

Distributed Coordination and Optimisation of
Network-Aware Electricity Prosumers

Paul Scott
November 2016

A thesis submitted for the degree of
Doctor of Philosophy of
The Australian National University

Declaration

Except where otherwise stated, this thesis is my own original work.

Paul Scott
November 2016

Acknowledgements

Firstly, I would like to thank to Sylvie Thiébaux for being a fantastic supervisor and mentor. Your advice and help has been invaluable and I greatly appreciate the time you invested. Thank you also to my supervisor Pascal Van Hentenryck and advisor Menkes van den Briel for your technical help, ideas and advice.

Thanks to Hassan Hijazi, Carleton Coffrin and Archie Chapman for sharing your expert technical knowledge and for the valuable discussions. To my fellow students Karsten, Josh, Jan, Fazlul, Hadi, Terrence and Boon Ping, I had a great time working with you and sharing in the experiences of PhD life.

Finally, I would like to acknowledge my family and friends, with a special thanks to Megan for all her encouragement, understanding and support.

Abstract

Electricity networks are undergoing a transformation brought on by new technologies, market pressures and environmental concerns. This includes a shift from large centralised generators to small-scale distributed generators. The dramatic cost reductions in rooftop solar PV and battery storage means that prosumers (houses and other entities that can both produce and consume electricity) will have a large role to play in future networks.

How can networks be managed going forward so that they run as efficiently as possible in this new prosumer paradigm? Our vision is to treat prosumers as active participants by developing a mechanism that incentivises them to help balance power and support the network. The whole process is automated to produce a near-optimal outcome and to reduce the need for human involvement.

The first step is to design an autonomous energy management system (EMS) that can optimise the local costs of each prosumer in response to network electricity prices. In particular, we investigate different optimisation strategies for an EMS in an uncertain household environment. We find that the uncertainty associated with weather, network pricing and occupant behaviour can be effectively handled using online optimisation techniques using a forward receding horizon.

The next step is to coordinate the actions of many EMSs spread out across the network, in order to minimise the overall cost of supplying electricity. We propose a distributed algorithm that can efficiently coordinate a network with thousands of prosumers without violating their privacy. We experiment with a range of power flow models of varying degrees of accuracy in order to test their convergence rate, computational burden and solution quality on a suburb-sized microgrid. We find that the higher accuracy model, although non-convex, converges in a timely manner and produces near-optimal solutions. We also develop simple but effective techniques for dealing with residential shiftable loads which require discrete decisions.

The final part of the problem we explore is prosumer manipulation of the coordination mechanism. The receding horizon nature of our algorithm is great for managing uncertainty, but it opens up unique opportunities for prosumers to manipulate the actions of others. We formalise this form of receding horizon manipulation and investigate the benefits manipulative

agents can obtain. We find that indeed strategic agents can harm the system, but only if they are large enough and have information about the behaviour of other agents. For the rare cases where this is possible, we develop simple privacy-preserving identifiers that monitor agents and distinguish manipulation from uncertainty.

Together, these components create a complete solution for the distributed coordination and optimisation of network-aware electricity prosumers.

Contents

1	Introduction	1
2	Background	9
2.1	Power Systems	9
2.1.1	Markets	10
2.1.2	Generation	11
2.1.3	Transmission	13
2.1.4	Distribution	13
2.1.5	Microgrids	15
2.1.6	Consumption	15
2.2	Distributed Energy Technologies	17
2.2.1	Generators	17
2.2.2	Storage	17
2.2.3	Controllable Loads	18
2.3	DET Coordination	19
2.3.1	Conventional Pricing	20
2.3.2	Direct Load Control	21
2.3.3	Dynamic Pricing	21
2.3.4	Coordination	22
3	Home Energy Management	25
3.1	Introduction	25
3.2	Stochastic Programming	26
3.3	Deterministic EMS Problem	28
3.3.1	Devices	29
3.3.2	Optimisation Problem	29
3.4	Device Models	30
3.5	Stochastic EMS Problem	34
3.5.1	Random Parameters	34
3.6	Online Stochastic Optimisation	35
3.6.1	Executives	35
3.6.2	Expectation Formulation	36
3.6.3	2-Stage Formulation	36

3.7	Parameter Stochastic Models	39
3.7.1	Generalised Additive Models	39
3.7.2	Markov Models	41
3.8	Experiments	42
3.8.1	Controller Comparison	43
3.8.2	Device Cost Contributions	45
3.8.3	Device Operation Examples	46
3.8.4	Horizon Length	47
3.9	Related Work	47
3.10	Conclusion and Future Work	50
4	Network-Aware Coordination	51
4.1	Introduction	51
4.2	Power Flows	53
4.2.1	Network Power Flow Equations	53
4.2.2	Lines	55
4.3	Distributed Optimisation	57
4.4	Multi-Period Optimal Power Flow	60
4.4.1	Connections	61
4.4.2	Components	61
4.4.3	Optimisation Problem	62
4.5	Component Models	62
4.6	Distributed Algorithm	64
4.6.1	Algorithm	65
4.6.2	Residuals and Stopping Criteria	67
4.7	Implementation	67
4.8	Test Microgrid	68
4.9	Impact of Power Flow Models	70
4.9.1	Power Flow Models	70
4.9.2	Convergence	71
4.9.3	Solution Quality	74
4.10	Discrete Decisions	75
4.10.1	Methods	75
4.10.2	Convergence	77
4.10.3	Solution Quality	77
4.11	Pricing Uncertainty	78
4.11.1	PV Generation Uncertainty	79
4.12	Related Work	80
4.13	Conclusion and Future Work	82
5	Receding Horizon Manipulation	83
5.1	Introduction	83
5.2	Mechanism Design	84
5.2.1	Incentive Compatibility	85

5.2.2	Vickrey-Clarke-Groves Mechanisms	85
5.3	Power Balancing Problem	87
5.3.1	Single Horizon Version	87
5.3.2	Receding Horizon Version	87
5.3.3	Terminating Receding Horizon Version	88
5.3.4	Agents	89
5.4	Power Balancing Mechanism	89
5.4.1	KKT Conditions	90
5.4.2	Payments	91
5.4.3	Convexity	91
5.4.4	Limits	92
5.5	Manipulation	93
5.5.1	Examples	94
5.5.2	Strategies	96
5.6	Greedy Agent Strategy	97
5.6.1	Generator	99
5.6.2	Deferrable Load	100
5.6.3	Fixed Power Device	100
5.7	Greedy Agent Experiments	101
5.7.1	Problem Instances	101
5.7.2	Solving	102
5.7.3	Greedy Agent Results	102
5.8	Identifying Manipulative Agents	104
5.9	Inconsistency Identification	105
5.9.1	Revealed Preference Activity Rule	105
5.9.2	Cost Change Proxy	107
5.9.3	Inconsistency Identifier	108
5.10	Manipulation Identification	109
5.10.1	Cost Anticipation Test	109
5.11	Identification Experiments	109
5.11.1	Uncertain Problem Instances	110
5.11.2	Inconsistency Results	110
5.11.3	Anticipated Cost Results	111
5.12	Related Work	113
5.13	Conclusion and Future work	115
6	Conclusion	117
A	AC Power	121
A.1	Steady-State AC	121
A.1.1	Balanced 3-Phase	123
B	Proofs	125

Nomenclature	129
List of Publications	135
Bibliography	137

Chapter 1

Introduction

Electricity networks are undergoing a huge transformation due to the emergence of new technologies, market pressures and environmental concerns. The most significant trend is the transition from large centralised to small distributed forms of generation. The adoption of rooftop solar is shifting the ownership and control of generation away from large companies and into the hands of consumers. Such *prosumers* (houses and other entities that can both produce and consume electricity) have the opportunity to meet significant portions of their own energy needs and to sell excess power.

Distributed generation both positively and negatively affects electricity networks. It can reduce losses and allow for leaner, cheaper network designs; however, the weather dependent nature of distributed generation like solar photovoltaics (PV) means that supply does not always align with demand. In addition to this, uncontrolled or uncoordinated distributed generation can create voltage, peak power flow and ramping problems for networks.

Solar PV is already causing problems in some networks, and it will only get worse with time. For example, in some areas of the Australian network, the amount of rooftop PV has reached the point where utilities have to limit or reject new proposals if the impacts are assessed to be adverse [AEMO, 2012, CAT Projects and ARENA, 2015]. The only other option utilities currently have, network augmentation, is cost-prohibitive.

Other types of distributed energy technologies (DETs), including battery storage, electric vehicles (EVs) and smart appliances, pose similar opportunities but also similar risks to the network as PV [Ramchurn et al., 2012, Kok et al., 2009].

This thesis explores a new way forward for prosumers and utilities, where they work together to operate the network as efficiently and safely as possible. It addresses the question:

How can DETs be operated so that the benefits to prosumers are maximised whilst also supporting the network?

In answering this question, we will enable networks to embrace higher levels of DETs, most importantly rooftop PV and EVs. This will allow networks to operate more efficiently, with lower emissions and at lower cost. Prosumers will benefit through lower network costs, greater flexibility, rewards for actively supporting the network, and better utilisation of their DET assets.

We identify three parts that are necessary for a technical solution to the above problem, which form the motivation for the work in this thesis:

1. having autonomous controllers for DETs that can minimise prosumer energy costs in an uncertain environment;
2. coordinating the actions of multiple prosumer controllers to reduce system-wide costs and satisfy network constraints; and
3. aligning the motivations of self-interested prosumers with what is favourable for the wider network.

When combined, these parts deliver a technical solution for the *distributed coordination and optimisation of network-aware electricity prosumers*.

Home Energy Management

The first part of this thesis recognises the need to automatically control DETs so that the amount of human interaction is minimised. Prosumers do not want to invest time in manually controlling their devices, and even if they did, they would likely make poor decisions. Instead, what we propose is an automated energy management system (EMS) that acts on behalf of the prosumer to optimise their energy needs and costs, in a way that is transparent and does not encroach on their freedom.

We focus on the development of an EMS for household prosumers (see figure 1.1), although the same techniques could be applied to commercial, industrial or other types of prosumers. Houses are the most numerous (and interesting) type of network participant, accounting for around 25% of the total electricity consumption in countries like Australia [Vivid Economics, 2013].

The EMS communicates with devices present in the house and with the network. The house occupants provide simple, high-level preferences and constraints to the EMS, which then schedules the operation of the devices in a way that minimises energy costs.

What makes this especially challenging is the presence of uncertainty, which comes from a variety of sources including the weather, occupant behaviour and dynamic electricity prices. The EMS can make poor decisions if it does not take this uncertainty into account. For example, it is typically




Figure 1.1: An EMS controls household DETs including batteries, EVs, space heating/cooling and smart appliances.

beneficial to align consumption within the house with the solar PV output, which is not possible without a good prediction of solar production.

We design an EMS that can make sensible decisions in the face of uncertainty, using ideas from model predictive control (MPC), receding horizon control [Morari and Lee, 1999] and online stochastic optimisation [Van Hentenryck and Bent, 2006]. The EMS schedules the operation of devices over a forward horizon, using stochastic models for the uncertainty. This is crucial because the operation of state-based household devices (e.g., battery charge) needs to be planned out over multiple time steps. The EMS also operates online, updating its stochastic models and schedule as it progresses, enabling it to recover from poor earlier decisions and to take advantage of new opportunities.

We develop models for the devices and stochastic models for the uncertain aspects of the problem. Two different online algorithms are developed that use mixed-integer programming, and are compared to more reactive forms of control in an uncertain home environment—something that is missing from the existing literature.

We learn that for household devices and uncertainty, a simple approach based on a receding horizon and an expectation of the future produces solutions that are close to what can be obtained by an oracle. The approach can quickly compute solutions, making it suitable for real-time use. We also learn that more complicated techniques such as 2-stage stochastic programming provide only very small additional returns, and impose a much higher

computational burden.

These results show that, in practice, the uncertainty present in houses can be managed with computationally efficient methods and that the case for automating energy management is strong. The EMS can make local control decisions on behalf of occupants, providing the building blocks for the wide-scale control of DETs.

Network-Aware Coordination

The next step is to expand beyond a single prosumer by considering how multiple EMSs impact the network. When prosumers are left to optimise their own costs in an uncoordinated manner, the combination of their decisions can create undesirable new peaks in supply and demand [Ramchurn et al., 2011]. This effect, known as herding, has the potential to cause price volatility, and large voltage, ramping and line limit problems [Farhoodnea et al., 2013, de Hoog et al., 2013, Cain et al., 2010, Roozbehani et al., 2012].

To illustrate this problem, consider a network where prosumers own EVs and are subject to peak/off-peak pricing. If the off-peak period starts at 10 pm, then many of the EVs will be scheduled to begin charging at 10 pm, which could create a large spike in demand. On the other hand, spikes in supply can be produced in sunny conditions by PV systems when the overall demand for electricity is low. We cannot consider agents in isolation if we want to ensure the safe operation of the network and the efficient utilisation of available assets.

This motivates the next step where we take a network of prosumers with EMSs and look at the larger problem of jointly optimising prosumer and network costs. We take a *network-aware* approach, which means that we explicitly model the flow of power on the network to correctly account for the network's behaviour and limitations. Conventional networks already perform a similar task for large generating units, minimising dispatch costs whilst obeying transmission network constraints. When this is extended to allow prosumer participation, a significantly different approach is required.

Firstly, the scale of the problem is much larger than the conventional dispatch problem. Instead of coordinating a few hundred generating units, now many thousands or even millions of participants need to be considered (see figure 1.2). Secondly, the distributed nature of the problem means that localised distribution network constraints become significant, where before they could be ignored. Finally, prosumers have a more diverse range of behaviours and sensitivities which need to be considered.

We propose the use of distributed optimisation algorithms to handle this huge increase in problem size. Distributed algorithms also enable greater privacy for prosumers as their preferences and device behaviour are kept local. The algorithm we use, the alternating direction method of multipliers



Figure 1.2: A small sample of the many thousands of connected prosumers and distributed generators that need to be coordinated to balance the network.

(ADMM) [Boyd et al., 2011], works by locally exchanging price and power consumption/production information between prosumer EMSs and neighbouring network components. The algorithm negotiates a set of dynamic nodal prices and power allocations for the network and its participants. The prices accurately represent the marginal cost of supplying power to that part of the network, and are used to charge/pay prosumers for their actions. As power allocations for each prosumer are negotiated ahead of time, it prevents the herding effect exhibited by other pricing methods that have no two-way negotiation.

The handling of network and prosumer non-convexity in distributed algorithms is a significant unresolved problem that we tackle. Tractable distributed algorithms do not provide any guarantees of optimality when the underlying problem is non-convex. Unfortunately, the equations that model the physical flow of power on the network are non-convex, and the operation of certain household devices requires discrete decisions to be made. We investigate how such non-convexities can be managed in a tractable manner without sacrificing the quality of solutions.

What we find in our experiments is that the non-convex network models can be used directly with little to no loss in solution quality and performance. We identify and compare a number of approaches that can be used to handle the discrete decisions required to model shiftable smart appliances in houses. We also investigate how pricing structures can be used to incentivise prosumers to meet their commitments in the presence of uncertainty.

The end result is to show that a distributed approach to optimising and coordinating EMSs is not only possible, but highly effective.

Receding Horizon Manipulation

The final contribution is to ensure that the incentives offered to prosumers are aligned with the social objective of minimising the total system costs. Our distributed algorithm minimises the social objective, but only if agents act truthfully. By lying about their costs and constraints, prosumers can reduce their own costs, at the detriment of others prosumers and the social objective. We want to reduce the ability of agents to perform such *manipulation* and its impact on the system.

Existing electricity markets are already vulnerable to manipulation, with generating units adjusting their bidding strategies to maximise their advantage. Beyond this, the receding horizon nature of our algorithm creates unique opportunities for manipulation. The receding horizon is great for dealing with uncertainty, but the added flexibility also allows more opportunity for agents to manipulate others.

To illustrate this form of manipulation, take two neighbours, Alice and Bob, who both want to charge their EVs at 9 pm (see figure 1.3). The social solution is for them to both charge at 9 pm. Alice realises that if it were just her charging at 9 pm then the price will be lower. She artificially and temporarily raises the price at 9 pm by claiming to require more power than she actually needs. To avoid the price Bob charges at 8 pm. After Bob has finished charging, Alice can start telling the truth, allowing the price at 9 pm to reduce to a value lower than in the social solution.




Figure 1.3: Two homes competing to charge their electric vehicles.

It is difficult to correctly identify that Alice actually manipulated Bob, as we do not have access to her private information, and she could pretend that she was uncertain about her consumption requirements. There is a conflict between allowing agents to recover from uncertainty while also preventing manipulation.

We investigate receding horizon manipulation in detail. We develop a *greedy* agent that uses bilevel optimisation [Migdalas et al., 1998] to calculate strategies for manipulating other agents. This is used in experiments to investigate how much a manipulative agent can achieve and how much it

harms the system. We then develop non-invasive techniques which can be used to detect this form of manipulation. If successfully identified, agents can be fined or be subject to further auditing processes.

We find that it is possible for large agents or a large coalition of agents to gain an advantage and harm the overall social objective. We manage to successfully apply our identifiers to the actions of these agents, and distinguish their actions as being manipulative as opposed to just being honestly uncertain. It does not completely eliminate manipulation, but it acts as a strong deterrent, and makes it more difficult to calculate beneficial manipulative strategies. Our framework allows new identifiers to be added over time if agents develop new strategies.

Summary

To summarise, the key research question we tackle is how to operate distributed energy technologies in a way that maximises the benefits to prosumers whilst also supporting the network. We address this question in three parts that come together to form a complete solution: the development of autonomous home EMSs that can operate in an uncertain environment; the coordination of these EMSs across the network to reduce system-wide costs while satisfying network constraints; and the development of identifiers that deter strategic agents from manipulating the system.

We contribute to knowledge in each of these areas:

- We show how online stochastic optimisation based on a receding horizon can be used to efficiently control devices within a house, and that it produces solutions that compare favourably with a perfect information case. It was found to have significant advantages over reactive forms of control.
- We establish the suitability of different power flow models in a distributed optimisation setting and, how the non-convex power flow equations can be used directly to produce near-optimal results. Similarly, we show how discrete shiftable house loads in practice can be included in the distributed algorithm, and still achieve near-optimal results.
- We determine in the worst case how much strategic agents have to gain and how much they can harm the performance of the system. We formalise receding horizon manipulation in this power exchange setting, and develop and demonstrate identifiers that can distinguish strategic agents from uncertain agents.

Updating our power systems is crucial for the future of our environment and the productivity of our societies. The development of renewable energy and cheap battery storage, and the rapid electrification of the world requires

new flexible and efficient solutions. Our work develops a key component of the solution, by enabling prosumers to participate actively in the balancing and management of the network.

Thesis Outline

This thesis is structured as follows:

Chapter 2 provides a non-technical introduction to power systems, looking at topics such as energy markets, generation and power transmission/distribution. It also dedicates sections to describing the different types of DETs and the different approaches that have been used to incentivise and coordinate them.

