demand patterns of the houses were generated through simulations of the (proposed) houses [139]. The results of the analysis indicate that a CHP unit with a maximum output of 60 kW thermal suffices to supply the neighbourhood for all of the considered scenarios. As the typical ratio between heat and electricity

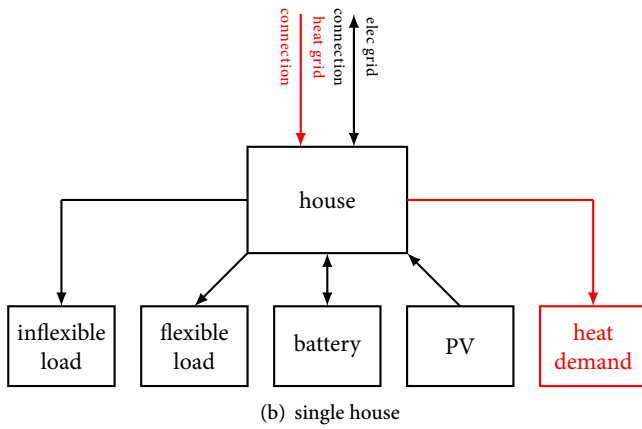
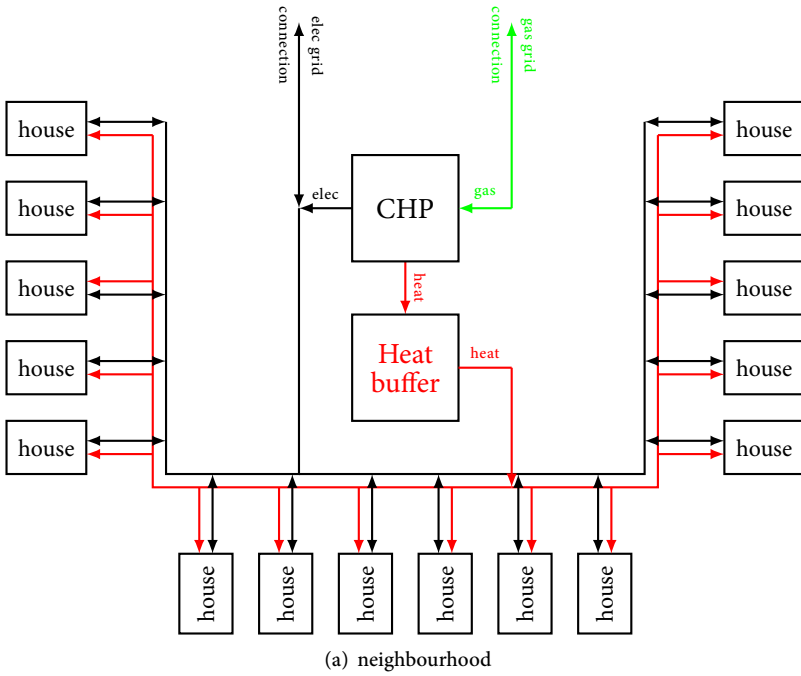


Figure 5.9: Schematic overview of the 16 houses case; (a) gives an overview of the grid and (b) an overview for a single house.



output of a CHP unit is about 2:1, we used this ratio in our simulation. As the size of the attached heat buffer has quite some impact on the flexibility of the CHP unit, we varied this size. While the size is set to 50 kWh for the standard case to reflect a realistic scenario in current practice, the buffer size is increased in case of smart control to 250 kWh to increase the flexibility provided by the CHP unit.

The size of the PV systems and the batteries was chosen based on the goal to ensure that the neighbourhood supplies its own energy as much as possible. This implies that the PV systems are sized to cover the daily remaining electricity demand of the neighbourhood after subtraction of the CHP production. The most challenging week for the considered combined system is in autumn, when PV production is low and heat demand is also relatively low, implying that the production of electricity from both PV and CHP is low. When analyzing data of a full year, week 43 turned out to have the lowest combined energy production from PV and the CHP unit. To ensure sufficient local production, the PV systems were sized to ensure enough production is available locally this week. Using panels of which 50% face south, 25% face west and 25% face east, all at an angle of 35 degrees, with an efficiency of 16% the required size of the system is 15 m<sup>2</sup> of PV panels per house.

To obtain the size of the batteries we first ran the simulation when no batteries are present. The batteries have to be sized such that they are able to store enough electrical energy from times of surplus production to provide any remaining demand drawn from the grid. From the resulting power flows obtained in the initial simulation we determined that a total storage capacity of 30 kWh is required. The maximum storage capacity requirement occurs in week 43, as this is the most difficult week for the considered system. To ensure the required storage capacity is met, we considered a two kWh battery for each house. Furthermore, note that the required storage capacity is relatively low. This is a consequence of the fact that already a lot of flexibility in electricity production from the central CHP unit with the (large) heat buffer is present in the neighbourhood.

The electric load for the houses is based on smart meter data obtained from houses in the Dutch town of Lochem (see [67]). For the smart appliances we used data on the load profiles and flexibility from [121]. For the base case no smart control is applied to the WMs and DWs. This means these devices run as soon as they are available. Furthermore, the CHP unit is steered in such a way that it meets the thermal demand of the neighbourhood with minimal changes in operational modes (i.e., between different production levels). For the control case, where the smart devices are utilized, we simulated both the case with and without the batteries. We use our profile steering approach for the control cases where we steer towards a zero profile, i.e., we minimize the import and export of electricity from the neighbourhood. Note that we only steer the use of electricity while only requiring the heat production to cover the demand. Furthermore, we assume the gas requirement of the CHP unit can always be satisfied. Finally, note that we use the algorithms for the continuous case (i.e., Algorithm 5.1) for the control of the CHP unit combined with the heat buffer and for the batteries.



For the simulations we consider three different weeks; week 6 in winter, week 26 in summer and week 43 in autumn. We present the simulation results for these weeks in Figures 5.10, 5.11, and 5.12 respectively. We see that the CHP unit does not switch often between different production levels in the base case, as is expected. This results in a rather erratic behaviour of the load profile of the considered system, resulting in peaks in both export and import of electricity. The CHP follows the electricity need of the neighbourhood as much as possible when control is applied, meaning that it mostly runs during the evenings, when demand for electricity is high. This causes the electric load requirements to be flattened out over the days for as much as the included storage allows. Without battery storage a slight import from the grid is required in both the summer and autumn week. When the batteries are included, the entire load profile can be smoothed enough to ensure that there is never a remaining net demand.

Finally, note that we focus on the potential for so-called soft-islanding in this study, where a connection to the main grid is still required and beneficial to ensure balance on smaller time scales and to export the surplus of energy observed through the year.

## 5.5 CONCLUSION

In this chapter we extended the model for EV charging introduced in the previous chapter to include discharging and SoC constraints for every time interval. This new model allows us to address cases when the EV can also discharge energy to the grid. Furthermore, this model can be applied to a stand-alone battery or a heat generator, such as a HP or CHP unit, combined with a heat buffer.

The extension of the model transforms the related optimization problem to include so-called cumulative bounds. We showed that we can efficiently solve the continuous version of this optimization problem. To do so we solve an instance of Problem *EVC*, presented in the previous chapter, and apply a divide and conquer strategy where we split the instance at the time interval where the cumulative bound is violated the most. For the discrete variant of the problem we considered the same modifications as in the previous chapter, allowing convex combinations of the feasible operational levels for every time interval. We solved this discrete variant using a greedy approach and argued that we expect that solutions in practice use convex combinations to switch between operational levels only rarely.

As an application of the developed algorithms, we considered a case study where sixteen houses in the Netherlands with a central CHP unit and rooftop PV are soft-islanded, i.e., as much of the demand as possible is supplied by local production. After the sizing of the system we applied our profile steering approach to two cases, one with batteries installed in the houses and one without these batteries. The results show that profile steering is capable of minimizing the need to import energy from the main grid even in difficult scenarios with low local production. Furthermore, with battery storage available the considered houses become a net exporter



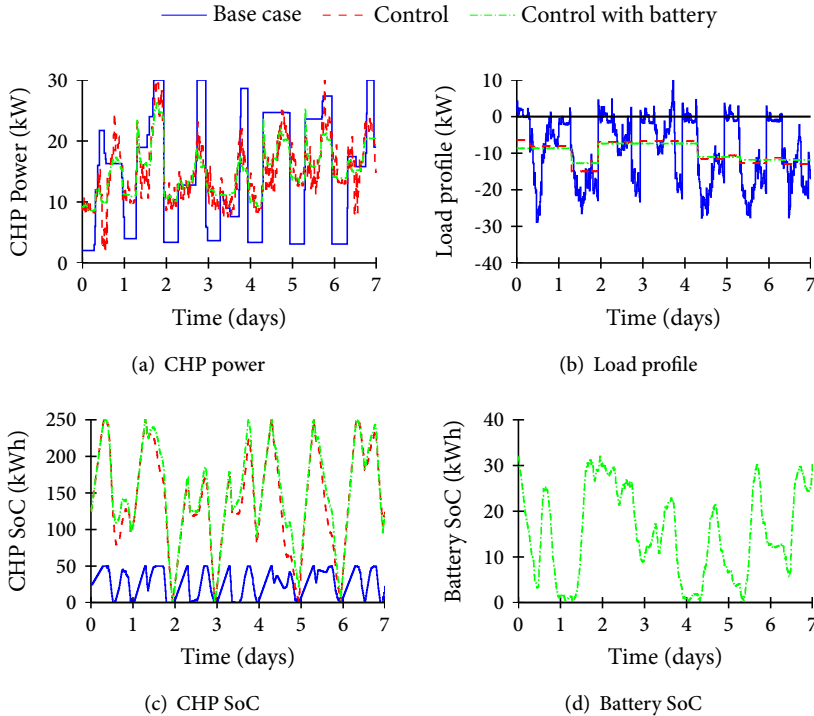


Figure 5.10: The results for the three different control cases for the 16 houses studied for week 6 (winter).

for all of the considered time intervals of fifteen minutes. We note that a connection to the main grid in such a scenario is still required to ensure balance also on a lower time-scale (i.e., sub-seconds and seconds) and to export the abundant energy which is available most of the time.

The models considered in the previous chapter could easily be adapted to include an efficiency factor for the charging of energy. As the same variable in this chapter covers both charging and discharging, this extension no longer applies when considering the models of this chapter. Furthermore, in the given models we did not consider losses depending on the SoC of the considered buffer. These losses are negligible for most batteries on the considered time-scales (in the order of less than one percent [53]) However, they are of importance when considering thermal storage. We consider this issue in the next chapter, where a model for heating, ventilation, and air conditioning is presented that incorporates such losses.

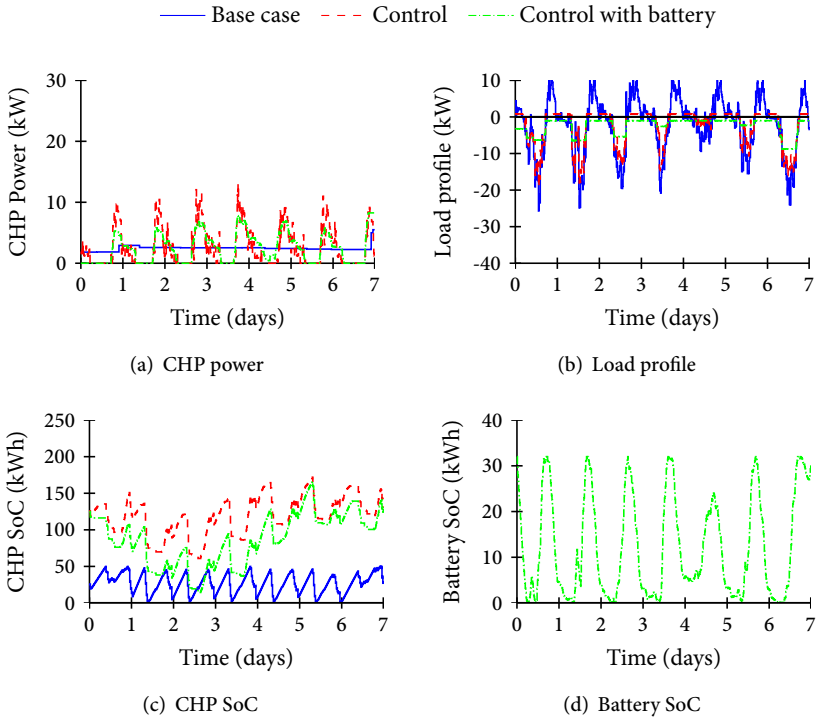


Figure 5.11: The results for the three different control cases for the 16 houses studied for week 26 (summer).

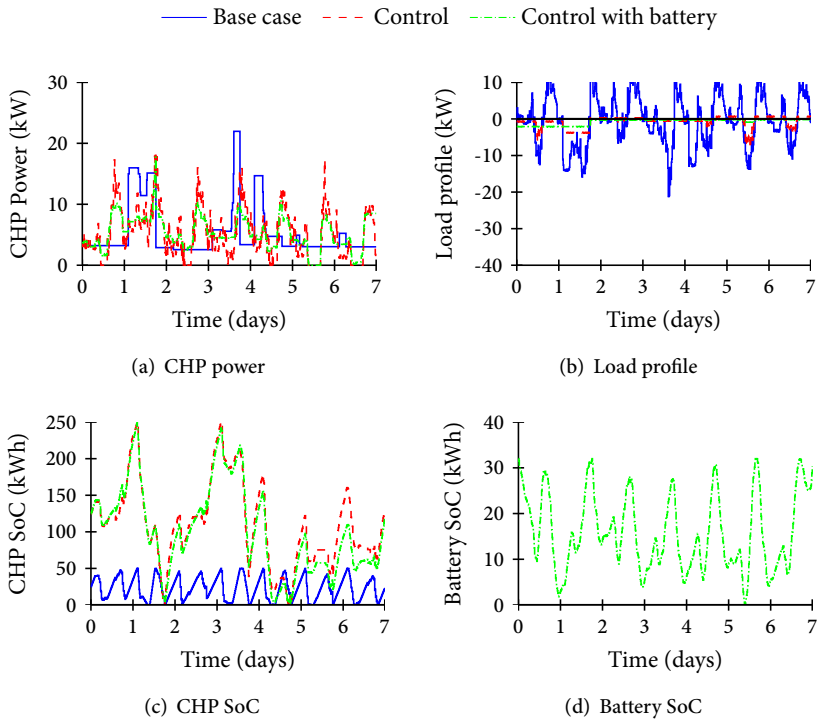


Figure 5.12: The results for the three different control cases for the 16 houses studied for week 43 (autumn).




# 6



## HEATING, VENTILATION, AND AIR CONDITIONING SYSTEMS


*ABSTRACT – A large share of our energy use is for thermal reasons, for hot tap water and to heat or cool our environment. To provide this heating or cooling, many heating systems exploit a thermal buffer. These systems often can be modelled as storage systems. In this chapter we extend the models from previous chapters to include an important aspect of heating and cooling systems: losses depending on the state of charge. We modify the solution approaches found in the previous chapter to be able to deal with the extended models. To test the solution approaches we use data from Austin, Texas. We use this data to determine the model coefficients and to study the effectiveness of profile steering when compared to traditional deadband control.*




This chapter is based on [TvdK:5].




■ space heating   ■ water heating   ■ space cooling   ■ appliances



(a) OECD average



(b) Texas



(c) Illinois

Figure 6.1: Distribution of the use of energy in households. The figures give a rough global overview (average of OECD countries) and the data for two states in the US with vastly different climates. Data taken from [75] and [133].

## 6.1 INTRODUCTION

In the previous chapters we discussed buffering devices and their importance in the future smart grid. We mostly discussed devices that currently have only low market penetration in the residential sector such as electric vehicles (*EVs*) and batteries. Furthermore, we neglected losses that depend on the state of charge (*SoC*) of the system. However, these losses play an important role in heating and cooling systems. Such systems are in general large consumers of energy, thus affecting the electricity grid in a significant way when run on electricity. However, they can also offer a large amount of flexibility to the grid. As such systems are already rather common in the residential sector, e.g., in the form of heat pumps (*HPs*) and heating, ventilation, and air conditioning (*HVAC*) systems, they can offer their flexibility already now rather than in the future. This means such systems can help to tackle the challenges existing with integrating renewables today and can increase the efficiency of our energy supply chain.

The heat demand of typical residential customers consists of two parts: space heating and (hot) tap water. We note that for warmer climates the space heating demand is replaced by a space cooling demand, especially during summer months. For an illustration of how the shares of energy consumption in households change with different climates see Figure 6.1. In Texas the space heating requirement is quite low, but a significant portion of the energy consumption of a household is instead used for cooling. We focus in this chapter on this cooling demand for the warmer climates and discuss a model for the flexibility that can be provided by HVAC systems used to cool residential buildings. However, most of the results in this chapter also apply to space heating requirements and hot tap water demands with only minor modifications.

To be able to use heating and cooling systems in a decentralized energy management (*DEM*) approach such as our profile steering approach, we first need to determine the available flexibility. In systems with a storage attached, this flexibility comes from the fact that the storage allows the system to consume electricity and transform it into heating or cooling energy before it is required. However, not all of these systems have such an explicit storage attached to it. In this case there is potential for flexibility by considering the building itself as a thermal storage unit. This allows the system to vary the indoor temperature, within user defined comfort limits, to adjust the electricity consumption of the cooling system when required. Most residential HVAC systems used today aim to keep the indoor temperature within pre-specified bounds, i.e., they use a deadband controller and thereby do not offer any flexibility. However, with more advanced control, such a system could exploit the flexibility by, e.g., switching on earlier to shift part of the energy requirement from a future time interval to the current one.

Of particular importance in heating and cooling systems are losses that depend on the *SoC* of the storage. For example, in a cooling scenario, houses heat up faster when the difference between indoor and ambient temperature is higher. In this





chapter we adapt the models from the previous chapters to include SoC dependent losses. A further aspect we discuss is based on our statement in Chapter 2 that an important step in a scheduling based DEM approach is *operational control* to account for differences between predictions and realizations. As exact predictions of thermal influences inside houses is near impossible, e.g., because it is dependent on user behaviour, we pay special attention to the operational control step in this chapter. We show that we can account for prediction errors, made by our model, using a system based on model predictive control (*MPC*).

The remainder of the chapter is outlined as follows. In the next section we briefly discuss related work on (residential) HVAC systems and their use in energy management (*EM*) approaches. Then, in Section 6.3 we introduce the thermal model used in this chapter and fit the model to data from the Pecan Street Inc. [73] dataset, detailing consumption data from houses in Austin, Texas. In Section 6.4 we formulate the resulting device level optimization problem first as a continuous problem and then as a discrete problem. For both variants we give an efficient solution method. As the thermal model we introduce may lead to significant prediction errors, we adapt the operational control step of the profile steering approach to account for these errors in Section 6.5. We use the solutions developed in this chapter in a simulation study in Section 6.6. Finally, we wrap up with a conclusion and discussion in Section 6.7.

## 6.2 RELATED WORK

For the control of HVAC systems it is important that the control maximizes user comfort. This is in itself already a difficult problem to solve. To this end many different control algorithms have been proposed and studied for HVAC systems (see, e.g., [5]). In this chapter we focus on MPC. In MPC a model of the controlled system, in this case a thermal model of a building or room, is used to predict the effect of possible control actions. The controller works in discrete time steps to determine a best schedule of control actions for a time horizon based on the used model and a predetermined objective. From this schedule only the first control action is used. The resulting changes in the modelled system are measured and a new schedule is made in the next time interval based on updated information and predictions of the model.

In [16] Bashash and Fathy studied HVAC systems from a control theory perspective. They aggregate the models of individual systems and show that the (simplified) aggregate model is bilinear. This allows the authors to develop corresponding control algorithms. They design a state-space feedback controller, which they verify using numerical simulations. Their results indicate that the controller is capable of getting the controlled HVAC systems to follow a desired load profile.

Cole *et al.* [32] consider both centralized and decentralized MPC for peak shaving using HVAC systems. They obtain their model using actual measurements obtained in Texas. For these measurements they develop a reduced order model





[31]. Their objective is to reduce the costs of the energy used for cooling. The decentralized control is capable of achieving results very similar to the centralized control, without requiring detailed privacy-sensitive information. Their results are mainly based on pre-cooling, i.e., the HVAC systems are turned on when energy is cheaper so the house is cooler when energy becomes more expensive, allowing the system to consume less at expensive times.

Aswani *et al.* [9] use learning-based MPC to reduce the overall energy consumption of an HVAC system used to cool a computer lab. Their thermal model is linear but nevertheless shows good results when used to predict the room temperature depending on different operations of the HVAC system. Their focus is on energy saving, thus reducing costs for the owner/operator of the system. Nevertheless, their models should also be applicable in the setting of an EM approach when slightly modified.

### 6.3 THERMAL MODEL

The buffer models for the cooling systems we consider in this chapter have a similar structure as in the previous chapter. We consider the whole structure, including internal air volume, of the house as a ‘buffer’, that needs to be kept between temperature bounds specified by the user (i.e., (one of) the occupants). This buffer heats up due to (thermal) energy leaking into the house from the environment (inflow) and the HVAC system cools the buffer by extracting this (thermal) energy. An important difference is that the outflow process is now controlled, as the HVAC system takes (thermal) energy from the room or house. Furthermore, the inflow process is uncontrolled and (partially) depends on the SoC of the buffer, i.e., the temperature inside the house. By this we mean that more energy flows into the building with a higher difference between indoor and outdoor temperature, i.e., when the buffer has a lower SoC. A similar model applies to many heating systems, where more energy dissipates to the environment when the temperature difference is higher.

To integrate HVAC systems in profile steering, we need to find local feasible schedules for the system. For this, we need to know the constraints for these schedules. For HVAC systems, these constraints are user comfort constraints. In a residential setting, these constraints are typically given by a temperature set point  $\bar{T}_t$  and an allowed maximum deviation from this set point  $D_t$  for every time interval. When considering a system denoted by  $m$ , these constraints lead to a local feasible set  $X^m$  given by:

$$X^m := \{\mathbf{x}^m \mid \bar{T}_t - D_t \leq T_t(\mathbf{x}^m) \leq \bar{T}_t + D_t\}, \quad (6.1)$$

where  $T_t(\mathbf{x}^m)$  is the indoor temperature, which depends on the schedule  $\mathbf{x}^m$  for the system, i.e., when the HVAC system is switched on. We assume the values of  $\bar{T}_t$  and  $D_t$  are known as they can, for example, be set by the user on a thermostat or home energy management system (HEMS). In order to determine the feasible set




Figure 6.2: Schematic overview of the interaction between HVAC unit energy consumption ( $\mathbf{x}^m$ ) and indoor temperature ( $T_t(\mathbf{x}^m)$ ).

$X^m$  of schedules we need to determine *how* the indoor temperature  $T_t$  depends on the schedule  $\mathbf{x}^m$ , i.e., how the energy consumption of the HVAC system influences  $T_t$ . In Figure 6.2 we give an example of how the energy consumption of an HVAC system influences the indoor temperature. Note that we depict the use of a (perfect) deadband controller (i.e., it switches the system on when the temperature reaches the upper bound and switches it off again as soon as the temperature reaches the lower bound).

In the next subsection we first specify the model we use for  $T_t$  and then fit it to data obtained in Austin, Texas. Then, we verify the model by comparing its output with measurements from the dataset in Subsection 6.3.2.

### 6.3.1 MODEL DETERMINATION

For our thermal model, we consider a linear discrete-time model. This agrees with the time steps used in profile steering and has been shown to work for small sized HVAC systems [9]. Furthermore, a linear model ensures scalability and tractability of our system for a larger number of houses. To determine the set of feasible schedules for the system, given by (6.1), we determine the indoor temperature  $T_{t+1}$  for the next time interval based on the current indoor temperature  $T_t$ , the HVAC average power consumption  $x_t$  and the outdoor temperature  $O_t$ :

$$T_{t+1} = aT_t + bx_t + cO_t + d_t, \quad (6.2)$$

where the coefficients  $a$ ,  $b$ ,  $c$ , and  $d_t$  are parameters of the model. Note that all the used variables are of the current time interval  $t$ . Furthermore, we assume that  $a \geq 0$ , meaning that the current temperature has a positive influence on the next time interval. Note that parameter  $d_t$  varies over time and can be used to model thermal gains and losses not captured by the other parts of the model, e.g., thermal gains from solar radiation and human occupancy/behaviour. This is similar to the model used for residential HVAC systems in [31]. To determine the constant  $d_t$



from historic data, we assume that it is invariant for the same time interval on different days, i.e.,  $d_t = d_{t+96}$  when using fifteen minute time intervals.

Note that we may assume that the indoor temperature does not change if the HVAC system is unused, the indoor and outdoor temperature are equal, and the other thermal gains/losses are zero (i.e.  $x_t = d_t = 0$  and  $T_{t+1} = T_t = O_t$ ). Thus we can assume that:

$$T_t = aT_t + cT_t, \quad (6.3)$$

from which we can conclude that  $a + c = 1$ . This allows us to rewrite (6.2) to

$$T_{t+1} - T_t = bx_t + c(O_t - T_t) + d_t, \quad (6.4)$$

relating the indoor temperature *change* between two consecutive time intervals to the HVAC system power consumption, the difference between indoor and outdoor temperature, and other thermal gains and losses.

To obtain the coefficients of the thermal model given in (6.2), we use data from the Pecan Street Inc. [73] dataset, which contains detailed electricity consumption data for a large body of houses predominantly in Austin, Texas. We combine this data with openly accessible weather data from Austin, Texas [2]. As the focus of this work is on cooling, we consider only data obtained in the summer of 2015, between the 1st of June and the 31st of September. We identified a total of ten households for which the required data is available. For each of these houses, we fit the model given in (6.2) using linear regression.

### 6.3.2 MODEL VERIFICATION

To verify the validity of the fitted model, we are mainly interested in the predictive power of the model, as this is exactly the purpose for which we designed the model. To this end we split the data of each house into two sets. For the first set, the training set, we have chosen to use the odd numbered days and fit the coefficients of the model using this set. The second set consists of the remaining days, which we use to predict indoor temperatures. These indoor temperatures are then compared to the measurements in the set, to determine the accuracy of predictions made by our model.