Chapter 3 introduces the residential energy management problem and develops online stochastic optimisation techniques for scheduling DETs in the presence of uncertainty. These algorithms are compared with reactive controllers and a clairvoyant controller on a household test case.

Chapter 4 extends this to the wider problem of how to coordinate EMSs over the network. A distributed algorithm is developed to perform the coordination and is experimented with on a test microgrid for different power flow models and methods for handling discrete decisions.

Chapter 5 investigates the potential for manipulation of the receding horizon mechanism. It formalises receding horizon manipulation, experiments with a greedy strategic agent and develops identifiers to detect this type of manipulation.

Chapter 6 concludes our findings and combined approach, and discusses future developments.

Chapter 2

Background

This chapter provides background material that sets the context and motivation for our problem. First we introduce power systems and what changes distributed energy technologies (DETs) are causing. Then we introduce the different types of DETs and the types of methods that have been proposed for controlling/coordinating them. In later chapters we develop our distributed algorithm with these DETs in mind, and perform more detailed comparisons to existing approaches.

Where relevant, later chapters include more topic-specific background material. Chapter 3 has a section on stochastic programming, chapter 4 contains sections on power flows and distributed optimisation, and chapter 5 has a section on mechanism design.

2.1 Power Systems

Electric power systems (also known as grids or networks) have been used for over a century to supply energy to our industries and homes, playing a central role in enabling the technological and quality-of-life improvements over this time. This section provides an overview of conventional power systems and discusses the pressures they are experiencing from new technologies and environmental concerns.

To help explain power systems, we refer to the National Electricity Market (NEM) in Australia. The NEM network spans the east coast of Australia, supplying some 19 million residents [AEMO, 2015a]. Many parallels can be drawn between the NEM and other deregulated electricity markets such as those in the USA.

Power systems consist of three main components: generators, consumers and the network of conductors that transport power between the two. The network is further divided into the transmission and distribution networks (see figure 2.1). Conventional power systems are operated by central authorities which manage the dispatch of generation and control of network

components.




Figure 2.1: The main components of a power system.

2.1.1 Markets

In deregulated power systems, markets play a key role in facilitating the supply of electricity and coordinating operational decisions. These centralised markets focus on coordinating the dispatch of large generation units and managing the transmission network. In chapter 4 we construct a market of our own, but one that coordinates prosumers and other network participants in a decentralised way on the distribution network.

Wholesale energy supply markets ensure that there is always enough bulk supply dispatched to meet forecast demand. Frequency regulation markets recruit generators to balance any wholesale market supply and demand prediction errors in real time.

The NEM is made up of one wholesale energy supply market (the spot market) and 8 frequency control markets. Generators bid into the market the amount of power they can supply for a particular price. The NEM dispatch engine clears the market every 5 minutes by selecting the lowest cost bids to match the forecast demand. The spot price is a 30 minute average of these dispatch prices, and it is the price at which all dispatched generators get paid.

Retailers buy electricity at the spot market price and onsell it to consumers. The retail tariffs offered to consumers vary between regions and depend on the type of customer (e.g., residential, commercial or industrial). Residential tariffs are typically either constant or time-of-use (TOU) where they change for different times of the day. Section 2.3.3 will discuss other more dynamic forms of pricing.

Retail tariffs recover wholesale market costs for the retailer, but also other costs to do with providing and maintaining the network. In the NEM at least, transmission and distribution costs overwhelm the wholesale market costs as shown in figure 2.2.

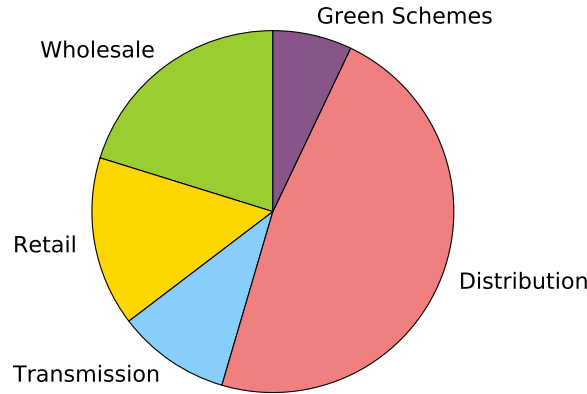


Figure 2.2: Breakdown of NEM residential electricity bill costs in 2014 [AER, 2014].

2.1.2 Generation

Generators get their energy from either renewable (hydro, wind and solar) or non-renewable (coal, gas, oil and nuclear) sources. In most countries the majority of electricity still comes from non-renewable sources. Table 2.1 shows how coal supplies around 75% of the NEM electricity.

Table 2.1: NEM capacity and production by generator type for 2014/15 financial year [AER, 2015]. Only includes registered generators (no rooftop PV).

Type	Capacity (%)	Production (%)
Black Coal	39.2	50.8
Brown Coal	14.3	25.7
Gas	20.1	11.6
Hydro	16.5	6.6
Wind	6.6	4.8
Other	3.2	0.5

Generators are dispatched so that supply always matches demand on the network. The network itself cannot store electricity, so any mismatch will cause fluctuations in frequency and voltage, which can ultimately damage equipment or trigger a blackout. This dispatch problem is becoming more challenging as networks transition away from large centralised forms of generation to small distributed generators, which is our primary motivation for using DETs to help balance supply and demand.

The physical limitations and running costs of different generator types influences what role they play in the network. Coal power stations have relatively low fuel costs, but their starting and stopping procedures can be lengthy and expensive. Gas turbines on the other hand can start, stop and change their output quickly, but have high fuel costs. Renewables, such

as wind and solar, have effectively zero marginal running costs, but their availability is weather dependent.

As such, coal power stations provide a steady base-load supply and gas turbines are typically peaking plants that get dispatched during peak load or high ramping rate conditions. The near-zero running costs of renewables means that they get dispatched whenever they are available. These behaviours are observed in table 2.1 when the normalised amount of energy each generator type *could* provide (capacity) is compared to the normalised amount of energy it *actually* produced (production). For example, although gas has 20.1% of the total capacity, it is only used to supply 11.6% of the actual load.

In most networks the vast majority of electricity comes from a relatively small number of very large power plants. For example, the NEM has less than 400 registered generating units, with the 100 largest having around 80% of the total network capacity [AEMO, 2015b] (see figure 2.3). These 100 generating units are from around 40 individual power plants, which is a tiny number when compared to the near 9 million metered consumers that the NEM serves.



Figure 2.3: The Yallourn W brown coal power station alone supplies around 8% of the NEM’s electricity (“Yallourn Power Station” by CSIRO, CC BY 3.0).

Large generators often have to be situated far from where the energy is needed for public safety reasons and/or the need to be collocated with a geographical feature (e.g., a coal mine, water source or hydro dam). The transmission network serves the role of transporting high volumes of electricity from the generators to where it is needed.

Solar, wind and other renewables are expected to replace non-renewable generation in the coming decades, both for environmental and economic reasons. For example, studies estimate that since 2013 wind has been cheaper than new build coal power in Australia [Bloomberg New Energy Finance, 2013]. A report from the same year estimated that PV would be the cheapest form of generation within the next 25 years [Syed, 2013]. The reductions PV

costs have outdone expectations, with recent studies estimating that utility-scale PV is already competitive with conventional generation in countries such as the US [Lazard, 2015].

These cost comparisons are for generators in the wholesale market, but renewables like PV are able to compete in the retail market by being installed in homes behind the meter. As shown previously in figure 2.2, the wholesale prices contribute only around 20% of the total retail costs in the NEM. This makes PV much more competitive installed in homes than it would otherwise be in the wholesale market, primarily because it bypasses the expensive network altogether by supplying power directly where it is needed. Such behind-the-meter PV makes up the majority of the installed PV, which was around 2.5% of the total electricity consumed in the NEM in 2014 [Johnston et al., 2015].

2.1.3 Transmission

Transmission networks provide high volume transport of electricity between generators and load centres (e.g., cities). Figure 2.4 shows the transmission network for the NEM which spans the whole Australian east coast connecting generators, cities and regional towns. This thesis focuses on what happens further downstream on the distribution network, as this is where DETs will have most impact; however, the same techniques can be applied to manage the transmission network.

Structurally, transmission networks are meshed (contain cycles) and are often designed and operated to be $N-1$ reliable, which requires that the failure of any one line, generator or network component will not bring the whole system down. Most networks are 3-phase alternating current (AC) with a small number of high voltage direct current (DC) links where economically viable.

Transmission networks operate at high voltages (typically in range 100kV to 500kV in the NEM) and are much more heavily instrumented, monitored and managed than distribution networks because of their more critical nature. The operators of transmission networks are concerned with issues of voltage regulation, managing congestion and stability.

2.1.4 Distribution

Distribution networks take power from the transmission network and distribute it out to all the loads in the local area. They are typically structurally radial, but some networks, especially in cities, can be meshed. In chapter 4 we develop a technique for coordinating DETs over a distribution network.

Substations (see figure 2.5) connect the transmission and distribution networks together by lowering the high transmission voltages to medium voltages (11kV in much of the NEM), which can more safely be transported

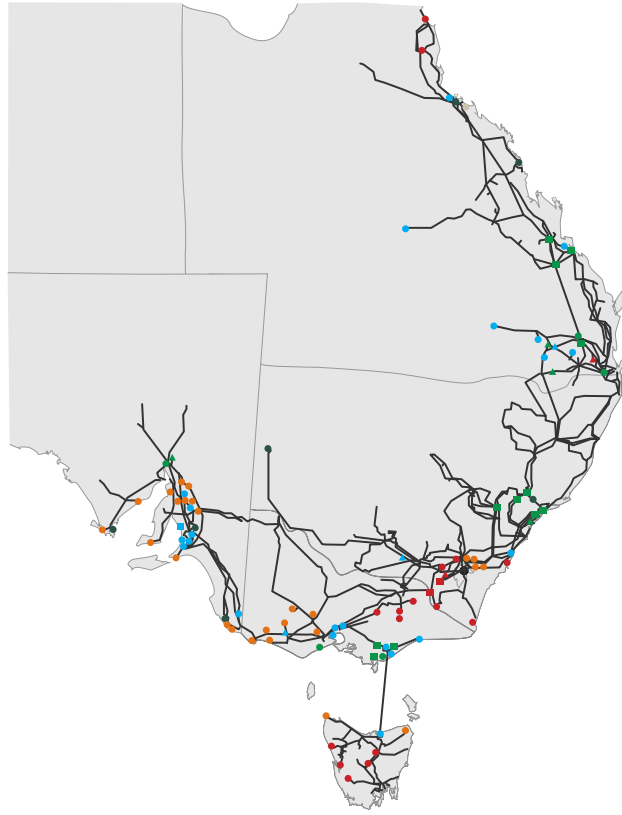


Figure 2.4: NEM transmission lines and generators (image adapted with permission from AER [2014]).

through built-up areas. These medium voltage lines, called feeders, connect to distribution transformers close to consumers. Distribution transformers bring the voltage down to the final voltage for consumption (400V three phase or 230V single phase for residential customers in the NEM).

Meters are used to record how much power enters the distribution network from the transmission network and how much leaves the network at the point of customer connection. AC power is split up into two components, *real power* which does actual work, and *reactive power* which oscillates back and forth between active electrical components (see appendix A for more details). Residential customers are typically only billed for real power consumption, while larger customers are also billed for their peak consumption of real and reactive power.

As shown in figure 2.2, distribution network investment and maintenance accounts for nearly half of all retail electricity costs. This is because distribution networks are designed so that all conductors and transformers can meet peak load requirements. These events might only occur for a handful of hours in a year, with the extra capacity (and investment) going under-

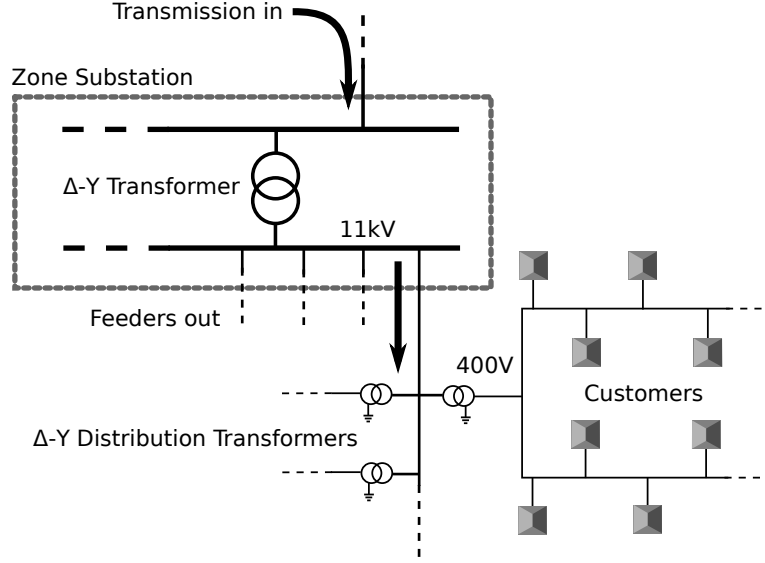


Figure 2.5: The layout of a distribution network. Distribution transformers have 4 conductor outputs, one for each phase and a neutral. Customers connect to either a single phase or to three phases.

utilised for the rest of the time. The techniques we develop in this thesis for the intelligent control of DETs can reduce these peaks and hence greatly reduce distribution network costs.

2.1.5 Microgrids

Microgrids are an emerging technology which hold particular promise in developing nations. Microgrids are small-scale networks (ranging in size from a small community all the way up to an entire city) that can disconnect from and operate independently of the rest of the network [Glover et al., 2011]. Microgrids provide a number of advantages including better resilience to faults on the network, and more organic and gradual electrification of remote or rural areas.

Like more conventional centralised networks, microgrids need to ensure that power is balanced, frequency is maintained and the network is safely operated. These events take place on a smaller scale, which means that there is more uncertainty and less room for error. In chapter 4 we test our distributed algorithm on a community-scale microgrid.

2.1.6 Consumption

In the NEM, the residential, commercial and manufacturing sectors each consume around 25% of all electricity, with the remaining quarter shared between public utilities, mining, agriculture and transport [Vivid Economics,

2013]. In this thesis we focus on residential customers; however, the same techniques we develop could be applied to participants in the other sectors. This focus is because residential customers:

- in aggregate have a large share of the overall consumption;
- are the most numerous;
- are the early adopters of many DETs such as PV, EVs and battery storage; and
- are widely dispersed throughout the network.

Figure 2.6 shows the consumption pattern for one house over a single day. The load profile has large peaks caused by the operation of different appliances. The power profile tends to smooth out as the loads of many houses are combined, but there are still events which are correlated between homes (e.g., working schedules, air conditioner usage and PV production). The aggregate profile tends to exhibit peaks in the morning and evening as residents leave for and return from work. By scheduling the operation of DETs (e.g., a smart dishwasher), the load profile of both a single house and the aggregate of many houses can be shaped. Chapter 3 investigates in detail how to optimise the operation of DETs within a single house (prosumer).

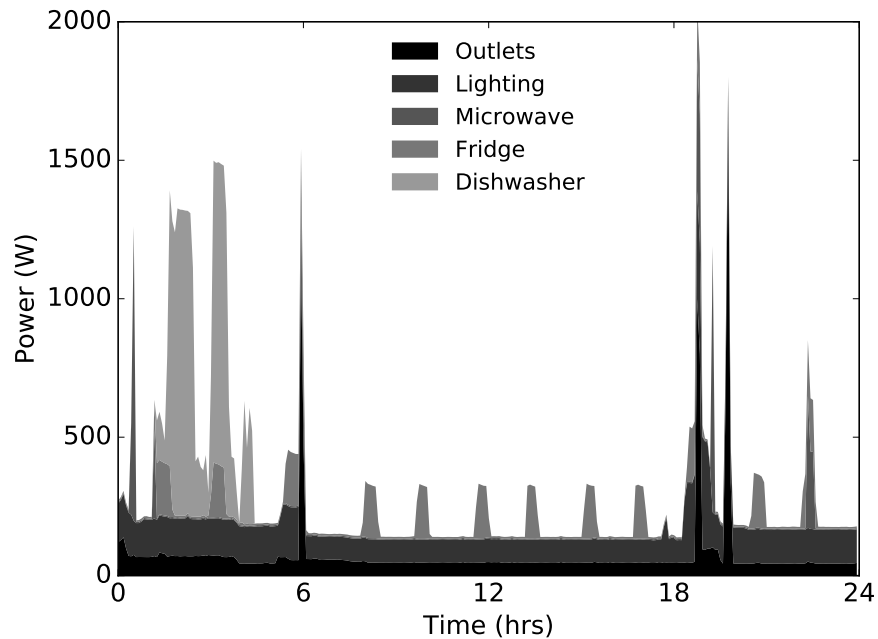


Figure 2.6: Loads in a house over a day (data from [Kolter and Johnson, 2011]).

2.2 Distributed Energy Technologies

DETs are any generators or controllable loads that are of a small enough scale that they can be connected directly to the distribution system or behind the meters of customers. The rapid development of such technologies is challenging the conventional notion of how networks are designed and operated.

As discussed previously, DETs are shifting generation away from large centralised power stations and are enabling individuals or communities to have more control over their own energy requirements. For example, DETs enable the formation of community owned microgrids, which can either operate stand-alone or trade energy with the wider network.

This thesis will often refer to consumers (houses or any agents) that have control and ownership over a set of DETs as *prosumers*. This captures the idea that they can both consume and produce energy, and that they are more involved in managing their own energy requirements.

The following sections describe some of the DETs that are available to houses, and which we will be working with in later chapters. We discuss their behaviour and what they can be used for. They range from batteries that can be charged or discharged on command, to less controllable solar PV generators.

2.2.1 Generators

Distributed generation ranges from the renewable-sourced solar PV and wind generators which have limited controllability to the dispatchable fuel-based micro gas turbine, diesel and fuel-cell generators. Solar PV is safe, low maintenance, unobtrusive and affordable, making it ideal for residential use. The other distributed generators are let down in one or more of these areas, which means that they are typically only viable in larger-scale grid-connected applications.

PV systems connect to an inverter, which converts DC power from the panels to AC power at the grid frequency. At the inverter the output real power can be controlled between zero and a maximum for the current sunlight intensity. Some inverters can also be configured to supply or sink reactive power. For most inverters these controls are not utilised, instead they are configured to output real power at maximum availability, at a fixed power factor (which fixes a ratio between real and reactive power).

2.2.2 Storage

With dramatic cost reductions in recent years, battery storage is starting to be a viable investment for residential customers. Storage enables prosumers to better utilise their rooftop PV systems by storing any excess generation

for later use. If they are exposed to time-varying pricing, storage also enables them to take advantage of the cheapest prices.

Like PV, batteries connect to an inverter, but this inverter can also consume real power to charge the battery. The maximum charge and discharge rates are a function of the battery state of charge and physical limits. Batteries deteriorate with use based on the number of charge/discharge cycles they undergo, the rate and depth at which this occurs, and the ambient temperature. These factors should be taken into consideration when controlling the battery, as the savings from a particular use might be outweighed by the resulting reduction in battery life.

The batteries in EVs can also be used for storage, but this use is complicated by the EVs' interaction with household occupants. EV batteries are only available when they are at home plugged in and their state of charge changes through driving. In addition to this, occupants expect a certain level of charge at different times of the day so that they have enough range to complete trips. Together these factors limit how EV batteries can be operated.

2.2.3 Controllable Loads

Not all conventional household loads can be controlled as they must be available on-demand for use by occupants. Such loads include lighting, cooking and entertainment. They can be made more energy efficient (e.g., lights that sense when an occupant enters a room) but for the most part they have no useful actions for control.

Loads that are controllable include dish washers, washing machines and water heating. These are typically more complicated than other DETs, as the occupants have expectations about their operation which add additional constraints and costs associated with particular decisions. Below we describe some of these technologies:

Smart Appliances— Smart appliances are regular appliances which have a communication and control interface built into them [van den Briel et al., 2013]. Relevant appliances include dish washers, washing machines and clothes dryers. When operating such a device, an occupant will typically be asked a completion deadline. The smart appliance can then be scheduled within this limit to take advantage of the cheapest electricity. Depending on the type of appliance it might be possible to delay the start time, pause and resume it part way through, or run it at different power settings.

Appliances such as clothes dryers and dish washers are ideal candidates as they can be energy intensive and occupants are often flexible with their use. Unfortunately, smart appliances are not always available, as they first require a occupant to issue a run request. Therefore, while smart appliances can play an important part, on their own they will not be enough to balance

a network that is supplied by high levels of renewable generation.

Space Heating/Cooling— Heating and cooling of houses is an energy intensive load. Occupants desire temperatures to be within comfortable bounds at certain times of the day. These bounds allow for some power consumption flexibility. Greater flexibility is available when the house has high levels of thermal mass or insulation.

In addition to this, occupants may allow the temperature to move outside the comfortable zone for short periods if it means that they will make a reasonable saving on their electricity bills. In cases where the house does not have much flexibility, control of space heating and cooling will only be beneficial if the electricity prices are dynamic and volatile.

Water Heating— Water heating is a flexible operation which needs to ensure that at all times enough hot water is stored to meet the requirements of occupants. In many networks water heaters are already controlled so that they take advantage of off-peak prices. This control can also take place when prices are dynamic or where excess solar generation is available during the day.

Pool Pumps— Swimming pools are relatively common in countries such as Australia. Water needs to be pumped through the pool's filtration or heating system. Typically a minimum volume of water needs to be pumped through the system each day, which is often a flexible operation.

2.3 DET Coordination

The prosumer and DET coordination we develop in chapter 4 has similar goals to demand response (DR) and demand-side management (DSM). DR and DSM are general terms used for schemes that influence or coordinate customer consumption on the demand side of the network. They are often interpreted as focussing just on loads when many of the techniques can also interact with prosumers, distributed generation, batteries and other grid connected DETs. Another point of confusion is that in microgrid settings, the distinction between the demand side and the supply side of the network is less apparent. For these reasons, we favour the use of prosumer or DET coordination when referring to our work.

The strength of our approach lies in its network modelling, consumer focus, near optimality, scalability, robustness to uncertainty, automation and resilience to manipulation. Other approaches have the advantage of being significantly less complex, but they do not achieve the long-term vision of managing high levels of DETs in a network supplied primarily by renewable resources. In later chapters we will compare our approach in more detail as it is explained and experimented with.

As an introduction, this section categorises and explains some of the existing techniques for controlling DETs (DR, DSM or otherwise) that fulfil a role analogous to conventional generator dispatch. Techniques that use DETs in a frequency regulation capacity are not considered, although in general they can be designed to work alongside DET dispatch.

2.3.1 Conventional Pricing

Utilities have used several different pricing schemes to influence the consumption and production behaviour of customers on the demand side of the network. The goal is either to shift consumption away from particular times or reduce peaks. Peak/off-peak metering has been used for more than 50 years in Australia [UOW, 2014], to shift some of the daytime loads to the night. Originally this involved separate peak and off-peak meters with separate wiring to designated peak or off-peak loads. Hot water heating is commonly used as an off-peak load, with the utilities using ripple control to remotely switch the system on during off-peak hours.

With modern smart meters the need for separate wiring is removed, thereby allowing all loads to take advantage of cheaper off-peak prices. They also enable more pricing intervals and a distinction between weekdays and weekends. The general name for such pricing structures is time-of-use (TOU) pricing. The idea is for these prices to reflect the wholesale market prices for a typical day, but to do so in a way that is transparent and not overly complicated for customers. The prices and time intervals are locked in when customers sign up for an electricity contract, or only infrequently updated by the retailer (e.g., annually).

More recent pricing developments have been in reply to the popular uptake of solar PV. PV first became financially viable because of government rebates and legislation that mandated retailers buy any excess PV generation from customers. The price that customers get paid at is called a feed-in tariff (FIT). In many markets the FITs were originally set at several times the average retail price of electricity, creating an asymmetry between consumption and production in order to encourage adoption.

FITs have reduced over time as PV technology has become cheaper. Subsequently, in many markets retailers are now free to set their own FIT prices. This has often resulted in the FIT being set much lower than the average retail price. This encourages houses to self consume their excess solar generation instead of exporting it back to the network.

These existing schemes encourage behavioural changes in customers, so that over time their consumption habits better align with the costs of generating and supplying electricity. What they cannot do is react to dynamic changes in the wholesale market and the distribution of electricity. The network will become increasingly more dynamic as more renewable generators come online, and as more solar PV and electric vehicles are adopted by cus-

tomers. Advances in DETs, communications and automated systems make it possible to design schemes that can handle more dynamic networks, with minimal disruption to customer behaviour.

2.3.2 Direct Load Control

In direct load control (DLC) schemes customers allow utilities or other third parties to take control of some of their devices, with customers paid some form of ongoing or per-use compensation as part of the agreement. DLC makes most sense for use with large industrial loads that can be reliably predicted. An example of this is aluminium smelters in Australia, which in extreme network situations can be shed by the network operator.

Optimisation techniques have been applied to assist networks with making DLC decisions when they have a portfolio of large industrial loads that can be controlled. For example, Pedrasa et al. [2009] use binary particle swarm optimisation to assist a utility in deciding when DLC should be performed, and which participant should be selected. The participants can have different load shedding costs and constraints in their contracts with the utility.

DLC has also been proposed for residential customers. Guo et al. [2008] consider DLC where the set point temperatures of household air conditioners are directly controlled by the utility. They use adaptive genetic algorithms to decide on the optimal set point allocation for each house, whilst considering temperature comfort bounds. Each house can override the utility signal, but by doing so they forfeit any incentives.

The problem with DLC is often scalability and privacy concerns, as data from each participant needs to be communicated to a central body which then solves a large problem centrally. DLC is more suitable for a network that only rarely needs to intervene, for example, just to reduce summers peaks. We focus on delivering a solution that on a daily basis seamlessly controls household devices to balance renewable generation.

2.3.3 Dynamic Pricing

As we have discussed, TOU pricing is fixed months or years in advance so cannot react to dynamic events. Dynamic pricing, also known as real-time pricing (RTP), allows for a continuously changing price signal, where the price is known at most only a short time in advance. Dynamic pricing comes the closest to exposing customers to the actual wholesale market, but the signal may be modified to reduce volatility and to incorporate network costs.

Dynamic pricing has the potential to achieve more efficient solutions for the network and unlike DLC does not require the utility to model and control every individual participant. On the negative side it exposes customers

to more complexity, and can be more difficult to predict the response of customers.

The utility has the task of setting a dynamic price based on what the current market and network conditions are, and based on how it expects participants to respond to the price. If this is not done carefully it can result in unanticipated herding behaviour as explained by Reddy and Veloso [2012]. Customers exposed to a dynamic price have to decide what their best response is. This can either be the occupants themselves being conscious of the price and manually changing their behaviour or, as we propose, having an automated home energy management system make decisions on behalf of the occupants.

The dynamic prices can also have a spatial dependence. For example, locational marginal prices represent the cost of serving electricity to a particular part of the network at a particular time. The advantage of this is that they can account for network losses and network congestion.

2.3.4 Coordination

Several methods that are not based on prices have been proposed to avoid the herding problems that naive pricing methods can experience [van den Briel et al., 2013, Shinwari et al., 2012]. A probability distribution which represents an “ideal” load curve is used to randomly select the start time of loads. Although this avoids having to select dynamic prices, the problem of how to choose an ideal load curve that obtains an optimal solution in expectation remains. At times this may be just as challenging as selecting sensible dynamic prices.

Ramchurn et al. [2011] and Reddy and Veloso [2012] reduce herding behaviour by providing agents with adaptive controllers. Using slow time constants to reduce the responsiveness of controllers and adding randomness to some decisions allows for a reduction in herding behaviour, but with a possible increase in the solution costs. Caron and Kesidis [2010] also use stochastic policies in order to reduce peaks in a cooperative game for scheduling loads, and investigate how the quality of solutions changes when sharing different amounts of private information.

A greater degree of coordination between participants is required to prevent the herding-type problems and to ensure quality solutions. Two-way communication between participants and the network (or amongst participants) can be used to provide feedback. For houses such communication will typically be performed by energy management systems on behalf of occupants. As with DLC, privacy and scalability can also be major issues.

Auctions, aggregators and distributed optimisation have been proposed to overcome these challenges. In auctions, participants submit bids to either consume or produce electricity, which are then centrally cleared by the market. The bids themselves can range from simple price and power pairs, to

utility functions that represent the demand curve of the participant, and they may be submitted iteratively [Ygge and Akkermans, 1996, Vytelingum et al., 2010, Chapman and Verbic, 2015] or all at once [Kok et al., 2009]. One of the challenges of these auction-based approaches is to develop appropriate bidding strategies for participants.

Virtual power plants and aggregators have been proposed as a means of enabling small distributed prosumers to participate in the wholesale market. Participants join an aggregator which can sell services to the wholesale market on behalf of the collective [Chalkiadakis et al., 2011]. The aggregator has to decide how to reward or incentivise its participants so that they achieve the desired outcome [Akasiadis and Chalkiadakis, 2013]. There remains a coordination problem, but at least it is smaller than the original and easier to integrate with existing electricity markets.

The coordination problem can be viewed as a distributed optimisation problem. Distributed optimisation algorithms provide a structured method of solving the problem with well-defined participant subproblems and a clear interface for communications [Mohsenian-Rad et al., 2010, Gatsis and Giannakis, 2012, Kraning et al., 2014]. Care must be taken to ensure the problem converges and that agents cannot cause significant harm to the system by providing misleading information.

Some of the approaches discussed take into consideration or model the network, but for the most part the network is overlooked. Modelling the power flows is necessary in order to accurately account for losses and network constraints such as voltage and current limits. We adopt a distributed optimisation approach that explicitly models the network in chapter 4.

Chapter 3

Home Energy Management

3.1 Introduction

For many residential electricity customers, DETs, including PV, EVs, battery storage and smart appliances, are now within financial reach. We develop a residential energy management system (EMS) for such a residential prosumer that minimises electricity costs by automatically operating local DETs (which we simply call devices). This provides value to prosumers that are exposed to time-varying pricing or who want to self consume solar generation, and is a building block for the coordination of network-aware prosumers that we develop in chapter 4.

The EMS must consider both occupant comfort and energy costs, for example, when controlling space heating/cooling. These two objectives are often conflicting, so we provide a means for occupants to indicate how they value comfort against cost savings. This combined objective is then used by the EMS when optimising the device operation.

One approach is for the EMS to have simple reactive device control policies, but due to the state-based nature of many devices this is often too short-sighted. Instead, what we propose is for the EMS to schedule the actions of devices over a forward time horizon, to better account for the implications of an action. This raises its own problems, because while the EMS might have a good idea of the current state of the system, there are many external influences which it cannot know exactly in advance including electricity prices, the weather and occupant behaviour.

To overcome this the EMS develops stochastic models of the uncertain external processes and takes these models into account when optimising its actions. The EMS can learn and tune these models over time as it collects more data. The EMS also makes use of an online algorithm to reduce the impact of uncertainty. We use a *receding horizon*, like in model predictive control, where the EMS only acts on decisions for one time step before the horizon shifts forward and a new optimisation is conducted. This

ensures that the algorithm always takes into account the most up-to-date information before making a decision.

We investigate the benefits of online optimisation for an EMS that is exposed to dynamic pricing in an uncertain environment, making two primary contributions: one conceptual and one algorithmic.

At the conceptual level, we present a compositional architecture for EMS optimisation, where each device can be modelled independently in terms of a collection of functions that encapsulate its behaviour. These devices are then assembled into a model of a home, from which optimisation problems for the EMS can derive.

At the algorithmic level, we present a comprehensive study of the value of EMS optimisation when future prices, occupant behaviour and environmental conditions are uncertain. The formulation uses models representative of physical devices and stochastic models trained on real weather and network demand data. These device and stochastic models are used in two online optimisation algorithms which are compared to simple controllers that use reactive policies.

The experimental results not only show the value of stochastic information, but also that online optimisation provides solutions that are close to the clairvoyant solutions which have perfect knowledge of the future. The online stochastic algorithms using mixed-integer linear programming (MILP) technology are fast and produce significantly better solutions than the reactive controllers. Also of interest is the comparison between the two online stochastic algorithms, and an experiment that investigates the optimal receding horizon size.

In section 3.2 we provide background material on stochastic programming which is a basis for one of the EMS optimisation strategies we employ. We then formalise a deterministic version of the residential EMS problem and develop device models for our experiments in sections 3.3–3.4. This is followed by the stochastic version of the problem, our online optimisation strategies and the stochastic models used in our experiments in sections 3.5–3.7. We provide a experimental comparison between our different EMS techniques in section 3.8, followed by a discussion of the related work in section 3.9 and finally our conclusion in section 3.10.

3.2 Stochastic Programming

Stochastic programming is a framework for modelling optimisation problems that have uncertain parameters [Shapiro et al., 2009]. In these problems some decisions need to be made before the values of the uncertain parameters are known. By taking into consideration the probability distributions of these parameters, stochastic programming aims to make decisions that are feasible for all possible scenarios, whilst also minimising the cost function (in

expectation). It fits naturally with scheduling problems where uncertainty is revealed over time [Van Hentenryck and Bent, 2006].

One of the online algorithms we develop for the EMS uses stochastic programming. Accounting for the stochastic nature of the problem can reduce the likelihood of making decisions that will later be regretted, and can open up new opportunities that would otherwise not have been possible.

In stochastic programming the problem is split up into two or more stages which represent consecutive time periods. Each stage has a set of decision variables which must be settled before the problem can progress to the next stage. The values of uncertain parameters are revealed between each stage. The simplest formulation to describe, and the one that we utilise in the EMS, is 2-stage stochastic programming.

In a 2-stage stochastic problem the variables are split up into two stages, where all uncertainty is assumed to be revealed after the first stage. Let the vectors x and y be the first and second stage variables respectively. Let the vector s represent a particular realisation of the stochastic parameters, also called a scenario. The set of all possible scenarios is S , only one of which will come true.

For a problem with cost function $f(x, y, s)$ and constraint $g(x, y, s) \leq 0$, the 2-stage stochastic problem is formulated as:

$$\min_x \mathbb{E}[q(x, s)] \quad (3.1)$$

$$q(x, s) := \min_y f(x, y, s) \quad (3.2)$$

$$\text{s.t. } g(x, y, s) \leq 0 \quad (3.3)$$

where $\mathbb{E}[\cdot]$ is the expectation over all scenarios S and $q(x, s)$ is the minimum cost achievable for a given scenario s and first-stage decision x .

The problem can be transformed into an equivalent deterministic form when the set of all possible scenarios S is finite. This effectively involves solving $|S|$ versions of the problem with common first-stage variables x . Each scenario s has its own copy of the second-stage variables y_s . The constraint is applied once to each scenario and the objective function of each scenario is weighted by the probability of the scenario occurring p_s . This equivalent deterministic formulation is:

$$\min_{x, y_s \forall s \in S} \sum_{s \in S} p_s f(x, y_s, s) \quad (3.4)$$

$$\text{s.t. } g(x, y_s, s) \leq 0 \quad \forall s \in S \quad (3.5)$$

Standard deterministic solvers (e.g., an LP solver if f and g are linear in x and y) can be used to solve the problem in this form.

If the number of scenarios is infinite (e.g., when there uncertain parameters are continuous random variables) or just very large, then a sample

average approximation (SAA) can be used to form an approximate but tractable version of the problem. This involves sampling a finite number of scenarios, say m scenarios, from the joint distribution, and giving them an equal weighting: $p_s = \frac{1}{m}$.

3.3 Deterministic EMS Problem

This section introduces the problem and formalises a deterministic version of it. The objective of the EMS is to minimise the cost of electricity and maximise occupant comfort through the control of household devices. To make the problem more concrete we have to make some assumptions, the first of which relates to the objectives.

Multi-objective problems can be interpreted and solved in a number of different ways [Collette and Siarry, 2013, Ehrgott and Gandibleux, 2000]. Our approach is to apply a linear scalarisation to the objectives, combining them into a single objective. This assumes that occupant discomfort can be weighted so that it is directly comparable to a monetary cost. Taking thermal comfort as an example, this assumes that it is possible for occupants to associate a price they are willing to pay to keep the temperature in their house within a certain range. To make this simple for occupants, we expect the EMS interface to present a “slider” which occupants can adjust between maximum comfort on one end and maximum frugality on the other. Behind the scenes this is interpreted as varying the weighting between the two objectives. Occupants can move the slider about until it matches their expectation through experience.

The next design choice we make is to discretise time, in effect taking a quasi-steady-state approximation of the control problem. This is appropriate because we are focused on the longer time scale device actions that require planning in advance. Faster transients and power system frequency control are better handled by myopic fast acting controllers, for example, dynamic-demand devices [Angeli and Kountouriotis, 2012] which can work alongside the EMS. As long as time is discretised into sufficiently small steps, the results will be close to what would be achievable with a fully continuous control signal.

For this discretisation, the index t represents the t -th time step, $\Delta\tau_t$ is the duration of the t -th time step in seconds and τ_t is the time at the end of the t -th time step in seconds, where $\tau_t > \tau_{t-1}$ and $\Delta\tau_t = \tau_t - \tau_{t-1}$ (see figure 3.1). This somewhat complicated formulation allows us to have variable step sizes over the forward scheduling horizon which, as we will discuss later, will allow us to focus computational resources where they are most needed.




Figure 3.1: The timings over a forward horizon with T time steps. The time steps (represented by the index t) range from 1 to T . The times τ are used to mark the times at which the steps begin and end, as the steps might have varying lengths.