The mean and standard deviation of the differences between the predicted values and measured values in the validation set for each house are given in Table 6.1. The average error made by the model when predicting indoor temperature changes is nearly 0. However, the standard deviation varies from house to house and indicates that the errors can be quite large. This is not surprising, as human behaviour can be a large factor in the thermal gains and losses of a household [9, 45]. Furthermore, human behaviour is in general hard to predict and erratic on the level of households. Therefore, we believe it is not realistic to pursue a model capable of accurately deriving the thermal gains and losses due to variations in human behaviour. Hence,



Table 6.1: Statistics on the errors made by the thermal model given in (6.2) for ten houses in Austin, Texas. Furthermore, we use the model to predict the daily average HVAC system use and compare this to the data.

House	Mean (°C)	Standard deviation (°C)	Auto- correlation	MAPE daily HVAC use (%)
1	$5.67 \times 10^{-5}$	0.17	-0.07	7.1
2	$-2.80 \times 10^{-3}$	0.15	-0.28	15.6
3	$1.06 \times 10^{-4}$	0.11	-0.38	8.4
4	$1.38 \times 10^{-3}$	0.31	0.17	31.8
5	$3.65 \times 10^{-3}$	0.17	0.26	25.1
6	$-4.21 \times 10^{-3}$	0.22	-0.21	14.3
7	$9.22 \times 10^{-4}$	0.12	-0.27	17.4
8	$5.62 \times 10^{-4}$	0.11	-0.18	13.2
9	$7.98 \times 10^{-4}$	0.38	-0.03	312.6
10	$-3.92 \times 10^{-3}$	0.23	-0.31	21.0

reliable EM approaches that plan and predict the use of HVAC systems must be able to account for these errors through other means. Later on in this chapter we discuss how we can deal with these errors within the operational control step of the profile steering approach.

To ensure that the model is suitable for use in profile steering, we investigate if the model gives reasonable HVAC power consumption prediction values compared to measured data. Using data from our validation set and (6.4) we estimate the expected required power consumption of the HVAC system for every time interval. As could be expected, due to the, sometimes large, errors made by the model, this sometimes leads to inaccurate or even infeasible values. However, as the average error made by the model is nearly zero, we consider an entire day instead, i.e., in the case of fifteen minute intervals we estimate the HVAC energy consumption for 96 intervals. Summing (6.4) over 96 intervals and rewriting it results in:

$$\sum_{i=t}^{t+96} x_i = \sum_{i=t}^{t+96} \frac{T_{i+1} - T_i - c(O_i - T_i) - d_i}{b}. \quad (6.5)$$

We use data from our validation set and (6.5) to estimate the HVAC power consumption for complete days and compare this estimate to the measured values. The mean absolute prediction error (MAPE) between the predicted consumption and the measured consumption for each day in the validation set is given in column five of Table 6.1. Generally, the prediction error made by our model is about ten to twenty-five percent. As we use only a linear model and human behaviour is very



hard to predict, we do not expect better predictions are possible using models of similar complexity. The very large value for house nine stands out. We believe it is due to the fact that the HVAC system in this house only runs sporadically. Furthermore, the daily HVAC power consumption for this house does not vary much over the two months of data in the validation set, indicating that the HVAC system is mainly used to compensate for thermal gains caused by other factors than outside temperature. As these are not explicitly modelled in our thermal model, it is not surprising that the model does not do very well in predicting the HVAC usage.

Based on the large standard deviation found for the distribution of the errors made by the models of houses four and nine, we deemed these models unfit and excluded them from the simulation study detailed later on in this chapter. As mentioned before, for house nine this is believed to be due to outdoor temperatures having very little effect on the indoor temperature. Furthermore, while the model of house five seemed to be a decent fit, it turned out that the HVAC system in this house has to run almost constantly to keep the building at a desirable temperature. This indicates that the HVAC system itself is undersized for its actual use and hence offers nearly no flexibility to any EM approach. Hence we also excluded this house from our simulation study.

Finally, we determined if a correlation exists between the various model parameters and the error made by the model. Hereby, no significant correlation was found. Also, we determined the autocorrelation between the time series of errors made by the models for each of the houses. The autocorrelation with a lag of a single time step is given in column four of Table 6.1. As this value is negative for each of the houses for which we deemed the model decent fits, it follows that we expect a (large) positive error to cause the error for the next time step to be negative.

## 6.4 HVAC SCHEDULING PROBLEM

In the previous section we determined a model, given in (6.2), which relates how the temperature inside a building or house changes depending on the current temperature, on the electricity consumption of the HVAC system and on the outdoor temperature. Furthermore, we argued that the schedule made by an HVAC system must maintain the temperature within a deadband around a given set point. The system now needs to locally determine an optimal schedule based on the received steering signal and these local constraints. This results in the following optimization problem:

**Problem 6.1 (HVACS).**

$$\min_{\mathbf{x}} f(\mathbf{x}), \quad (6.6a)$$

$$\text{s.t. } \overline{T_{t+1}} - D_{t+1} \leq aT_t + bx_t + cO_t + d_t \leq \overline{T_{t+1}} + D_{t+1} \quad t = 1, 2, \dots, T, \quad (6.6b)$$

$$\mathbf{x}^{\min} \leq \mathbf{x} \leq \mathbf{x}^{\max}. \quad (6.6c)$$



We call this optimization problem the HVAC scheduling (*HVACS*) problem. Using the same substitution as used in the previous chapters we can again assume that  $\mathbf{x}^{min} = 0$ . We note that in Constraints (6.6b) the bounds  $\overline{T}_{t+1} - D_{t+1}$  and  $\overline{T}_{t+1} + D_{t+1}$  are for the next time interval. This is because we use a discrete thermal model, meaning that the value  $T_t$  can be interpreted as both the temperature at the start of time interval  $t$  and at the end of time interval  $t - 1$ . Thus, the constraints only pose limits on the temperature at the beginning/end of time intervals and we assume the temperature in between is well-behaved enough to not violate user comfort.

When considering Constraint (6.6b) we note that  $T_1$  is the temperature at the start of the optimization horizon, which we assume to be known. Furthermore, we assume we have reasonable predictions of the outdoor temperature  $O_t$  for  $t = 1, 2, \dots, T$ . Now, using (6.2) to substitute  $T_t$  we obtain that (6.6b) is equivalent to:

$$\overline{T}_{t+1} - D_{t+1} \leq a^2 T_{t-1} + b(ax_{t-1} + x_t) + c(aO_{t-1} + O_t) + ad_{t-1} + d_t \leq \overline{T}_{t+1} + D_{t+1} \quad (6.7)$$

Now, again using (6.2) we recursively substitute  $T_{t-1}, T_{t-2}, \dots, T_2$  to obtain that (6.6b) is equivalent to:

$$\overline{T}_{t+1} - D_{t+1} \leq \sum_{t'=1}^t [a^{t-t'} bx_{t'} + a^{t-t'} cO_{t'} + a^{t-t'} d_{t'}] + a^t T_1 \leq \overline{T}_{t+1} + D_{t+1} \quad (6.8)$$

Defining

$$B_t := \frac{\overline{T}_{t+1} - D_{t+1} - a^t T_1 - \left( \sum_{t'=1}^t a^{t-t'} cO_{t'} + a^{t-t'} d_{t'} \right)}{b}, \quad (6.9a)$$

$$C_t := \frac{\overline{T}_{t+1} + D_{t+1} - a^t T_1 - \left( \sum_{t'=1}^t a^{t-t'} cO_{t'} + a^{t-t'} d_{t'} \right)}{b}, \quad (6.9b)$$

we can rewrite (6.8) to

$$B_t \leq \sum_{t'=1}^t a^{t-t'} x_{t'} \leq C_t \quad t = 1, 2, \dots, T. \quad (6.10)$$

If we now substitute (6.10) for (6.6b) in Problem *HVACS*, we obtain a problem very similar to Problem *BC* we discussed in the previous chapter. We note that, similar to the assumption we made in the previous chapter, we can assume that  $B_T = C_T$ . Furthermore, Constraints (6.10) can be interpreted as SoC constraints on a leaky buffer. To calculate the SoC at time interval  $t$ , given by  $\sum_{t'=1}^t a^{t-t'} x_{t'}$ , the charging done in interval  $t'$ , given by  $x_{t'}$  is discounted by a factor  $a$  for every time interval between  $t'$  and  $t$ . In other words, a percentage of the energy stored into the buffer in time interval  $t'$  is lost during every time interval between  $t'$  and  $t$ .



In Subsection 6.4.2 we derive an efficient solution method for Problem *HVACS*, which is similar to the solution we obtained for Problem *BC*. To do so, we first study the problem we obtain when we drop Constraints (6.10) except for the last time interval. This problem can be transformed into Problem *EVC*, as studied in Chapter 4, using a substitution and can be solved using the same approaches. We then solve Problem *HVACS* using a divide and conquer approach very similar to the approach we used for Problem *BC*.

#### 6.4.1 HVACS WITHOUT CUMULATIVE BOUNDS

We first study Problem *HVACS* where we drop Constraints (6.10) for all but the last time interval. This problem is given by:

**Problem 6.2** (*rHVACS*).

$$\min_{\mathbf{x}} f(\mathbf{x}), \quad (6.11a)$$

$$\text{s.t. } B_T \leq \sum_{t=1}^T a^{T-t} x_t \leq C_T, \quad (6.11b)$$

$$\mathbf{x}^{\min} \leq \mathbf{x} \leq \mathbf{x}^{\max}. \quad (6.11c)$$

We call this problem the relaxed HVAC scheduling (*rHVACS*) problem. If we substitute  $y_t = a^{T-t} x_t$  and  $C = C_T$ , where we note that we assumed that  $f(\mathbf{x}) = \sum_t f_t(x_t)$  and  $B_T = C_T$ , we obtain

$$\min_{\mathbf{y}} \sum_{t=1}^T f_t\left(\frac{y_t}{a^{T-t}}\right), \quad (6.12a)$$

$$\text{s.t. } \sum_{t=1}^T y_t = C \quad (6.12b)$$

$$x_t^{\min} \leq \frac{y_t}{a^{T-t}} \leq x_t^{\max} \quad t = 1, 2, \dots, T. \quad (6.12c)$$

We observe that the function  $g_t(y_t) := f_t\left(\frac{y_t}{a^{T-t}}\right)$  is convex since  $f_t$  is convex. Furthermore, by defining

$$y_t^{\min} := a^{T-t} x_t^{\min}, \quad (6.13a)$$

$$y_t^{\max} := a^{T-t} x_t^{\max}, \quad (6.13b)$$

and substituting this together with  $g_t$  into (6.12) we obtain

$$\min_{\mathbf{y}} \sum_{t=1}^T g_t(y_t), \quad (6.14a)$$

$$\text{s.t. } \sum_{t=1}^T y_t = C, \quad (6.14b)$$

$$\mathbf{y}^{\min} \leq \mathbf{y} \leq \mathbf{y}^{\max}, \quad (6.14c)$$



which is an instance of Problem *EVC* (see Chapter 4). We note that, with the substitution given above, the structure of the objective function often does not change significantly. For instance, for quadratic objective functions, discussed in the previous chapters, the substitution of  $y_t = a^{T-t}x_t$  into such objective functions keeps the objective quadratic, meaning that the solution methods discussed in Chapter 4 remain applicable.

For Problem *EVC* we have the optimality conditions stated in Lemma 4.1, from which we obtain the following optimality conditions for Problem *rHVACS*.

**Lemma 6.1.** *A solution  $\mathbf{x}$  to Problem *rHVACS* is optimal if and only if it is feasible and there exists a multiplier  $\lambda$  such that, for all  $t$ :*

$$x_t^{\min} < x_t < x_t^{\max} \quad \Rightarrow \quad \frac{f_t^-(x_t)}{a^{T-t}} \leq \lambda \leq \frac{f_t^+(x_t)}{a^{T-t}}, \quad (6.15a)$$

$$x_t = x_t^{\min} \quad \Rightarrow \quad \frac{f_t^+(x_t)}{a^{T-t}} \geq \lambda, \quad (6.15b)$$

$$x_t = x_t^{\max} \quad \Rightarrow \quad \frac{f_t^-(x_t)}{a^{T-t}} \leq \lambda. \quad (6.15c)$$

*Proof.* The proof follows immediately from substituting  $x_t = \frac{y_t}{a^{T-t}}$  and  $g_t(y_t) = f_t(\frac{y_t}{a^{T-t}})$  and applying Lemma 4.1.  $\square$

These optimality conditions intuitively state that, in an optimal solution, no two time intervals  $t$  and  $t'$  and  $\delta > 0$  can be found such that increasing  $x_t$  by  $\frac{\delta}{a^{T-t}}$  and decreasing  $x_{t'}$  by  $\frac{\delta}{a^{T-t'}}$  decreases the objective while keeping  $\mathbf{x}$  feasible. Note that this modification to  $\mathbf{x}$  keeps  $\sum_{t=1}^T a^{T-t}x_t$  the same, i.e., it always keeps  $\mathbf{x}$  feasible with respect to Constraint (6.11b).

The above shows that, with minor modifications, we can solve Problem *rHVACS* by applying any solution method applicable to Problem *EVC* (see, e.g., the solution methods presented in Chapter 4). We note that instances of Problem *rHVACS* do not often match to the constraints in reality. This is because we do not consider constraints on the SoC on other time intervals besides the last. While we could do this for the EV in Chapter 4 because we could assume that the SoC was a non-decreasing function of time, this is no longer the case here as there are losses involved with the current SoC for every interval. For example, consider an instance where it is undesirable to draw any energy from the grid in the last few time intervals. Then, an optimal solution to *rHVACS* will charge more than a total of  $C$  on the first intervals, after which only energy is discharged through the losses until the final SoC is equal to  $C$ . This potentially violates the capacity constraint of the buffer.





### 6.4.2 SOLUTION APPROACH FOR HVACS

In order to solve Problem *HVACS* we use an approach similar to the divide and conquer strategy formulated for Problem *BC* in the previous chapter. The approach is based around a slight modification of the result given in Lemma 5.1 for Problem *BC*. In our solution approach to Problem *BC* we first construct a candidate solution by solving an instance of Problem *EVC*. Then we take the maximum violation of the SoC-constraints in the found candidate solution and split the instance into two separate instances by setting the lower and upper bound to the same value for the time interval where the maximum violation occurs. The reason to take the maximum violation is that fixing the cumulative bound at larger violations can already make it such that the solution also becomes feasible for smaller violations of the bounds. This follows from the fact that the total charging must remain the same, i.e., from the constraint that  $\sum_{t=1}^T x_t = C$ . In Problem *HVACS* this constraint now includes a discount factor  $a$ . This factor implies that to reduce the total charging, i.e., the SoC, on time interval  $t$  by  $\delta$  the charging done at an interval  $t' < t$  needs to be decreased by more than  $\delta$ . In other words, reducing  $x_t$  by  $\delta$  reduces  $\sum_{t'=1}^{t^*} x_{t'}$ , with  $t^* > t$ , by less than  $\delta$ . To take this effect into account, we need to scale the violations obtained in our candidate solutions to determine at which time interval the cumulative bound of an optimal solution is tight.

**Lemma 6.2.** *Consider an instance of Problem HVACS with  $C_T = B_T$  and let  $\mathbf{y}$  be an optimal solution to the instance of Problem *rHVACS* obtained by ignoring the cumulative bounds for all indices except the last. Assume that  $\mathbf{y}$  is not feasible for the considered instance of Problem *HVACS* and define*

$$k := \arg \max_t \left\{ a^{T-t} \left( \sum_{t'=1}^t a^{t-t'} y_{t'} - C_t \right), a^{T-t} \left( B_t - \sum_{t'=1}^t a^{t-t'} y_{t'} \right) \right\}. \quad (6.16)$$

*Then, there is an optimal solution  $\mathbf{x}$  to the considered instance of Problem HVACS such that, if  $\sum_{t=1}^k a^{k-t} y_t > C_k$  then  $\sum_{t=1}^k a^{k-t} x_t = C_k$  and, on the other hand, if  $\sum_{t=1}^k a^{k-t} y_t < B_k$  then  $\sum_{t=1}^k a^{k-t} x_t = B_k$ .*

*Proof.* The proof is similar to that of Lemma 5.1 given in the previous chapter. Let  $\mathbf{x}$  be an optimal solution to the considered instance of Problem *HVACS* and assume that  $\sum_{t=1}^k a^{k-t} y_t > C_k$  and  $\sum_{t=1}^k a^{k-t} x_t \neq C_k$ . Since  $\mathbf{x}$  is feasible, it follows that  $\sum_{t=1}^k a^{k-t} x_t < C_k$ . Let  $l$  be the last index before  $k$  for which the upper cumulative bound is met by  $\mathbf{x}$  with equality, i.e.  $l := \max\{t < k \mid \sum_{t'=1}^t a^{t-t'} x_{t'} = C_t\}$  and  $l := 0$  if this set is empty. Furthermore, let  $m$  be the first index after  $k$  for which the upper cumulative bound is met by  $\mathbf{x}$  with equality, i.e.  $m := \min\{t > k \mid \sum_{t'=1}^t a^{t-t'} x_{t'} = C_t\}$ . Note that  $\sum_{t=1}^T a^{T-t} x_t = C_T$  by assumption, hence  $m$  is well defined and  $m \leq n$ . We wish to show that there exist indices  $n$  and  $n^*$  such that  $l < n \leq k < n^* \leq m$  and  $y_n > x_n$  and  $y_{n^*} < x_{n^*}$ .

We start by showing that such an index  $n$  exists. We know, by construction of  $k$ , that



$$a^{T-l} \left( \sum_{t=1}^l a^{l-t} y_t - C_l \right) \leq a^{T-k} \left( \sum_{t=1}^k a^{k-t} y_t - C_k \right). \quad (6.17)$$

Dividing both sides by  $a^{T-k}$  we obtain

$$a^{k-l} \left( \sum_{t=1}^l a^{l-t} y_t - C_l \right) \leq \sum_{t=1}^k a^{k-t} y_t - C_k. \quad (6.18)$$

From this we obtain that

$$\begin{aligned} C_k - a^{k-l} C_l &\leq \sum_{t=1}^k a^{k-t} y_t - \sum_{t=1}^l a^{k-l} a^{l-t} y_t \\ &= \sum_{t=1}^k a^{k-t} y_t - \sum_{t=1}^l a^{k-t} y_t \\ &= \sum_{t=l+1}^k a^{k-t} y_t. \end{aligned} \quad (6.19)$$

On the other hand, by assumption on  $\mathbf{x}$  and construction of  $l$ , we have that:

$$\sum_{t=1}^k a^{k-t} x_t < C_k, \quad (6.20a)$$

$$\sum_{t=1}^l a^{l-t} x_t = C_l. \quad (6.20b)$$

Multiplying (6.20b) by  $a^{k-l}$  and combining it with (6.20a) gives

$$C_k - a^{k-l} C_l > \sum_{t=1}^k a^{k-t} x_t - \sum_{t=1}^l a^{k-t} x_t = \sum_{t=l+1}^k a^{k-t} x_t. \quad (6.21)$$

If we now combine (6.19) with (6.21) we obtain

$$\sum_{t=l+1}^k a^{k-t} y_t > \sum_{t=l+1}^k a^{k-t} x_t. \quad (6.22)$$

Since  $a > 0$  this implies that there indeed exists an  $n$  with  $l < n \leq k$  such that  $y_n > x_n$ .

Next we show that an  $n^*$  with  $k < n^* \leq m$  and  $y_{n^*} < x_{n^*}$  exists. Similar to the derivation of (6.19) we obtain, by construction of  $k$ , that



$$C_m - a^{m-k} C_k \geq \sum_{t=k+1}^m a^{m-t} y_t \quad (6.23)$$

Also, similar to the derivation of (6.21) and by assumption on  $\mathbf{x}$  and construction of  $m$  we obtain

$$C_m - a^{m-k} C_k < \sum_{t=k+1}^m a^{m-t} x_t. \quad (6.24)$$

Combining (6.23) and (6.24) we obtain

$$\sum_{t=k+1}^m a^{m-t} x_t > \sum_{t=k+1}^m a^{m-t} y_t, \quad (6.25)$$

which implies that there indeed exists an  $n^*$  with  $k < n^* \leq m$  such that  $y_{n^*} < x_{n^*}$ .

The optimality conditions stated for  $rHVACS$  in Lemma 6.1 apply to  $\mathbf{y}$ . Therefore, we obtain that  $\frac{f_n^-(y_n)}{a^{T-n}} \leq \frac{f_{n^*}^+(y_{n^*})}{a^{T-n^*}}$ . Furthermore, by the convexity of  $f_n$  and  $f_{n^*}$  respectively it follows that  $\frac{f_n^+(x_n)}{a^{T-n}} \leq \frac{f_n^-(y_n)}{a^{T-n}}$  and  $\frac{f_{n^*}^+(y_{n^*})}{a^{T-n^*}} \leq \frac{f_{n^*}^-(x_{n^*})}{a^{T-n^*}}$ . Thus we obtain

$$\frac{f_n^+(x_n)}{a^{T-n}} \leq \frac{f_n^-(y_n)}{a^{T-n}} \leq \frac{f_{n^*}^+(y_{n^*})}{a^{T-n^*}} \leq \frac{f_{n^*}^-(x_{n^*})}{a^{T-n^*}}. \quad (6.26)$$

By construction of  $l$  and  $m$  it follows that if we increase  $x_n$  by  $\frac{\epsilon}{a^{T-n}}$  and decrease  $x_{n^*}$  by  $\frac{\epsilon}{a^{T-n^*}}$  for some small  $\epsilon > 0$  the solution remains feasible. Furthermore, from (6.26) it follows that doing so does not increase the objective value for sufficiently small  $\epsilon$ . On the other hand, doing so increases  $\sum_{t=1}^k a^{k-t} x_t$ . We can repeat the above process until  $\sum_{t=1}^k a^{k-t} x_t = C_k$ . This shows that  $\sum_{t=1}^k a^{k-t} y_t > C_k$  implies that  $\sum_{t=1}^k a^{k-t} x_t = C_k$ . The proof for the case that  $\sum_{t=1}^k a^{k-t} y_t < B_k$  and  $\sum_{t=1}^k a^{k-t} x_t > B_k$  is symmetric.  $\square$

The above lemma implies that, with minimal modifications, we can apply the same solution approach we used for Problem  $BC$  to Problem  $HVACS$ . These modifications entail updating  $k$  to reflect the definition as given in Lemma 6.2 and replacing the call to an algorithm solving instance of Problem  $EVC$  to a call to an algorithm solving instance of Problem  $rHVACS$ . Note that we argued that instance of  $rHVACS$  can be solved by an algorithm for Problem  $EVC$  after a linear substitution. For completeness sake we list the modified approach of Algorithm 5.1 in Algorithm 6.1, with  $optHVACS$  denoting a call to an algorithm that solves an instance of  $rHVACS$  to optimality.




---

**Algorithm 6.1** Recursive algorithm *optHVACS* for problem *HVACS*


---

```

1:  $\mathbf{x}_{1 \rightarrow T}$  = Function optBC( $\mathbf{f}_{1 \rightarrow T}$ ,  $\mathbf{x}_{1 \rightarrow T}^{max}$ ,  $\mathbf{B}_{1 \rightarrow T}$ ,  $\mathbf{C}_{1 \rightarrow T}$ )
2:  $\mathbf{y}_{1 \rightarrow T} :=$  optrHVACS( $\mathbf{f}_{1 \rightarrow T}$ ,  $\mathbf{u}_{1 \rightarrow T}$ ,  $\mathbf{C}_T$ )
3: if  $\mathbf{y}_{1 \rightarrow T}$  is feasible then
4:    $\mathbf{x}_{1 \rightarrow T} := \mathbf{y}_{1 \rightarrow T}$ 
5: else
6:    $k := \arg \max_t \left\{ a^{T-t} \left( \sum_{t'=1}^t a^{t-t'} y_{t'} - C_t \right), a^{T-t} \left( B_t - \sum_{t'=1}^t a^{t-t'} y_{t'} \right) \right\}$ 
7:   if  $\sum_{t=1}^k a^{k-t} y_t > C_k$  then
8:      $B_k = C_k$ 
9:      $B_t = B_t - C_k$  and  $C_t = C_t - C_k$  for  $t = k + 1, k + 2, \dots, T$ 
10:   else
11:      $C_k = B_k$ 
12:      $B_t = B_t - B_k$  and  $C_t = C_t - B_k$  for  $t = k + 1, k + 2, \dots, T$ 
13:   end if
14:    $\mathbf{x}_{1 \rightarrow k} =$  optHVACS( $\mathbf{f}_{1 \rightarrow k}$ ,  $\mathbf{x}_{1 \rightarrow k}^{max}$ ,  $\mathbf{B}_{1 \rightarrow k}$ ,  $\mathbf{C}_{1 \rightarrow k}$ )
15:    $\mathbf{x}_{k+1 \rightarrow T} =$  optHVACS( $\mathbf{f}_{k+1 \rightarrow T}$ ,  $\mathbf{x}_{k+1 \rightarrow T}^{max}$ ,  $\mathbf{B}_{k+1 \rightarrow T}$ ,  $\mathbf{C}_{k+1 \rightarrow T}$ )
16: end if
17: Return  $\mathbf{x}_{1 \rightarrow T}$ 

```

---

### 6.4.3 DISCRETE VARIANT OF HVACS

Similar to discrete variants of the EV and battery problems we formulate a discrete variant of Problem *HVACS*. The discussion and results below follow to a large extent from techniques and approaches already presented in the previous chapters. We add them for completeness sake.

Again, the true discrete variant is  $\mathcal{NP}$ -hard, as it is an extended version of Problem *dEVC*, for which we proved  $\mathcal{NP}$ -hardness in Chapter 4. Therefore, we directly state the variant using convex combinations.