3.3.1 Devices

At a high level, a device has a vector of variables $x_t \in \mathbb{R}^M$ and a vector of parameters¹ $r_t \in \mathbb{R}^L$ for each time step t . The variables represent device actions or states and the parameters represent external factors which impact the operation of the device (e.g., occupant usage requests or ambient temperatures). Devices have an operation cost function $f : \mathbb{R}^M \times \mathbb{R}^L \mapsto \mathbb{R}$ and a power function $h : \mathbb{R}^M \times \mathbb{R}^L \mapsto \mathbb{R}$ which take these vectors as inputs. The operation cost represents any comfort, fuel, deterioration or other cost associated with the operation of the device, and the power function returns the power that the device either consumes (+ve) or produces (−ve). Finally, devices have a vector-valued constraint function $g : \mathbb{R}^M \times \mathbb{R}^L \times \mathbb{R}^M \times \mathbb{R}^L \mapsto \mathbb{R}^N$ which applies to the variables and parameters in consecutive time steps and which is satisfied when the component-wise inequality $g(x_t, r_t, x_{t-1}, r_{t-1}) \leq 0$ holds. This generic constraint function can be used to do anything from placing bounds on variables to constraining state transitions.

3.3.2 Optimisation Problem

A house is simply a set of devices \mathcal{D} , together with bounds \underline{P} and \bar{P} on the instantaneous amount of power that the house can transfer to or from the grid. The EMS controls the decision variables for each of these devices in order to minimise the costs for the overall house.

We make use of a deterministic formulation of the EMS optimisation problem as a building block for the stochastic formulations. In this formulation we have a forward horizon of T time steps, over which the values of all parameters are known. The objective is to choose device actions to reduce the total cost over this future horizon. Inputs include the device initial states $x_{d,0}$, electricity price λ_t , house background power usage P_t^b (the aggregation of uncontrollable electrical consumption, e.g., lighting, entertainment and cooking), and device parameters $r_{d,t}$. The variables at each time step include the device variables and the total house power consumption P_t .

¹To simplify the fomulation we assume $\Delta\tau_t$ is one of the parameters in this vector.

The deterministic optimisation problem is as follows:

$$\min_{x_{d,t}, P_t} \sum_{t=1}^T \Delta\tau_t \lambda_t P_t + \sum_{d \in \mathcal{D}} \sum_{t=1}^T f_d(x_{d,t}, r_{d,t}) \quad (3.6)$$

$$\text{s.t. } P_t \in [\underline{P}, \bar{P}] \quad \forall t \in \{1, \dots, T\} \quad (3.7)$$

$$P_t = P_t^b + \sum_{d \in \mathcal{D}} h_d(x_{d,t}, r_{d,t}) \quad \forall t \in \{1, \dots, T\} \quad (3.8)$$

$$g_d(x_{d,t}, r_{d,t}, x_{d,t-1}, r_{d,t-1}) \leq 0 \quad \forall d \in \mathcal{D}, t \in \{1, \dots, T\} \quad (3.9)$$

The appropriate solver for this optimisation problem depends on the form of the device functions. For example, if all functions are linear, then an linear programming (LP) solver can be used, or if they are convex an interior point convex solver. The device models presented in the next section result in a mixed-integer linear program (MILP), which in our experiments is solved with Gurobi [Gurobi Optimization, Inc., 2014].

3.4 Device Models

What makes this problem particularly interesting is the behaviour of the different types of devices that the EMS needs to control. Before moving onto the stochastic part of the problem, we take a look at the some of these devices and the models that will be used in the experiments.

These descriptions are less formal in order to make them easier to understand. With reformulation they all fit into the general device definition of the previous section which is more useful at a higher level of abstraction. Unique symbols are used for variables and parameters instead of dealing directly with the vectors x_t and r_t for a device. Variables for device power $P_t = h(x_t, r_t)$ and operation cost $c_t = f(x_t, r_t)$ are used in place of the associated functions. Finally, individual constraints are provided instead of explicitly writing the vector-valued constraint function g .

The physical behaviour of devices has been approximated by linearising their physical equations. Only significant steps of this process are mentioned in the device descriptions. Parameters used in the experiments were selected to be representative of typical devices. For example, the EV battery capacity is equivalent to that of a Nissan Leaf, and the house floor area is typical for an average Australian house. Some parameters were difficult to source so had to be estimated, for example, the efficiency of the EV battery.

Battery— A battery has a stored energy state $E_t \in [0, \bar{E}]$ and action variables that represent the rate of charge/discharge: $P_t^c \in [0, \bar{P}^c]$, $P_t^d \in [0, \bar{P}^d]$. The battery has a fixed efficiency $\eta \in [0, 1]$, and the stored energy transitions according to the following constraint:

$$E_t = E_{t-1} + \Delta\tau_t (\eta P_t^c - P_t^d) \quad (3.10)$$

The overall power consumption of the battery is $P_t = P_t^c - P_t^d$. A battery lifetime cost c_t is associated with power that is discharged from the battery through a lifetime price ψ : $c_t = \Delta\tau_t\psi P_t^d$.

Electric Vehicle— An EV is the same as the above battery, but with a few additional constraints. Firstly the parameter $u^h \in \{0, 1\}$ indicates whether the EV is home, and the battery can only be charged/discharged when this is the case:

$$u_t^h = 0 \implies P_t = 0 \quad (3.11)$$

The parameter $P_t^r \in \mathbb{R}_+$ represents the power drawn from the battery whilst driving. This modifies the state update function as follows:

$$E_t = E_{t-1} + \Delta\tau_t \left(\eta P_t^c - P_t^d - P_t^r \right) \quad (3.12)$$

The final constraint is on the amount of energy stored in the battery. The house occupants provide an parameter $E_t^m \in [0, \bar{E}]$ that represents the minimum energy that the EV battery should have in it at each time step. This value represents how much energy the occupant expects to need if they drive away in the car at a given time. This is not a hard constraint as the draw from driving can bring the battery charge below this limit, but it ensures that if the battery power does fall below, then it charges back up as fast as possible (at the maximum charging rate \bar{P}^c).

$$u_t^h = 1 \implies E_t \geq \min [E_{t-1} + \Delta\tau_t \eta \bar{P}^c, E_t^m] \quad (3.13)$$

Hot Water Heating— The hot water system is made up of a storage tank and an electric heating element. We ignore the details of the interaction between hot and cold water in the tank and consider the state of the tank as being the amount of energy $E_t \in [0, \bar{E}]$ it contains above the inlet cold water temperature. The tank is considered empty of hot water when this value is zero. The action variable is the power setting of the electric heater $P_t \in [0, \bar{P}]$ at each time step. The parameter $P_t^d \in \mathbb{R}_+$ is the amount of power drawn from the tank to meet occupant demand. The energy state update function is given by:

$$E_t = E_{t-1} + \Delta\tau_t \left(P_t - P_t^d - P_t^l + P_t^u \right) \quad (3.14)$$

The variable $P_t^l \in \mathbb{R}_+$ represents thermal losses from the tank to the outdoor environment. The rate of loss depends on how full the tank is and the difference in temperature between the water set point $T^s \in \mathbb{R}$ and the outdoor temperature $T_t^o \in \mathbb{R}$ through a thermal resistance $R \in \mathbb{R}_+$:

$$P_t^l = \frac{1}{R} \frac{E_t}{\bar{E}} (T^s - T_t^o) \quad (3.15)$$

The variable $P_t^u \in \mathbb{R}_+$ is a recourse variable that is used to indicate the amount of hot water demand which goes unmet, i.e. water drawn from the

tank when it is empty. This is heavily penalised as a cost c_t through an unmet demand price ψ : $c_t = \Delta\tau_t\psi P_t^u$.

The hot water system has a minimum stored energy level $E^m \in [0, \bar{E}]$, much like the electric vehicle. If drawn water brings the energy level of the tank below this value then the heater must work as hard as possible to bring the energy back up. Occupants can adjust this parameter to reduce the likelihood of running out of hot water.

$$E_t \geq \min \left[E_{t-1} + \Delta\tau_t \left(\bar{P} - P_t^d - P_t^l + P_t^u \right), E^m \right] \quad (3.16)$$

Heating and Cooling— The house temperature is controlled by a heat pump that heats/cools water, which is then pumped through piping embedded in the floor of the house. The temperatures of the floor and the air in the room $T_t^f, T_t^a \in \mathbb{R}$ are state variables. The action variables are the amount of thermal energy used to heat or cool the floor of the house $P_t^h, P_t^c \in \mathbb{R}_+$. This is limited by the heat pump electrical power consumption $P_t \in [0, \bar{P}]$ through heating and cooling coefficients of performance (COPs) $\eta_t^h \in [\underline{\eta}^h, \bar{\eta}^h], \eta_t^c \in [\underline{\eta}^c, \bar{\eta}^c]$:

$$P_t = \frac{1}{\eta_t^h} P_t^h + \frac{1}{\eta_t^c} P_t^c \quad (3.17)$$

The COPs depend on the temperatures of the two thermal wells between which the heat pump is operating. We assume that the internal thermal well is at a constant temperature and that the external well is at the outdoor temperature $T_t^o \in \mathbb{R}$. We approximate the COPs as linear functions of T_t^o for some constants $a^h, a^c \in \mathbb{R}_+$ and $b^h, b^c \in \mathbb{R}$, with hard upper and lower limits:

$$\eta_t^h = \min \left[\max \left[a^h T_t^o + b^h, \underline{\eta}^h \right], \bar{\eta}^h \right] \quad (3.18)$$

$$\eta_t^c = \min \left[\max \left[-a^c T_t^o + b^c, \underline{\eta}^c \right], \bar{\eta}^c \right] \quad (3.19)$$

Heat can transfer between the floor and the outdoor environment $P_t^{fo} \in \mathbb{R}$, the floor and the air in the room $P_t^{fa} \in \mathbb{R}$, and the air in the room and the outdoor environment $P_t^{ao} \in \mathbb{R}$. We use simple lumped thermal resistances $R^{fo}, R^{fa}, R^{ao} \in \mathbb{R}_+$ to govern these heat flows:

$$P_t^{fo} = \frac{1}{R^{fo}}(T_t^f - T_t^o), \quad P_t^{fa} = \frac{1}{R^{fa}}(T_t^f - T_t^a), \quad P_t^{ao} = \frac{1}{R^{ao}}(T_t^a - T_t^o) \quad (3.20)$$

The temperature state update is given by:

$$T_t^f = T_{t-1}^f + \frac{\Delta\tau_t}{m^f \kappa^f} \left(P_t^h - P_t^c - P_t^{fo} - P_t^{fa} + A^f I_t \right) \quad (3.21)$$

$$T_t^a = T_{t-1}^a + \frac{\Delta\tau_t}{m^a \kappa^a} \left(P_t^{fa} - P_t^{ao} + P_t^g \right) \quad (3.22)$$

where $m^f, m^a, \kappa^f, \kappa^a \in \mathbb{R}_+$ are the floor and air, mass and specific heat capacity coefficients respectively. Sunlight enters through the windows at an irradiance $I_t \in \mathbb{R}_+$ and lands on a floor area $A^f \in \mathbb{R}_+$. The input $P_t^g \in \mathbb{R}_+$ is thermal power generated by occupant metabolisms and background electric appliances which contributes to heating the air in the room.

The final relation we have is for the comfort cost c_t which depends on the distance of the air temperature from an occupant specified set point temperature $T_t^s \in \mathbb{R}$. The occupants also specify two time-varying comfort prices ψ^a, ψ^b , one of which is only included after a threshold temperature difference ΔT^{th} :

$$c_t = \begin{cases} \Delta\tau_t(\psi_t^a + \psi_t^b)|T_t^a - T_t^s| & \text{if } |T_t^a - T_t^s| \geq \Delta T^{th} \\ \Delta\tau_t\psi_t^a|T_t^a - T_t^s| & \text{otherwise} \end{cases} \quad (3.23)$$

This is a flexible way of allowing occupants to represent their thermal comfort preferences. They can have relatively low costs associated with temperatures close to their ideal temperature, but then have a much higher cost take over once the threshold is exceeded.

The ASHRAE [2013] handbook describes thermal comfort as a complicated function of not just the air temperature, but also other factors including the clothing worn by occupants, humidity, radiative heat transfers and the season. Many of these factors will be slow changing with the seasons or a function of the house design itself. Our simple but flexible temperature set point and temperature bounds can be adjusted by occupants as the seasons change. Alternatively, something like the Nest thermostat can be used to learn the best temperature set point as a function of a whole range of other factors, from intuitive occupant feedback.

Shiftable Load— Shiftable loads are devices that need to run once within a time window. An occupant sets two parameters: a start time t^s and a last allowed start time t^l , between which the controller must schedule the device to run. Examples of this kind of device include washing machines, clothes dryers and dish washers. We model non-preemptive shiftable loads which can have time varying power consumptions.

The variables $u_t \in \{0, 1\}$ are used to indicate when the shiftable load starts. A value of $u_t = 1$ indicates that the run starts at time step t . A shiftable load has a function $\chi : \mathbb{R}_+ \mapsto \mathbb{R}_+$ that represents the cumulative energy consumed by the device. It is a monotonic function which takes a run duration and returns the cumulative amount of energy that the device has consumed up to that duration. Constraints on the start time indicator u_t and the device powers $P_t \in \mathbb{R}_+$ are given by:

$$\sum_{k=t^s}^{t^l} u_k = 1, \quad P_t = \sum_{k=t^s}^t u_k \frac{\chi(\tau_t - \tau_{k-1}) - \chi(\tau_{t-1} - \tau_{k-1})}{\Delta\tau_t} \quad (3.24)$$

Photovoltaics— PV panels have no action variables as the amount of electricity they generate is purely determined by the solar irradiance parameter. We model a PV system ignoring temperature and shading effects and by assuming the panels lay on a horizontal surface. The generated electric power $P_t \in \mathbb{R}_-$ is then a simple function of the panel area $A \in \mathbb{R}_+$, efficiency $\eta \in [0, 1]$ and global irradiance parameter $I_t \in \mathbb{R}_+$: $P_t = -\eta AI_t$.

3.5 Stochastic EMS Problem

The deterministic EMS formulation requires perfect knowledge about what will happen over the time horizon, while in practice the values of some of the device parameters will be unknown in advance. These include the outdoor temperature, solar irradiance, internal heat generation, hot water demand, EV usage and shiftable load requests. Future electricity prices and background house power consumption are also unknown. The values of these parameters only become known to the EMS in real time (e.g., a measurement of temperature), or in some cases a small amount of time in advance (e.g., dynamic prices published half an hour in advance).

This is a problem because the state-based nature of devices requires the EMS to plan ahead in order to make sensible decisions. The approach we take is to model the parameters as random variables, and then use online stochastic optimisation [Van Hentenryck and Bent, 2006] in the EMS to exploit these statistical models in order to make the best decisions on average. The EMS can learn these models over time using data it collects.

3.5.1 Random Parameters

We use the term random parameter in place of random variable to avoid confusing them with the decision variables in our optimisation problem. We introduce a vector $W_t = (W_{t,1}, W_{t,2}, \dots)$ for each time step which contains the random parameters for the electricity price λ_t , background house power P_t^b and device parameters $r_{d,t}$. A random parameter may depend on random parameters within the same or previous time steps, but not on future time steps.

A value of the random parameter $W_{t,k}$ is represented by $w_{t,k}$ (a random variate). An asterisk $*$ is used to indicate a known value of a parameter. For example, if $W_{t,k}$ represents the random parameter for outdoor temperature at time step t , then $w_{t,k}^*$ denotes its actual measured value.

At a real world time of τ , the indices of all known parameters for time step t is given by the set $K_{\tau,t}$ (recall that a parameter becomes known a short time in advance or only in real-time). These sets monotonically grow in size as τ increases.

3.6 Online Stochastic Optimisation

Online stochastic optimisation algorithms make decisions for one step at a time using stochastic information for any unknowns. After each time step more parameter values become known and the result of actions is revealed. Decisions for the next period are then computed and the process is repeated. Online stochastic optimisation has been used successfully on a wide variety of problems (e.g., see Powell et al. [2012], Van Hentenryck and Bent [2006]).

Our algorithms make use of a receding horizon, as illustrated in figure 3.2. Optimisation is performed for each horizon using stochastic information for any unrevealed parameters and then the actions for the first time step are executed in the real world. In the formulations that follow we assume that the horizon is T time steps long and that the first step in the horizon corresponds to $t = 1$. The online algorithm performs its optimisation for the horizon when the current time is equal to the time at the start of the horizon, in this case $\tau = \tau_0$.




Figure 3.2: Receding horizon for 3 consecutive iterations.

The next section discusses the executives which are used to actually implement the scheduled decisions, and then in the following sections two approaches to solving the stochastic optimisation part of the problem for each horizon: an expectation and a 2-stage algorithm.

3.6.1 Executives

Some parameters might still be uncertain in the first time step, which means that it might not be possible to execute the decisions made by the optimisation directly as given. For example, if the hot water heater is scheduled to run at full power, but the demand for hot water turns out to be less than expected, then the tank might be heated beyond its limits.

In order to manage this problem, we develop simple executives which take in the scheduled actions and modifies them based on what actually occurs in real-time as the uncertainty is revealed. Our EMS therefore is comprised of two parts: an online scheduling algorithm, and an executive controller that implements the schedule decisions for the next time step using very simple policies that can run in real-time. When shorter time steps are used in the online scheduling, the intervention of an executive will become less necessary.

The core action variables for the devices are: the power into the device for the battery, EV and hot water tank; the heating and cooling power of the underfloor heat pump; and the start time of the shiftable loads. The executive modifies these actions based on what happens during the first time step in real-time, and calculates the resulting values of all other state or auxiliary variables at the end of the time step. Using the hot water heating example from above: if the tank reaches its upper limit during the first time step, then the executive cuts back the scheduled power so that the tank is not overheated.

3.6.2 Expectation Formulation

The expectation online algorithm takes the conditional expected value of any unknown parameters in the optimisation horizon, and solves the deterministic version of the problem given in equations (3.6–3.9). We use the term expected value loosely because, in truth, the expected value is used only where it makes sense, which is typically for continuous inputs. For the rest of the inputs, the most likely value is calculated instead. For example, the expected value is used for outdoor temperatures and the most likely value is used for the washing machine requests.

Both of these calculations are performed using the joint distribution for any unknown parameters in the horizon, conditioned on any known parameters in the horizon and prior to it. For example, assuming that the random parameters have a dependence that stretches at most Q time steps into the past, the probabilities of the unknown random parameters in the horizon are given by:

$$\begin{aligned} \mathcal{P} \big(& W_{t,k} = w_{t,k} \quad \forall k \notin K_{\tau_0,t} \quad \forall t \in \{1, \dots, T\} \\ & | W_{t,k} = w_{t,k}^* \quad \forall k \in K_{\tau_0,t} \quad \forall t \in \{-Q, \dots, T\} \big) \end{aligned} \quad (3.25)$$

3.6.3 2-Stage Formulation

In this algorithm, 2-stage stochastic programming is used within each horizon. This provides an approximation to a full multi-stage stochastic program which are, in general, known to be extremely computationally challenging [Shapiro, 2006].

The 2-stage algorithm uses more information from the random parameter probability distributions by working directly with samples instead of the expectation. Figure 3.3 provides a selection of samples and the expected value for outdoor temperature over 12 hours. This highlights how a collection of samples can capture the variance of the random parameter.

Traditionally, in 2-stage stochastic programming there is no uncertainty in the first stage [Shapiro et al., 2009]. The most natural choice of a first stage is the first time step but, as we have discussed, the first time step may have uncertainty in it. There are two ways forward.




Figure 3.3: Five conditional samples and expected value (dashed line) for outdoor temperatures over 12 hours.

The first approach is to treat the first time step as fully deterministic, by taking the expected values of uncertain parameters in the first time step (as done in the previous section). The first stage includes time step 1, and the second stage time steps $2, \dots, T$. The scenarios in the second stage are then sampled from the joint distribution of random parameters in the second stage, conditioned on any known parameters in and prior to the second stage, and the expected values that were calculated for the unknown first stage parameters. Using $w_{1,k}^\dagger$ to represent the expected values in the first stage, the joint distribution to sample from is:

$$\begin{aligned} \mathcal{P} \big(& W_{t,k} = w_{t,k} \quad \forall k \notin K_{\tau_0,t} \quad \forall t \in \{2, \dots, T\} \\ & | \quad W_{t,k} = w_{t,k}^* \quad \forall k \in K_{\tau_0,t} \quad \forall t \in \{-Q, \dots, T\}, \\ & W_{1,k} = w_{1,k}^\dagger \quad \forall k \notin K_{\tau_0,1} \quad \big) \end{aligned} \quad (3.26)$$

The sample average approximation is used to limit the problem to m scenarios, which are given by the set $S := \{1, \dots, m\}$. We subscript variables and parameters by an $s \in S$ to indicate which scenario it belongs to. The

2-stage optimisation problem is:

$$\min_{x_{d,t,s}, P_{t,s}} \frac{1}{m} \sum_{s \in S} \left[\sum_{t=1}^T \Delta \tau_t \lambda_{t,s} P_{t,s} + \sum_{d \in \mathcal{D}} \sum_{t=1}^T f_d(x_{d,t,s}, r_{d,t,s}) \right] \quad (3.27)$$

$$\text{s.t. } x_{d,1,s_1} = x_{d,1,s_2} \quad \forall d \in \mathcal{D}, s_1, s_2 \in S \quad (3.28)$$

$$P_{1,s_1} = P_{1,s_2} \quad \forall d \in \mathcal{D}, s_1, s_2 \in S \quad (3.29)$$

$$P_{t,s} \in [P, \bar{P}] \quad \forall t \in \{1, \dots, T\}, s \in S \quad (3.30)$$

$$P_{t,s} = P_{t,s}^b + \sum_{d \in \mathcal{D}} h_d(x_{d,t,s}, r_{d,t,s}) \quad \forall t \in \{1, \dots, T\}, s \in S \quad (3.31)$$

$$g_d(x_{d,t,s}, r_{d,t,s}, x_{d,t-1,s}, r_{d,t-1,s}) \leq 0 \quad \forall d \in \mathcal{D}, t \in \{1, \dots, T\}, s \in S \quad (3.32)$$

Equations (3.28) and (3.29) tie together the variables in the first time step (first stage) so that they share a common value. As discussed above, expected values were used for all uncertain parameters in the first time step, which means that for a device d and any two scenarios s_1 and s_2 : $r_{d,1,s_1} = r_{d,1,s_2}$.

The second 2-stage approach is to represent the behaviour of the executives for the first time step within the scheduler. The first stage then only represents the first time step device action variables. The second stage covers the whole horizon, where the scenarios are sampled from the joint distribution given by (3.25). The first stage variables are linked with the first time step in the second stage through constraints that implement the device executives.

For example, for the hot water heating (with reference to its model in section 3.4), we link the first stage action variable P' to the second stage variable $P_{1,s}$ in the first time step for scenario s with the relation:

$$P_{1,s} = \min(\max(P', \underline{P}_s), \bar{P}_s) \quad (3.33)$$

The parameters \underline{P}_s and \bar{P}_s are calculated for the particular scenario of interest, depending on the values of $P_{1,s}^d$ and $T_{1,s}^o$, and the amount of energy in the tank at the beginning of the horizon E_0 . They represent the minimum and maximum amount of power that can go toward heating the tank in order to satisfy the minimum tank energy heating trigger, and so that the tank does not exceed its maximum energy. The relation above can be implemented as a piecewise-linear constraint between the two variables $P_{1,s}$ and P' .

We do not formalise the second approach here or provide its results in detail because, as we will discuss in section 3.8.2, it does not produce results that are much different from the first approach.

3.7 Parameter Stochastic Models

In this section we present the stochastic models that we developed for our experiments. The focus is not to provide high accuracy models of the underlying processes, which in itself is a significant research task, but rather for them to be representative enough so that a reliable comparison of the EMS strategies can be made.

A particular requirement of the 2-stage EMS strategy is the generation of not just the expectation of the random parameters into the future, but also the generation of possible scenarios for these parameters conditioned on past events. We investigated a number of different model types before settling on generalised additive models for the continuous parameters like temperature, and Markov models for the more discrete occupant driven behaviours such as shiftable load requests.

3.7.1 Generalised Additive Models

A generalised additive model (GAM) [Hastie and Tibshirani, 1990] is a hybrid statistical model which combines properties of a generalised linear model (GLM) and an additive model (AM). For a GAM a probability distribution from the exponential family is selected for the response random variable Y . The expectation of the response is related to the predictor variables x_i through a link function g that has been appropriately selected for the chosen distribution:

$$g(\mathbb{E}[Y]) = \beta_0 + f_1(x_1) + f_2(x_2) + \dots \quad (3.34)$$

where the functions f_j are some parametric or non-parametric smooth functions. The backfitting algorithm is one such algorithm that can be used for performing estimation.

The `mgcv` R package [Wood, 2011] was used for fitting the models and generating scenarios for the continuous stochastic parameters. Through experimentation, we found a Gaussian distribution produced the best results for each model, which in effect reduces the model from a GAM to an AM. However, we make use of some of the more advanced features in the `mgcv` tool, such as using multivariate smooth functions, which are absent from the standard AM formulation.

We will now describe some of the models that were used for the continuous random parameters. The models were constructed to predict the value of a stochastic parameter r_t at the upcoming time step, using the value of the parameter at the previous time step r_{t-1} , the time of the day *tod*, the day of the week *dow* and other quantities as predictors.

Weather forecast data was used as predictors for the parameters temperature, global irradiance, wind and total network electricity demand (used to

estimate the electricity price). These forecasts and actual data for the parameters relevant to the states of NSW and ACT in Australia were obtained from the Bureau of Meteorology (BOM) and the Australian Energy Market Operator (AEMO) and used to train the models and in EMS experiments. The forecasts include daily predictors of minimum and maximum temperatures $tmin, tmax$ and morning and afternoon cloud cover $cldam, cldpm$ and wind speeds $windam, windpm$.

The mgcv models are presented below in mgcv pseudo-code for the continuous stochastic parameters. The identify link function was used for each model, in line with the use of Gaussian distributions. The s functions represent smoothing functions which optionally have their basis dimension restricted by setting k .

$$\mathbb{E}[\text{temp}_t] = s(\text{temp}_{t-1}) + s(\text{tod}, tmax, tmin, k = 20) \quad (3.35)$$

$$\mathbb{E}[\text{global}_t] = s(\text{global}_{t-1}) + s(\text{tod}, cldam, cldpm, k = 20) \quad (3.36)$$

$$\mathbb{E}[\text{wind}_t] = s(\text{wind}_{t-1}) + s(\text{tod}, windam, windpm, k = 20) \quad (3.37)$$

$$\mathbb{E}[\text{demand}_t] = \text{demand}_{t-1} + s(\text{tod}, dow, tmax, k = 20) \quad (3.38)$$

From top to bottom these are: the outdoor temperature, global irradiance, wind and aggregate network load. The wind and aggregate network load are used in predictions of the dynamic price as we will discuss shortly.

Once the models have been fit to training data, the expected values of the continuous variables over a forward horizon (for the distribution 3.25) can be calculated by repeatedly calling the mgcv predict function. Each time it is conditioned on the known predictors and previous values. In order to generate a scenario, a sample is drawn from a normal distribution with standard deviation equal to the result returned by the original fit, and then added to the result of the prediction step. This value is then used to condition the predictor for the next prediction, and the process is repeated until enough values for the number of time steps of interest have been generated.

Figure 3.4 shows five scenarios sampled from the outdoor temperature model. The temperatures are sampled for 10 days into the future, conditioned on the current temperature and the forecast minimum and maximum temperatures.

The best approach to implementing dynamic pricing in retail markets is still an open question. It will likely have a shape that is representative of the wholesale market, but with less volatility. We designed our dynamic price to be a quadratic function (which represents an increasing marginal supply price [Ramchurn et al., 2011]) of the amount of power that fossil fuel sources must supply to meet total network load. This is the total network demand (estimated with a GAM) minus the generation from renewable sources such as wind and solar. The dynamic price is only revealed to the EMS 30 minutes in advance.




Figure 3.4: Five scenarios for outdoor temperature, generated from the forecast minimums and maximums (dashed lines) over ten days.

The resulting prices have a similar magnitude to wholesale market prices in the NEM (see section 2.1). This is only a part of the overall residential retail prices, which also include network tariffs and other retail fees, but it is the part that has the strongest time dependence. This means that the costs presented in our experiments section should not be treated as absolute values as they do not include these additional retail charges and fixed costs.

3.7.2 Markov Models

A Markov model uses a random variable to model the state of the system, where the distribution of the variable in future states only depends on the current state, not on the states preceding it. In particular we make use of a semi-Markov models, which allows the distribution to also rely on the amount of time that the model has been in the current state.

Semi-Markov models were used to capture the behaviour of four occupants of a particular house in the ACT in Australia. These models provide the consumption patterns and profiles for the stochastic parameters such as hot water demand, uncontrollable energy consumption, shiftable load requests and EV usage. Each occupant had its own model that captured key activities, for example, watching TV, taking a shower and driving to work. Two dedicated models for clothes washing/drying (see figure 3.5) and dish washing were also developed. In addition to the states, the models also

specify the probability of transitioning from one activity to the next within certain time windows.

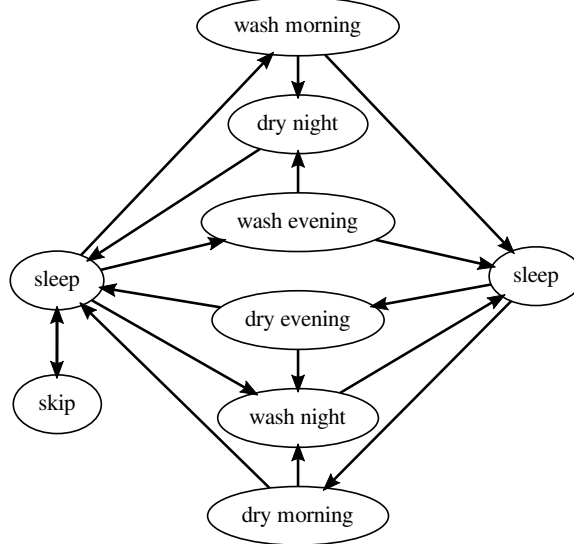


Figure 3.5: Model of washing machine and clothes dryer usage.

A state transition can trigger the use of a resource. A simple example is the *cooking* activity/state which can trigger use of the oven resource. Another one is the *work* activity which removes the EV resource from the home and turns the lighting off. These resources have a direct impact the parameters in our optimisation problem. Conditional sampling through these models is used to generate scenarios.

Whilst this scheme was convenient for our experiments, other more data-driven options are possible: we could simply gather and use a database of raw scenarios, or learn model parameters from disaggregated demand data [Kolter et al., 2010, Parson et al., 2012].

3.8 Experiments

We implemented the expectation (Expect) and 2-stage (Stage) online algorithms in Python using Gurobi as a backend to solve the MILPs for each horizon. The devices in section 3.4 were implemented and included in the experimental house, and conditional samplers were created for the uncertain parameters in section 3.7. We developed a simulator that uses the same physical equations as the optimisation to simulate the execution of actions in the real world. We compare the performance of the expectation and 2-stage controllers with two purely reactive controllers (Naive and Simple), and a controller that has perfect information (Perfect).

The *naive* reactive controller represents a household that either has no ability or no desire to respond to a dynamic price. It starts shiftable devices as soon as a request is received, fills up the hot water tank in off-peak hours, charges the EV only if it is below the requested minimum level, maintains the room at the set point temperature and never uses the battery bank. This corresponds to what people do now, in a setting without an EMS.

The *simple* reactive controller uses simple device action policies to decide how to respond to changes in price. It delays running a shiftable device until it reaches either a cheap price or the last available start time; uses thresholds about a moving average of the price to decide when to charge or discharge the battery, EV and hot water system; and maintains the room at the set point temperature like the naive controller.

The *perfect* controller has perfect foresight about what will happen in the future. It optimises the deterministic problem in equations (3.6–3.9) over the whole experiment duration with full knowledge of w_t^* . This controller (which is not possible in a real setting) produces a lower bound on the objective, for comparison with the other controllers.

3.8.1 Controller Comparison

Nine problem instances were generated for the controllers to be tested on. Each instance spans 7 days and is made up of known parameter values that are typical of a week in February in the ACT Australia. The online algorithms were set up to use 16 hour optimisation horizons (the gains are minimal beyond this length as seen in section 3.8.4), with 15 minutes for the first two time steps and 30 minute time steps for the remainder of the horizon. This was done to demonstrate the ability of the algorithm to focus the computation on the more immediate actions, where there is more certainty about random parameters. The reactive and perfect controllers had 15 minute time steps. The 2-stage algorithm sampled 60 scenarios for each second stage. The amount of time spent optimising in Gurobi per day was on average 1 second for the expectation algorithm and 4 minutes for the 2-stage algorithm (using a single core of an Intel i7-2600 3.4GHz CPU). Whilst the 2-stage is much slower, its computation time is still small when spread out over a day.

The controller costs are plotted in figure 3.6 for each problem instance. These results are adjusted to account for any energy that remains in the battery, EV, or hot water system at the end of the experiment. This is done by valuing the left-over energy at the average price for the last 24 hours. Without this adjustment it would not be a fair comparison since any controller that anticipates the need to store energy for a future purpose would perform poorly if it does so just before the experiment ends. This is an artefact of the finite length of our experiments. With longer durations this problem is less significant.

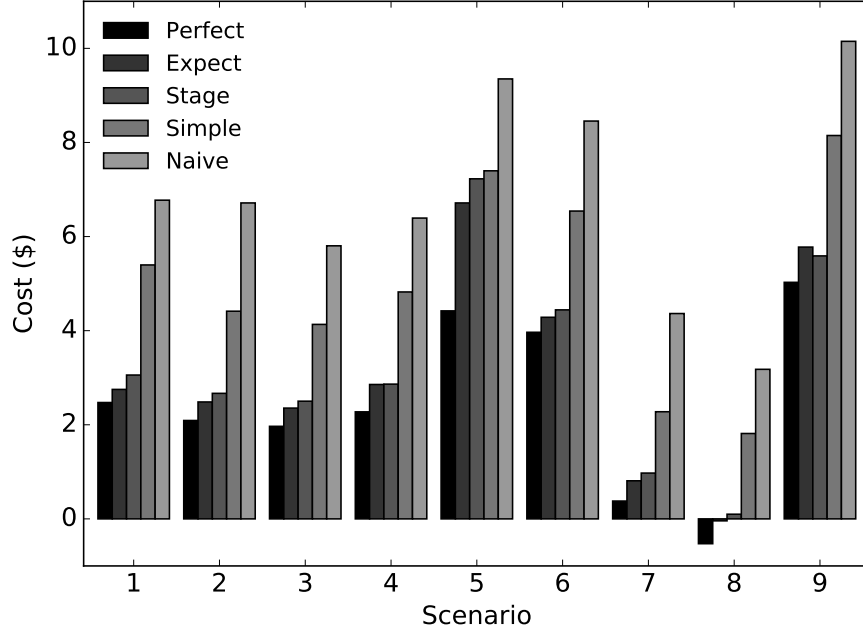


Figure 3.6: Comparison of controller costs for each problem instance.

The results show that the expectation and 2-stage algorithms get quite close to the performance of the controller with perfect foresight (except for instance 5) and they achieve significant cost reductions over the two reactive controllers (approximately 35% and 52% less costs than the simple and naive controllers). This justifies the assumption we made earlier that planning EMS actions ahead of time can produce much better outcomes than purely reactive control. It also highlights how an EMS can significantly reduce costs in a dynamic pricing environment, as all the controllers greatly improve over the naive results (which represents a house with no EMS).

To test the robustness of the 2-stage algorithm, we reran the experiments 4 times using different random seeds for the scenario generator. The standard deviation of the results was typically less than 2 cents for each instance, except instance 5 which had a standard deviation of 42 cents (caused by a volatile hot water consumption scenario that will be discussed shortly).

The expectation controller slightly outperforms the 2-stage controller on average and in each individual instance except 9. This tells us that the nature of the stochastic EMS problem is handled well by the simpler online algorithm that just looks at the expectation of unknown parameters. We anticipated that the 2-stage approach would do at least as well as the expectation approach, but in fact it appears to be let down by the volatility of some of the unknown parameters and the limited number of samples taken in its second stage, which we will discuss in the next section.

3.8.2 Device Cost Contributions

Figure 3.7 shows the costs attributed to each device, averaged across all instances. The costs for the expectation and 2-stage algorithms are very similar, and also close to the perfect algorithm. The devices that gain the most benefit from using these receding-horizon algorithms over the reactive controllers are the battery, EV and underfloor heating/cooling system (HVAC). The shiftable devices also gain a benefit, but the overall affect on cost for these is much smaller.

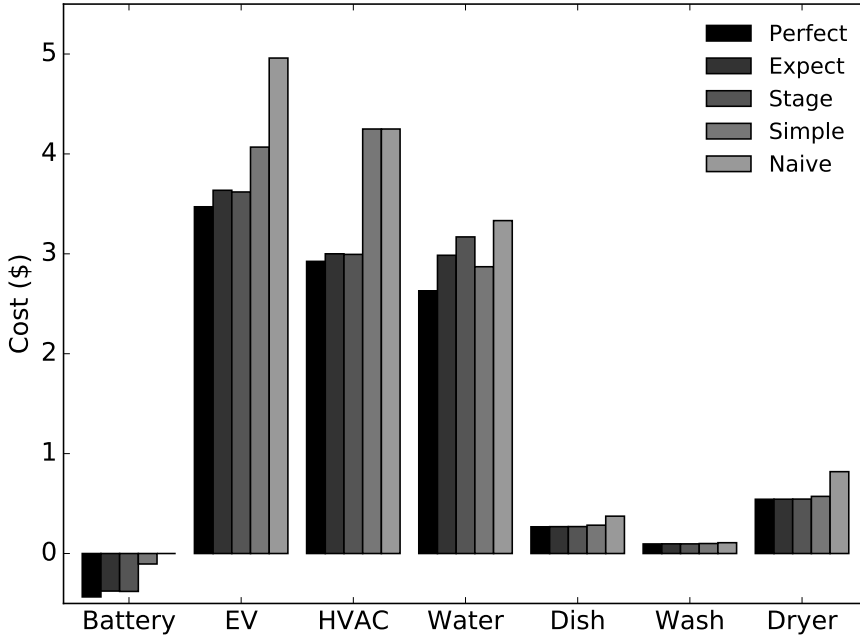


Figure 3.7: Costs for each device averaged over all problem instances.