**Problem 6.3** (*dHVACS*).

$$\min_{\mathbf{y}} \sum_{t=1}^T F_t(x_t) = \sum_{t=1}^T \sum_{j=0}^{m_t} y_t^j f_t(z_t^j), \quad (6.27a)$$

$$\text{s.t. } B_t \leq \sum_{t'=1}^t a^{t-t'} x_{t'} \leq C_t \quad t = 1, 2, \dots, T, \quad (6.27b)$$

$$x_t = \sum_{j=0}^{m_t} y_t^j z_t^j \quad t = 1, 2, \dots, T, \quad (6.27c)$$

$$\sum_{j=0}^{m_t} y_t^j = 1 \quad t = 1, 2, \dots, T, \quad (6.27d)$$

$$y_t^j \geq 0 \quad t = 1, 2, \dots, T; j = 0, 1, \dots, m_t. \quad (6.27e)$$



Note that we use Constraints (6.27b) similarly to Constraints (6.10), which are equivalent to Constraints (6.6b) for Problem *HVACS* with the appropriate choice of  $B_t$  and  $C_t$  respectively. Furthermore, recall that  $z_t^j$  denote the feasible operation levels of the device for time interval  $t$  for  $j = 0, 1, \dots, m_t$ . Also,  $F_t$  can be rewritten as a piecewise linear function obtained from  $f_t$  by linearizing  $f_t$  between each of  $z_t^j$ 's for  $j = 0, 1, \dots, m_t$ . Note that we assume that the losses on the SoC of the buffer are independent of the order in which the different operation levels are used within a time interval.

To solve Problem *dHVACS* we combine previous results. Similar to instances of Problems *dBC* and *rdEVC*, we can rewrite an instance of Problem *dHVACS* to an instance of Problem *HVACS* with piecewise linear objective. To solve this instance of *HVACS* we can then apply the approach described above, i.e., we can solve this instance by applying Algorithm 6.1. The time complexity of this approach is  $O(TM)$ , where  $M$  denotes the total number of breakpoints of the piecewise linear objective functions.

As discussed in previous chapters, this complexity result relies on an algorithm using median find, which is rather inefficient for small to medium sized inputs. For practical applicability, we pursue a direct approach for instances of Problem *HVACS* with piecewise linear objective functions below. As Problem *HVACS* closely resembles problem *BC*, we modify the greedy approach we found for Problem *BC* with piecewise linear objective. One of the important differences is that now, to satisfy the lower cumulative bound  $B_t$  for interval  $t$ , increasing  $x_{t'}$  by  $\delta$  for intervals  $t' \leq t$  no longer has the same effect on the cumulative sum for every  $t'$ . In particular, if we increase  $x_{t'}$  by  $\delta$  for  $t' < t$ , then the cumulative sum for  $t$  is increased by  $a^{t-t'}\delta$ . Now, to determine which piece gives the lowest increase of the objective value per unit of increase of the cumulative sum for  $t$ , we need to discount the slopes of the pieces for time interval  $t' < t$  with a factor  $\frac{1}{a^{t-t'}}$ . Note that this is exactly the result we obtained in Lemma 6.1.

Similar to the approach given in Algorithm 5.2, we iteratively want to ensure that the lower bounds  $B_1, B_2, \dots, B_T$  are met. To ensure that  $\sum_{t'=1}^k a^{k-t'} x_{t'} = B_k$  in iteration  $k$ , we only increase the values of  $x_1, x_2, \dots, x_k$  until the equality is met. Note that if we increase  $x_t$  with  $1 \leq t \leq k$  by  $\delta$  in iteration  $k$ , then this cumulative sum  $\sum_{t'=1}^k a^{k-t'} x_{t'}$  increases by  $a^{k-t}\delta$ . For the choice of the index  $t$  for which we increase  $x_t$  we use the piece with the lowest slope (after discounting as described above). Thus, for iteration  $k$ , we track which of the currently active pieces for time intervals  $1, 2, \dots, k$  has the lowest discounted slope (for the definition of an active piece see Subsection 4.5.3). In the algorithm we track this by using an ordered set  $S$ . The set  $S$  contains the discounted slopes of the currently active pieces for the intervals  $1, 2, \dots, k$ . We note that in some cases all pieces for a time interval may have already been used and hence no piece is active for this time interval. Furthermore, if for some  $t' \geq t$  the upper cumulative bound  $C_{t'}$  is met, the pieces for time interval  $t$  can no longer be used.



When moving from iteration  $k - 1$  to  $k$  the pieces in  $S$  need to be discounted by an additional factor  $\frac{1}{a}$ . Since all the pieces are discounted by this same factor, the ordering of  $S$  does not change. Furthermore, we need to add the slope of the first piece of  $f_k$  to  $S$ . Note that we do not need to discount the slope of this first piece of  $f_k$ . However, if in iteration  $k$  we completely use a piece of  $f_t$  with  $t < k$  then we need to add the next piece of  $f_t$  to  $S$ . We note that this piece does need to be discounted first by a factor  $a^{k-t}$  before it can be added to the (sorted) set  $S$ .

The above procedure takes care of the lower bounds  $B_t$ . However, we also need to ensure that the upper cumulative bounds, given by the values  $C_1, C_2, \dots, C_T$ , are respected. In Algorithm 5.2 we ensured these bounds are respected through a book-keeping variable  $V_t$ , which tracks how much we can increase  $x_t$  without violating the cumulative upper bounds that include  $x_t$ . In the following we describe how we ensure that these bounds are respected for Problem *HVACS* with piecewise linear objective. When increasing  $x_t$  for some  $t$  by  $\delta$ , the cumulative sum for interval  $t^*$ , with  $t^* > t$  is increased by  $a^{t^*-t}\delta$ . Hence, the upper cumulative bound  $C_{t^*}$  is respected if

$$\delta \leq \frac{1}{a^{t^*-t}} \left( C_{t^*} - \sum_{t'=1}^{t^*} a^{t^*-t'} x_{t'} \right). \quad (6.28)$$

Thus, to ensure that all cumulative upper bounds are met, we need to ensure that the increase  $\delta$  of  $x_{t^*}$  is bounded by

$$\delta \leq \min_{t \geq t^*} \left\{ \frac{1}{a^{t^*-t}} \left( C_{t^*} - \sum_{t'=1}^{t^*} a^{t^*-t'} x_{t'} \right) \right\}. \quad (6.29)$$

Similar to the approach in Algorithm 5.2 we define

$$V_t = \min_{t \geq t^*} \left\{ \frac{1}{a^{t^*-t}} \left( C_{t^*} - \sum_{t'=1}^{t^*} a^{t^*-t'} x_{t'} \right) \right\}. \quad (6.30)$$

Note that we have

$$V_{t-1} = \min \left\{ C_{t-1} - \sum_{t'=1}^{t-1} a^{t-1-t'} x_{t'}, \frac{1}{a} V_t \right\}, \quad (6.31)$$

for  $t < T$  and

$$V_T = C_T - \sum_{t'=1}^T a^{T-t'} x_{t'}. \quad (6.32)$$

Hence we can iteratively calculate the value of  $V_t$ , starting with interval  $T$  and working backwards, if we know the cumulative sum for each interval. To this end we define variables  $\bar{C}_t$  to track the value of  $C_t - \sum_{t'=1}^t a^{t-t'} x_{t'}$  during the approach and use this to update  $V_t$  after we increase  $x_t$  for some  $t$ . The approach is summarized in Algorithm 6.2.




---

**Algorithm 6.2** Greedy approach *pwlHVACS* for *HVACS* with piecewise linear convex objective

---

```

1:  $\mathbf{x}_{1 \rightarrow T} = \text{Function } pwlHVACS(\mathbf{f}_{1 \rightarrow T}, Z, \mathbf{x}_{1 \rightarrow T}^{max}, \mathbf{B}_{1 \rightarrow T}, \mathbf{C}_{1 \rightarrow T})$ 
2:  $S := \emptyset$ , and  $\overline{C}_t := C_t$ , and  $x_t := 0$  for all  $t$ 
3:  $V_T := \overline{C}_T$ 
4: for  $l = T - 1, T - 2, \dots, 1$  do
5:    $V_l := \min\{\overline{C}_l, \frac{1}{a} V_{l+1}\}$ 
6: end for
7: for  $k = 1, 2, \dots, T$  do
8:   Multiply the elements of  $S$  by  $\frac{1}{a}$ 
9:   Insert  $s_k^1$  into  $S$  such that  $S$  remains ordered non-decreasingly
10:   $R := B_k - \sum_{t'=1}^k a^{k-t'} x_k$ 
11:  while  $R > 0$  do
12:    Take  $t$  and  $j$  such that  $s_t^j$  is the first piece from  $S$ 
13:     $\delta := \min\{R, V_t, z_t^j - z_t^{j-1}\}$ 
14:     $x_t = x_t + \delta$ 
15:     $R = R - a^{k-t} \delta$ 
16:    for  $l = t, t + 1, \dots, T$  do
17:       $\overline{C}_l = \overline{C}_l - a^{k-t} \delta$ 
18:    end for
19:     $V_T = \overline{C}_T$ 
20:    for  $l = T - 1, T - 2, \dots, 1$  do
21:       $V_l = \min\{\overline{C}_l, \frac{1}{a} V_{l+1}\}$ 
22:    end for
23:    for  $j', t'$  such that  $s_{t'}^{j'} \in S$  and  $V_{t'} = 0$  do
24:       $S = S \setminus \{s_{t'}^{j'}\}$ 
25:    end for
26:    if  $\delta = z_t^j - z_t^{j-1}$  then
27:       $S = S \setminus \{s_t^j\}$ 
28:      if  $j < m_t$  and  $V_t > 0$  then
29:        Insert  $a^{k-t} s_t^{j+1}$  into the ordered vector  $S$ 
30:      end if
31:    else if  $V_t > 0$  then
32:       $z_t^{j-1} = z_t^{j-1} + \delta$ 
33:    end if
34:  end while
35: end for
36: Return  $\mathbf{x}_{1 \rightarrow T}$ 

```

---



**Lemma 6.3.** *Algorithm 6.2 solves an instance of Problem HVACS with piecewise linear objective functions to optimality in time  $O(TM)$ , with  $M$  the total number of breakpoints of the objective functions.*

*Proof.* The feasibility of the algorithm follows from the fact that pieces are added in order to the set  $S$  and the lower and upper cumulative bounds are respected, per the discussion above.

To prove optimality of the algorithm, consider an instance of HVACS with piecewise linear convex objective functions. Let  $\mathbf{y}$  be an optimal solution, and let  $\mathbf{x}$  be the solution produced by Algorithm 6.2. Furthermore, let  $t$  be the smallest index for which  $x_t \neq y_t$ . We consider two cases.

*Case 1*  $x_t > y_t$ : Let  $t^* > t$  be the smallest index for which  $x_{t^*} < y_{t^*}$ . This index must exist since  $\sum_{t'=1}^T a^{T-t'} x_{t'} = B_T = \sum_{t'=1}^T a^{T-t'} y_{t'}$ . Furthermore, let  $s_t$  and  $s_{t^*}$  be the slopes of the pieces on which  $x_t$  and  $x_{t^*}$  lie respectively. In case  $x_t \in Z_t$ , we pick the piece with  $x_t$  as endpoint and in case  $x_{t^*} \in Z_{t^*}$  we pick the piece with  $x_{t^*}$  as begin point. By the fact that the right and left derivatives of  $f_t$  and  $f_{t^*}$  are non-decreasing it follows that  $s_t \geq f_t^+(y_t)$  and  $f_{t^*}^-(y_{t^*}) \geq s_{t^*}$ . Also, by the optimality of  $\mathbf{y}$  it follows from Lemma 6.1 that  $\frac{f_t^+(y_t)}{a^{T-t}} \geq \frac{f_{t^*}^-(y_{t^*})}{a^{T-t^*}}$ . Note that, since both  $\mathbf{x}$  and  $\mathbf{y}$  are feasible, we can find a  $\delta > 0$  such that we can decrease  $x_t$  by  $a^{T-t}\delta$  and increase  $x_{t^*}$  by  $a^{T-t^*}\delta$ , until either  $x_t$  is equal to  $y_t$  or  $x_{t^*}$  is equal to  $y_{t^*}$ , without violating feasibility. Finally, note that doing so does not increase the objective value.

*Case 2*  $x_t < y_t$ : This case can be treated symmetrical to the previous case.

The above process can be repeated until  $\mathbf{x} = \mathbf{y}$  without increasing the objective value. This shows that  $\mathbf{x}$  is indeed optimal.

Finally, to show that the time complexity is  $O(TM)$ , we note that the while loop in Algorithm 6.2 is executed at most  $O(M)$  times, as was also the case for Algorithm 5.2. As the steps in Algorithm 6.2 still can be executed in time  $O(T)$ , the complexity remains  $O(TM)$ .  $\square$

## 6.5 OPERATIONAL CONTROL OF HVACs

In Section 6.3 we discussed a thermal model of houses in Austin, Texas and in Section 6.4 we showed how we can schedule the HVAC systems using this thermal model. However, as mentioned, the thermal model suffers from inaccuracies resulting, e.g., from human behaviour. Thus, the resulting schedules are not guaranteed to satisfy the user thermal comfort constraints at all times.

In this section we discuss how we can adapt the schedules in the operational control step of the profile steering DEM approach (see Section 2.3) to ensure that these constraints are met. Furthermore, in Subsection 6.5.2, we discuss how most residential HVAC systems are currently operated and how we can simulate this behaviour in a simulation using discrete time intervals. We will use this simulation of the current



behaviour of HVAC systems as a reference case for the simulation study we perform in the next section.

### 6.5.1 PREDICTION ERRORS IN PROFILE STEERING

To be able to deal with prediction errors made by our thermal model, we assume that we can accurately track the indoor temperature. Thus, the system can react whenever the measured temperature deviates from the predicted value. The schedule made by the profile steering approach specifies energy consumption per time interval of, e.g., fifteen minutes. The thermal model predicts the resulting temperature values for these consumption values. However, it is possible that a prediction error causes the real temperature to deviate from this predicted value and that the system violates the thermal bounds set by the user for the next time interval. This can happen when, for example, (much) more energy is needed for cooling than originally predicted. In such cases we assume that the system has to switch its state to prevent the loss of (too much) user comfort. Consequently the system uses a different amount of energy than scheduled in the current time interval. This can occur quite frequently if the system operates near its thermal limits, which usually happens if the full flexibility of the system is exploited. To prevent this issue we suggest that the thermal limits used in the scheduling phase of the profile steering approach are taken to be stricter than the actual limits set by the user, i.e., the maximum allowed deviation from the set point used by our model is taken to be lower than specified by the user.

Above we discussed how we can ensure that not too much user comfort is lost in the current time interval due to prediction errors for the next time interval(s). However, such prediction errors also impact the available flexibility for these next time intervals. Below, we discuss how we can adapt the schedule to account for this change in flexibility. In Section 6.3 we mentioned that our model cannot accurately predict human behaviour and that we assume that it is not realistic to expect that such accurate predictions exist. To ensure that the system nevertheless behaves correctly we propose to use an MPC like approach. By this we mean that new schedules are made at every time interval, e.g., every fifteen minutes. For such a schedule, only the first time interval is used to control the HVAC system. For the profile steering approach this implies that a new schedule has to be made for every device, since the available flexibility constantly changes between time intervals. These changes in available flexibility impact the best use of flexibility of other devices, hence a full profile steering scheduling phase has to be executed for every time interval.

If the execution of a scheduling phase of profile steering is fast enough it is possible to implement the above procedure. This is however a costly operation from both a computational and communication perspective. We note that particularly the computational issue plays an important role when performing simulation studies, which are commonly done for the analysis of EM approaches. As an alternative, we propose to use the hierarchy of the profile steering approach to reduce the computational complexity. Instead of rescheduling the flexibility of all devices





in the system for every time interval for which prediction errors are made, we propose that this is only done, e.g., for predetermined time intervals or whenever the prediction errors are judged too large. In other steps where prediction errors occur we propose that these errors are only taken into consideration locally through the rescheduling of the energy use of the device (in this case the HVAC system). For the device, a schedule can be made, using the updated parameters, that best fits the steering signal the device received.

### 6.5.2 BASE CASE

Above we argued how we can integrate HVAC systems into our profile steering approach using the thermal model developed earlier in this chapter. To investigate the potential of such an EM approach, we need a reference case to compare our results with. While we have a dataset available from Austin, Texas detailing HVAC energy consumption values and indoor temperatures, this dataset does not give us any information on the most vital parameter; the available flexibility. In order to determine this flexibility using the developed thermal model we require both a temperature set point  $\bar{T}_t$  and an allowed deviation from this set point  $D_t$ .

To tackle this problem we implemented a basic form of standard control for HVAC systems. Most residential HVAC systems are equipped with a thermostat having a deadband controller [31, 45]. The thermostat turns the HVAC system on to maximum power when the temperature reaches the upper limit and turns it off once it reaches the lower limit. However, when implementing such a controller in our system using fifteen minute time intervals, the granularity of the intervals is too coarse to capture the dynamics of the system if decisions can only be made at the beginning of every time interval. This is especially the case if the room temperature exceeds the thermal limits shortly after the start of a time interval.

To alleviate this issue we use linear interpolation to determine when the system switches on during a time interval. At the beginning of such a time interval  $t$  we know if the HVAC is running or not. Furthermore, we calculate the temperature  $T_{t+1}$  at the start of the next time interval using our thermal model (6.2). If this value exceeds the thermal comfort bounds specified we use linear interpolation to determine when the HVAC system switches on or off. For example if the system is turned off at the start of time interval  $t$  and  $T_{t+1} > \bar{T}_{t+1} + D_{t+1}$  we estimate the fraction  $\tau$  of the time interval the system stays off by:

$$T_t + \tau(T_{t+1} - T_t) = \bar{T}_{t+1} + D_{t+1}. \quad (6.33)$$

The total energy consumption of the system for time interval  $t$  can now be calculated by taking  $(1 - \tau)$  times the consumption of energy if the system would run for a full interval. This value for the energy consumption can subsequently be used in (6.2) to calculate the resulting temperature at the start of the next time interval. The procedure for determining when the system switches off if the lower thermal limit is reached is similar.



## 6.6 SIMULATION STUDY

In this section we discuss the results of a simulation study to investigate the potential of HVAC systems using profile steering. To do this, we compare the results from a simulation using our profile steering approach to those of a simulation with deadband control. Furthermore, we investigate the potential of profile steering if the errors made by our thermal model are assumed to be known a priori, i.e., we consider profile steering where we assume the thermal model gives perfect predictions. The latter gives us an upper bound on the increase in effectiveness of our approach with a ‘perfect’ thermal model.

Within our simulation study we have to incorporate errors which are made by our thermal model. In a real-life implementation of an EM approach, we assume that these errors can easily be observed from local measurements, e.g., from a thermostat. However, we do not have access to such data in a simulation. To still be able to incorporate the errors made by our thermal model, we used the Pecan Street Inc. dataset, described above. We split this data set in two parts, a training and a validation part. We used the latter part to verify the obtained model. Furthermore, we use it also to obtain errors made by our thermal model.

We construct the errors used in our simulation study as follows. To validate our model, we compared measured temperatures in the validation part of the dataset with predictions made by our model. The differences between measured and predicted values give us a set of prediction errors. We use these differences to construct a sample set for the prediction error made by our model. In every time step of the simulation we sample an error from the sample set and add it to the temperature calculated by our model. We do this for both the profile steering approach and the base case with deadband control. Note that it is possible that the resulting temperatures in either scenario violate the thermal limits, particularly if we add a large disturbance term. In such cases the disturbance is too large to be accounted for by the used system and we allow the temperature to violate the thermal constraints. We believe such situations mostly occur in cases with a large influence due to human behaviour, for example a door or window being left open for a long period of time, for which it is not feasible to expect the system to operate within its limits.

For the simulation study we consider the seven houses with reasonable thermal models, described in Section 6.3. For these houses we also have the (uncontrollable) load of the other devices besides the HVAC system available. For the simulation we chose a set point of  $23^{\circ}\text{C}$  with a maximum allowed deviation of  $0.5^{\circ}\text{C}$ . We apply the profile steering approach to the group of these seven houses for a full week with the goal to flatten the load profile. Due to the prediction errors discussed in Section 6.5 we used  $D_t = 0.3$  for the profile steering approach. Furthermore, for the profile steering approach we plan twelve hours ahead using fifteen minute time intervals. We adjust the local schedule after every time step to account for prediction errors in the thermal model. Finally, we use a full scheduling phase of the approach every six hours, i.e., after half the scheduling horizon has elapsed.



The results are given in Figure 6.3 (a) and Table 6.2. Note that in the figure we give load duration curves, i.e., we sort the load values non-decreasingly, to accentuate the difference between the approaches. The table is used to indicate the minimum and maximum values of power attained over the scheduling horizon. The results show a significant improvement in the load profile of the considered houses when profile steering is applied. Furthermore, the results show that a perfect thermal model does not significantly increase the ability of profile steering to flatten the load profile during peak hours. However, the curve stays above zero, indicating that all locally produced energy is consumed locally when perfect predictions are used.

To study the effect of increasing the number of houses involved we studied two additional cases, one with fourteen houses and another with twenty-one houses. For these cases we replicated the thermal models of the seven houses, using the uncontrollable load data of different weeks. In Figures 6.3 (b) and (c) and Table 6.2 the results for these cases are shown. They show that our methodology is able to significantly flatten the aggregated profile of the houses. However, in both cases the peak reduction can be improved upon even more when perfect predictions are available. We believe that the discrepancy between these cases and the case with seven houses is caused by the fact that the major peaks in the case of seven houses are large but short, allowing the HVAC system to compensate, even in the case of significant errors in the thermal models. These results indicate that improvements in the thermal model combined with good human behaviour forecasting techniques can further increase the potential of DSM with residential HVAC systems.

Finally, we studied the effect of varying the maximum allowed deviation from the set point temperature in the seven houses case. The results for this case and the used deviations  $D_t$  are given in Figure 6.4 and Table 6.2. We note that for deadband control we only consider the scenario where we allow a deviation of  $0.5^\circ\text{C}$  as the other scenarios (with higher allowed deviations) give nearly identical results. The results show that increasing the allowed deviation improves the ability for the system to level the load duration curve. However, further increasing the allowed deviations gives limited improvement. We believe that this is because the capability of reducing the HVAC system energy consumption during times of high energy consumption of the other loads is already rather high for the studied seven houses with an allowed deviation from the set point of  $0.5^\circ\text{C}$ .

## 6.7 CONCLUSION

In this chapter we considered a thermal model to determine the flexibility available from HVAC systems to be used by EM approaches. We used a linear thermal model to keep the complexity of the resulting optimization problem for the device low. We fit the model to data from the Pecan Street inc. dataset, gathered in Austin, Texas. The prediction errors of the model are mainly due to human behaviour. As such behaviour is hard to predict in general, we believe that the model can be used for

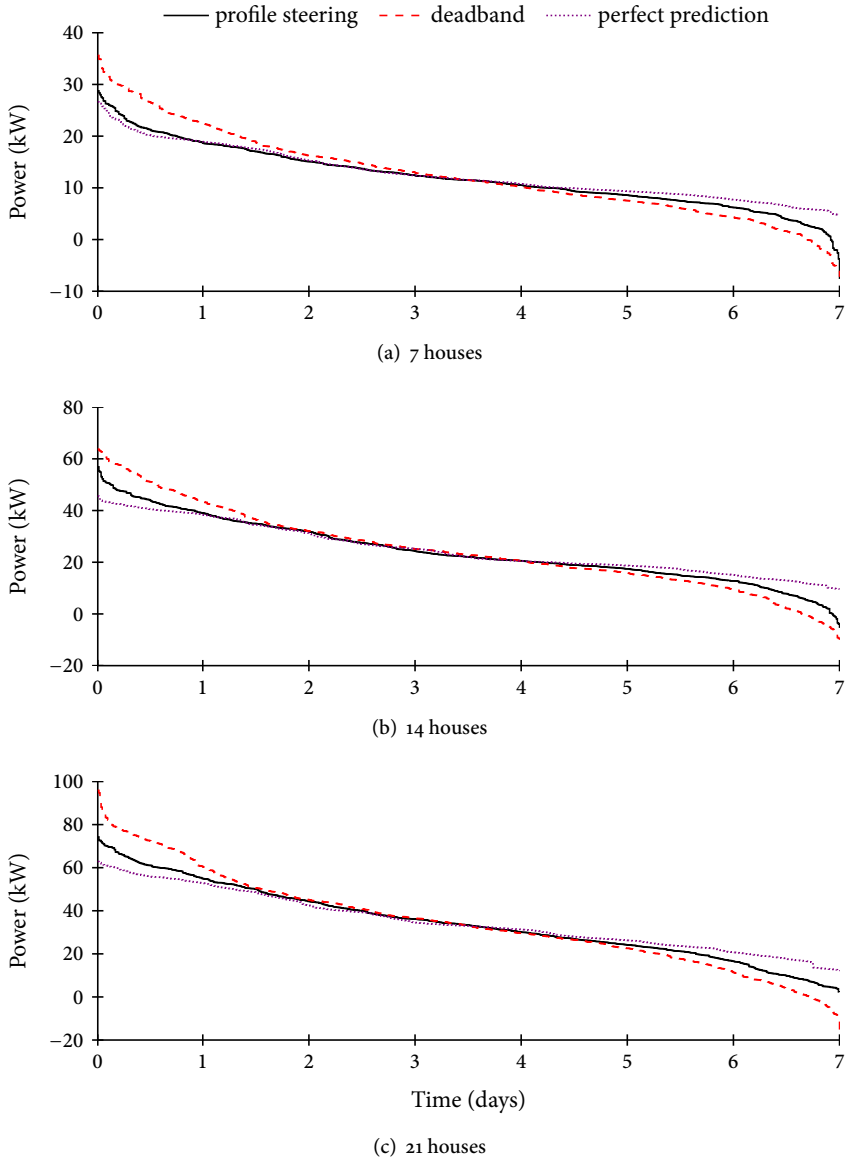


Figure 6.3: The load duration curve for (a) 7, (b) 14, and (c) 21 houses for which we control the HVAC system energy consumption to flatten the overall profile.

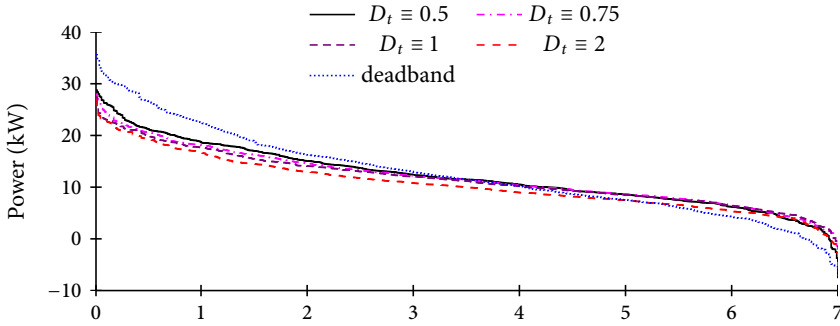


Figure 6.4: Load duration curve when controlling the HVAC system of seven houses with different allowed deviations from the set point.

Table 6.2: The minimum and maximum power values for the simulated cases in Figures 6.3 and 6.4.

# houses	case	min power	max power
7	profile steering	-7.6	28.8
7	deadband	-7.2	35.6
7	perfect predictions	4.8	26.6
14	profile steering	-5.4	56.9
14	deadband	-10.0	63.8
14	perfect predictions	9.5	45.6
21	profile steering	2.3	74.2
21	deadband	-17.2	95.9
21	perfect predictions	12.0	62.8
7	deadband	-7.2	35.6
7	$D_t \equiv 0.5$	-7.6	28.8
7	$D_t \equiv 0.75$	-5.7	27.9
7	$D_t \equiv 1$	-1.6	27.3
7	$D_t \equiv 2$	-2.7	27.0

steering HVAC systems in profile steering. Note that the model used in this chapter is equivalent to a resistance capacitor (RC) model with a single resistance and a single capacitor. These RC models are often used in the literature (see, e.g., [13, 138]). However, for many buildings a higher order model, i.e., a model equivalent to a circuit with more resistances and/or capacitors, is required to capture the full dynamics of the systems. The integration of such higher order models into our profile steering approach is outside the scope of this thesis and is therefore left for future work.

We note that in any case it is important to properly deal with prediction errors made on the input parameters of our profile steering approach. In Subsection 6.5.1 we gave a method to deal with such prediction errors. We believe that further research



on methods to deal with these prediction errors can increase the capabilities of our approach. Furthermore, we specifically considered the case of HVAC systems. However, the model should also be applicable to other systems where losses based on the SoC of the storage system play an important role, such as heating systems.

Next, we considered the resulting device level optimization problem called *HVACS*. This problem has a lot of similarities with the problems considered in previous chapters. We first formulated a relaxed problem wherein the cumulative bounds for all but the last time interval are dropped. We showed that this problem, after a transformation, is equivalent to Problem *EVC*, as studied in Chapter 4. Furthermore, we showed that a similar result to that found for Problem *BC* (Chapter 5) holds for Problem *HVACS*. This result allowed us to apply, after minor modifications, the divide and conquer algorithm obtained for *BC* to *HVACS*. As the problems previously studied in this thesis are also applicable to other fields outside of EM, we believe the problems studied in this chapter could also be of value to other fields. However, investigating if the results of this chapter can be applied elsewhere is outside the scope of this thesis.

As mentioned, the thermal model developed in this chapter suffers from significant prediction errors. To ensure that our profile steering approach produces feasible schedules when considering HVAC systems, we need to ensure the operational control step is robust against these errors. We argued that we can ensure that the temperature is generally kept between bounds within the discrete time steps by setting stricter bounds on the temperature than allowed by the user. Furthermore, we deal with prediction errors by rescheduling the energy consumption locally at the start of every time interval. Next to the local rescheduling, we also reschedule all HVAC systems together, using a full profile steering scheduling step, at pre-determined intervals, e.g., every few hours. Further improvements could potentially be obtained if a rescheduling step involving all HVAC systems together is not done at pre-determined time intervals, but whenever the schedules of the individual HVAC systems deviate too much from their original plan. We believe an investigation of how the system can track the deviations by devices from their original schedules, and when and how to respond to such deviations can increase the robustness of the profile steering approach.

We applied the approach developed in this chapter in a simulation study that compares profile steering using HVAC systems with deadband control. For this simulation study we used the same houses we used to obtain a thermal model in Section 6.3. The results show that profile steering can significantly flatten the load profile even if only seven houses are present. The results do not significantly improve with perfect predictions. However, as an improvement is observed for the cases with fourteen and twenty-one houses when perfect predictions are used, we believe this to be an artefact of the considered scenario. A further investigation into this effect can result in a better understanding of the dynamics of the presented profile steering approach in the presence of prediction errors.

Simulations with a higher allowed deviation from the set point, at least for seven



houses, indicate that the improvements in flattening of the load profile suffer from diminishing returns. From this we conclude that it is likely that the flexibility provided by HVAC systems scales sub-linearly with the increase in allowed temperature deviations. An investigation if this is always the case can provide useful results in determining where we can (easily) obtain the most flexibility in the current and future grid.

Summarizing, the model developed in this chapter allows us to incorporate losses that depend on the state of charge of the system into device level models in profile steering. This is of particular importance in, for example, heating and cooling applications. The devices used in these applications often have another important constraint associated with minimum run and off times. Such constraints, while important, are out of the scope of this work. Furthermore, the results we obtained when simulating fourteen and twenty-one houses indicate that improvements can be made by improving the thermal model. This indicates that further research into this area can potentially lead to significant improvements in our profile steering approach. We believe that this is a non-trivial problem, as a large influence on the energy consumption in the residential sector comes from human behaviour, which is hard to model.








# 7



## TRANSFORMER AGEING AND EM

*ABSTRACT – One of the important aspects of the costs of our (future) energy supply chain are the costs of the infrastructure. In the electricity grid these costs are to a large extent maintenance and investment costs. As one of the key elements in the energy transition is a new energy management approach, we study the impact such an energy management approach can have on the ageing of power transformers in this chapter. For this study we use a transformer model developed by the IEEE, and we derive charging schedules for electric vehicles that minimize the ageing of transformers using this model. We compare the obtained results to several other charging strategies. The obtained results indicate that flattening the load profile of a neighbourhood of houses, which can be achieved through, e.g., profile steering, also minimizes ageing of the transformer supplying this neighbourhood.*



This chapter is based on [TvdK:6].




Figure 7.1: The investments made by Enexis in their gas and electricity distribution grids over the period 2011-2015. Data taken from [44].




Figure 7.2: The operational costs made by Enexis for their distribution grids over the period 2011-2015. Data taken from [44].




Figure 7.3: The expected asset value for the DSO PPL in the UK over the period 2015-2020. Note that the growth is constant, implying a near constant and significant investment requirement. Data taken from [111].

## 7.1 INTRODUCTION

In Chapter 3 we introduced our profile steering energy management (*EM*) approach. We argued that this approach should assist in keeping the total costs of our energy supply chain reasonable while we transition towards a system based on clean, renewable resources. An important part of the total costs is the cost of the infrastructure. For electricity this part entails the cost of maintenance and of investment in the electricity grid. With significant penetration levels of (renewable) generation on the customer side, the stress put on the grid, and in particular the distribution grid, increases significantly [102]. This implies that the associated maintenance and investment costs increase to levels much above the current values. As an example, a total investment requirement of 49 billion euros up to 2032 is expected for the German electricity distribution grid [19].

To get a better insight in the magnitude of the costs, we give the maintenance and investment costs for the Dutch distribution system operator (*DSO*) Enexis in Figures 7.1 and 7.2. Enexis serves about 31% of the Dutch customers and is the third largest DSO in the Netherlands. The costs have been rather stable over the last few years. Furthermore, the costs are expected to remain stable in coming years (see Figure 7.3 for a projection in the UK with a stable growth of asset value). However, these expectations are based on the assumption that the transition towards an energy supply chain based on renewables does not incur heavy reinforcement costs. To ensure this assumption holds EM approaches are generally expected to assist in the better management of flexibility in the (distribution part of the) future smart grid [120, 142]. Also, the maintenance and investment costs are relatively high, meaning that even small percentile savings imply large reductions in absolute costs. In particular, novel EM approaches can unlock such reductions in costs.

In most studies on EM the grid is implicitly assumed to benefit from the applied approaches, while in some cases conflicts of interest can exist between the different stakeholders [104]. Most studies on EM use the assumption that increasing local use of locally produced energy and flattening out the demand profile decreases stress on grid assets (see, e.g., [39, 122]). In this chapter we aim to investigate what the best use of flexibility provided by (residential) customers actually is from a grid perspective.

To investigate the impact of EM on the grid, we study the ageing of distribution transformers, as these are among the most expensive assets in the distribution grids [122]. Note that also stress on (underground) cables is an important issue, as their replacement is costly. However, if we reduce stress on the transformer we generally also reduce stress on the cables. To this end we study a model of the ageing of transformers and use this model to determine what the best schedules for the charging of electric vehicles (*EVs*) are to minimize this ageing. More specifically, we use the transformer ageing model provided by the IEEE in their C57.91-2011 standard [4]. This model was originally designed to assist system operators in determining the size of transformers before they are installed in (new) grids. For





this the models relates the use of the transformer, i.e., the energy flowing through the transformer, to the ageing of the asset. As this is also the relation we require to determine the impact of EM approaches on the asset's ageing, we believe the IEEE model to be suitable for our application.

The remainder of this chapter is outlined as follows. In the next section we briefly discuss some related work to provide some background. Then, in Section 7.3 we introduce the used transformer ageing model. As the full model is highly complex, we introduce a simplifying assumption that ensures that we can formulate an optimization problem in Section 7.4 that is efficiently solvable. This optimization problem determines the charging of EVs such that minimal ageing occurs in the transformer. Afterwards, in Section 7.5, we validate that our simplifying assumption does not alter the results significantly. Furthermore, we compare the minimization of transformer ageing to other scenarios, in particular to our profile steering approach, in this section. Finally, we finish with some conclusions in Section 7.6.

## 7.2 RELATED WORK

In this section we discuss some related work on the impact of EM on transformer ageing. As mentioned above, few studies explicitly incorporate the ageing of grid assets in their analysis. Most of these studies use the the model from the IEEE standard C57.91-2011 [4], which we also use in this chapter. Below we discuss a few publications that use this model as background to the study we perform in this chapter.

Jargstrof *et al.* [78] study the effects of EM on transformer ageing in a Belgian case. They use the IEEE model as a basis and linearize part of the model to obtain a mixed-integer quadratic program, which they solve using a commercial solver. The flexible devices they consider are smart white goods (dish washer, washing machine, and tumble dryer) and electric heat generation systems combined with hot water buffers. They conclude that the increase in transformer ageing is more pronounced when the transformer is operated near or at its limits without EM applied. We note that in this case study the available flexibility is rather limited.

In [70], Humayun *et al.* study a Finnish scenario, which focusses on load shifting offered by smart white goods, HVAC systems, and refrigerators and freezers. They steadily increase the load on the system to study the potential of EM to reduce asset ageing when the asset is under significant stress. In particular they aim to keep the temperature of the winding inside the transformer below a specific threshold, as this threshold is given by manufacturers as the maximum operating temperature under normal conditions. Similar to Jargstrof *et al.* they simplify the IEEE model to obtain a mixed-integer quadratic program. The authors conclude that the overall load on transformers can be significantly increased without sacrificing transformer life if an EM approach is used. They observed that the capability to increase transformer life proportionally increased with the amount of flexibility available on the demand side.



Moghe *et al.* [93] consider the impact of electric driving on transformer ageing. They use a Monte Carlo simulation with data from two major cities in the US (Phoenix, Arizona and Seattle, Washington). They compare uncontrolled charging (i.e., at full power upon arrival) and controlled charging of the EVs. Their control strategy only allows a vehicle to charge when the total load on the transformer does not exceed a predetermined threshold. Within each of the simulated (random) scenarios they compare the resulting transformer lifetime in both cases with the lifetime when no EVs are present. They base their results on a model using differential equations obtained from the IEEE standard. Their results indicate that, while the introduction of (large amounts of) EVs has adverse effects on transformer ageing, this effect can be significantly mitigated through smart charging. We note that their controlled charging strategy can be improved upon, potentially leading to a further decrease in asset ageing.

### 7.3 TRANSFORMER AGEING MODEL

In this section we study the ageing model of transformers given by IEEE in standard C57.91-2011 [4]. First, we give the details of the model. As this full model is complex it is difficult to determine the best use of flexibility offered by distributed energy resources (*DERs*) with respect to ageing. To arrive at a computationally tractable problem we simplify the model. This makes the resulting optimization problem that schedules the use of flexibility to minimize transformer ageing convex. In turn, this allows us to use, e.g., a convex solver to calculate the optimal use of flexibility for minimizing transformer ageing.

#### 7.3.1 THERMAL AGEING MODEL

The full model from the IEEE standard is based on the relation between the temperature of the transformer and its degradation. In turn, the temperature of the transformer is related to the ambient temperature and the load on the transformer. The ageing is determined through the calculation of ageing factors for a set of time intervals. These ageing factors indicate how much the transformer aged compared to normal operation (which is usually at 110°C). For a given time interval  $t$ , the ageing factor  $F_t^{AA}$  can be calculated using:

$$F_t^{AA} = e^{\frac{15000}{383}} e^{-\frac{15000}{\Theta_t^H + 273}}, \quad (7.1)$$

where  $\Theta_t^H$  is the so called hottest spot temperature in °C. This is the temperature of the hottest spot on the transformer windings, i.e., the point where the transformer is assumed to degrade the quickest.

As an example, if the ageing factor calculated over a day is 0.5, then the expected remaining lifetime, under normal operation, of the transformer is reduced by half a day instead of a full day. On the other hand, if the ageing factor is 3, then the expected remaining lifetime, under normal operation, of the transformer is reduced



by three days instead. In Figure 7.4 two scenarios of the loading of a transformer are plotted. The top plots give the load on the transformer (as a fraction of the load at which the transformer was rated) together with the hottest spot temperature (we come back to this later). The bottom plots give the ageing factor calculated for each of the hours as well as the cumulative ageing factor. This cumulative ageing factor is around 25 for the mild overloading case, indicating that the transformer degrades about 25 times faster than while operating under a constant load at which it was rated. On the other hand, the cumulative factor is about 450 for the severe overloading case, indicating that such a scenario degrades the transformer very fast (if the transformer is expected to last 20 years under normal operating conditions it only lasts about 16 days under severe overloading conditions).

To calculate the ageing for longer periods of time, with multiple time intervals involved, the equivalent ageing factor  $F^{EQA}$  is used. This factor is given by:

$$F^{EQA} = \frac{\sum_t F_t^{AA} \delta_t}{\sum_t \delta_t}, \quad (7.2)$$

where  $\delta_t$  is the length of time interval  $t$ . Note that, in case all time intervals are of equal length,  $F^{EQA}$  is equivalent to the average of the ageing factors of the considered time intervals. Alternatively, if the time intervals are of equal length, one can also use the sum of the individual ageing factors (as done in Figure 7.4)

The hottest spot temperature  $\Theta_t^H$  is calculated as:

$$\Theta_t^H = O_t + \Delta\Theta_t^{TO} + \Delta\Theta_t^H, \quad (7.3)$$

where  $O_t$  is the ambient temperature,  $\Delta\Theta_t^{TO}$  is the temperature *rise* of the oil, used as coolant, over the ambient temperature and  $\Delta\Theta_t^H$  is the temperature *rise* of the hottest spot on the winding over the temperature of the oil, both for time interval  $t$ . This temperature is also plotted in Figure 7.4.  $\Delta\Theta_t^{TO}$  is calculated using:

$$\Delta\Theta_t^{TO} = (\Delta\Theta_t^{TO,ult} - \Delta\Theta_t^{TO,ini}) \left(1 - e^{-\frac{\delta_t}{\tau^{TO}}}\right) + \Delta\Theta_t^{TO,ini}, \quad (7.4)$$

where  $\Delta\Theta_t^{TO,ini}$  is the initial temperature rise at the start of the interval,  $\tau^{TO}$  is a parameter given in minutes and  $\Delta\Theta_t^{TO,ult}$  is the ultimate temperature rise at the given load for the time interval, i.e., the temperature rise the oil will reach if the transformer has to supply a load  $x_t$ , given for time interval  $t$ , for a very long time (up to infinity). If several consecutive time intervals are considered,  $\Delta\Theta_{t+1}^{TO,ini}$  is assumed to be equal to  $\Delta\Theta_t^{TO}$ , the temperature rise of the previous time interval. Similarly  $\Delta\Theta_t^H$  is calculated as:

$$\Delta\Theta_t^H = (\Delta\Theta_t^{H,ult} - \Delta\Theta_t^{H,ini}) \left(1 - e^{-\frac{\delta_t}{\tau^w}}\right) + \Delta\Theta_t^{H,ini}, \quad (7.5)$$

with the various parameters similarly defined as in (7.4). While the parameter  $\tau^w$  can be assumed to be independent of the temperature of the winding, the parameter



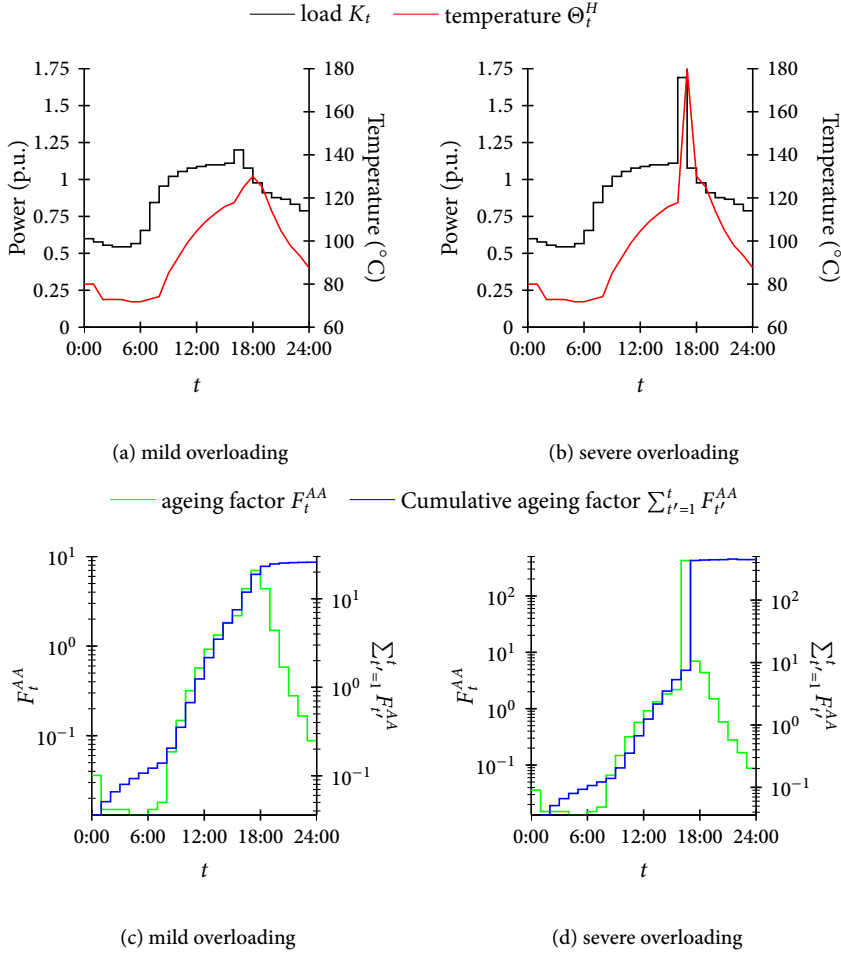


Figure 7.4: Example of the behaviour of the hottest spot temperature for a transformer with (a) mild and (b) severe overloading and the ageing factor for (c) mild and (d) severe overloading. Example data taken from [4].

$\tau^{TO}$  does depend on the initial temperature rise  $\Delta\Theta_t^{TO,ini}$  and ultimate temperature rise  $\Delta\Theta_t^{TO,ult}$  by:

$$\tau^{TO} = \tau^{TO,R} \frac{\left( \frac{\Delta\Theta_t^{TO,ult}}{\Delta\Theta_t^{TO,R}} \right) - \left( \frac{\Delta\Theta_t^{TO,ini}}{\Delta\Theta_t^{TO,R}} \right)}{\left( \frac{\Delta\Theta_t^{TO,ult}}{\Delta\Theta_t^{TO,R}} \right)^{\frac{1}{n_1}} - \left( \frac{\Delta\Theta_t^{TO,ini}}{\Delta\Theta_t^{TO,R}} \right)^{\frac{1}{n_1}}}, \quad (7.6)$$

where  $\tau^{TO,R}$  is the value of parameter  $\tau^{TO}$  for the load value at which the trans-



former was rated,  $\Delta\Theta_t^{TO,R}$  is the ultimate temperature rise of the oil for the load value at which the transformer was rated, and  $n_1$  is an exponent based on the type of cooling used. Note that  $n_1$  typically takes values between 0.8 and 1, where in the case that  $n_1 = 1$ ,  $\tau^{TO}$  is assumed to be equal to  $\tau^{TO,R}$  for any initial and ultimate temperature rises.

Finally,  $\Delta\Theta_t^{TO,ult}$  and  $\Delta\Theta_t^{H,ult}$  can be calculated as:

$$\Delta\Theta_t^{TO,ult}(K_t) = \Delta\Theta^{TO,R} \left( \frac{(K_t)^2 R + 1}{R + 1} \right)^{n_1}, \quad (7.7a)$$

$$\Delta\Theta_t^{H,ult}(K_t) = \Delta\Theta^{H,R} (K_t)^{2n_2}, \quad (7.7b)$$

where  $K_t$  is the ratio between the load for interval  $t$  and the load at which the transformer is rated,  $\Delta\Theta_t^{H,R}$  is the hottest spot temperature rise over the oil at the load at which the transformer is rated, and  $n_2$  is an exponent similar to  $n_1$ , also typically taking values between 0.8 and 1.

### 7.3.2 SIMPLIFYING ASSUMPTION

The parameter  $\tau^{TO}$  for the top oil temperature rise, calculated by (7.6) depends, through a complex relation, on the initial and ultimate temperature rise. Furthermore, the standard states that the relation given in (7.6) is chosen such that both the initial rate of change of the temperature rise of the oil as well as the final temperature rise of the oil are correctly approximated by the model. However, intermediate values of the oil temperature rise might deviate. As we are exactly interested in these intermediate temperatures we make the following assumption.

**Assumption 7.1.** *The effect on the ageing factors of different values of the parameter  $\tau^{TO}$  for different initial and ultimate temperature rises of the oil over the ambient temperature can be neglected. Thus, we take  $\tau^{TO} = \tau^{TO,R}$ .*

The model stated in Subsection 7.3.1, i.e., without Assumption 7.1 we call the ‘full model’. On the other hand, the model we obtain below, using Assumption 7.3.1, we denote by the ‘simplified model’. Using Assumption 7.1, we have a closer look at (7.4). As noted, for consecutive time intervals we have  $\Delta\Theta_t^{TO,ini} = \Delta\Theta_{t-1}^{TO}$ . Furthermore, we assume that each time interval is of equal length, i.e., we take  $\delta_t = \delta$  for all  $t$ .



Now (7.4) can be rewritten as:

$$\begin{aligned}
 \Delta\Theta_t^{TO} &= \left(1 - e^{-\frac{\delta}{\tau^{TO}}}\right) \Delta\Theta_t^{TO,ult} + e^{-\frac{\delta}{\tau^{TO}}} \Delta\Theta_t^{TO,ini} \\
 &= \left(1 - e^{-\frac{\delta}{\tau^{TO}}}\right) \left(\Delta\Theta_t^{TO,ult} + e^{-\frac{\delta}{\tau^{TO}}} \Delta\Theta_{t-1}^{TO,ult}\right) + e^{-\frac{2\delta}{\tau^{TO}}} \Delta\Theta_{t-1}^{TO,ini} \\
 &= \left(1 - e^{-\frac{\delta}{\tau^{TO}}}\right) \sum_{t'=t-1}^t \left[ e^{\frac{(t'-t)\delta}{\tau^{TO}}} \Delta\Theta_{t'}^{TO,ult} \right] \\
 &\quad + e^{-\frac{2\delta}{\tau^{TO}}} \left(1 - e^{-\frac{\delta}{\tau^{TO}}}\right) \Delta\Theta_{t-2}^{TO,ult} + e^{-\frac{3\delta}{\tau^{TO}}} \Delta\Theta_{t-2}^{TO,ini} \\
 &= \left(1 - e^{-\frac{\delta}{\tau^{TO}}}\right) \sum_{t'=t-2}^t \left[ e^{\frac{(t'-t)\delta}{\tau^{TO}}} \Delta\Theta_{t'}^{TO,ult} \right] \\
 &\quad + e^{-\frac{3\delta}{\tau^{TO}}} \left(1 - e^{-\frac{\delta}{\tau^{TO}}}\right) \Delta\Theta_{t-3}^{TO,ult} + e^{-\frac{4\delta}{\tau^{TO}}} \Delta\Theta_{t-3}^{TO,ini}.
 \end{aligned} \tag{7.8}$$

Continued backwards substitution allows us to rewrite (7.4) further into:

$$\Delta\Theta_t^{TO} = \left(1 - e^{-\frac{\delta}{\tau^{TO}}}\right) \sum_{t'=1}^t \left[ e^{\frac{(t'-t)\delta}{\tau^{TO}}} \Delta\Theta_{t'}^{TO,ult} \right] + e^{-\frac{t\delta}{\tau^{TO}}} \Delta\Theta_1^{TO,ini}, \tag{7.9}$$

where  $\Delta\Theta_1^{TO,ini}$  is the temperature rise of the oil over ambient at the start of the considered time horizon. Similarly we can rewrite (7.5) into:

$$\Delta\Theta_t^H = \left(1 - e^{-\frac{\delta}{\tau^H}}\right) \sum_{t'=1}^t \left[ e^{\frac{(t'-t)\delta}{\tau^H}} \Delta\Theta_{t'}^{H,ult} \right] + e^{-\frac{t\delta}{\tau^H}} \Delta\Theta_1^{H,ini}, \tag{7.10}$$

where  $\Delta\Theta_1^{H,ini}$  is the temperature rise of the hottest spot over the oil at the start of the considered time interval. Note that the above equations for  $\Theta_t^{TO}$  and  $\Theta_t^H$  are affine maps of  $\Delta\Theta_{t'}^{TO,ult}$  and  $\Delta\Theta_{t'}^{H,ult}$  for  $t' = 1, 2, \dots, t$  respectively. This observation allows us to proof a convexity result for the simplified model later on.