The 2-stage algorithm performs fractionally better than the expectation algorithm on all devices except the hot water heater. In fact, its poor performance on the hot water heater outweighs the small gains it makes on the other devices. Also, interestingly, the simple reactive controller actually performs better on the hot water heater than both of the online receding horizon algorithms.

If we recall the hot water heater model, it has a recourse variable that represents the unmet hot water demand. This recourse variable has a very high cost which represents the dissatisfaction associated with a cold shower. This along with the high volatility of hot water use (e.g., the large spikes in consumption for shower usage), means that mistakes in predictions can have a significant cost associated with them.

For this reason the reactive controllers are designed to be quite conservative, and so do not take much advantage of price differences. Unfortunately

the expectation and 2-stage controllers get it wrong often enough that they actually do worse than the simple reactive controller. This is not too surprising for the expectation algorithm, but the 2-stage algorithm performs even worse than this.

The first thing we checked was whether or not this was just a statistical anomaly. We extended the problem instances from 9 to 36, but still experience the same trend. We also thought this might be caused by our treatment of the first time step unknowns. To completely remove all unknowns from the first stage we implemented the second approach described in section 3.6.3, where we explicitly model the executives in the optimisation. This only improved the results by 1.5% on average relative to the original 2-stage approach.

We tried to increase the number of second stage scenarios from 60 up to 100, but this did not significantly improve the results either. More samples beyond this point would start to become computationally prohibitive in both experimental and real world settings.

The best explanation we have is that the intrinsically volatile nature of hot water usage makes it difficult to apply stochastic programming directly to the probability distribution with only a limited number of samples. Extra outside rules (e.g., to keep a certain minimum charge level at certain times) can be used to make up for this lack of predictive power. For example, we see that the simple reactive controller that uses simple rules can produce good results for the hot water system.

A combination of the 2-stage approach with the reactive control rules for the hot water system would produce the best overall outcome. The expectation approach with the reactive hot water rules would only be a fraction worse off. From a practical point of view, the expectation online algorithm is faster and much simpler to implement. The stochastic models are also simpler to implement because they do not require a diverse range of scenarios to be generated.

3.8.3 Device Operation Examples

An example of the control provided by the 2-stage algorithm for the under-floor heating/cooling operation is provided in figure 3.8 for problem instance 1. This figure shows the room air temperature, outdoor temperature and the set point temperature. The white sections indicate when temperature control is required. The figure shows how the controller pre-cools the house before occupants arrive home on hot afternoons (for example, around the times 3.7, 4.7 and 5.7).

Figure 3.9 gives an example of the power exchanged between the house and the grid for one day, along with the price. As expected, most consumption occurs when the price is low and then sold back to the grid from the battery, EV and PV when it is high. The expectation and 2-stage con-




Figure 3.8: Example of underfloor heating/cooling control.

trollers follow the general trend of the perfect controller with some small divergences.

3.8.4 Horizon Length

Figure 3.10 shows the results of an experiment that investigates how the horizon length affects performance. This plot shows the performance of the perfect controller running as an online algorithm where it is restricted to only having perfect foresight a certain distance into the future (the horizon length). The experiment is performed on instance 1 for a number of different horizon lengths and the results are compared to the original perfect controller that can see the full 7 days. The results show that there is little to be gained by looking any further into the future than 20 hours.

3.9 Related Work

Much of the existing literature on residential energy management and DR focuses on deterministic formulations over a fixed horizon where the scheduler has perfect foresight [Ramchurn et al., 2011, Gatsis and Giannakis, 2012]. The work that has considered uncertainty in the problem typically focus on just one aspect (e.g., dynamic pricing) [Mohsenian-Rad and Leon-Garcia, 2010], or use simple stochastic models (e.g., no correlation between time steps) [Tischer and Verbic, 2011, Jacomino and Le, 2012]. Model-predictive

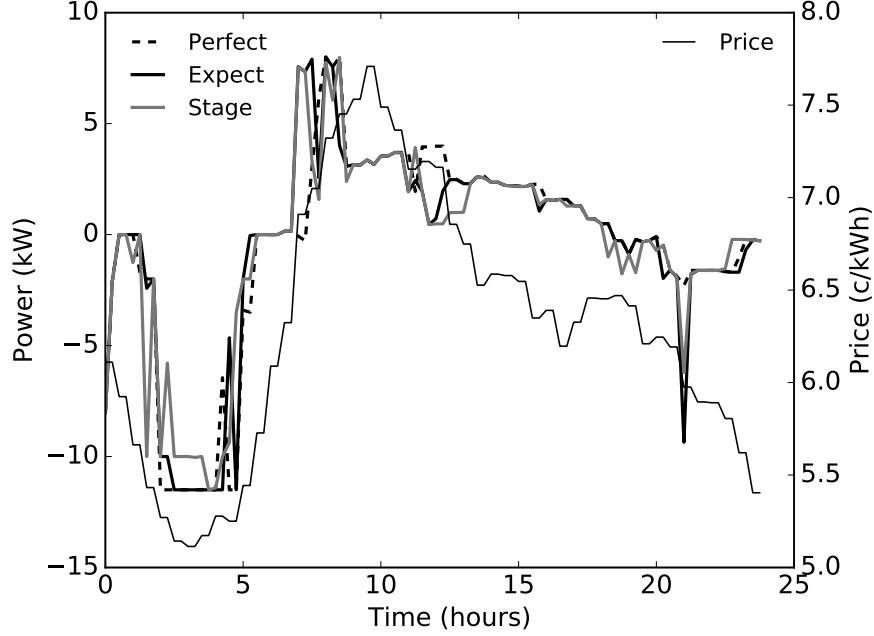


Figure 3.9: House power profiles for one day (inverse power is plotted: $-P_t$).

control has been used with uncertainty in device model parameters and measurement noise [Yu et al., 2012], but not the uncertainty used for the underlying processes.

Dynamic programming [Tischer and Verbic, 2011, Kim and Poor, 2011] and Q-learning [Levorato et al., 2010] have been used in conjunction with Markov decision process (MDP) formulations of the residential load scheduling problem, to generate policies that allocate power to each device. Dynamic programming has also been used to develop optimal policies for batteries exposed to uncertain PV production and uncertain price signals [Carvalho et al., 2012, Grillo et al., 2012]. MDP approaches suffer from severe scalability issues, especially since the state space needs to be discretised. Moreover, MDPs seem somewhat excessive for our problem, given that uncertainty does not depend on the decisions taken. Our online approach is more scalable and natural in the presence of exogenous uncertainty.

Tischer and Verbic [2011] found that acting on the basis of an optimal dynamic programming solution did not provide any benefit over acting on expectations. Our results extend this finding to cover more diverse uncertainty models, as we found in most instances that our expectation algorithm performs just as well as the 2-stage approach we developed.

The work closest to ours (which was done independently) compares robust optimisation against 2-stage stochastic programming for scheduling residential loads [Chen et al., 2012]. Uncertainty is restricted to the dynamic pricing which is known for the first stage but becomes uncertain




Figure 3.10: Costs for different horizon lengths compared to the cost achieved by the perfect controller.

thereafter. The objective includes minimising expected cost and the probability mass of “risky” scenarios whose price exceeds a certain threshold. Comfort is handled by imposing hard constraints under which appliances must run, rather than by inclusion into the objective. In this setting, 2-stage stochastic programming was observed to provide benefits over robust scheduling.

Industry is now offering products for home automation. While general energy management solutions are not yet widely available, there are well developed solutions for specific problems. A well known example is the Nest smart thermostat. Over time it learns occupant preferences from manual temperature adjustments, so that it can perform these adjustments automatically at different times of day and under different external factors. This has the potential to improve occupant comfort and to improve energy efficiency. It is more focused on these aspects than responding to wider electricity network pricing signals.

Reposit provides an EMS product that optimises the combined operation of solar PV and battery storage in homes. It takes into consideration the electricity tariffs, solar generation and household energy use to work out the best utilisation of a battery system. This EMS could be expanded to work with smart appliances and other controllable DETs in homes, as we have presented in this chapter.

The scope of our analysis goes significantly beyond these results, by exploring uncertainty from a large range of sources and by identifying the value of stochastic information. We enable richer sources of uncertainty

to be considered in our framework, by allowing inputs to be revealed at arbitrary points in time.

3.10 Conclusion and Future Work

The EMS we have developed in this chapter is crucial for enabling households to efficiently operate DETs, and through them modify their pattern of consumption. Creating this flexibility in household consumption is what is needed to help networks with balancing renewable energy and keeping the network within its operating limits.

To the best of our knowledge, this is the first work to produce a scalable and accurate solution in the presence of uncertainty about future prices, occupant behaviour and environmental conditions. Using models representative of physical devices and random processes, we have shown the monetary and comfort cost savings that can be achieved by using online stochastic algorithms over reactive control, and the comparison of performance between a 2-stage approach and acting on expectations. Studies such as the one provided in this chapter are important for rallying industry and customers towards more effective energy management schemes.

Further research is needed to investigate how closely reality can be modelled with random processes, and if in turn they are suitable for online learning. We also need to further investigate how time step sizes influence performance, and to conduct more experiments for different months of the year to get a broader understanding of the value of such technologies. The experimental setup we have developed can be used to experiment with and compare different pricing schemes, for example: TOU pricing, RTP and prices that change depending on whether the house is buying or selling electricity.

Chapter 4

Network-Aware Coordination

4.1 Introduction

The EMS presented in the previous chapter successfully reduces costs for a home exposed to dynamic electricity prices by automatically controlling local DETs. From the point of view of a utility, dynamic pricing can be used to coordinate the combined actions of many thousands of such homes; however, if not done carefully, this can cause undesirable herding or other unexpected impacts. In this chapter we present a method of calculating these prices, and hence for coordinating multiple EMSs, so that the combined result is the most efficient for the overall system.

Conventional networks use centralised markets to solve the similar problem of dispatching generators (see section 2.1.1). These markets, to one degree or another, seek to achieve an optimal power flow (OPF) [Glover et al., 2011, chapter 12], which is a traditional power systems optimisation problem that is concerned with minimising operational costs. The goal is to dispatch generators in a way that minimises costs, so that all loads are met and without overloading the network. Such markets and the traditional OPF problem were not designed to operate in the new world of prosumers where every customer is potentially an active participant. This massive increase in scale, the time-coupled behaviour of DETs and the unique preferences of prosumers means that a different approach is needed.

We envisage a future where network operators provide a competitive electricity market that anyone can participate in, and where the distinction between generators, loads and prosumers is removed. The overall operation of this market is managed through distributed algorithms and with each EMS calculating their own small local part of the overall problem. This will be of particular importance for the operation of microgrids, which require more finesse to ensure that demand and supply are balanced and that the

network is in a safe operating state in each instance.

Several works have applied distributed solving techniques to the problem of coordinating many participants [Kraning et al., 2014, Gatsis and Giannakis, 2012, Mohsenian-Rad et al., 2010]. These distributed algorithms greatly parallelise the problem and help to preserve the privacy of participants. As a by-product, they provide a natural market mechanism for allocating payments between participants. Theoretically, these algorithms require the problem to be convex in order to guarantee convergence to a globally optimal solution. However, the behaviour of many household loads are discrete in nature [Ramchurn et al., 2011], and the equations that govern how power physically flows on the network are non-convex.

We show that these theoretical problems can in practice be dealt with, specifically in the context of microgrids where the problem of balancing supply and demand is more challenging because individual participants have more influence. We show that for a distributed algorithm in a microgrid, exact non-convex power flow models perform well compared to inexact convex models, which makes them a valuable candidate in practice. Secondly, we find that the non-convex nature of discrete house loads to be a non-issue, and that in practice simple approaches to handling these discrete loads are effective at the microgrid level. By solving these problems, we show that the use of distributed algorithms for managing the balance of power in a microgrid is in practice not only possible, but also highly effective.

We formulate our EMS coordination problem as a multi-period OPF problem to account for multiple time steps over a day, which can be used as part of a day-ahead pricing scheme or, as we propose, a receding horizon control algorithm. We solve the multi-period OPF problems in a distributed manner by adapting the alternating direction method of multipliers (ADMM) approach presented in [Kraning et al., 2014]. We experiment with a range of power flow models of varying degrees of accuracy, to compare their relative behaviour in a distributed algorithm. We then introduce and compare several approaches layered on top of ADMM which manage the introduction of discrete variables into the problem. Technically, our contributions can be summarised as:

- A comprehensive experimental comparison of the convergence of five commonly used power flow models when used for distributed OPF in a microgrid context.
- The identification that the exact non-convex power flow model in practice not only converges in this context, but also finds near-optimal solutions in a timely fashion relative to other models.
- The introduction and comparison of three simple but effective approaches to managing the discrete shiftable loads that are typically found within houses.

Combined, these results show that distributed EMS coordination using ADMM can achieve near optimal solutions in a time frame that is practical for receding horizon control in this challenging microgrid setting, even though the theoretical results do not extend this far. This work brings distributed DR closer to the point where it can be deployed in real networks, including microgrids.

In the next two sections we provide some background on the different power flow models and distributed optimisation. In sections 4.4–4.6 we formulate the multi-period OPF problem and present the distributed algorithm we use to solve it. The test microgrid is introduced in section 4.8 before presenting our results in sections 4.9–4.10 on power flows and discrete decisions.

4.2 Power Flows

In the majority of cases, alternating current (AC) is used for transporting electrical energy across networks. AC power flow equations model the physical behaviour of such systems, making them useful in simulations and in decision making processes.

In this section we introduce network power flows and the different power flow equation variations that we experiment with later in this chapter. We assume a balanced 3-phase network, for which we use a single-phase AC equivalent. We also make a quasi-steady-state approximation of the system, which ignores transient effects and discretises time. Appendix A provides an introduction to AC power, which might be worth reading before continuing here.

4.2.1 Network Power Flow Equations

The networks we investigate are made up of three main categories of components: buses, lines and participants (which includes traditional generators, loads and prosumers). Participants connect to buses, and lines connect between buses.

We use some non-standard diagrams and notation throughout this thesis. This is done to make it easier to understand and formalise the problem for use with a distributed algorithm. A compact matrix representation of the power flow problem is typically used in power systems literature, whereas we use a more explicit and flexible component-based representation.

Figure 4.1 illustrates a common line diagram and a component-based diagram of the network. In the component-based form, a network is made up of buses (small circles), lines (rectangles) and participants (large circles). Components connect to buses through their terminals (small solid squares).

In power systems, network equations are typically formulated in terms of voltages and power flows instead of voltages and currents. Power flows




Figure 4.1: A 3-bus 3-line network with 4 participants (2 generators and 2 loads), shown as a line diagram (left) and a component-based diagram (right). A single-phase network with implied ground connections is used to represent an equivalent balanced 3-phase network.

are more natural to use, as many traditional generators and loads can be treated as constant power sources or sinks. Current formulations are more appropriate where there are only current sources and constant impedance components. Below we derive a power flow equivalent for Kirchhoff's current law (KCL).

Let each bus n have a complex voltage \mathbf{V}_n . Let \mathcal{T}_n represent the terminals connected to bus n , and \mathbf{I}_i the complex current entering the bus from terminal $i \in \mathcal{T}_n$. At each bus KCL must hold. It can be transformed into an equivalent power form by taking the complex conjugate of (A.8), and multiplying it by the voltage at the node. For a bus n :

$$0 = \sum_{i \in \mathcal{T}_n} \bar{\mathbf{I}}_i \mathbf{V}_n = \sum_{i \in \mathcal{T}_n} \mathbf{S}_i \quad (4.1)$$

In rectangular form this breaks into two equations which conserve power:

$$\sum_{i \in \mathcal{T}_n} P_i = 0 \quad (4.2)$$

$$\sum_{i \in \mathcal{T}_n} Q_i = 0 \quad (4.3)$$

P_i and Q_i represent the real and reactive power (relative to ground) entering the bus from connected terminal i . Components need to relate the quantities of \mathbf{V} and \mathbf{S} at their terminals. We have seen that for the bus this is given by a common voltage for all connected terminals, and the conservation of power. Additional constraints such as bus voltage bounds can be provided to represent operating limits. In the next section we will discuss the relations for a line, which for the power flow formulation is where much of the computational complexity comes in.

4.2.2 Lines

Models of varying degrees of accuracy exist for overhead and underground power lines (see Glover et al. [2011, chapter 5]). One of the simplest approaches is to model a balanced 3-phase power line as a series impedance. Higher accuracy finite element models can be used, but the series impedance is accurate enough for our problem where we are primarily interested in distribution lines which are relatively short. Using a and b to indicate the two terminals (as shown in figure A.1), Ohm's law provides the current flowing in a series impedance:

$$(\mathbf{V}_a - \mathbf{V}_b) = \mathbf{Z}\mathbf{I}_{ab} \quad (4.4)$$

The power consumed by the line is:

$$\mathbf{S}_{ab} = \bar{\mathbf{I}}_{ab}(\mathbf{V}_a - \mathbf{V}_b) \quad (4.5)$$

The power “flowing into” terminal a of the line is:

$$\mathbf{S}_a = \bar{\mathbf{I}}_{ab}\mathbf{V}_a \quad (4.6)$$

Cancelling the current term by using Ohm's law:

$$\mathbf{S}_a = \bar{\mathbf{Y}}(\bar{\mathbf{V}}_a - \bar{\mathbf{V}}_b)\mathbf{V}_a \quad (4.7)$$

$$= \bar{\mathbf{Y}}(V_a^2 - V_a V_b \angle(\theta_a - \theta_b)) \quad (4.8)$$

$$= (G - iB)(V_a^2 - V_a V_b \cos(\theta_a - \theta_b) - iV_a V_b \sin(\theta_a - \theta_b)) \quad (4.9)$$

Breaking it up into components:

$$P_a = GV_a^2 - GV_a V_b \cos(\theta_a - \theta_b) - BV_a V_b \sin(\theta_a - \theta_b) \quad (4.10)$$

$$Q_a = -BV_a^2 + BV_a V_b \cos(\theta_a - \theta_b) - GV_a V_b \sin(\theta_a - \theta_b) \quad (4.11)$$

The equations for terminal b can be obtained from the above by switching a and b where they appear. The end result is a set of 4 nonlinear equations which relate the terminal voltages in polar form to the powers in rectangular form. If the admittance is considered constant, then we have a set of 8 real variables and 4 equations.

Lines have an upper bound on how much current they can carry before they get too hot and either destroy something or deform and short out. This is called the line thermal limit. This is often approximated by setting a bound on apparent power. If S_{th} is the line thermal limit then:

$$P_a^2 + Q_a^2 \leq S_{th}^2 \quad (4.12)$$

$$P_b^2 + Q_b^2 \leq S_{th}^2 \quad (4.13)$$

We refer to these 4 equations and 2 inequalities as simply just the AC power flow model or sometimes the *exact* model.