## 7.4 USING EM TO REDUCE TRANSFORMER AGEING

As mentioned, the ageing of a transformer depends on the energy flowing through it. As such, it can be influenced by an EM approach which is able to change the energy profile of the neighbourhood behind the transformer. In this chapter we focus on only a part of the load behind the transformer, namely the residential charging of EVs, as these are novel, large loads that are expected to have a large impact on the distribution grid [141]. We aim to optimize the charging of the EVs such that the ageing of the transformer is minimized. We consider a set  $\mathcal{M}$  of EVs which offer their flexibility to the EM approach. We schedule the charging of the EVs over a time horizon of  $T$  time intervals, which results for each EV  $m \in \mathcal{M}$  in a schedule  $\mathbf{x}^m = (x_1^m, x_2^m, \dots, x_T^m)$ . For the schedules  $\mathbf{x}^m$  we assume that the same local (i.e., device specific) constraints apply as we used in Chapter 4, i.e., we assume



$$\sum_{t=1}^T x_t^m = C^m, \quad (7.11a)$$

$$0 \leq \mathbf{x}^m \leq \mathbf{x}^{max,m}, \quad (7.11b)$$

with  $C^m$  the required charge of EV  $m$  and  $\mathbf{x}^{max,m}$  the vector of upper bounds on the charging allowed for EV  $m$  over the horizon. Based on the above we define the local constraint set  $X^m$  to be given by:

$$X^m = \{\mathbf{x}^m \mid \mathbf{x}^m \text{ satisfies (7.11)}\}. \quad (7.12)$$

We neglect grid losses within the distribution grid, implying that the energy profile flowing through the transformer of the neighbourhood for time interval  $t$  can be calculated by adding the load of the EVs to the base load of the neighbourhood, i.e., the load of the devices that cannot be controlled. The ratio of load at interval  $t$  to the load at which the transformer is rated can then be calculated by:

$$K_t = \frac{\beta_t + \sum_m x_t^m}{\mathcal{R}}, \quad (7.13)$$

where  $\beta_t$  is the load of the uncontrolled devices for time interval  $t$  and  $\mathcal{R}$  is the load at which the transformer is rated. Since we consider equal length time intervals, minimizing transformer ageing is equivalent to minimizing the average ageing factor over the considered time horizon. This is in turn equivalent to minimizing the sum of the ageing factors. This leads to the following optimization problem:

**Problem 7.1 (MA).**

$$\min_{\mathbf{x}} \sum_t F_{AA}^t, \quad (7.14a)$$

$$\text{s.t. } \mathbf{x}^m \in X^m \quad \forall m \in \mathcal{M}, \quad (7.14b)$$

$$\text{and (7.1), (7.3), (7.9), (7.10), (7.13)}. \quad (7.14c)$$

We call this problem the minimize ageing (MA) problem and note that it uses the simplified model. Problem MA is a nonlinear optimization problem. However, as we show below, it is a convex optimization problem. To obtain this convexity result, we first show that the ultimate temperature rise of both the oil and the hottest spot is convex in the ratio of the load of any time interval to the load at which the transformer was rated.

**Lemma 7.1.** *Both  $\Delta\Theta_t^{H,ult}$  and  $\Delta\Theta_t^{TO,ult}$  are convex in the ratio  $K_t$  of the load for time interval  $t$  and the load at which the transformer was rated.*

*Proof.* We first consider  $\Delta\Theta_t^{H,ult}$ , which is given by (7.7b). The second derivative of (7.7b) with respect to  $K_t$  is given by:



$$2n_2(2n_2 - 1)\Delta\Theta_t^{H,R}(K_t)^{2n_2-2}, \quad (7.15)$$

which is positive since  $n_2 \in [0.8, 1]$  and  $(K_t)^{2n_2-2}$  is positive for  $K_t \geq 0$ .

Next we consider  $\Delta\Theta_t^{TO,ult}$ , which is given by (7.7a). The second derivative of (7.7a) with respect to  $K_t$  is given by:

$$\frac{2n_1R((K_t)^2R + 1)^{n_1-2}((2n_1 - 1)(K_t)^2 + 1)}{(R + 1)^{n_1}}, \quad (7.16)$$

which is positive because  $n_1 \in [0.8, 1]$ .  $\square$

Lemma 7.1 allows us to show that Problem MA is a convex problem.

**Theorem 7.1.** *Problem MA is convex for load values that keep the hottest spot temperature below 7227 °C, i.e., for all realistic load values.*

*Proof.* For each  $m \in \mathcal{M}$ , the set  $X^m$  of locally feasible schedules is convex since (7.11a) is linear and (7.11b) is a bounding box. Thus, it remains to show that the ageing factor is a convex function of  $x_t^m$ . For any given  $t$ ,  $K_t$  is given by an affine map of the  $x_t^m$ 's for all  $m \in \mathcal{M}$  (see (7.13)). Hence,  $\Delta\Theta_t^{H,ult}$  and  $\Delta\Theta_t^{TO,ult}$  are convex in each  $x_t^m$  by Lemma 7.1. Furthermore,  $\Delta\Theta_t^H$  and  $\Delta\Theta_t^{TO}$  are given by increasing affine maps of each  $\Delta\Theta_{t'}^{H,ult}$  and  $\Delta\Theta_{t'}^{TO,ult}$  respectively, for  $t' \leq t$  (see (7.9) and (7.10)). Hence these are convex in each of the  $x_t^m$ 's. Finally, note that  $\Theta_t^H$  is an increasing affine map of  $\Delta\Theta_t^H$  and  $\Delta\Theta_t^{TO}$  (see (7.3)), hence it is convex in each of the  $x_t^m$ 's.

Next we consider  $F_t^{AA}$  which is given by (7.1). The first derivative with respect to  $\Theta_t^H$  is given by:

$$e^{\frac{15000}{383}} e^{\frac{15000}{\Theta_t^H + 273}} \frac{15000}{(\Theta_t^H + 273)^2}. \quad (7.17)$$

The second derivative is given by:

$$15000e^{\frac{15000}{383}} e^{\frac{15000}{\Theta_t^H + 273}} \left( \frac{15000}{(\Theta_t^H + 273)^4} - \frac{2}{(\Theta_t^H + 273)^3} \right). \quad (7.18)$$

Note that the first derivative is strictly positive for any value of  $\Theta_t^H$  and the second derivative is strictly positive for

$$\frac{15000}{(\Theta_t^H + 273)^4} > \frac{2}{(\Theta_t^H + 273)^3}, \quad (7.19)$$

which is true for  $\Theta_t^H < 7227$ , i.e., for all temperatures resulting from realistic energy flows. Thus, the ageing factors  $F_t^{AA}$  are indeed convex in  $\Theta_t^H$  for any realistic loading pattern. The theorem now follows from the fact that  $\Theta_t^H$  is a convex function of each of the  $x_t^m$ 's and  $F_t^{AA}$  is a convex and increasing function of  $\Theta_t^H$ .  $\square$

Theorem 7.1 allows us to solve the Problem *MA* using a standard convex solver. Using this, we can compare the obtained load profiles of different charging strategies for the EVs.



## 7.5 SIMULATION STUDY

In this section we study the potential of an EM approach to reduce the ageing of transformers through a simulation study using the models developed in the previous section. We use load data from a real world neighbourhood transformer in the town of Lochem in the Netherlands to construct instances of Problem *MA* that represent scenarios with significant penetration levels of EVs. We compare the optimization results, obtained using the simplified model, with the temperature profile and ageing factors calculated using the full model, i.e., the model described in Section 7.3. Furthermore, we compare the results of the optimization to an implementation of our profile steering EM approach, where we steer towards flattening the load profile. For this comparison we calculate the ageing of the transformer when the profile steering approach is applied. Finally, we calculate the ageing when no coordination between the charging of the EVs is applied. For this uncoordinated charging we consider two cases: the EVs are either charged as quickly as possible upon arrival or their charging is equally spread over the period that the vehicle is available for charging.

### 7.5.1 CONSIDERED SCENARIO

As mentioned, we consider a transformer in the town of Lochem for which we have detailed load measurements available. This transformer supplies a neighbourhood of residential customers. The transformer is rated at 400 kVA with an average winding rise of 65 °C. As a test case we use data from November 3rd until November 9th in 2014, for which the base load profile is given in Figure 7.5 (a) (denoted by ‘Base’). From the available data we calculate the fifteen minute averaged load values. The temperature profile of the transformer resulting from this base load is given in Figure 7.5 (b).

The values of the required parameters forming the base of calculating the ageing are given in Table 7.1. Cooling of the transformer is done through natural convection, known as ONAN, cooling. For this type of cooling, [4] suggests that the values of the exponents  $n_1$  and  $n_2$  are set to 0.8. For the initial temperature of the top oil and hottest spot we use an equivalent load  $B$  to the load profile of the previous day (for details see [4]).

For the other parameters of the transformer we used values from [3] for a 65 °C degrees average winding rise transformer. Finally, for the ambient temperature we used measurement data from a weather station in Eefde, near Lochem.

Next we specify the data used for the EVs in the considered scenarios. We assume there are 50 EVs in the neighbourhood. For each EV we assume that it arrives at

Table 7.1: Used parameter values within problem *MA*.

Parameter	Value	Parameter	Value
$\Delta\Theta^{TO,R}$	55 °C	$\Delta\Theta^{H,R}$	20 °C
$\tau^{TO}$	210 min	$\tau^w$	5 min
$B$	97897 W	$R$	2.7
$\Delta\Theta_1^{TO,ini}$	21.0 °C	$\Delta\Theta_1^{H,ini}$	1.53 °C
$n_1$	0.8	$n_2$	0.8

a randomly selected time between 17:00 and 19:00. Furthermore we assume that each EV has to be fully charged by a time randomly selected between 6:00 and 8:00 the next day. The total required charge and maximal charging power of each EV is varied over the scenarios. This is done to simulate the effect of a higher penetration of EVs as doubling the required charging and maximal charging power of each EV is nearly the same as modelling twice as many vehicles from the transformer perspective. As a basis, we use the scenario where each EV can charge with a maximal power of 3.8 kW and has to charge a random amount between 10 and 15 kWh every night. We call this scenario *low penetration*. Furthermore, we performed simulations for the case that the required charging and maximal charging power are either doubled or tripled. These scenarios are called *medium penetration* and *high penetration* respectively.

### 7.5.2 OPTIMIZING TRANSFORMER LIFETIME

To optimize the transformer lifetime by controlling the charging of the EVs, we implemented Problem *MA* in the AIMMS modelling software [1]. Within AIMMS we used the CONOPT solver to produce a solution of Problem *MA*. While CONOPT only finds a local optimum, the convexity result, given in Theorem 7.1, ensures that this is in fact a global optimum.

As discussed above, we used the simplified model in order to be able to derive our results. To verify that our results do not deviate too far from the full model we compare the temperatures and resulting ageing for the load profile found by the simplified model with those obtained when using the full model. The results are listed in Table 7.2.

For the considered scenarios and found charging profiles of the EVs using the simplified model gives slightly lower temperature values and hence gives slightly lower ageing. However, for the considered cases, the differences are small enough to assume that the found charging profiles give transformer ageing results that are almost identical with the minimum that would be obtained when using the full model. The resulting load profiles can be found in Subfigure (a) of Figures 7.5-7.7 for the low, medium and high penetration scenarios respectively, where the curve is denoted by ‘Opt’. Furthermore, the temperature profiles of each of these scenarios can be found in Subfigure (b) of Figures 7.5-7.7. When minimizing transformer ageing we observe a slight increase in the total load profile at the end of the night.



Table 7.2: Differences between outcomes of the ageing model with the simplified model and the full model for optimized EV charging for the scenarios.

Scenario	Maximum absolute temperature difference (°C)	Ageing difference (%)
light penetration	0.02	0.04
medium penetration	0.03	0.17
heavy penetration	0.09	0.65

This increase can be explained by a lower ambient temperature during these time intervals. Finally, the increase of transformer ageing for each of the scenarios is listed in Table 7.3. This table gives the increase in the ageing factor compared to the ageing under the base load (i.e., with no EVs present). An increase of 100% means that the transformer is expected to last only half as long compared to a scenario where no EVs are present.

### 7.5.3 PROFILE STEERING

For the simplified model of transformer ageing, the temperature rise of both the oil and hottest spot is an increasing convex function of the load on the transformer and the ageing of the transformer is an increasing convex function of the hottest spot temperature. Therefore, flattening the load profile should intuitively reduce the ageing of the transformer significantly. As a comparison to the presented optimization directly towards transformer ageing, we implemented our profile steering EM approach, with as goal to achieve a flat profile. As the objective we used the Euclidean norm and a desired profile of zero energy usage. To derive the schedules of the EVs we used the algorithms from Chapter 4. The resulting temperature and load profiles can be found in Figures 7.5-7.7 and Table 7.3 where they are listed as ‘PS’. The profile steering approach spreads the charging of the EVs such that the resulting profile is as flat as possible. The influence on ageing and temperature is discussed in Subsection 7.5.5.

### 7.5.4 UNCONTROLLED CHARGING

For comparison, we simulated the scenarios also without a coordination mechanism for the EV charging. While currently most EVs start charging at maximum power directly when plugged-in until they are fully charged, several studies have shown that this can cause a tremendous amount of stress on the grid [67]. Hence, in the future, EV owners might be incentivized to spread the charging of their own vehicle equally over the time period for which the EV is available for charging. To assess the potential of our optimization approach we implemented both cases. The resulting temperature and load profiles for both cases in each of the scenarios can be found also in Figures 7.5-7.7 respectively, where ‘Max’ denotes the case of maxi-



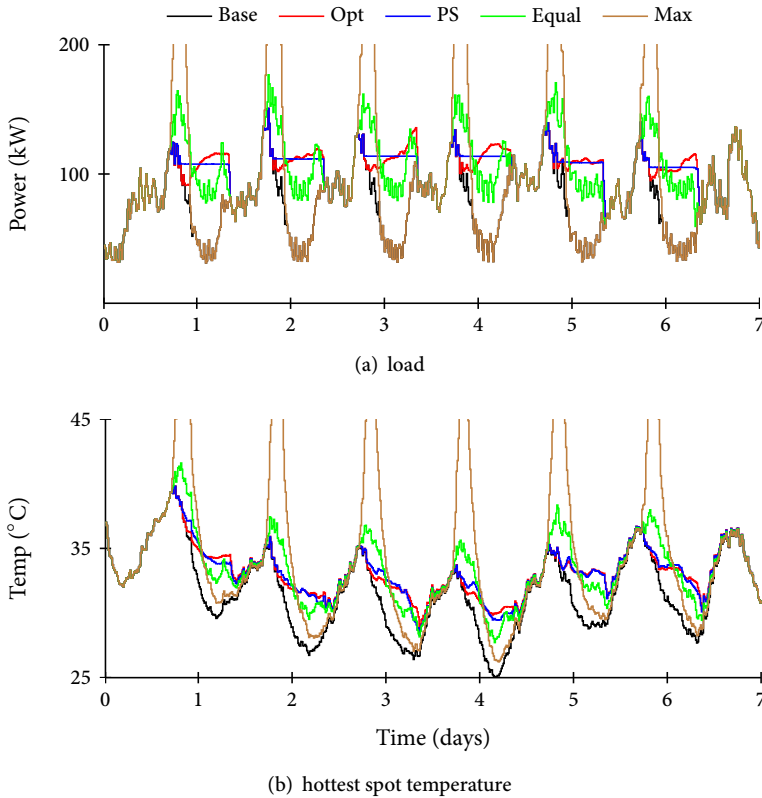


Figure 7.5: Results for low penetration case (note that the ‘Max’ case is clipped).

mal charging upon arrival and ‘Equal’ denotes the case of equal loading over the entire charging interval. The increase in ageing over the base load can be found in Table 7.3 in the rows Max and Equal.

#### 7.5.5 COMPARISON OF THE CASES

The results of the simulation study show a very large reduction in transformer ageing when smart charging is adopted. Specifically in the high penetration case, the transformer ages rapidly when the vehicles are fully charged upon arrival. This is mainly caused by the high peaks in this case that overload the transformer, causing high temperatures and extreme wearing. While flat charging already significantly decreases transformer ageing, this can be further improved by applying an EM approach steering the EV charging.

When comparing directly minimizing transformer ageing with the profile steering approach, we conclude that the differences in temperature and ageing are minimal.




Figure 7.6: Results for low penetration case (note that the ‘Max’ case is clipped).

For all considered scenarios, transformer ageing minimization slightly favours time intervals later during the night for charging. This can be explained by a lower ambient temperature for these time intervals, which allows a slightly higher load on the transformer with the same resulting temperature of the hottest spot and thus the same ageing. For the considered scenarios however, an EM approach that steers towards flattening the load profile of the neighbourhood provides adequate results with respect to minimizing transformer ageing.

## 7.6 CONCLUSION

In this chapter we studied the ageing of an important and costly asset in the electricity (distribution) grid and how this can be influenced by an EM approach. This research was motivated by the fact that the energy transition is expected to lead to an increase in stress on the electricity grid, and in particular the distribution grid. This increased stress mainly results from the electrification of our energy




Figure 7.7: Results for low penetration case (note that the ‘Max’ case is clipped).

Table 7.3: The increase (in percentages) of the equivalent ageing factor when comparing the base load and the additional load caused by the different EV loading strategies.

	light penetration (%)	medium penetration (%)	heavy penetration (%)
Opt	20.6	76.5	200.9
PS	20.9	77.6	202.1
Equal	27.6	92.2	243.9
Max	161.0	4763.5	22306.9



use and the introduction of local (renewable) generation. The increased stress can cause a large increase in investments in the grid especially if the energy transition is not handled properly. EM approaches can help to keep the energy transition economically feasible by using flexibility to reduce the stress on the grid.

To study the impact of the electrification on our electricity grid we study the ageing of a (distribution) transformer in this chapter. For this we use an ageing model developed by IEEE. This model relates energy flows through the transformer to ageing of the asset. The relation is, in its base, nonlinear. We showed that the given relation is convex under a mild assumption. This allowed us to formulate a convex optimization problem for the case that EV charging can be steered to minimize the ageing of the transformer.

We compared the results of this optimization towards decreased transformer ageing to other charging strategies in a simulation study using different levels of EV penetration in a Dutch neighbourhood. The results of the simulation study indicate that charging EVs maximally upon arrival causes the transformer to degrade rapidly. Significant improvements are made by spreading the charging of each EV over its complete charging horizon. Further improvements are obtained if the EV charging is coordinated with the charging of the other EVs and the profiles of the other loads in the neighbourhood. Finally, the results obtained when directly minimizing transformer ageing and when applying our profile steering approach with a flat profile as the target profile are very similar. This indicates that steering towards a flat profile is beneficial for the (local) grid. As the computational complexity of directly optimizing towards transformer ageing is much higher than that of the profile steering approach, the above indicates that profile steering offers a good approach when asset ageing is of importance and the computational power is limited.

The model used in this chapter is based on thermal properties of the transformer. The used transformer model was originally designed to help size transformers before they are installed in (new) grids. With this goal in mind, some of the assumptions made in the model are on the pessimistic side. Nevertheless, we expect a similar (convex) relation to hold between energy flows and degradation for other assets, such as cables. The results in this chapter should extend to the assets for which such a similar (convex) relation holds. Furthermore, we note that the results indicating that profile steering gives near optimal results regarding transformer degradation may also be extended to other EM approaches that flatten the profile of a neighbourhood.





# 8



## SUMMARY AND CONCLUSION

*ABSTRACT – This chapter is divided into three parts. First the main results of the thesis are summarized (Section 8.1). The main results are the introduction of an EM approach called profile steering, the development of device level models and solution methods and a brief discussion on asset ageing in the smart grid and the multi-objective nature of our future energy system, shared by different stakeholders. In the second section we recap on the research questions posed in Chapter 1 and summarize the important results that answer these questions (Section 8.2). Finally, we give some recommendations on potential ways to improve the approach in future work to make profile steering more viable for practical implementation in the future smart grid (Section 8.3).*






Table 8.1: Lowest obtained complexity for device level solution methods with quadratic objective functions for both the continuous version and the relaxation of the discrete version.  $T$  denotes the number of time intervals and  $M$  denotes the total number of pieces in the piecewise linear approximation of the discrete versions of the problems.

Problem	Complexity (continuous)	Complexity (discrete)
Electric vehicle (Chapter 4)	$O(T)$	$O(M)$
Storage devices (Chapter 5)	$O(T^2)$	$O(TM)$
HVAC systems (Chapter 6)	$O(T^2)$	$O(TM)$

Table 8.2: Overview of our requirements on an EM approach, as listed in Chapter 2, and a description of how profile steering satisfies them.

Requirement	Description of the solution in profile steering
Scheduling	By design, the approach solves scheduling problems to determine when and how to use the flexibility available in the various DERs.
Scalability	Through the decentralized approach most of the computation is done on a local level, allowing for heavy use of parallelism.
Heterogeneity	Because device level problems are solved locally, heterogeneity only implies different DERs need to be able to make their own schedule, which we showed is possible for several device classes.
Feasibility	Device schedules are made on the local level, where constraints are known. The natural mapping between the control hierarchy and the physical grid leads to an approach where grid limitations are taken into account.
Privacy	Privacy sensitive information is only required on a local level to make the device level schedules. As only the resulting energy profile is communicated upwards and aggregated on various levels, the approach keeps privacy sensitive information very local.



## 8.1 SUMMARY OF THE OBTAINED RESULTS

With the shift towards an energy system based on clean renewable sources, the system is changing, particularly in our electricity grid. An important trend is that flexibility is lost on the production side. Traditional EM approaches depend on a large amount of flexibility on the production side. Furthermore, these traditional approaches do not scale to the large amount of DERs envisioned in the future smart grid. Therefore the traditional approaches are no longer applicable in the future system. Thus, a new EM approach is needed to enable the energy transition towards a sustainable energy system.

In Chapter 3 of this thesis we propose the profile steering EM approach. It is a novel EM approach that is applicable to the future (smart) grid. Profile steering is capable of managing the available flexibility from DERs through a decentralized control approach. In this decentralized approach, the decisions how and when to use flexibility are made locally (i.e., by the devices). The coordination between the use of flexibility from different sources is achieved through steering signals. These signals are sent by a central controller and indicate to the devices when it is useful for the overall system goals to use their flexibility.

In Chapters 4 to 6 we studied how to make the local decisions for the devices, particularly in the profile steering approach. We assumed that the steering signals used in our profile steering approach result in convex and separable objective functions for the local problems. We showed that the resulting problems are resource allocation problems. In particular we studied buffering devices, which utilize an internal storage to fulfil their primary function. In Chapter 4 we considered a device that can only be charged, e.g., an EV. For such a device we discussed efficient solution approaches for both the continuous and discrete variants of the problem.

In Chapter 5 we extended the buffering devices model to include discharging to capture, e.g., stand-alone batteries and heat generators combined with heat storage. We showed that the resulting problems can be solved using a divide and conquer strategy in the continuous case. This resulted in an efficient solution approach using the methods developed in Chapter 4. Furthermore, we extended the approach from Chapter 4 to the discrete case of the model including discharging. Then, in Chapter 6 we included state of charge (SoC) dependent losses to the model. We showed that the results from Chapter 5 also apply to this case after minor modifications. The best obtained complexity for each of the device level problems is given in Table 8.1.

Finally, in Chapter 7 we studied a model for transformer ageing. We showed that, under a mild assumption, we can formulate the objective of minimizing transformer ageing when steering the charging of EVs as a convex optimization problem. This allowed us to calculate optimal charging strategies for the EVs with respect to transformer ageing and compare them to other EM approaches. In particular, the results indicate that flattening load profiles gives near optimal results with respect to transformer ageing.



## 8.2 CONCLUSION & DISCUSSION

In this thesis the central research question we studied is:

*How can we effectively and efficiently manage the flexibility provided by (future) energy resources in the electricity grid to facilitate the changes occurring due to the energy transition?*

In particular we clarified effective and efficient to mean that:

*An EM approach is required to employ **scheduling** in a **scalable** manner to use flexibility offered by a large **heterogeneous** set of resources in a **feasible** manner at the appropriate time while **respecting user privacy**.*

To answer this question we first consider the sub-questions from the introduction.

*How can we effectively manage the coordination of flexibility use **between** devices?*

To effectively coordinate the use of flexibility between different DERs in the future smart grid we introduced the profile steering approach. In this approach the devices decide on the local level when and how to use flexibility. The devices make these decisions based on steering signals. These steering signals are communicated through a control hierarchy, where the centralized controller sends steering signals based on the system goals. The control hierarchy reflects the physical structure of the grid in a tree structure. We achieve this by subsequently dividing the DERs over controllers on the different levels of the tree to achieve different groups on each level in the control hierarchy. For example, devices in an LV grid can be grouped according to their house, phase, and feeder. Each group is managed by a different controller on the corresponding hierarchical level, where this controller adapts the steering signals to ensure the devices it governs use their flexibility when this best fits the system goals specified by the steering signals.

Using the strong resemblance between the control hierarchy and the physical grid, we incorporated grid constraints into the approach. This is achieved through limitations set on the allowed schedules of when and how flexibility is used by the group of DERs that a controller governs. These limits correspond to the physical limits of the corresponding asset, e.g., the total load on a cable is kept between bounds by a controller governing all the DERs connected to this cable (or by a controller that is situated higher in the control hierarchy).

Throughout this thesis we have shown, by means of simulations, that the profile steering approach is capable of steering the combined energy profile of different neighbourhoods to better fit a desired profile of the central controller. One of the strengths of the approach is that it is not limited to a single goal or signal. Rather, a different aspect of the obtained energy profile can be emphasized and considered by the approach by using different steering signals. Note, that the results in this thesis hold for any convex separable objective functions used as steering signals. These functions cover a broad range of desired objectives in EM, e.g., cost minimization, peak shaving, and local self consumption.





*What are the local decision problems in future EM approaches and how can we solve them on the local level?*

In profile steering, the devices make their own local decisions on when and how to use flexibility. As scalability is an important requirement for any future EM approach, we believe it is likely that any future EM approach requires some level of decision making on the local level. As mentioned, the studied objective functions cover a broad range of goals in EM and should therefore be applicable in many situations.

We mentioned in Chapter 2 that DERs fall into four different classes, of which we considered two. The first class we considered, the time-shiftables, can generally be solved by an exhaustive search through the state space of potential solutions. This means that it is possible to consider all different ways of scheduling the flexibility use of the device and to determine the best possible schedule through comparison. This is no longer true when considering the second class of DERs we studied; the buffers. Assuming that the steering signals results in convex and separable objective functions, the resulting optimization problems fall into the class of resource allocation problems. Several problems in the class of resource allocation have been well studied in literature.

For the buffer class of DERs we considered three types of problems which we believe cover many different (future) devices in this class. In particular we considered a model of the charging of an EV. We extended this model to include discharging, such that it covers vehicle-to-grid capabilities and stand-alone batteries. We further expanded upon the extension by including SoC dependent losses. These losses play an important role in many heating and cooling systems.

For the model and each of its extensions we considered a continuous case where the amount of energy consumption in a time interval is limited by a lower and an upper bound. Also, for each case we considered a discrete variant, where the feasible amounts of energy consumption/production for a time interval are given by a finite set. Determining if a feasible solution for the discrete variant of the EV model exists is already  $\mathcal{NP}$ -hard and, hence, this is also the case for the considered extensions. To circumvent this problem we instead consider a relaxation where we allow convex combinations of energy values given in the discrete set.

For the resulting optimization problems on the device level we formulated efficient solution approaches, whereby the efficiency of our solution approaches is two fold. First, we show that we can solve these problems with approaches that have a low asymptotic complexity (see Table 8.1), particularly in the case that the steering signals results in quadratic local objectives. This implies that the solution approaches discussed in this thesis are scalable to large input sizes if needed. Furthermore, for the studied problems we obtained solution approaches that have both a low asymptotic complexity and allow an efficient implementation in practice. Therefore, we believe we showed that the developed solution approaches can be implemented on local and low-cost controllers, ensuring that indeed the local problems can be solved locally. Furthermore, while not detailed in this thesis, initial experiments



showed that our solution methods for the local problems require much less time and resources than an implementations of the same problems in (commercial) solvers available for these problems does. This indicates that our solution methods are indeed suitable for implementation on the (very) low cost controllers expected to be used in devices/on the household level in the future smart grid.

*Can an EM approach assist in realizing goals of different stakeholders in the (future) smart grid?*

The examples throughout this thesis show that it is indeed possible for an EM approach to realize multiple goals simultaneously. In particular, in the simulation study performed in Chapter 5, both the total amount of energy that the group of houses imports/exports as well as the transportation peaks are minimized. This is mainly caused by the fact that both objectives are captured in the used objective function (the squared Euclidean norm), combined with the assumption that storage does not incur losses. Another example of such a result is given in Chapter 7. The obtained results indicate that steering towards a flat profile assists with minimizing asset ageing in the local grid. Furthermore, ensuring the energy profile of a neighbourhood is flat or follows a predetermined desired profile, as is also possible in the profile steering approach, is beneficial for energy suppliers, as they can generally supply such an energy profile with less costs than an arbitrary profile (with peaks).

The profile steering approach with our developed local algorithms can also be used to combine different objectives under the following condition. It is required that the different objectives can be captured together into a single objective function. This can be achieved, for example, through weighting the different objectives. Furthermore, the resulting objective needs to be separable and convex in order for our solution approaches to the local problems to still be applicable.

All in all we believe that profile steering, combined with the device level solution methods developed in this thesis, is capable of efficiently and effectively managing flexibility provided in the future grid. Thus, the approach is capable of assisting in facilitating the changes caused by the energy transition. Below we quickly recap on the individual requirements that entail this efficient and effective management, as we listed above (for an overview see also Table 8.2 on page 180).

- » By design the profile steering approach is based around *scheduling* the use of flexibility. In other words, the approach inherently looks ahead to determine when it is best to use the available flexibility. This results in an approach that is more robust than approaches that do not look ahead.
- » The proposed approach is *scalable*, because we distribute most of the computation to the device level. Furthermore, we showed that the device level problems allow efficient solution methods, demonstrating that such solution methods can run on local, low-cost hardware. This in turn implies that the approach is capable of exploiting parallelism in the computation of local schedules by distributing the computation to the device level.



- » In this thesis we discussed four different classes of DERs and gave solution approaches for devices in both the time-shiftable and buffering class. Because these device classes were designed to include devices with similar characteristics, we believe the obtained results cover a large share of the (future) devices in these classes. Furthermore, because of the decentralized nature, the approach is agnostic to the specific device details; it only requires a scheduled energy profile from the device. Thus, devices for which the solution methods we presented do not apply can be incorporated as long as they are capable of making local decisions. This implies the profile steering approach can be applied to a large set of *heterogeneous* DERs.
- » The *feasibility* of our approach is guaranteed on both the device and the grid level. First, the devices make their own decisions. We believe it is safe to assume that the constraints on the operation of a device are known locally, hence the decentralized approach ensures locally feasible schedules. Second, the control hierarchy used in the approach closely resembles the structure of most residential grids. This allowed us to incorporate constraints on the use of grid assets (i.e., minimum and maximum power flows) into the approach.
- » The *privacy* issue is partially solved through the decentralized approach with the proposed control hierarchy. By design, the privacy sensitive information required to make (locally) feasible schedules is only required on the local level. While some information can still be deduced from the resulting energy profiles, we believe the information can be sufficiently hidden through the aggregation process used in the various levels of the control hierarchy. Furthermore, we note that less privacy sensitive information is available to the central controller than in many other approaches found in the literature.

### 8.3 RECOMMENDATIONS

In this final section of my thesis (disregarding the Appendix), I spend some words on my recommendations for future work (some of which I'll hopefully work on myself). Looking back I think I have a great deal *more* questions than at the time I started over four years ago, but I was told this is usual for PhD students. I believe that the work I did made a contribution to the advancement of (decentralized) energy management, and that part of it uncovered where more work is required. Thus, below, a non-exhaustive list of my recommendations in no particular order is given.

The control hierarchy for the profile steering approach as proposed in Chapter 3 closely resembles the physical grid. Experiments have shown that the number of levels in the tree used for the control structure can influence the efficiency of the implementation. In particular, it can be beneficial to terminate the iterative phase of profile steering aggressively by, e.g., limiting the total amount of allowed iterations. Furthermore, additional benefits can potentially be obtained by searching for more efficient implementations of the iterative phase of the profile steering approach. For example, it could be beneficial to allow multiple devices to update



their current schedule to their proposed schedule in the iterative phase of profile steering, potentially resulting in faster convergence. For a further discussion on this and some results see [126].

The aforementioned relation between the control hierarchy and the physical grid leads to an approach that is capable of incorporating grid (capacity) constraints directly into the distributed approach. Next to these capacity constraints, also constraints on the delivered power quality are of importance (e.g., the voltage level [99]). A further integration of such constraints into decentralized energy management (*DEM*) approaches is the topic of on-going work of a second PhD student in our group working on the same project, parts of which he published in, e.g., [66].

The device level problems studied in Chapter 4 to 6 model the behaviour and available flexibility of devices in the buffer class of DERs. While the models presented in this thesis should sufficiently represent several devices in this class, for some devices the models potentially do not capture all important aspects. As an example, heating devices often take some time to ramp up to their desired output and need a period of time to achieve their maximum efficiency. Also, some devices might be limited in how often their operational level can be switched, to prevent wearing of these devices. How to incorporate such characteristics into the presented models and how this affects the solution approaches is an important topic for future research to bring EM approaches closer to reality.

The device level problems studied in this thesis fall into the class of resource allocation problems. The solution methods in this thesis use results from previous work in this area. As resource allocation is applicable in many fields (see Subsection 4.3.1), the presented solution methods are potentially also of interest in these fields. Therefore, it should be interesting to study other potential applications of the found solution methods, for example, in the field of green computing. Furthermore, the solution approaches presented in this thesis have a asymptotic complexity at or near the complexity of state-of-the-art approaches. However, as mentioned in Section 4.4 for example, the practical running time of some of these approaches is rather high for small to medium sized instances, specifically for approaches using a linear time version of median find [18]. An interesting topic of research is a detailed comparison between the various presented approaches herein, other state-of-the-art techniques, and (commercial) solvers when solving practical instances of the local problems encountered in the EM setting.

Chapter 7, and other examples throughout this thesis, touched upon the subject of different goals for different stakeholders. An EM approach that is capable of bringing together the goals of all the stakeholders in the future (smart) grid is of higher value to these stakeholders and is therefore more likely to be adopted. For this reason, it is of particular interest to study these different stakeholders and their goals to see if they can be aligned in a single EM approach such as profile steering. Note that the barrier for cooperation between stakeholders is not always of a technical nature. For example, in many countries discharging a battery to the grid is seen as the production of energy, which network operators are by law not

allowed to do. This results in the prohibition of network operators from properly using a battery to support the grid. This highlights that a study into cooperation between stakeholders should not only be of a technical nature, but should also include the legislative aspect.



### 8.3.1 DEALING WITH UNCERTAINTY

In Chapter 2, we mentioned that an EM approach based on scheduling naturally requires a prediction and operational control phase too. The prediction phase is required to be able to obtain good results in the scheduling phase. As some parameters in the EM setting are hard to predict, e.g., human behaviour on a small scale, development of good prediction techniques together with an investigation of how prediction errors influence the results is an important future step. However, it should be evident that no prediction is ever going to be perfect. Hence, it is important to consider ways to mitigate the effects of uncertain predictions.

With efficient solution methods it is in theory possible to use rescheduling to deal with prediction errors, e.g., in case a significant deviation between predictions and realizations is observed. We note that we briefly touched on this in Chapter 6. This rescheduling means that whenever a significant deviation between a realization and a prediction is observed, the system reschedules the flexibility use to account for this error. We believe two aspects of such a rescheduling approach are important to consider.

First, it is important to determine when a deviation from a prediction is significant enough that it needs to be accounted for. Small errors generally do not cause a change in the system's overall performance that outweighs the burden of rescheduling the use of flexibility. Also, some errors potentially cancel out on higher levels in the control hierarchy.

The second important aspect of rescheduling is that in order to deal with prediction errors as they are detected, sufficient flexibility must be available to correct them by rescheduling devices. A potential way to ensure enough flexibility is available is to use past data to give an estimation (possibly pessimistic) of the required flexibility to deal with prediction errors. Based on this, the total flexibility incorporated in the initial scheduling procedure may be limited to ensure that still enough flexibility is available to deal with prediction errors later on.

Next to rescheduling, uncertainty can be dealt with in other ways. One example is to use an online approach that only requires input parameters that are easy to predict (see, e.g., [57]). Another potential strategy is to incorporate techniques from robust or stochastic optimization into the scheduling approach to incorporate the uncertainty of parameters directly into the approach. Finally, on-going research in our group is showing that an event-based approach (similar to that presented in [36]) using the local device algorithms presented in this thesis has the potential to deal with prediction errors without the need to reschedule the use of devices for which this is not preferable.



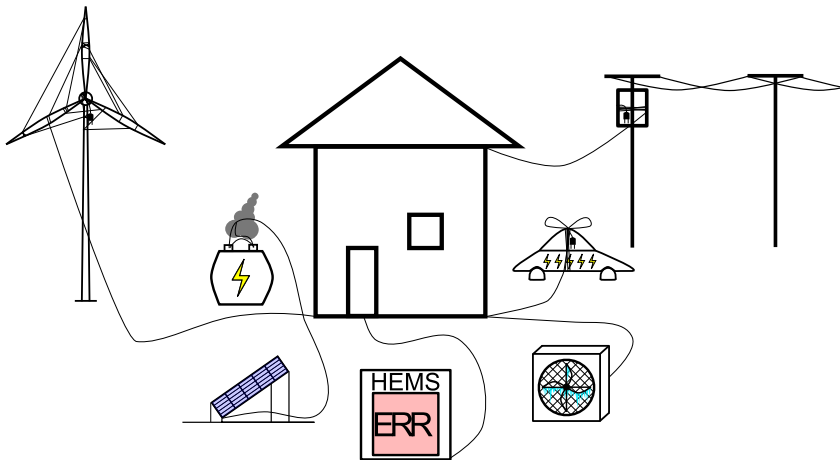


# A



## MATHEMATICAL BACKGROUND

*ABSTRACT – In this appendix we briefly introduce some mathematical background to the problems studied in this thesis. Central to the studies in this thesis is the field of convex optimization, as many problems we study are convex optimization problems. We briefly introduce several properties of convex functions and convex optimization problems that are used throughout proofs in this thesis. Furthermore, we consider the complexity of algorithms. To do so we very briefly introduce the notation and the concepts of polynomial time complexity and the complexity classes  $\mathcal{P}$  and  $\mathcal{NP}$ .*



The content of this appendix is deeply rooted in literature, therefore we omit most of the citations and instead refer the interested reader to textbooks such as [114] for convex analysis and [8] for complexity theory.

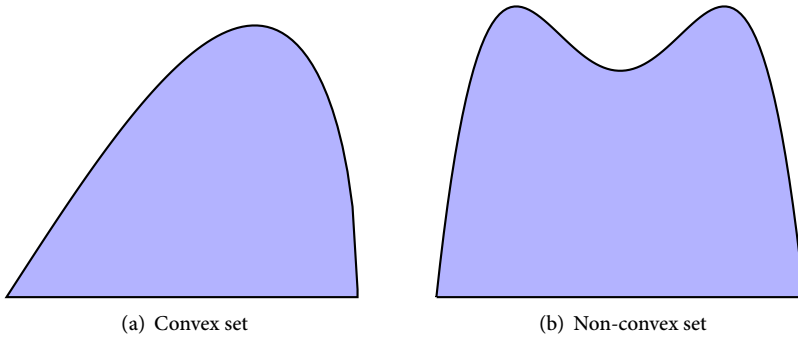


Figure A.1: Examples of a convex set and a non-convex set.




Figure A.2: Examples of a convex function and a non-convex function.




Figure A.3: Example of a convex function of two variables.

## A.1 CONVEX OPTIMIZATION

Below we give a short introduction to convex optimization. For more details we refer the interested reader to one of the many textbooks in the field, e.g., [114].

### A.1.1 CONVEX SETS

Convex sets form the basis of convex analysis. In  $\mathbb{R}^n$ , the convex sets are the sets for which convex combinations of points inside the set also lie in this set. More formally:

**Definition A.1.** A set  $S \subset \mathbb{R}^n$  is called convex if for every  $x, y \in S$  and every  $\lambda \in [0, 1]$  the point  $z := \lambda x + (1 - \lambda)y \in S$ .

Examples of both a set that is convex and a set that is not convex are given in Figure A.1.

### A.1.2 CONVEX FUNCTIONS

A function  $f : D \rightarrow \mathbb{R}$  with  $D \subseteq \mathbb{R}^n$  is called convex if the set of points on and above the graph of  $f$  is a convex set. This is equivalent to requiring that the line between any two points  $f(x)$  and  $f(y)$  on the graph lies above the graph. More formally, this gives:

**Definition A.2.** A function  $f : D \rightarrow \mathbb{R}$  with  $D \subseteq \mathbb{R}^n$  is called convex if, for every  $x, y \in D$  and every  $\lambda \in [0, 1]$ , it holds that

$$f(\lambda x + (1 - \lambda)y) \leq \lambda f(x) + (1 - \lambda)f(y). \quad (\text{A.1})$$

Examples of both a function that is convex and a function that is not convex, both from  $\mathbb{R}$  to  $\mathbb{R}$ , are given in Figure A.2. An example of a convex function from  $\mathbb{R}^2$  to  $\mathbb{R}$  is given in Figure A.3

Convexity of functions is preserved by several operations, of which we list several used throughout this thesis below (without proofs)

- » If  $f$  and  $g$  are convex with the same domain, then so is  $h(x) := f(x) + g(x)$ .
- » If  $f$  and  $g$  are convex and  $g$  is monotonically non-decreasing, then  $h(x) := g(f(x))$  is convex.
- » If  $f$  is convex with domain  $D \subset \mathbb{R}^m$  and  $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$  is an affine map given by  $g(x) = Ax + b$  with  $A \in \mathbb{R}^{m \times n}$ , then  $h(x) := f(g(x)) = f(Ax + b)$  is convex.

If  $f$  is a function from  $\mathbb{R}$  to  $\mathbb{R}$ , we have the following. If  $f$  is differentiable then  $f$  is convex if and only if its derivative  $f'$  is a monotonically non-decreasing function. Furthermore, if  $f$  is convex and  $f'$  exists,  $f'$  is continuous. Finally, if  $f''$  exists then  $f$  is convex if and only if  $f''$  is non-negative.



Next we consider the case that  $f$  is convex but not necessarily differentiable. Furthermore, we assume that the domain of  $f$  is an interval  $[a, b]$  of  $\mathbb{R}$ . For this function  $f$  the left and right derivatives at the point  $x$  are defined as the limit of  $f(x+h)/h$  as  $h$  respectively increases or decreases towards zero. Formally this gives:

**Definition A.3.** For a function  $f : [a, b] \rightarrow \mathbb{R}$ , the left derivative  $f^-(x)$  and the right derivative  $f^+(x)$  are given by

$$f^-(x) = \lim_{h \uparrow 0} \frac{f(x+h)}{h}, \quad (\text{A.2a})$$

$$f^+(x) = \lim_{h \downarrow 0} \frac{f(x+h)}{h}. \quad (\text{A.2b})$$

The derivative at  $x$  exists if and only if  $f^-(x)$  and  $f^+(x)$  exist and are equal. As mentioned before, if the derivative of a convex function  $f$  exists, then it is continuous. This does not necessarily hold for the left and right derivatives, however, they are what is called left-continuous and right-continuous respectively. Here a function is left-continuous, respectively right-continuous, in  $x$  if the limit of the function value increasing, respectively decreasing, towards  $x$  is the same as  $f(x)$ . More formally this gives:

**Definition A.4.** A function  $f : [a, b] \rightarrow \mathbb{R}$  is called left-continuous if, for every  $y \in (a, b]$ , it holds that

$$\lim_{x \uparrow y} f(x) = f(y). \quad (\text{A.3})$$

Furthermore,  $f$  is called right-continuous if, for every  $y \in [a, b)$ , it holds that

$$\lim_{x \downarrow y} f(x) = f(y). \quad (\text{A.4})$$

We note that a function  $f : [a, b] \rightarrow \mathbb{R}$  is continuous in  $x \in (a, b)$  if and only if it is both left-continuous and right-continuous in  $x$ .

For a convex function  $f : [a, b] \rightarrow \mathbb{R}$  the following properties hold (stated without proofs).

**Lemma A.1.** Let  $f : [a, b] \rightarrow \mathbb{R}$  be a convex function. Then the following properties hold:

- » The left derivative  $f^-(x)$  exists on the interval  $(a, b]$ . Furthermore, it is monotonically non-decreasing and left-continuous on this interval and it is finite on  $(a, b)$ , though possibly  $\infty$  at  $b$ .
- » The right derivative  $f^+(x)$  exists on the interval  $[a, b)$ . Furthermore, it is monotonically non-decreasing and right-continuous on this interval and it is finite on  $(a, b)$  though possibly  $-\infty$  on  $a$ .



» For  $a \leq x < y \leq b$  we have that

$$f^+(x) \leq \frac{f(y) - f(x)}{y - x} \leq f^-(y). \quad (\text{A.5})$$

» For any  $x$  with  $a < x < b$  we have that

$$f^-(x) \leq f^+(x). \quad (\text{A.6})$$

### A.1.3 CONVEX OPTIMIZATION

In the field of mathematical optimization, problems are often formulated as

**Problem A.1 (Opt).**

$$\begin{aligned} \min_x \quad & f(x), \\ \text{s.t.} \quad & x \in X. \end{aligned}$$

In Problem *Opt*,  $f$  is called the objective function and  $X$  the feasible set. If  $f$  is a convex function and  $X$  a convex set, then Problem *Opt* is called a convex problem. In this thesis we often consider problems written as

**Problem A.2 (CO).**

$$\begin{aligned} \min_x \quad & f(x), \\ \text{s.t.} \quad & g_i(x) \leq 0 \quad \forall i, \end{aligned}$$

where  $f$  and the  $g_i$ 's are convex. We note that such problems are indeed convex problems, as the feasible set  $X$ , given by

$$X = \{x \mid g_i(x) \leq 0 \quad \forall i\}, \quad (\text{A.7})$$

is a convex set.

For a convex problem a local minimizer is also a global minimizer. More formally this can be stated as (without proof).

**Lemma A.2.** *Let  $f$  be a convex function with convex domain  $D$  and  $x_0 \in D$ . Furthermore, let  $x_0$  be a local minimizer of  $f$ , i.e., there is a  $\delta > 0$  such that for all  $x$  with  $\|x_0 - x\|_2 \leq \delta$  it holds that  $f(x_0) \leq f(x)$ . Then  $x_0$  is also a global minimizer, i.e., for every  $x \in D$  it holds that  $f(x_0) \leq f(x)$ .*

The above lemma implies that to find a global optimum of a convex problem it is sufficient to search for a local optimum. This often means that more efficient solution approaches exist as local minima are generally easier to find. For example, for a differentiable function the local minima can be found by checking the points for which the derivative equals zero.

## A.2 COMPLEXITY

In mathematics and (theoretical) computer science, the complexity of algorithms is often of importance. To study the complexity we are interested in the number of elementary computational steps that the algorithm requires. This complexity is generally defined as the *worst case* number of (elementary) computational steps the algorithm requires as the size of the input asymptotically grows to infinity. The reason that the asymptote is considered is because we are often interested if an algorithm scales to large instances. For such large instances other effects, such as overhead or initialization, have a negligible effect. We note that the introduction given below is very brief and refer the interested reader to a textbook on the subject (e.g., [8]) for more information.

### A.2.1 NOTATION

In order to talk about the (time) complexity of an algorithm, the big-O notation is often used. A function  $f$  is called big-O of  $g$  (denoted  $O(g)$ ) if the growth of  $f$  is asymptotically bounded by that of  $g$ . Formally this gives

**Definition A.5.** For two functions  $f, g : \mathbb{R} \rightarrow \mathbb{R}$  one writes

$$f(n) = O(g(n)) \tag{A.8}$$

if and only if there exist constants  $k$  and  $n_0$  such that

$$|f(n)| \leq k|g(n)| \quad \forall n \geq n_0. \tag{A.9}$$

As an example the function  $2n^2 + 3n + 5$  is  $O(n^2)$  and also  $O(n^3)$ ,  $O(n^4)$ , etc., but not  $O(n)$ . This allows us to formally define the complexity of an algorithm.

**Definition A.6.** Let  $F(n)$  define the worst case number of elementary computational steps required by an algorithm for instances with input size  $n$ . In other words,  $F(n)$  is an upper bound for the number of steps the algorithm requires for an instance of size  $n$ . Then we say that the algorithm has a complexity of  $O(g(n))$  if  $F(n) = O(g(n))$ .

For example, the complexity of binary searching in a sorted array of length  $n$  is  $O(\log n)$  and the complexity of finding the smallest element in an unsorted array of length  $n$  is  $O(n)$ .

Complementary to the big-O notation, the big-Omega notation is often also used. This is formally defined as

**Definition A.7.** For two functions  $f, g : \mathbb{R} \rightarrow \mathbb{R}$  one writes

$$f(n) = \Omega(g(n)) \tag{A.10}$$

if and only if there exist constants  $k$  and  $n_0$  such that

$$|f(n)| \geq k|g(n)| \quad \forall n \geq n_0. \tag{A.11}$$





In other words,  $f(n) = \Omega(g(x))$  implies that  $f$  grows asymptotically at least as fast as  $g$ .

### A.2.2 POLYNOMIAL TIME COMPLEXITY

An algorithm is said to be of polynomial (time) complexity if its worst case required number of (elementary) computational steps is bounded by some polynomial in the input size. More formally this gives

**Definition A.8.** *An algorithm with worst case required number of (elementary) computational steps  $F(n)$ , where  $n$  is the size of the input, is said to have polynomial time complexity if  $F(n) = O(f(n))$  for some fixed polynomial  $f(n)$ .*

Problems for which a solution can be found by an algorithm of polynomial time complexity are often considered ‘nice’. This is because the time to find a solution is assumed not to grow too ‘fast’ in the input size. Note however that an algorithm with a complexity of  $O(n^{100})$  probably still takes far too long to run for any practical size of the input. The class of problems for which an algorithm of polynomial time complexity exists is denoted by  $\mathcal{P}$ . For instance, finding the smallest element of unsorted array is a problem that is in  $\mathcal{P}$ .

Another important class of problems is the set  $\mathcal{NP}$ . This class contains all problems for which it is easy to verify if a solution is valid. More formally, a problem belongs to  $\mathcal{NP}$  if there exists an algorithm of polynomial time complexity that can check if a given solution is valid. We consider as an example the subset sum problem.

**Definition A.9.** *The subset sum problem is the problem that asks, given a set of  $n$  integers  $S = \{s_1, s_2, \dots, s_n\}$  and a goal  $A$ , if there exists a subset  $S' \subset S$  that sums to  $A$ , i.e.  $\sum_{s_i \in S'} s_i = A$ .*

The subset sum problem belongs to the class  $\mathcal{NP}$  because it can be checked in a linear number of steps if a given subset  $S' \subset S$  sums to  $A$ .

We note that any problem that belongs to  $\mathcal{P}$  also belongs to  $\mathcal{NP}$ . A natural question is to ask if the classes are in fact equal (i.e., if every problem in  $\mathcal{NP}$  also belongs to  $\mathcal{P}$ ). This is one of the fundamental problems in current day mathematics and computer science that has puzzled many scientists for the past decades.

Many important and practically relevant problems in  $\mathcal{NP}$  seem easy at first hand. In fact, many people who first encounter the subset sum problem intuitively believe an efficient solution approach must exist. Nonetheless, so far no polynomial time algorithm has been found for subset sum, nor for many other seemingly easy problems in the class  $\mathcal{NP}$ . In fact, in 1971 Cook proved that some problems in this class are so fundamental that, if they are of polynomial complexity, then all problems in  $\mathcal{NP}$  are [34]. These problems are called  $\mathcal{NP}$ -complete problems. The list of these  $\mathcal{NP}$ -complete problems has since that time been extended by a lot of other problems, including the subset sum problem.



The general consensus is that some of the problems in  $\mathcal{NP}$  are so difficult that they cannot be of polynomial complexity, implying that any  $\mathcal{NP}$ -complete problem cannot be of polynomial complexity. Hence, if a problem is shown to be  $\mathcal{NP}$ -complete, researchers often stop searching for a solution approach with polynomial complexity. Instead, they turn to other techniques to try and tackle the problem, such as approximations, heuristics or relaxations of the problem.