Power flow equations are used in many network optimisation problems, including optimal power flow (OPF) [Coffrin et al., 2015, Momoh et al., 1999], optimal transmission switching (OTS) [Coffrin et al., 2014, Fisher et al., 2008], capacitor placement [Aguiar and Cuervo, 2005, Huang et al., 1996] and expansion planing [Taylor and Hover, 2011]. Unfortunately, the equations presented above are nonlinear and non-convex, which in general makes it difficult for optimisation techniques to find and/or prove globally optimal solutions to the above problems.

For these reasons, the equations are often either approximated or relaxed before being used in network optimisation problems. Approximations attempt to simplify the equations without losing too much accuracy. Convex relaxations of the equations ensure that when used in an optimisation context, the resulting objective is a lower bound on what is possible in the original problem. We present some of the alternative forms that we test with our distributed algorithm.

Linear DC (LDC)— This approximation is a common linearisation known as the linear DC model [Schweppe and Rom, 1970, Stott et al., 2009]. By assuming $V_a \approx V_b \approx 1$, $R \ll X \implies G \approx 0$, and $|\theta_a - \theta_b|$ is small, the equations simplify to:

$$P_a = -B(\theta_a - \theta_b) \quad (4.14)$$

$$Q_a = 0 \quad (4.15)$$

This reduces the number of variables from 8 to 4: P_a , P_b , θ_a and θ_b . Because $P_a = -P_b$, the model is lossless.

Dist-Flow (DF)— The dist-flow relaxation [Farivar et al., 2011, Baran and Wu, 1989] and an equivalent SOCP relaxation [Jabr, 2006, Bose et al., 2014] provide a convex relaxation of the line model which ignores voltage phase angles. By manipulating Ohm's law and relaxing two equalities (equations (4.18–4.19)) the following relations can be produced:

$$P_a + P_b = RI_{ab}^2 \quad (4.16)$$

$$Q_a + Q_b = XI_{ab}^2 \quad (4.17)$$

$$P_a^2 + Q_a^2 \leq V_a^2 I_{ab}^2 \quad (4.18)$$

$$P_b^2 + Q_b^2 \leq V_b^2 I_{ab}^2 \quad (4.19)$$

$$V_b^2 - V_a^2 + 2(RP_a + XQ_a) = (R^2 + X^2)I_{ab}^2 \quad (4.20)$$

The voltages and currents only appear as squared terms, which can be replaced by new variables, resulting in a convex problem with linear equalities and second-order cone constraints. The voltages at the buses must also be replaced with their squared versions.

Quadratic Constraint (QC)— The quadratic constraint model developed by Coffrin et al. [2015] and Hijazi et al. [2014] is also a convex relaxation. It is a strengthening of the DF model, which reintroduces the voltage phase angles. It introduces intermediate variables, takes relaxations of the sine and cosine terms in the exact model and sequentially applies McCormick relaxations to multilinear terms. The full formulation is omitted due to its length, but can be seen in Hijazi et al. [2014].

Quadratic Approximation (QA)— A quadratic approximation (QA) was proposed by Kraning et al. [2014]. While it does not appear to be in common use, it is provided here for comparison because of its use in a similar distributed optimisation setting. By assuming that $V_a \approx V_b \approx 1$ and that $|\theta_a - \theta_b|$ is small, the exact equations can be simplified to:

$$P_a - P_b = -2B(\theta_a - \theta_b) \quad (4.21)$$

$$P_a + P_b = \frac{1}{4G}(P_a + P_b)^2 + \frac{G}{4B^2}(P_a - P_b)^2 \quad (4.22)$$

A convex relaxation of this approximate model (not the original exact model) can be obtained by changing the equality in equation (4.22) to an inequality:

$$P_a + P_b \geq \frac{1}{4G}(P_a + P_b)^2 + \frac{G}{4B^2}(P_a - P_b)^2 \quad (4.23)$$

Others Models— Other common models include a semi-definite program [Bai et al., 2008] and linearisation that incorporates reactive power and voltage, the LPAC model [Coffrin and Van Hentenryck, 2014].

4.3 Distributed Optimisation

Distributed computation can greatly speed up certain calculations through parallelisation. It can also be used in circumstances where input data for the computation cannot be brought together to one central location, for example, when there are privacy concerns surrounding the data. These are some of the reasons why we apply a distributed algorithm to the prosumer coordination problem. In this section we introduce some of the different distributed optimisation techniques and provide sources for more information on the topic. This is not a comprehensive review of the literature, but rather a gentle introduction to some of the most significant approaches.

Bertsekas and Tsitsiklis [1989] present a range of distributed numerical techniques for solving different problems, from systems of equations to dynamic programming and constrained optimisation. Not all problems gain a speed advantage from being solved in a distributed manner, for example, when there is a tight coupling between the different parts of the problem.

To distribute an optimisation problem it is first decomposed into smaller subproblems. These subproblems can then be solved in parallel on the

same machine or on distributed infrastructure. Some problems are trivially decomposable while others may require a significant reformulation. The subproblem results are communicated either amongst each other or with a master problem (see figure 4.2). In all but the most simple problems, multiple solving and communication iterations are required before a solution is obtained for the overall problem. A master node can be used to coordinate the actions of subproblems (e.g., terminate the algorithm on reaching a solution), whereas a distributed approach is required when the problem is fully decentralised.




Figure 4.2: Decomposition into subproblems with and without master problem coordination.

Specialised algorithms have been developed for different types of optimisation problems. Some algorithms can provide global optimality guarantees whereas others only attempt to find a local optimum or approximate solution. The theory behind distributed optimisation is well developed for convex optimisation problems, with efficient algorithms that converge to a globally optimal solution.

For convex problems two decomposition methods are commonly used: primal and dual decompositions. They differ on whether the master problem solves a primal or a dual version of the original problem. We provide examples from Boyd et al. [2007] of these decompositions for the unconstrained optimisation problem:

$$\min_{x_1, x_2, y} f_1(x_1, y) + f_2(x_2, y) \quad (4.24)$$

Here y is a complicating variable as it prevents us from trivially separating the functions f_1 and f_2 . For the primal decomposition the problem is split up into a master problem and two subproblems:

$$\min_y \phi_1(y) + \phi_2(y) \quad (4.25)$$

$$\phi_1(y) := \min_{x_1} f_1(x_1, y) \quad (4.26)$$

$$\phi_2(y) := \min_{x_2} f_2(x_2, y) \quad (4.27)$$

For the dual decomposition we first reformulate the original problem by

duplicating the complicating variables and tying them together with a constraint:

$$\min_{x_1, x_2, y_1, y_2} f_1(x_1, y_1) + f_2(x_2, y_2) \quad (4.28)$$

$$\text{s.t. } y_1 - y_2 = 0 \quad (4.29)$$

The Lagrangian for this problem is:

$$\mathcal{L}(x_1, x_2, y_1, y_2, \nu) = f_1(x_1, y_1) + f_2(x_2, y_2) + \nu(y_1 - y_2) \quad (4.30)$$

Where ν is a Lagrangian multiplier. Finally, the dual decomposition consists of a master problem and two subproblems based on this Lagrangian:

$$\max_{\nu} \phi_1(\nu) + \phi_2(\nu) \quad (4.31)$$

$$\phi_1(\nu) := \inf_{x_1, y_1} f_1(x_1, y_1) + \nu y_1 \quad (4.32)$$

$$\phi_2(\nu) := \inf_{x_2, y_2} f_2(x_2, y_2) - \nu y_2 \quad (4.33)$$

The master problem for the primal (dual) problem is solved with an iterative algorithm that fixes the value y (ν), solves the subproblems for this fixed value, and then updates y (ν) to a new value. The update of the primal (dual) values can be done using gradient, subgradient or quasi-newton methods depending on the form of f_1 and f_2 .

We show an example of the subgradient method applied to the dual decomposition. For the k -th iteration, given a value ν^{k-1} for the Lagrangian multiplier at the previous iteration, the algorithm proceeds as follows:

$$x_1^k, y_1^k := \arg \min_{x_1, y_1} f_1(x_1, y_1) + \nu^{k-1} y_1 \quad (4.34)$$

$$x_2^k, y_2^k := \arg \min_{x_2, y_2} f_2(x_2, y_2) - \nu^{k-1} y_2 \quad (4.35)$$

$$\nu^k := \nu^{k-1} + \rho^k (y_1^k - y_2^k) \quad (4.36)$$

where ρ^k is a positive step parameter. The subproblems given by equations (4.34) and (4.35) can be solved in parallel, which is then followed by the master problem dual variable update in equation (4.36).

The distributed algorithm we use is the alternating direction method of multipliers (ADMM) [Boyd et al., 2011, Douglas and Rachford, 1956, Gabay and Mercier, 1976]. It is a distributed version of the method of multipliers, which has improved convergence on problems that are not strictly convex compared to dual gradient and subgradient methods [Boyd et al., 2011].

There are a whole range of other distributed algorithms that are designed to improve convergence over the more basic methods. These include distributed Newton-like methods [Jadbabaie et al., 2009, Zargham et al., 2011], Nesterov-like methods [Jakovetic et al., 2012] and distributed interior

point algorithms [Pakazad et al., 2013]. These can greatly speed up convergence, but are often specialised to particular problems and result in a greater subproblem computation and communication burden.

There are several other common decomposition methods and algorithms that are worth mentioning. For LPs, Dantzig-Wolfe decomposition [George B. Dantzig, 1960] and Benders' decomposition [Benders, 1962] can solve very large problems with particular structures. Distributed constraint optimisation (DCOP) has been used to solve constraint programming problems in a distributed manner [Faltings, 2006, Modi et al., 2005].

4.4 Multi-Period Optimal Power Flow

Chapter 3 has shown how online algorithms, in particular receding horizon algorithms, are effective at making decisions in the presence of uncertainty. We adopt a receding horizon algorithm here, but instead of minimising the costs for a single house in each horizon, we minimise the costs for the whole network by solving a multi-period OPF problem. We are primarily interested in the new aspects of the problem in this chapter, so we just focus on the performance of the distributed algorithm within a single horizon.

In this section we formalise the multi-period OPF problem for a single horizon in a way that is readily decomposable and can be used with the ADMM algorithm. This horizon has a length of T time steps.

We adopt a component-based representation of the network similar to that presented in section 4.2. The difference is that we treat buses just like any other component, providing them with terminals and only allowing connections between pairs of terminals. An illustration of this is provided in figure 4.3 where small solid squares represent terminals and solid lines between terminals represent connections. This enables us to decompose the problem by breaking the network up at connections.




Figure 4.3: Component-based representation.

In our model a network consists of a set of components $\mathcal{D} := \{1, \dots, |\mathcal{D}|\}$,

a set of terminals $\mathcal{T} := \{1, \dots, |\mathcal{T}|\}$ and a set of connections $\mathcal{E} \subseteq \mathcal{T} \times \mathcal{T}$ which are pairs of terminals. Each component $d \in \mathcal{D}$ (e.g., bus, line, generator, load, prosumer) has a set of terminals $\mathcal{T}_d \subseteq \mathcal{T}$ which can be connected to the terminals of other components, where the \mathcal{T}_d sets partition \mathcal{T} .

4.4.1 Connections

Connections exist between the terminals of two different components. We use the quantities of real power, reactive power, voltage and voltage phase angle, $P_i, Q_i, V_i, \theta_i \in \mathbb{R}^T$ respectively, to model the flow of power into a component through a terminal $i \in \mathcal{T}$. These are vectors in order to capture each time step in the horizon. For convenience, we use a vector $y_i \in \mathbb{R}^{4T}$ to represent all variables for a terminal $i \in \mathcal{T}$, where $y_i := (P_i, Q_i, V_i, \theta_i)^\top$.

When two terminals are connected together, $(i, j) \in \mathcal{E}$, we pose the following *connection constraints*:

$$P_i + P_j = 0, \quad Q_i + Q_j = 0, \quad V_i - V_j = 0, \quad \theta_i - \theta_j = 0 \quad (4.37)$$

The first two constraints ensure that any power leaving one terminal enters the other. The second two ensure that the connected terminals have the same voltage and phase angle. This duplication of variables is necessary in order to decompose the problem for our distributed algorithm. To avoid confusion, recall that connections and terminals are different from lines and buses (see figures 4.3 and 4.1).

We rewrite these constraints as $y_i + \mathcal{A}y_j = 0$ for the terminal vectors y_i and y_j , where \mathcal{A} is the appropriate $4T \times 4T$ diagonal matrix that reproduces (4.37). For convenience we also define the *connection function* $h : \mathbb{R}^{4T} \times \mathbb{R}^{4T} \mapsto \mathbb{R}^{4T}$ as the LHS of this constraint: $h(y_i, y_j) := y_i + \mathcal{A}y_j$.

4.4.2 Components

Each component $d \in \mathcal{D}$ has a vector of variables $x_d \in \mathbb{R}^{M_d}$ which implicitly includes the variables for its terminals (y_i for all $i \in \mathcal{T}_d$) plus any additional internal variables. Internal variables can be used, for example, to model the start time of shiftable loads or a battery's state of charge (section 3.4 offers other examples). Components also have an objective function $f_d : \mathbb{R}^{M_d} \mapsto \mathbb{R}$ and a constraint function $g_d : \mathbb{R}^{M_d} \mapsto \mathbb{R}^{N_d}$, where $g_d(x_d) \leq 0$.

The objective function is used to model any costs or preferences that a component may have other than the direct payments they make to the market for their consumption. Examples include generator fuel costs and thermal comfort preferences in houses (as discussed in chapter 3). Components such as lines and buses, will typically have no objective function.

Component constraints relate the component terminal variables to each other and to any internal variables. They provide hard limits on what power the component can exchange with its connected neighbour components.

4.4.3 Optimisation Problem

The objective of the multi-period OPF problem for one horizon is to minimise the sum of all component cost functions, subject to internal component constraints and external terminal connection constraints. This is a utilitarian view of the problem, where we minimise the total combined cost of the system. Formally for a network this is:

$$\min_{x_d \forall d \in \mathcal{D}} \sum_{d \in \mathcal{D}} f_d(x_d) \quad (4.38)$$

$$\text{s.t. } g_d(x_d) \leq 0 \quad \forall d \in \mathcal{D} \quad (4.39)$$

$$h(y_i, y_j) = 0 \quad \forall (i, j) \in \mathcal{E} \quad (4.40)$$

4.5 Component Models

In the following section we describe the models used for the components in our experiments. We avoid writing the constraint functions g_d explicitly, but instead write the component constraints in a way that is easy to understand. The constraints we present for each component, including the equalities, can be transformed into a function of the form g_d , which is more convenient when formulating the optimisation problem at a higher level.

When necessary, we use $t \in \{0, \dots, T\}$ to index vectors by time, otherwise we imply standard vector operations. The index where $t = 0$ is used to represent the value of the variable at the beginning of the current horizon, which we assume is a known fixed value.

Bus— A bus has a variable number of terminals which depends on how many other components connect to it. For example, a bus might be connected to a generator, a load and 3 lines for a total of 5 terminals. Regardless of the number of terminals, the constraints take the form:

$$\sum_{i \in \mathcal{T}_d} P_i = 0 \quad \sum_{i \in \mathcal{T}_d} Q_i = 0 \quad (4.41)$$

$$V_i = V_j, \quad \theta_i = \theta_j \quad \forall i, j \in \mathcal{T}_d \quad (4.42)$$

The first two constraints are an expression of KCL as described in section 4.2.1. The remaining constraints ensure that all terminal voltages and phase angles are the same. Recall that these constraints can be reworked into constraint functions g_d for bus components.

Line— A line is a two terminal component which transports power from bus to bus. In our experiments we use the various line models discussed in section 4.2.2. We use a maximum apparent power $S \in \mathbb{R}_+$ to represent a thermal limit. In addition we have voltage magnitude limits \underline{V}, \bar{V} and a phase angle

difference limit $\bar{\theta}$ to ensure safe operating conditions. Let $\{i, j\} = \mathcal{T}_d$ be the line's terminals, then these limits translate to the following constraints:

$$V_{i,t}, V_{j,t} \in [\underline{V}, \bar{V}] \quad \forall t \in \{1, \dots, T\} \quad (4.43)$$

$$|\theta_{i,t} - \theta_{j,t}| \leq \bar{\theta} \quad \forall t \in \{1, \dots, T\} \quad (4.44)$$

Generator— A generator is a single terminal component which produces real and reactive power. Let i be the generator's terminal. In our formulation the generator has a floating phase angle and voltage. A generator has lower and upper real and reactive power limits $\underline{P}, \bar{P}, \underline{Q}, \bar{Q}$, a ramping rate $P^r \in \mathbb{R}_+$ and a quadratic cost function f :

$$f(x_d) = P_i^\top \Psi P_i - \psi^\top P_i \quad (4.45)$$

$$|P_{i,t} - P_{i,t-1}| \leq P^r \quad \forall t \in \{1, \dots, T\} \quad (4.46)$$

$$P_{i,t} \in [\underline{P}, \bar{P}], \quad Q_{i,t} \in [\underline{Q}, \bar{Q}] \quad \forall t \in \{1, \dots, T\} \quad (4.47)$$

where the diagonal matrix $\Psi \in \mathbb{R}_+^{T \times T}$ and vector $\psi \in \mathbb{R}_+^T$ are price coefficients for the horizon. More advanced generator models with non-convex start up costs and minimum outputs can be modelled in this framework but are not considered here.

House— Houses are the final type of component we consider. A house can be further decomposed into components, e.g., one for each device, but instead we stop decomposing the problem once we reach the house EMS. As in chapter 3, the EMS centrally controls the rest of the house devices. At the scale of a single house the problem is tractable, so there is little benefit in further distributing the problem. Also the privacy advantages of distributed control are not nearly as important within a single house as they are between house and the rest of the network.

Houses have a single terminal and just like for generators we leave their voltages and phase angles floating. As in chapter 3, a house is made up of a set of devices and limits on its terminal powers, but now we also model the reactive power of devices. We use a hat to distinguish house devices \hat{d} and the device set $\hat{\mathcal{D}}$ from components d and the component set \mathcal{D} . This notation highlights how the decomposition into components could have extended to house devices.

If a component d is a house with terminal i , then the power limit on the house is $P_{i,t}^2 + Q_{i,t}^2 \leq S_i^2$. The house's real and reactive power is the sum of any uncontrollable background power consumption P^b and Q^b and the power consumed by its devices $\hat{\mathcal{D}}_d$:

$$P_i = P^b + \sum_{\hat{d} \in \hat{\mathcal{D}}_d} P_{\hat{d}} \quad (4.48)$$

$$Q_i = Q^b + \sum_{\hat{d} \in \hat{\mathcal{D}}_d} Q_{\hat{d}} \quad (4.49)$$

Any of the devices modelled and tested in chapter 3 could be included in the house. Within this framework we implemented models for shiftable loads, batteries and solar PV. Most of our experiments are focused around houses with just shiftable loads, as they are a source of non-convexity that we are interested in.