As a final note, the classes  $\mathcal{P}$  and  $\mathcal{NP}$  and the notion of  $\mathcal{NP}$ -completeness were originally formulated for decision problems, which are problems formulated as a yes/no question. These classes are less suitable when considering optimization problems. This is because, to determine if a given solution to an optimization problem is optimal, it is often necessary to solve the considered optimization problem. For optimization problems the notion of  $\mathcal{NP}$ -hardness is often used. For this notion the corresponding decision problem to an optimization problem is considered. Such a corresponding decision problem does not ask for an optimal solution but rather if a solution with an objective value of at most (or at least, in case of maximization) a given value exists. The optimization problem is said to be  $\mathcal{NP}$ -hard if the corresponding decision problem is  $\mathcal{NP}$ -complete. Note that, in case the range of the objective value of the optimization problem is small enough, this optimization problem can be solved by a bisection search on its range using its corresponding  $\mathcal{NP}$ -complete decision problem. This implies that finding an efficient solution approach for many  $\mathcal{NP}$ -hard optimization problems is the same as finding an efficient solution approach for (any of) their corresponding decision problems.







# ACRONYMS



C	CC-CV	constant-current constant-voltage
	CHP	combined heat and power
D	DEM	decentralized energy management
	DER	distributed energy resource
	DSM	demand side management
	DSO	distribution system operator
	DVFS	dynamic voltage and frequency scaling
	DW	dish washer
E	EF-Pi	Energy Flexibility Platform and interface
	EM	energy management
	EV	electric vehicle
H	HEMS	home energy management system
	HP	heat pump
	HV	high voltage
	HVAC	heating, ventilation, and air conditioning
I	ICT	information and communication technologies
K	KKT	Karush-Kuhn-Tucker
L	LV	low voltage
M	MPC	model predictive control
	MV	medium voltage
P	PHEV	plug-in hybrid electric vehicle
	PV	photovoltaic
S	SoC	state of charge
T	TE	transactive energy
	ToU	time-of-use
	TSO	transmission system operator
U	UCP	unit commitment problem
V	V2G	vehicle-to-grid
W	WM	washing machine

200



# LIST OF SYMBOLS



$\beta_t$	Base load for time interval $t$
$\Delta\Theta_t^H$	Temperature rise of hottest spot over oil
$\Delta\Theta_t^{H,ini}$	Initial temperature rise of hottest spot
$\Delta\Theta_t^{H,R}$	Ultimate temperature rise of hottest spot at rated load
$\Delta\Theta_t^{H,ult}$	Ultimate temperature rise of hottest spot
$\Delta\Theta_t^{TO,ini}$	Initial temperature rise of oil
$\Delta\Theta_t^{TO,R}$	Ultimate temperature rise of oil at rated load
$\Delta\Theta_t^{TO,ult}$	Ultimate temperature rise of oil
$\Delta\Theta_t^{TO}$	Temperature rise of oil over ambient
$\delta_t$	Length of time interval $t$
$\ell$	Length of a program of a time-shiftable
$\hat{\mathbf{x}}$	Candidate Energy schedule
$\mathbf{p}$	Energy profile of a program of a time-shiftable
$\mathbf{x}$	Scheduled energy profile, potentially aggregated
$\mathbf{x}^{max}$	Upper bound on the scheduled energy use
$\mathbf{x}^{min}$	Lower bound on the scheduled energy usage
$\mathcal{M}$	Set of devices
$\mathcal{R}$	Rated load of transformer
$\mathcal{T}$	Set of time intervals
$\overline{T}_t$	Temperature set point for interval $t$
$\tau^w$	Time parameter for hottest spot temperature change
$\tau^{TO,R}$	Value of $\tau^{TO}$ for rated load
$\tau^{TO}$	Time parameter for oil temperature change
$\Theta_t^H$	Hottest spot temperature
$B_t$	Lower cumulative bound (SoC bound)



$C$	Required charging of EV
$C_t$	Upper cumulative bound (SoC bound)
$D_t$	Allowed deviation from set point for interval $t$
$f$	Objective function
$F^{EQA}$	Equivalent ageing factor for multiple time intervals
$F_t$	Objective of Problem $rdEVC$
$F_t^{AA}$	Ageing factor for time interval $t$
$K_t$	Ratio of load on transformer to rated load
$m$	Device/index denoting the devices
$m_t$	Index of the last (largest) discrete feasible point for $x_t$
$n_1$	Exponent for oil temperature rise based on the type of cooling used
$n_2$	Exponent for hottest spot temperature rise based on type of cooling used
$O_t$	Outdoor temperature
$s_t^j$	Slope of the piece between breakpoints $z_t^{j-1}$ and $z_t^j$
$t$	Time interval/index denoting the time interval
$t^a$	Arrival time interval of a time-shiftable
$t^d$	Deadline for starting of a time-shiftable
$t^s$	Start time of the program of a time-shiftable
$T_t$	Indoor temperature
$T_t(\mathbf{x}^m)$	Indoor temperature for interval $t$
$X$	Set of grid constraints
$X^m$	Set of local (i.e., technical and user) constraints
$y_t^j$	Convex multiplier for interval $t$ and point $z_t^j$
$z_t^j$	Feasible discrete point for $x_t$
$Z_t$	Set of the feasible points for $x_t$
$BC$	Battery charging problem (Problem 5.1)
$dBC$	Discrete battery charging
$dEVC$	Discrete EV charging problem
$EVC$	EV charging problem

<i>HVACS</i>	HVAC scheduling problem
<i>rdEVC</i>	Relaxed discrete EV charging problem
<i>rHVACS</i>	Problem <i>HVACS</i> with the cumulative bounds dropped except for the last time interval







# BIBLIOGRAPHY



- [1] Aimms software. [online] [www.aimms.com](http://www.aimms.com) last accessed on 25-01-2017. (Cited on page 171).
- [2] wunderground. [online] <http://www.wunderground.com/> last accessed on 25-01-2017. (Cited on page 135).
- [3] IEEE guide for loading mineral-oil-immersed overhead and pad-mounted distribution transformers rated 500 kVA and less with 65 degrees C or 55 degrees C average winding rise. *ANSI/IEEE Std C57.91-1981*, 1981. doi: 10.1109/IEEESTD.1981.81014. (Cited on page 170).
- [4] IEEE guide for loading mineral-oil-immersed transformers and step-voltage regulators - Redline. *IEEE Std C57.91-2011 (Revision of IEEE Std C57.91-1995) - Redline*, March 2012. (Cited on pages 161, 162, 163, 165, and 170).
- [5] A. Afram and F. Janabi-Sharifi. Theory and applications of HVAC control systems – a review of model predictive control (MPC). *Building and Environment*, 72:343 – 355, 2014. ISSN 0360-1323. doi: 10.1016/j.buildenv.2013.11.016. (Cited on pages 37 and 132).
- [6] M. R. Alam, M. St-Hilaire, and T. Kunz. Computational methods for residential energy cost optimization in smart grids: a survey. *ACM Comput. Surv.*, 49(1):2:1–2:34, Apr. 2016. ISSN 0360-0300. doi: 10.1145/2897165. (Cited on pages 27 and 34).
- [7] Arbeitsgruppe Erneuerbare Energien - Statistik. Zeitreihen zur entwicklung der erneuerbaren energien in deutschland. Technical report, Bundesministerium für Wirtschaft und Energie, 2016. [Online] [http://www.erneuerbare-energien.de/EE/Redaktion/DE/Downloads/zeitreihen-zur-entwicklung-der-erneuerbaren-energien-in-deutschland-1990-2015.pdf?\\_\\_blob=publicationFile&v=7](http://www.erneuerbare-energien.de/EE/Redaktion/DE/Downloads/zeitreihen-zur-entwicklung-der-erneuerbaren-energien-in-deutschland-1990-2015.pdf?__blob=publicationFile&v=7) last accessed 25-01-2017. (Cited on page 2).
- [8] S. Arora and B. Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009. (Cited on pages 189 and 194).
- [9] A. Aswani, N. Master, J. Taneja, D. Culler, and C. Tomlin. Reducing transient and steady state electricity consumption in HVAC using learning-based model-predictive control. *Proceedings of the IEEE*, 100(1):240–253, Jan 2012. ISSN 0018-9219. doi: 10.1109/JPROC.2011.2161242. (Cited on pages 133, 134, and 135).
- [10] I. Atzeni, L. G. Ordóñez, G. Scutari, D. P. Palomar, and J. R. Fonollosa. Demand-side management via distributed energy generation and storage optimization. *IEEE Transactions on Smart Grid*, 4(2):866–876, June 2013. ISSN 1949-3053. doi: 10.1109/TSG.2012.2206060. (Cited on page 34).



- [11] M. Babar, P. H. Nguyen, V. Cuk, and I. G. Kamphuis. The development of demand elasticity model for demand response in the retail market environment. In *IEEE PowerTech 2015*, pages 3231–3236, June 2015. doi: 10.1109/PTC.2015.7232789. (Cited on page 34).
- [12] F. Bach. Convex analysis and optimization with submodular functions: a tutorial. *arXiv preprint arXiv:1010.4207*, pages 1–41, 2010. (Cited on page 86).
- [13] P. Bacher and H. Madsen. Identifying suitable models for the heat dynamics of buildings. *Energy and Buildings*, 43(7):1511 – 1522, 2011. ISSN 0378-7788. doi: <http://doi.org/10.1016/j.enbuild.2011.02.005>. (Cited on page 154).
- [14] V. Bakker. *Triana: a control strategy for Smart Grids: Forecasting, planning & real-time control*. PhD thesis, University of Twente, Enschede, January 2012. (Cited on page 35).
- [15] A. Barbato and A. Capone. Optimization models and methods for demand-side management of residential users: a survey. *Energies*, 7(9):5787, 2014. ISSN 1996-1073. doi: 10.3390/en7095787. (Cited on page 34).
- [16] S. Bashash and H. K. Fathy. Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads. In *Proceedings of the 2011 American Control Conference*, pages 4546–4553, June 2011. doi: 10.1109/ACC.2011.5990939. (Cited on page 132).
- [17] Bloomberg Technology. Battery cost plunge seen changing automakers most in 100 years, 2016. [Online] <https://www.bloomberg.com/news/articles/2016-10-11/battery-cost-plunge-seen-changing-automakers-most-in-100-years> last accessed on 26-01-2017. (Cited on page 105).
- [18] M. Blum, R. W. Floyd, V. Pratt, R. L. Rivest, and R. E. Tarjan. Time bounds for selection. *Journal of Computer and System Sciences*, 7(4):448 – 461, 1973. ISSN 0022-0000. (Cited on pages 80, 93, 95, and 186).
- [19] BMWI. Moderne verteilernetze für deutschland (verteilernetzstudie. Technical report, BMWI, 2014. [Online] <http://www.bmwi.de/Redaktion/DE/Publikationen/Studien/verteilernetzstudie.html> last accessed on 30-01-2017. (Cited on page 161).
- [20] M. G. C. Bosman. *Planning in Smart Grids*. PhD thesis, University of Twente, Enschede, July 2012. (Cited on pages 34 and 35).
- [21] 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, 2012. ISSN 0377-2217. doi: 10.1016/j.ejor.2011.07.033. (Cited on page 25).
- [22] M. G. C. Bosman, A. Molderink, V. Bakker, G. Smit, and J. L. Hurink. Multilevel unit commitment in smart grids. In *Proceedings of the 1st International Conference on Operations Research and Enterprise Systems*, pages 361–370, 2012. ISBN 978-989-8425-97-3. doi: 10.5220/0003716903610370. (Cited on pages 21 and 23).



- [23] K. M. Bretthauer and B. Shetty. Quadratic resource allocation with generalized upper bounds. *Operations Research Letters*, 20(2):51 – 57, 1997. ISSN 0167-6377. (Cited on page 85).
- [24] K. M. Bretthauer and B. Shetty. A pegging algorithm for the nonlinear resource allocation problem. *Computers & Operations Research*, 29(5):505 – 527, 2002. ISSN 0305-0548. doi: 10.1016/S0305-0548(00)00089-7. (Cited on page 85).
- [25] Centraal Bureau voor de Statistiek. Hernieuwbare energie in nederland 2015. Technical report, Centraal Bureau voor de Statistiek, 2016. [Online] [https://www.cbs.nl/-/media/\\_pdf/2016/39/hernieuwbare-energie-in-nederland-2015.pdf](https://www.cbs.nl/-/media/_pdf/2016/39/hernieuwbare-energie-in-nederland-2015.pdf) last accessed on 25-01-2017. (Cited on page 2).
- [26] H. Chen, T. N. Cong, W. Yang, C. Tan, Y. Li, and Y. Ding. Progress in electrical energy storage system: a critical review. *Progress in Natural Science*, 19(3):291 – 312, 2009. ISSN 1002-0071. doi: 10.1016/j.pnsc.2008.07.014. (Cited on page 106).
- [27] F. Claessen, B. Claessens, M. Hommelberg, A. Molderink, V. Bakker, H. Toersche, and M. van den Broek. Comparative analysis of tertiary control systems for smart grids using the flex street model. *Renewable Energy*, 69:260 – 270, 2014. ISSN 0960-1481. doi: 10.1016/j.renene.2014.03.037. (Cited on page 38).
- [28] B. J. Claessens, S. Vandael, F. Ruelens, and M. Hommelberg. Self-learning demand side management for a heterogeneous cluster of devices with binary control actions. In *IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe) 2012*, pages 075:1–075:8, Oct 2012. doi: 10.1109/ISGTEurope.2012.6465679. (Cited on pages 27, 37, and 38).
- [29] B. J. Claessens, S. Vandael, F. Ruelens, K. D. Craemer, and B. Beusen. Peak shaving of a heterogeneous cluster of residential flexibility carriers using reinforcement learning. In *IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe) 2013*, pages 116–120, Oct 2013. doi: 10.1109/ISGTEurope.2013.6695254. (Cited on pages 37, 51, and 72).
- [30] K. Clement-Nyns, E. Haesen, and J. Driesen. The impact of charging plug-in hybrid electric vehicles on a residential distribution grid. *IEEE Transactions on Power Systems*, 25(1):371–380, Feb 2010. ISSN 0885-8950. doi: 10.1109/TPWRS.2009.2036481. (Cited on pages 24 and 69).
- [31] W. J. Cole, K. M. Powell, E. T. Hale, and T. F. Edgar. Reduced-order residential home modeling for model predictive control. *Energy and Buildings*, 74:69 – 77, 2014. ISSN 0378-7788. doi: 10.1016/j.enbuild.2014.01.033. (Cited on pages 133, 134, and 150).
- [32] W. J. Cole, J. D. Rhodes, W. Gorman, K. X. Perez, M. E. Webber, and T. F. Edgar. Community-scale residential air conditioning control for effective grid management. *Applied Energy*, 130:428 – 436, 2014. ISSN 0306-2619. doi: 10.1016/j.apenergy.2014.05.067. (Cited on page 132).
- [33] J. Cook, N. Oreskes, P. T. Doran, W. R. Anderegg, B. Verheggen, E. W. Maibach, J. S. Carlton, S. Lewandowsky, A. G. Skuce, S. A. Green, D. Nuccitelli, P. Jacobs, M. Richardson, B. Winkler, R. Painting, and K. Rice. Consensus on consensus: a synthesis of consensus estimates on human-caused global warming. *Environmental Research Letters*, 11(4):048002, 2016. (Cited on page 3).



- [34] S. A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the third annual ACM symposium on Theory of computing*, pages 151–158. ACM, 1971. (Cited on page 195).
- [35] Council of European Energy Regulators. CEER benchmarking report 5.2 on the continuity of electricity supply data update. Technical report, Council of European Energy Regulators, 2015. [Online] [http://www.ceer.eu/portal/page/portal/EER\\_HOME/EER\\_PUBLICATIONS/CEER\\_PAPERS/Electricity/Tab4/](http://www.ceer.eu/portal/page/portal/EER_HOME/EER_PUBLICATIONS/CEER_PAPERS/Electricity/Tab4/) last accessed on 25-01-2017. (Cited on pages 16 and 17).
- [36] K. D. Craemer, S. Vandael, B. Claessens, and G. Deconinck. An event-driven dual coordination mechanism for demand side management of phev. *IEEE Transactions on Smart Grid*, 5(2):751–760, March 2014. ISSN 1949-3053. doi: 10.1109/TSG.2013.2272197. (Cited on pages 37 and 187).
- [37] J. M. Danskin. The theory of max-min, with applications. *SIAM Journal on Applied Mathematics*, 14(4):641–664, 1966. (Cited on page 75).
- [38] G. J. H. de Goeijen, G. J. M. Smit, and J. L. Hurink. An integer linear programming model for an Ecovat buffer. *Energies*, 9(8):592, 2016. ISSN 1996-1073. doi: 10.3390/en9080592. (Cited on page 106).
- [39] R. Deng, Z. Yang, M. Y. Chow, and J. Chen. A survey on demand response in smart grids: mathematical models and approaches. *IEEE Transactions on Industrial Informatics*, 11(3):570–582, June 2015. ISSN 1551-3203. doi: 10.1109/TII.2015.2414719. (Cited on pages 34, 52, and 161).
- [40] Dr. Claw. De tweede hint, 2017. [Online] <https://www.google.nl/search?q=de+tweede+hint+is+dat+de+sleutels+prominent+aanwezig+zijn+beginnenden+op+pagina+1>. (Not cited).
- [41] B. Dunn, H. Kamath, and J.-M. Tarascon. Electrical energy storage for the grid: a battery of choices. *Science*, 334(6058):928–935, 2011. ISSN 0036-8075. doi: 10.1126/science.1212741. (Cited on page 106).
- [42] J.-P. Dussault, J. A. Ferland, and B. Lemaire. Convex quadratic programming with one constraint and bounded variables. *Mathematical Programming*, 36(1):90–104, 1986. ISSN 0025-5610. (Cited on page 75).
- [43] S. Dutta and T. J. Overbye. Optimal storage scheduling for minimizing schedule deviations considering variability of generated wind power. In *Power Systems Conference and Exposition (PSCE), 2011 IEEE/PES*, pages 598–604, March 2011. doi: 10.1109/P-SCE.2011.5772521. (Cited on page 107).
- [44] Enexis. Annual report 2015. Technical report, Enexis, 2016. [Online] <http://annualreport.enexis.nl/2015/annualreport/userfiles/pdf/Enexis-Annual-Report-2015.pdf> last accessed on 25-01-2017. (Cited on page 160).
- [45] V. L. Erickson and A. E. Cerpa. Occupancy based demand response HVAC control strategy. In *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, BuildSys '10, pages 7–12, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0458-0. doi: 10.1145/1878431.1878434. (Cited on pages 135 and 150).



- [46] T. Ericson. Direct load control of residential water heaters. *Energy Policy*, 37(9):3502 – 3512, 2009. ISSN 0301-4215. doi: 10.1016/j.enpol.2009.03.063. (Cited on page 35).
- [47] A. T. Ernst and G. Singh. Taming wind energy with battery storage. In *Operations Research Proceedings 2007: Selected Papers of the Annual International Conference of the German Operations Research Society (GOR) Saarbrücken, September 5–7, 2007*, pages 199–204, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg. ISBN 978-3-540-77903-2. doi: 10.1007/978-3-540-77903-2\_31. (Cited on page 107).
- [48] European Commission. European alternative fuels observatory, 2016. URL <http://www.eafo.eu>. Last Accessed on 25-01-2017. (Cited on page 68).
- [49] European Commission. 2020 climate & energy package, 2016. URL [http://ec.europa.eu/clima/policies/strategies/2020\\_en](http://ec.europa.eu/clima/policies/strategies/2020_en). Last Accessed on 25-01-2017. (Cited on page 3).
- [50] European Commission. 2030 energy strategy, 2016. URL <https://ec.europa.eu/energy/en/topics/energy-strategy/2030-energy-strategy>. Last Accessed on 25-01-2017. (Cited on page 3).
- [51] European Commission. Eurostat your key tot european statistics, 2016. URL <http://ec.europa.eu/eurostat/web/energy/data/main-tables>. Last Accessed on 25-01-2017. (Cited on pages 16 and 42).
- [52] European Commission Joint Research Centre. Smart metering deployment in the european union, 2016. [Online] <http://ses.jrc.ec.europa.eu/smart-metering-deployment-european-union> last accessed on 25-01-2017. (Cited on page 24).
- [53] A. Evans, V. Strezov, and T. J. Evans. Assessment of utility energy storage options for increased renewable energy penetration. *Renewable and Sustainable Energy Reviews*, 16(6):4141 – 4147, 2012. ISSN 1364-0321. doi: 10.1016/j.rser.2012.03.048. (Cited on pages 106 and 125).
- [54] H. L. Ferreira, R. Garde, G. Fulli, W. Kling, and J. P. Lopes. Characterisation of electrical energy storage technologies. *Energy*, 53:288 – 298, 2013. ISSN 0360-5442. doi: 10.1016/j.energy.2013.02.037. (Cited on page 106).
- [55] Flexiblepower Alliance Network. EF-Pi framework, 2016. URL <http://flexible-energy.eu/>. Last Accessed on 25-01-2017. (Cited on page 31).
- [56] L. Gan, U. Topcu, and S. H. Low. Optimal decentralized protocol for electric vehicle charging. *IEEE Transactions on Power Systems*, 28(2):940–951, May 2013. ISSN 0885-8950. doi: 10.1109/TPWRS.2012.2210288. (Cited on page 71).
- [57] M. E. T. Gerards and J. L. Hurink. Robust peak-shaving for a neighborhood with electric vehicles. *Energies*, 9(8):594, 2016. ISSN 1996-1073. doi: 10.3390/en9080594. (Cited on page 187).
- [58] M. E. T. Gerards, J. L. Hurink, and P. K. F. Hölzenspies. A survey of offline algorithms for energy minimization under deadline constraints. *Journal of Scheduling*, 19(1):3–19, 2016. ISSN 1099-1425. doi: 10.1007/s10951-015-0463-8. (Cited on page 74).



- [59] J. Gönsch and M. Hassler. Sell or store? an ADP approach to marketing renewable energy. *OR Spectrum*, 38(3):633–660, 2016. ISSN 1436-6304. doi: 10.1007/s00291-016-0439-x. URL 10.1007/s00291-016-0439-x. (Cited on page 106).
- [60] GridWise Architecture Council. Gridwise transactive energy framework, january 2015. URL [http://www.gridwiseac.org/pdfs/te\\_framework\\_report\\_pnnl-22946.pdf](http://www.gridwiseac.org/pdfs/te_framework_report_pnnl-22946.pdf). Last Accessed on 25-01-2017. (Cited on page 26).
- [61] GTM Research. U.S. energy storage monitor: Q4 2016 executive summary. Technical report, Wood Mackenzie Business, 2016. [Online] <https://www.greentechmedia.com/research/subscription/u.s.-energy-storage-monitor> last accessed on 25-01-2017. (Cited on page 104).
- [62] M. Hall. EU plans renewables-friendly grid expansion to 2050, 2013. URL [http://www.pv-magazine.com/news/details/beitrag/eu-plans-renewables-friendly-grid-expansion-to-2050\\_100010988/#axzz4MD1euAHF](http://www.pv-magazine.com/news/details/beitrag/eu-plans-renewables-friendly-grid-expansion-to-2050_100010988/#axzz4MD1euAHF). last accessed on 25-01-2017. (Cited on page 24).
- [63] R. E. Hebner, F. M. Uriatre, A. Kwasinski, A. L. Gattozzi, H. B. Estes, A. Anwar, P. Cairoli, R. A. Dougal, X. Feng, H.-M. Chou, L. J. Thomas, M. Pipattanasomporn, S. Rahman, F. Katiraei, M. Steurer, M. O. Faruque, M. A. Rios, G. A. Ramos, M. J. Mousavi, and T. J. McCoy. Technical cross-fertilization between terrestrial microgrids and ship power systems. *Journal of Modern Power Systems and Clean Energy*, 4(2):161–179, 2016. ISSN 2196-5420. doi: 10.1007/s40565-015-0108-0. (Cited on page 24).
- [64] D. Hochbaum and S.-P. Hong. About strongly polynomial time algorithms for quadratic optimization over submodular constraints. *Mathematical Programming*, 69(1-3):269–309, 1995. ISSN 0025-5610. (Cited on pages 80, 86, 93, and 114).
- [65] M. Hommelberg and B. Claessens. System for distributing electrical energy over a cluster of electrical devices, method for distributing electrical energy over a cluster of electrical devices of such a system and controller for use in such a system, June 27 2013. URL <https://www.google.com/patents/WO2013093119A1?c1=en>. WO Patent App. PCT/EP2012/076,882. (Cited on page 37).
- [66] G. Hoogsteen, A. Molderink, J. L. Hurink, and G. J. M. Smit. Managing energy in time and space in smart grids using triana. In *IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe) 2014*, pages 1408–1413, Oct 2014. doi: 10.1109/ISGTEurope.2014.7028973. (Cited on pages 34, 36, 98, 99, and 186).
- [67] G. Hoogsteen, A. Molderink, J. L. Hurink, G. J. M. Smit, F. Schuring, and B. Kootstra. Impact of peak electricity demand in distribution grids: a stress test. In *IEEE PowerTech 2015*, pages 1028–1034, June 2015. doi: 10.1109/PTC.2015.7232412. (Cited on pages 24, 28, 36, 54, 57, 123, and 172).
- [68] G. Hoogsteen, A. Molderink, J. L. Hurink, and G. J. M. Smit. Generation of flexible domestic load profiles to evaluate demand side management approaches. In *IEEE International Energy Conference (ENERGYCON) 2016*, pages 13–18, April 2016. doi: 10.1109/ENERGYCON.2016.7513873. (Cited on page 54).



- [69] W. Huang and Y. Wang. An optimal speed control scheme supported by media servers for low-power multimedia applications. *Multimedia Systems*, 15(2):113–124, 2009. ISSN 0942-4962. (Cited on page 111).
- [70] M. Humayun, M. Z. Degefa, A. Safdarian, and M. Lehtonen. Utilization improvement of transformers using demand response. *IEEE Transactions on Power Delivery*, 30(1): 202–210, Feb 2015. ISSN 0885-8977. doi: 10.1109/TPWRD.2014.2325610. (Cited on page 162).
- [71] L. M. Hvattum, I. Norstad, K. Fagerholt, and G. Laporte. Analysis of an exact algorithm for the vessel speed optimization problem. *Networks*, 62(2):132–135, 2013. ISSN 1097-0037. (Cited on pages 74 and 111).
- [72] S. Iacovella, F. Ruelens, P. Vingerhoets, B. Claessens, and G. Deconinck. Cluster control of heterogeneous thermostatically controlled loads using tracer devices. *IEEE Transactions on Smart Grid*, 8(2):528–536, March 2017. ISSN 1949-3053. doi: 10.1109/TSG.2015.2483506. (Cited on page 37).
- [73] P. S. Inc. Pecan street inc. dataport 2016. URL <https://dataport.pecanstreet.org>. last accessed on 25-01-2017. (Cited on pages 132 and 135).
- [74] International Energy Agency. Pvps annual report 2015. Technical report, Photovoltaic power systems technology collaboration programme, 2016. [Online] [http://www.iea-pvps.org/index.php?id=6&eID=dam\\_frontend\\_push&docID=3195](http://www.iea-pvps.org/index.php?id=6&eID=dam_frontend_push&docID=3195) last accessed on 25-01-2017. (Cited on pages 5 and 6).
- [75] International Energy Agency. What drives energy consumption in households?, 2016. [Online] <http://www.iea.org/newsroom/graphics/2014-07-22-what-drives-energy-consumption-in-households.html> last accessed on 25-01-2017. (Cited on page 130).
- [76] International Energy Agency, Clean Energy Ministerial, Electric Vehicles Initiative. Global ev outlook 2016. Technical report, Electric Vehicles Initiative, 2016. [Online] [https://www.iea.org/publications/freepublications/publication/Global\\_EV\\_Outlook\\_2016.pdf](https://www.iea.org/publications/freepublications/publication/Global_EV_Outlook_2016.pdf), last accessed on 25-01-2017. (Cited on page 68).
- [77] J. P. Iria, F. J. Soares, and R. J. Bessa. Optimized demand response bidding in the wholesale market under scenarios of prices and temperatures. In *IEEE PowerTech 2015*, pages 1121–1127, June 2015. doi: 10.1109/PTC.2015.7232428. (Cited on page 34).
- [78] J. Jargstorf, K. Vanthournout, T. D. Rybel, and D. V. Hertem. Effect of demand response on transformer lifetime expectation. In *IEEE PES Innovative Smart Grid Technologies (ISGT Europe) 2012*, pages 2011:1–2011:8, Oct 2012. doi: 10.1109/ISGTEurope.2012.6465805. (Cited on page 162).
- [79] H. G. J. Kamp. Tijdpad uitrol slimme meter en dynamische leveringstarieven. Technical report, Ministerie van Economische Zaken, 2016. [Online] <https://www.rijksoverheid.nl/documenten/kamerstukken/2016/03/04/kamerbrief-over-tijdpad-uitrol-slimme-meter-en-dynamische-leveringstarieven> last accessed on 25-01-2017. (Cited on page 24).



- [80] W. Karush. Minima of functions of several variables with inequalities as side conditions, 1939. Master Thesis, Department of Mathematics, University of Chicago, Chicago, Illinois. (Cited on page 75).
- [81] M. Kesler, M. C. Kisacikoglu, and L. M. Tolbert. Vehicle-to-grid reactive power operation using plug-in electric vehicle bidirectional offboard charger. *IEEE Transactions on Industrial Electronics*, 61(12):6778–6784, Dec 2014. ISSN 0278-0046. (Cited on pages 85, 88, and 96).
- [82] P. Kleinmann and R. Schultz. A simple procedure for optimal load dispatch using parametric programming. *Zeitschrift für Operations Research*, 34(3):219–229, 1990. ISSN 1432-5217. doi: 10.1007/BF01415985. (Cited on page 74).
- [83] J. Kok. *The PowerMatcher smart coordination for the smart electricity grid*. PhD thesis, July 2013. (Cited on page 36).
- [84] M. Koller, T. Borsche, A. Ulbig, and G. Andersson. Defining a degradation cost function for optimal control of a battery energy storage system. In *IEEE PowerTech 2013*, pages 1446–1451, June 2013. doi: 10.1109/PTC.2013.6652329. (Cited on page 106).
- [85] H. W. Kuhn and A. W. Tucker. Nonlinear programming. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 481–492, Berkeley, Calif., 1951. University of California Press. (Cited on page 75).
- [86] W.-C. Kwon and T. Kim. Optimal voltage allocation techniques for dynamically variable voltage processors. *ACM Trans. Embed. Comput. Syst.*, 4(1):211–230, Feb. 2005. ISSN 1539-9087. (Cited on page 88).
- [87] M. Li, A. C. Yao, and F. F. Yao. Discrete and continuous min-energy schedules for variable voltage processors. *Proceedings of the National Academy of Sciences of the United States of America*, 103(11):3983–3987, 2006. (Cited on pages 74 and 88).
- [88] B. Lydic. How California’s Rule 21 inverter requirements expand grid capacity, limit energy (revenue) generation, 2016. [Online] <http://solarbuildermag.com/featured/california-rule-21-inverters-explained/> last accessed on 25-01-2017. (Cited on page 31).
- [89] J. Marcos, I. de la Parra, M. García, and L. Marroyo. Control strategies to smooth short-term power fluctuations in large photovoltaic plants using battery storage systems. *Energies*, 7(10):6593, 2014. ISSN 1996-1073. doi: 10.3390/en7106593. (Cited on page 107).
- [90] H. Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952. ISSN 1540-6261. (Cited on page 75).
- [91] K. McKenna and A. Keane. Discrete elastic residential load response under variable pricing schemes. In *IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe) 2014*, pages 210–215, Oct 2014. doi: 10.1109/ISGTEurope.2014.7028769. (Cited on page 34).
- [92] MetalsTech Ltd. Market outlook, 2016. [Online] <http://www.metalsstech.net/market-outlook/> last accessed on 25-01-2017. (Cited on page 104).





- [93] R. Moghe, F. Kreikebaum, J. E. Hernandez, R. P. Kandula, and D. Divan. Mitigating distribution transformer lifetime degradation caused by grid-enabled vehicle (GEV) charging. In *Energy Conversion Congress and Exposition (ECCE), 2011 IEEE*, pages 835–842, Sept 2011. doi: 10.1109/ECCE.2011.6063857. (Cited on page 163).
- [94] A. H. Mohsenian-Rad, V. W. S. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia. Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid. *IEEE Transactions on Smart Grid*, 1(3): 320–331, Dec 2010. ISSN 1949-3053. doi: 10.1109/TSG.2010.2089069. (Cited on page 34).
- [95] A. Molderink. *On the three-step control methodology for Smart Grids*. PhD thesis, University of Twente, May 2011. (Cited on pages 34 and 35).
- [96] 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 *IEEE PES Innovative Smart Grid Technologies (ISGT) 2010*, pages 289–296, Jan 2010. doi: 10.1109/ISGT.2010.5434731. (Cited on page 35).
- [97] Y. Mou, H. Xing, Z. Lin, and M. Fu. Decentralized optimal demand-side management for PHEV charging in a smart grid. *IEEE Transactions on Smart Grid*, 6(2):726–736, March 2015. ISSN 1949-3053. (Cited on pages 81 and 101).
- [98] O. Mégel, J. L. Mathieu, and G. Andersson. Maximizing the potential of energy storage to provide fast frequency control. In *IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe) 2013*, pages 724–728, Oct 2013. doi: 10.1109/ISGTEurope.2013.6695380. (Cited on page 107).
- [99] NEN-EN 50160:2010. Voltage characteristics of electricity supplied by public distribution network, 2010. (Cited on pages 98 and 186).
- [100] I. Norstad, K. Fagerholt, and G. Laporte. Tramp ship routing and scheduling with speed optimization. *Transportation Research Part C: Emerging Technologies*, 19(5): 853 – 865, 2011. ISSN 0968-090X. (Cited on page 74).
- [101] J. Nutma. More comprehensive demand side management by the integration of the PowerMatcher and Triana, 2016. Master Thesis, Department of electrical engineering, mathematics, and computer science, University of Twente, Enschede, The Netherlands. (Cited on page 37).
- [102] S. Nykamp. PhD thesis, University of Twente, Enschede, October 2013. (Cited on pages 21, 26, 29, 107, and 161).
- [103] S. Nykamp, A. Molderink, J. L. Hurink, and G. J. Smit. Statistics for PV, wind and biomass generators and their impact on distribution grid planning. *Energy*, 45(1): 924 – 932, 2012. ISSN 0360-5442. The 24th International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy, ECOS 2011. (Cited on pages 3, 23, and 24).
- [104] S. Nykamp, M. G. C. Bosman, A. Molderink, J. L. Hurink, and G. J. M. Smit. Value of storage in distribution grids - competition or cooperation of stakeholders? *Smart Grid, IEEE Transactions on*, 4(3):1361–1370, Sept 2013. ISSN 1949-3053. doi: 10.1109/TSG.2013.2254730. (Cited on page 161).



- [105] S. Nykamp, A. Molderink, J. L. Hurink, and G. J. M. Smit. Storage operation for peak shaving of distributed pv and wind generation. In *Innovative Smart Grid Technologies (ISGT), 2013 IEEE PES*, pages 13–18, Feb 2013. doi: 10.1109/ISGT.2013.6497786. (Cited on page 106).
- [106] S. Nykamp, T. Rott, N. Dettke, and S. Kueppers. The project "ElChe" Wetrtingen: storage as an alternative to grid reinforcements - experiences, benefits and challenges from a DSO point of. In *International ETG Congress 2015; Die Energiewende - Blueprints for the new energy age; Proceedings of*, pages 3.1:1–3.1:6, Nov 2015. (Cited on pages 5, 23, 26, 28, 105, and 107).
- [107] O. Ongkiehong and Senternovem EOS. Electricity grids, description of the state under the dutch energy research program eos, 2006. [Online] <https://www.rvo.nl/sites/default/files/bijlagen/StatusReportElectricitygridsEOS.pdf> last accessed on 25-01-2017. (Cited on page 19).
- [108] M. Patriksson. A survey on the continuous nonlinear resource allocation problem. *European Journal of Operational Research*, 185(1):1–46, 2008. ISSN 0377-2217. (Cited on page 75).
- [109] M. Patriksson and C. Strömberg. Algorithms for the continuous nonlinear resource allocation problem: new implementations and numerical studies. *European Journal of Operational Research*, 243(3):703–722, 2015. ISSN 0377-2217. (Cited on page 75).
- [110] A. Poullikkas, S. Papadouris, G. Kourtis, and I. Hadjipaschalis. Storage solutions for power quality problems in Cyprus electricity distribution network. *AIMS Energy*, 2(20140101):1–17, 2014. ISSN 2333-8334. doi: 10.3934/energy.2014.1.1. (Cited on page 106).
- [111] PPL. PPL: a diverse pure-play electric utility company, 2016. [Online] [https://www.sec.gov/Archives/edgar/data/922224/000092222416000137/exhibit99\\_1.htm](https://www.sec.gov/Archives/edgar/data/922224/000092222416000137/exhibit99_1.htm) last accessed on 25-01-2017. (Cited on page 160).
- [112] D. Pudjianto, P. Djapic, M. Aunedi, C. K. Gan, G. Strbac, S. Huang, and D. Infield. Smart control for minimizing distribution network reinforcement cost due to electrification. *Energy Policy*, 52:76 – 84, 2013. ISSN 0301-4215. doi: 10.1016/j.enpol.2012.05.021. Special Section: Transition Pathways to a Low Carbon Economy. (Cited on page 24).
- [113] A. Rahimi, M. Zarghami, M. Vaziri, and S. Vadhva. A simple and effective approach for peak load shaving using battery storage systems. In *North American Power Symposium (NAPS), 2013*, pages 34–38, Sept 2013. doi: 10.1109/NAPS.2013.6666824. (Cited on page 106).
- [114] R. T. Rockafellar. *Convex analysis*. Princeton university press, 1970. (Cited on pages 189 and 191).
- [115] N. Rotering and M. Ilic. Optimal charge control of plug-in hybrid electric vehicles in deregulated electricity markets. *IEEE Transactions on Power Systems*, 26(3):1021–1029, Aug 2011. ISSN 0885-8950. doi: 10.1109/TPWRS.2010.2086083. (Cited on pages 70 and 71).



- [116] M. Rowe, T. Yunusov, S. Haben, W. Holderbaum, and B. Potter. The real-time optimisation of DNO owned storage devices on the LV network for peak reduction. *Energies*, 7(6):3537, 2014. ISSN 1996-1073. doi: 10.3390/en7063537. (Cited on page 106).
- [117] M. H. H. Schoot Uiterkamp. Robust planning of electric vehicle charging, 2016. Master Thesis, Department of Electrical Engineering, Mathematics and Computer Science, University of Twente, Enschede, The Netherlands. (Cited on page 69).
- [118] F. C. Schweppe, R. D. Tabors, J. L. Kirtley, H. R. Outhred, F. H. Pickel, and A. J. Cox. Homeostatic utility control. *IEEE Transactions on Power Apparatus and Systems*, PAS-99(3):1151–1163, May 1980. ISSN 0018-9510. doi: 10.1109/TPAS.1980.319745. (Cited on pages 33 and 37).
- [119] F. Shahnia, M. T. Wishart, A. Ghosh, G. Ledwich, and F. Zare. Smart demand side management of low-voltage distribution networks using multi-objective decision making. *IET Generation, Transmission Distribution*, 6(10):968–1000, October 2012. ISSN 1751-8687. doi: 10.1049/iet-gtd.2011.0781. (Cited on page 35).
- [120] P. Siano. Demand response and smart grids—a survey. *Renewable and Sustainable Energy Reviews*, 30:461 – 478, 2014. ISSN 1364-0321. doi: 10.1016/j.rser.2013.10.022. (Cited on pages 4, 34, and 161).
- [121] R. Stamminger. Synergy potential of smart appliances. Technical report, Rheinische Freidrich-Wilhelms-Universität Bonn, 2008. [Online] <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.378.4289&rep=rep1&type=pdf> last accessed on 21-01-2017. (Cited on pages 54 and 123).
- [122] G. Strbac. Demand side management: Benefits and challenges. *Energy Policy*, 36(12): 4419 – 4426, 2008. ISSN 0301-4215. (Cited on pages 4 and 161).
- [123] M. Swierczynski, D. I. Stroe, A. I. Stan, R. Teodorescu, R. Lærke, and P. C. Kjær. Field tests experience from 1.6mw/400kwh li-ion battery energy storage system providing primary frequency regulation service. In *IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe) 2013*, pages 225–229, Oct 2013. doi: 10.1109/ISGT-Europe.2013.6695277. (Cited on page 107).
- [124] W. Tang, S. Bi, and Y. J. . Zhang. Online coordinated charging decision algorithm for electric vehicles without future information. *IEEE Transactions on Smart Grid*, 5 (6):2810–2824, Nov 2014. ISSN 1949-3053. (Cited on page 74).
- [125] E. Telaretti, M. Ippolito, and L. Dusonchet. A simple operating strategy of small-scale battery energy storages for energy arbitrage under dynamic pricing tariffs. *Energies*, 9(1):12, 2016. ISSN 1996-1073. doi: 10.3390/en9010012. (Cited on pages 34 and 106).
- [126] H. A. Toersche. *Effective and efficient coordination of flexibility in smart grids*. PhD thesis, University of Twente, oct 2016. (Cited on pages 35, 36, 51, 53, 65, and 186).
- [127] H. A. Toersche, A. Molderink, J. L. Hurink, and G. J. M. Smit. Cascaded column generation for scalable predictive demand side management. In *IEEE International Energy Conference (ENERGYCON) 2014*, pages 1228–1235, May 2014. doi: 10.1109/ENERGYCON.2014.6850580. (Cited on page 35).



- [128] A. Toliyat, A. Kwasinski, and F. M. Uriarte. Effects of high penetration levels of residential photovoltaic generation: observations from field data. In *2012 International Conference on Renewable Energy Research and Applications (ICRERA)*, pages 37–42, Nov 2012. doi: 10.1109/ICRERA.2012.6477269. (Cited on page 3).
- [129] R. Tyagi and J. W. Black. Emergency demand response for distribution system contingencies. In *IEEE PES Transmission and Distribution Conference and Exposition 2010*, pages 922–925, April 2010. doi: 10.1109/TDC.2010.5484598. (Cited on pages 25 and 35).
- [130] United Nations Framework Convention on Climate Change. Adoption of the paris agreement, 2015. [Online] <http://unfccc.int/resource/docs/2015/cop21/eng/109r01.pdf> last accessed on 25-01-2017. (Cited on page 3).
- [131] United Nations Framework Convention on Climate Change. Report of the conference of the parties on its twenty-first session, held in Paris from 30 november to 13 december 2015, 2016. [Online] <http://unfccc.int/resource/docs/2015/cop21/eng/10.pdf> last accessed on 25-01-2017. (Cited on page 3).
- [132] United Nations Treaty Collection. UNTC depositary - status of treaties, 2016. [Online] [https://treaties.un.org/pages/ViewDetails.aspx?src=TREATY&mtdsg\\_no=XXVII-7-d&chapter=27&clang=\\_en](https://treaties.un.org/pages/ViewDetails.aspx?src=TREATY&mtdsg_no=XXVII-7-d&chapter=27&clang=_en) last accessed on 25-01-2017. (Cited on page 3).
- [133] US Energy Information Administration. State fact sheets on household energy use, 2016. [Online] [https://www.eia.gov/consumption/residential/reports/2009/state\\_briefs/](https://www.eia.gov/consumption/residential/reports/2009/state_briefs/) last accessed on 25-01-2017. (Cited on page 130).
- [134] US Energy Information Administration. Electric power annual, 2016. [Online] <http://www.eia.gov/electricity/annual/> last accessed on 25-01-2017. (Cited on page 23).
- [135] P. P. J. van den Bosch. Optimal static dispatch with linear, quadratic and non-linear functions of the fuel costs. *IEEE Transactions on Power Apparatus and Systems*, PAS-104(12):3402–3408, Dec 1985. ISSN 0018-9510. doi: 10.1109/TPAS.1985.318869. (Cited on page 74).
- [136] P. P. J. van den Bosch and F. A. Lootsma. Scheduling of power generation via large-scale nonlinear optimization. *Journal of Optimization Theory and Applications*, 55(2):313–326, 1987. ISSN 1573-2878. doi: 10.1007/BF00939088. (Cited on page 74).
- [137] R. van Gerwen and H. de Heer. Position paper flexibility value chain, October 2015. [Online] <https://usef.energy/News-Reads/Publications.aspx#> last accessed on 25-01-2017. (Cited on page 34).
- [138] R. P. van Leeuwen, J. B. de Wit, J. Fink, and G. J. M. Smit. House thermal model parameter estimation method for model predictive control applications. In *2015 IEEE Eindhoven PowerTech*, pages 1–6, June 2015. doi: 10.1109/PTC.2015.7232335. (Cited on page 154).



- [139] R. P. van Leeuwen, J. Fink, and G. J. M. Smit. Central model predictive control of a group of domestic heat pumps - case study for a small district. In *Proceedings of the 4th International Conference on Smart Cities and Green ICT Systems*, pages 136–147, 2015. ISBN 978-989-758-105-2. doi: 10.5220/0005434301360147. (Cited on page 121).
- [140] P. van Oirsouw. *Netten voor distributie van elektriciteit*. Phase to Phase, 2nd edition, 2012. ISBN 978-90-817983-1-0. (Cited on page 18).
- [141] S. Vandael, B. Claessens, M. Hommelberg, T. Holvoet, and G. Deconinck. A scalable three-step approach for demand side management of plug-in hybrid vehicles. *Smart Grid, IEEE Transactions on*, 4(2):720–728, June 2013. ISSN 1949-3053. doi: 10.1109/TSG.2012.2213847. (Cited on pages 37, 71, 72, and 167).
- [142] J. S. Vardakas, N. Zorba, and C. V. Verikoukis. A survey on demand response programs in smart grids: pricing methods and optimization algorithms. *IEEE Communications Surveys Tutorials*, 17(1):152–178, Firstquarter 2015. ISSN 1553-877X. doi: 10.1109/COMST.2014.2341586. (Cited on pages 33, 34, and 161).
- [143] T. Vidal, P. Jaillet, and N. Maculan. A decomposition algorithm for nested resource allocation problems, 2014. Working paper - MIT. (Cited on page 120).
- [144] H. M. Wagner. A two-phase method for the simplex tableau. *Operations Research*, 4(4):443–447, 1956. ISSN 0030-364X. doi: 10.1287/opre.4.4.443. (Cited on page 58).
- [145] S. Wogrin, D. Galbally, and J. Reneses. Optimizing storage operations in medium- and long-term power system models. *IEEE Transactions on Power Systems*, 31(4): 3129–3138, July 2016. ISSN 0885-8950. doi: 10.1109/TPWRS.2015.2471099. (Cited on pages 21 and 23).
- [146] F. Yao, A. Demers, and S. Shenker. A scheduling model for reduced CPU energy. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science, 1995*, pages 374–382, Oct 1995. (Cited on pages 74 and 120).
- [147] S. You, J. Hu, A. B. Pedersen, P. B. Andersen, C. N. Rasmussen, and S. Cha. Numerical comparison of optimal charging schemes for electric vehicles. In *IEEE Power and Energy Society General Meeting 2012*, pages 5152–5157, July 2012. doi: 10.1109/PESGM.2012.6345356. (Cited on pages 34, 69, and 71).
- [148] W. Yu, W. Rhee, S. Boyd, and J. M. Cioffi. Iterative water-filling for gaussian vector multiple-access channels. *IEEE Transactions on Information Theory*, 50(1):145–152, Jan 2004. ISSN 0018-9448. doi: 10.1109/TIT.2003.821988. (Cited on page 81).



# LIST OF PUBLICATIONS



- [TvdK:1] Thijs van der Klauw, Marco E. T. Gerards, and Johann L. Hurink. Resource allocation problems in decentralized energy management. *OR Spectrum*, pages 1–25, 2017. ISSN 1436-6304. doi: 10.1007/s00291-017-0474-2.
- [TvdK:2] Thijs van der Klauw, Johann L. Hurink, and Gerard J. M. Smit. Scheduling of electricity storage for peak shaving with minimal device wear. *Energies*, 9(6):465, 2016. ISSN 1996-1073. doi: 10.3390/en9060465.
- [TvdK:3] Thijs van der Klauw, Marco E. T. Gerards, Gerwin Hoogsteen, Gerard J. M. Smit, and Johann L. Hurink. Considering grid limitations in profile steering. In *IEEE International Energy Conference (ENERGYCON) 2016*, pages 968–973, April 2016. doi: 10.1109/ENERGYCON.2016.7514033.
- [TvdK:4] Thijs van der Klauw, Marco E. T. Gerards, Gerard J. M. Smit, and Johann L. Hurink. Optimal scheduling of electrical vehicle charging under two types of steering signals. In *IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe) 2014*, pages 85–91, Oct 2014. doi: 10.1109/ISGTEurope.2014.7028746.
- [TvdK:5] Thijs van der Klauw, Gerwin Hoogsteen, Marco E. T. Gerards, Johann L. Hurink, Xianyong Feng, and Robert E. Hebner. Assessing the potential of residential hvac systems for demand-side management. In *IEEE PES Innovative Smart Grid Technologies Conference (ISGT) 2016*, pages 0107:1–0107:5, Sept 2016. doi: 10.1109/ISGT.2016.7781193.
- [TvdK:6] Thijs van der Klauw, Marco E. T. Gerards, and Johann L. Hurink. Using demand-side management to decrease transformer ageing. In *IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe) 2016*, pages 1–6, October 2016. accepted for publication.
- [TvdK:7] Marco E. T. Gerards, Hermen A. Toersche, Gerwin Hoogsteen, Thijs van der Klauw, Johann L. Hurink, and Gerard J. M. Smit. Demand side management using profile steering. In *IEEE PowerTech 2015*, pages 536–541, June 2015. doi: 10.1109/PTC.2015.7232328.
- [TvdK:8] Gerwin Hoogsteen, Thijs van der Klauw, Albert Molderink, Johann L. Hurink, Gerard J. M. Smit, Xianyong Feng, and Robert E. Hebner. Balancing islanded residential microgrids using demand side management. In *IEEE PES Innovative Smart Grid Technologies Conference (ISGT) 2016*, pages 0057:1–0057:5, Sept 2016. doi: 10.1109/ISGT.2016.7781167.
- [TvdK:9] Krystian X. Perez, Michael Baldea, Thomas F. Edgar, Gerwin Hoogsteen, Richard P. van Leeuwen, Thijs van der Klauw, Bart Homan, Jirka Fink, and

Gerard J. M. Smit. Soft-islanding a group of houses through scheduling of CHP, PV and storage. In *IEEE International Energy Conference (ENERGYCON) 2016*, pages 605–610, April 2016. doi: 10.1109/ENERGYCON.2016.7513972.

220



## THIS THESIS

---

LIST OF PUBLICATIONS

```
@phdthesis{klauw2017:thesis,  
  author={van der Klauw, Thijs},  
  title={Decentralized Energy Management with Profile Steering -- Resource  
    Allocation Problems in Energy Management},  
  school={University of Twente},  
  address={PO Box 217, 7500 AE Enschede, The Netherlands},  
  year={2017},  
  month={may},  
  day={19},  
  number={CTIT Ph.D. thesis Series No. 17-424},  
  issn={1381-3617; CTIT Ph.D. Thesis Series No. 17-424},  
  isbn={978-90-365-4301-9},  
  doi={10.3990/1.9789036543019}  
}
```

---

BIB<sub>T</sub>E<sub>X</sub> of this thesis



## ABOUT THE AUTHOR



Thijs van der Klauw was born on July 31st, 1989. After finishing high school in 2007 he went to study mathematics at the University of Utrecht. He received his B.Sc. and M.Sc. degrees in 2010 and 2012 respectively. His focus during his studies was on operations research and logic. Wanting to switch to a more practically focused research group he came to the University of Twente to pursue his Ph.D. in a combined project on smart energy between the groups Discrete Mathematics and Mathematical Programming (DMMP) and Computer Architecture for Embedded Systems (CAES). His main research interests include scheduling and optimal usage of all forms of energy storage in smart (micro-)grids.





ISBN 978-90-365-4301-9



9 789036 543019