A shiftable load is a single terminal component used to model electrical loads like dish washers and clothes dryers. The starting time of the shiftable load is flexible, and the EMS will schedule them to minimise electricity costs. The shiftable load model used is a simplified version of that presented in section 3.4. We expect the results presented here will carry over to the more complicated model.

Shiftable loads must start running between an earliest $t^e \in \mathbb{N}$ and a latest $t^l \in \mathbb{N}$ start time. To model this we introduce binary variables $u \in \{0, 1\}^T$ for the horizon. A value of 1 indicates that the device starts at the corresponding time step. The device has a runtime of $t^r \in \mathbb{N}$ consecutive time steps, during which it consumes a (nominal) load of $P^{\text{nom}} \in \mathbb{R}$. For a shiftable device \hat{d} the following constraints apply:

$$P_{\hat{d},t} = P^{\text{nom}} \sum_{t'=t-t^r+1}^t u_{t'} \quad \sum_{t=t^e}^{t^l} u_t = 1 \quad (4.50)$$

$$u_t = 0 \quad \forall t \notin \{t^e, \dots, t^l\} \quad (4.51)$$

A convex relaxation of this component can be obtained by relaxing the integrality requirement, i.e. allowing $u \in [0, 1]^T$.

4.6 Distributed Algorithm

The next step is to show how this multi-period optimisation problem can be solved in a distributed manner. The end result is an iterative algorithm where each component (household, generator and network device) selfishly optimises its own consumption/production profile for the currently standing terminal prices. These profiles are then communicated amongst connected components and the terminal prices are updated in a way that helps drive agreement.

We use the alternating direction method of multipliers (ADMM) algorithm introduced in section 4.3 to distribute and solve the problem. It is a dual decomposition method that applies the augmented Lagrangian to the complicating constraints. For our problem the augmented Lagrangian applied to the connection constraints (4.40) is:

$$\mathcal{L} := \sum_{d \in \mathcal{D}} f_d(x_d) + \sum_{(i,j) \in \mathcal{E}} \left[\frac{\rho}{2} \|h(y_i, y_j)\|_2^2 + \lambda_i^\top h(y_i, y_j) \right] \quad (4.52)$$

where $\lambda_i \in \mathbb{R}^{4T}$ are dual variables for each connection $(i, j) \in \mathcal{E}$ and $\rho \in (0, \infty)$ is a penalty parameter.

The dual variables represent the locational marginal prices in our problem, or put another way, connection dependent dynamic prices. These prices are used to charge (or pay) components for the power that they exchange through their terminals. For example, a component with a terminal i (where $(i, j) \in \mathcal{E}$) will pay an amount equal to:

$$\lambda_{i,P}^\top P_i + \lambda_{i,Q}^\top Q_i + \lambda_{i,V}^\top V_i + \lambda_{i,\theta}^\top \theta_i \quad (4.53)$$

where we have split up the dual variables so that it is clear how they associate with each physical power quantity. These prices are based on not just the cost of generation, but also account for line losses and adjust to prevent congestion. They provide a natural market mechanism for the distribution of payments from consumers to producers.

4.6.1 Algorithm

A single iteration of the ADMM algorithm is broken down into 3 steps. The first two steps involve solving subproblems for two groups of components, and the third step updates the dual prices. It is this alternating between solving subproblems for different groups of components that gives the algorithm its name.

Each component is allocated to either group I or group II, as indicated by the sets \mathcal{D}_I and \mathcal{D}_{II} . The variables of the components in these groups are represented by $x_I := (x_d : d \in \mathcal{D}_I)$ and $x_{II} := (x_d : d \in \mathcal{D}_{II})$. Figure 4.4 shows two examples of how a network can be grouped.




Figure 4.4: Two different groupings for the same network. The solid regions indicate components in group I and the striped regions components in group II. Contiguous regions indicate a set of components that have intra-group connections. The example on the left shows a case where all buses are in a group of their own, which allows the network to be fully decomposed (no intra-group connections). The example on the right does not decompose as much, which makes subproblems harder to solve, but it may converge in fewer iterations.

The connections are split into three sets: the intra-group connections \mathcal{E}_I and \mathcal{E}_{II} , and the inter-group connections \mathcal{E}_{III} . The intra-group connections are the connections between components in the same group, while the inter-group connections are those between components in different groups:

$$\mathcal{E}_I := \{(i, j) : (i, j) \in \mathcal{E}, i \in \mathcal{T}_{d_1}, j \in \mathcal{T}_{d_2} : d_1, d_2 \in \mathcal{D}_I\} \quad (4.54)$$

$$\mathcal{E}_{II} := \{(i, j) : (i, j) \in \mathcal{E}, i \in \mathcal{T}_{d_1}, j \in \mathcal{T}_{d_2} : d_1, d_2 \in \mathcal{D}_{II}\} \quad (4.55)$$

$$\mathcal{E}_{III} := \{(i, j) : (i, j) \in \mathcal{E}, i \in \mathcal{T}_{d_1}, j \in \mathcal{T}_{d_2} : d_1 \in \mathcal{D}_I, d_2 \in \mathcal{D}_{II}\} \quad (4.56)$$

For simplicity, we assume that the first terminal in each inter-group connection belongs to a component in the first group. The augmented Lagrangian relaxation is applied to the inter-group connections, so we only need dual variables for the connections in the set \mathcal{E}_{III} . The intra-group connections are treated as hard constraints.

Superscripts such as $x^{(k)}$ are used to indicate the value of a variable x after the k -th algorithm iteration. At the start all inter-group connection terminal and dual variables are initialised to some values $y_i^{(0)}$ and $\lambda_i^{(0)}$. For the k -th iteration the ADMM algorithm proceeds as follows:

1. Minimize \mathcal{L} for x_I , holding x_{II} constant at its $k - 1$ value
2. Minimize \mathcal{L} for x_{II} , holding x_I constant at its k value
3. Update the dual variables λ

For our optimisation problem this becomes:

$$\begin{aligned} x_I^{(k)} = \arg \min_{x_I} \quad & \sum_{d \in \mathcal{D}_I} f_d(x_d) \\ & + \sum_{(i,j) \in \mathcal{E}_{III}} \left[\frac{\rho^{(k)}}{2} \|h(y_i, y_j^{(k-1)})\|_2^2 + \lambda_i^{(k-1)\top} h(y_i, y_j^{(k-1)}) \right] \\ \text{s.t.} \quad & g_d(x_d) \leq 0 \quad \forall d \in \mathcal{D}_I \\ & h(y_i, y_j) = 0 \quad \forall (i, j) \in \mathcal{E}_I \end{aligned} \quad (4.57)$$

$$\begin{aligned} x_{II}^{(k)} = \arg \min_{x_{II}} \quad & \sum_{d \in \mathcal{D}_{II}} f_d(x_d) \\ & + \sum_{(i,j) \in \mathcal{E}_{III}} \left[\frac{\rho^{(k)}}{2} \|h(y_i^{(k)}, y_j)\|_2^2 + \lambda_i^{(k-1)\top} h(y_i^{(k)}, y_j) \right] \\ \text{s.t.} \quad & g_d(x_d) \leq 0 \quad \forall d \in \mathcal{D}_{II} \\ & h(y_i, y_j) = 0 \quad \forall (i, j) \in \mathcal{E}_{II} \end{aligned} \quad (4.58)$$

$$\lambda_i^{(k)} = \lambda_i^{(k-1)} + \rho^{(k)} h(y_i^{(k)}, y_j^{(k)}) \quad \forall (i, j) \in \mathcal{E}_{III} \quad (4.59)$$

In the simple case when ρ is held constant, f_c and g_c are convex, and h is affine, ADMM converges to a global optimum [Boyd et al., 2011].

If a component has no intra-group connections, then it is separable in the subproblem for its group, and can therefore be solved in parallel. We adopt the partitioning scheme where \mathcal{D}_{II} contains all buses and \mathcal{D}_{I} the rest of the network (the example on the left in figure 4.4). This allows us to fully separate all components within groups, since buses will never connect to other buses ($\mathcal{E}_{\text{II}} = \emptyset$) and non-bus components will never connect to other non-bus components ($\mathcal{E}_{\text{I}} = \emptyset$). In this way each component acts as an independent agent and communicates only to other directly connected agents. As an additional benefit, some components are simple enough when separated that they have closed-form solutions that can be calculated at each iteration, instead of invoking an optimisation routine [Peng and Low, 2014]. We adopt such closed-form solutions for buses as proposed in [Kraning et al., 2014].

4.6.2 Residuals and Stopping Criteria

As in [Kraning et al., 2014], we use primal and dual residuals to define the stopping criteria for our algorithm. The primal residuals represent the constraint violations at the current solution. We combine the residuals of all connections into a single vector r_p . By indexing into the inter-group connections $\{(i_1, j_1), (i_2, j_2), \dots\} = \mathcal{E}_{\text{III}}$, the primal residuals are:

$$r_p^{(k)} := ((y_{i_1}^{(k)} + \mathcal{A}y_{j_1}^{(k)}), (y_{i_2}^{(k)} + \mathcal{A}y_{j_2}^{(k)}), \dots) \quad (4.60)$$

The dual residuals give the violation of the KKT stationarity constraint at the current solution. We collect the dual residuals for each connection into the vector r_d . For ADMM, the dual residuals are (see [Boyd et al., 2011] for derivation):

$$r_d^{(k)} := (\rho^{(k)} \mathcal{A}(y_{j_1}^{(k)} - y_{j_1}^{(k-1)}), \rho^{(k)} \mathcal{A}(y_{j_2}^{(k)} - y_{j_2}^{(k-1)}), \dots) \quad (4.61)$$

These residuals approach zero as the algorithm converges to a KKT point. We consider that the algorithm has converged when the scaled 2-norms of these residuals are smaller than a tolerance ϵ : $\frac{1}{\sqrt{n}} \|r_p^{(k)}\|_2 < \epsilon$, $\frac{1}{\sqrt{n}} \|r_d^{(k)}\|_2 < \epsilon$. Here n is the total number of inter-group terminal constraints $4T|\mathcal{E}_{\text{III}}|$ minus the number of terminal constraints that are trivially satisfied (e.g., floating voltages and phase angles for generators). This is used to keep the tolerance independent of problem size.

4.7 Implementation

In a fully distributed real-world implementation every house, generator, bus, line, and other component could have its own collocated computational node.

In practice it might make more sense to have the computational parts of the network located separately from their components and even grouped together. For example, all the buses and lines of a single distribution feeder could be managed by a single node, which communicates to downstream houses and the upstream substation. A whole range of practical issues around speed, communications, costs, maintenance and robustness would need to be considered before a decision could be made on the right architecture.

We developed an experimental implementation of the algorithm and models presented in previous sections. It is a sequential implementation of the ADMM algorithm so that it can be run on a single machine; however, we timed the slowest component at each iteration to get an idea of how long a fully distributed implementation would take. This implementation was designed with flexibility in mind, so that a wide range of experiments could be conducted.

The implementation is written in C++ using Gurobi [Gurobi Optimization, Inc., 2014] and Ipopt [Wächter and Biegler, 2006, HSL Archive, 2014] as backend solvers for subproblems. Gurobi is used for mixed-integer linear or quadratically constrained problems, and Ipopt for more general nonlinear problems. CasADi [Andersson, 2013] was used as a modelling and automatic differentiation front end to Ipopt. The experiments were run on machines with 2 AMD 6-Core Opteron 4184, 2.8GHz, 3M L2/6M L3 Cache CPUs and 64GB of memory.

4.8 Test Microgrid

Our experiments are based around a modified 70 bus 11kV benchmark distribution network [Das, 2006] (shown in Figure 4.5), which was chosen because it has a comparable size to that of an Australian suburb. We close all tie lines in the network in order to change it from a radial to a meshed configuration. Over time we expect microgrids to take on more of a meshed network structure to improve reliability and efficiency, and to better utilise distributed generation (this is currently the case in some city centres).

The benchmark comes with a static PQ load at each bus, which we replace with a number of houses (around 50 on average) that depends on the size of the static load. The houses are connected directly to the 11kV buses as we have no data on the low voltage part of the network. We assume that the power bounds we place on each house will be sufficient to prevent any capacity violation of the low voltage network.

For our experiments, each house includes an uncontrollable background power draw and two shiftable loads. Each house has an apparent power limit of $S = 10\text{kVA}$, which in practice would be chosen just below the point where the circuit breakers in each house would trip.




Figure 4.5: The 70 bus test network showing only the buses and lines. Each bus has on average around 50 houses connected to it, and generators/substations are located at the grey buses.

We develop a typical house load profile P_t^l by modifying an aggregate autumn load profile for the ACT region in Australia (data from the Australian Energy Market Operator (AEMO)). We assume that houses consume on average 20kWh per day. This provides the basis for all uncontrollable house background loads. For the purposes of these experiments, we assume that the static PQ loads in the benchmark were recorded when loads were at 75% of their peak. We divide the benchmark static real power at each bus by how much power a typical house consumes at 75% of its peak power (1.45kW). Rounding down this number gives us an estimate of the number of houses which would be located at a given bus. This approach produces a total of 3674 houses for the network, about the number in a typical Australian suburb.

We place two generators in the network where the distribution system connects to upstream substations. These can be thought of as either dispatchable microgenerators or as representing the cost of importing power into the microgrid.

We randomise some of the generator and house load parameters to produce different problem instances, as can be seen in table 4.1. The time horizon spans 24 hours with 15 minute time steps, which produces a problem instance with over 2 million variables per horizon. The experiments were run with a primal and dual stopping tolerance of $\epsilon = 10^{-4}$ and a fixed penalty parameter of $\rho = 0.5$. To improve numerical stability, we scale the system to a per-unit representation with base values at 11kV and 100kVA. This means that a real power residual of 10^{-4} translates to 10W for a connection, or about 1% of the average house load.

The starting values for the distributed algorithm are the same for all

Table 4.1: Component parameters.

Comp	Param	Value	Units
Gen	ψ_t	$\max(4, \sim \mathcal{N}(40, 8^2))$	\$/MWh
Gen	$\Psi_{t,t}$	$\max(1, \sim \mathcal{N}(10, 2^2))$	\$/MWhMW
Gen	\underline{P}, \bar{P}	$-S \times h/2, 0$	kW
Gen	\underline{Q}, \bar{Q}	$0.2\underline{P}, -0.2\bar{P}$	kvar
House	\bar{P}_t	$\sim \mathcal{N}(P_t^l, (0.2l_t)^2)$	kW
House	\bar{Q}_t	$0.3P_t$	kvar
Shift 1	t^r	$\max(15, \sim \mathcal{N}(90, 18^2))$	min
Shift 1	P^{nom}	$\max(0.3, \sim \mathcal{N}(3, 0.6^2))$	kW
Shift 2	t^r	$\max(15, \sim \mathcal{N}(60, 12^2))$	min
Shift 2	P^{nom}	$\max(0.1, \sim \mathcal{N}(1, 0.2^2))$	kW
Shift	t^e, t^l	$0, T - t^r$	

terminals and all time steps. All are zero except for the voltage magnitudes $V_t = 1$ and the real power connection constraint dual variables $\lambda_t^P = 5$, which translates to a price of \$200/MWh. This is a naive (or cold) starting point as it uses no information about the particular network instance.

In addition to the microgrid, we also ran a series of experiments on randomly generated networks, similar to those described in [Kraning et al., 2014]. These randomly generated networks ranged in size from 20 to 2000 buses, and were designed to be highly congested. We will occasionally mention some of the results from these random networks when they differ from those of our 70 bus microgrid.

4.9 Impact of Power Flow Models

In this section we investigate how the ADMM method performs with different power flow models. We assess 5 different models, of varying degrees of accuracy and complexity, in order to establish the relative trade-offs when used as part of a distributed algorithm.

4.9.1 Power Flow Models

As discussed in section 4.2.2, the non-convex AC power flow equations are often either relaxed or approximated before being used in power systems optimisation problems. The convex relaxations provide a lower bound on the globally optimal solution while the approximations can produce results with a cost higher or lower than the true global optimal. The relaxations and approximations often produce solutions that are not feasible for the exact model; however, they are often much simpler to compute, and their solutions can be used as a heuristic or to provide bounds.

Some of these models have been used with distributed algorithms, which we will discuss in greater detail in section 4.12. What is lacking is a comparison of the relative strengths and weaknesses between the different models when used in this context. In this section we compare how the distributed ADMM algorithm performs when using the AC, QC, DF, LDC, and QA line models. We compare the differences in the solution quality, feasibility, processing time and number of iterations for our test network. What we find is that even though the AC equations are non-convex, in practice they converge and perform well compared to the other approaches. We also find that there is the potential to obtain faster convergence using the QA model, but at the expense of accuracy.

We generate 60 random instances of our test microgrid with the binary variables for the shiftable loads relaxed. These are then solved using the distributed algorithm presented in section 4.6, for each of the 5 different power flow models. Below we discuss the convergence of the algorithm and then the quality and accuracy of the solutions.

4.9.2 Convergence

For all 60 instances and all 5 power flow models the algorithm converged. This was expected for all the convex models, but we had no guarantee for the non-convex AC model. This gives us confidence that the exact AC model, even though non-convex, can in practice be used within distributed algorithms, at least to find a local optimum.

Table 4.2 provides the number of iterations and time taken to converge in the form of means and standard deviations for the 60 instances. The parallel solve time is the amount of time required to solve the problem in a fully distributed implementation. This was measured by summing together the time of the slowest component for each group in each iteration. In absolute terms, the parallel solve time is relatively small despite the fact that our implementation was designed with flexibility in mind, not performance. The QA model is much faster than the other models. It converges in half the number of iterations required by the next fastest model, but as we will see in the next section, it gives us an inaccurate result.

Table 4.2: Iterations and parallel solve time for line models.

Model	Iterations (std.)	Time in sec (std.)
AC	1945 (17)	148 (12)
QC	1951 (14)	546 (33)
DF	1933 (26)	110 (8)
LDC	4140 (50)	244 (8)
QA	1027 (52)	15 (1)

The congested random networks produce similar results. One difference is that for a number of instances the QA model was infeasible (would not converge) due to its tendency to exaggerate line losses, where we had a valid AC solution. It is expected that the LDC model, and the relaxations to some degree, will exhibit the reverse effect: returning a solution when there is no feasible solution for the exact model. However, we did not come across such a scenario in our experiments with the microgrid and congested random networks.

Figure 4.6 shows an example of the primal residual (see section 4.6.2) convergence for different line models (the dual residuals are similar). The AC, QC and DF model curves are indistinguishable from one another.




Figure 4.6: Convergence of primal residuals. AC, QC and DF overlap.

One result that might be unintuitive is the fact that the LDC model converges poorly when it is in fact a very simple linear model. Large oscillations build up across the network during the solution process which slows the rate of convergence. This appears to be due to the fact that the LDC line model is not strictly convex and that it ignores line losses. This means it can be very sensitive to changes in dual variable prices. The quadratic term in the augmented Lagrangian is supposed to aid convergence in such cases, but it still struggles with the LDC line model.

The “stronger” convexity of the QA line model and the fact that it overestimates the line losses provides it with a fast convergence devoid of oscillations. The AC, QC and DF models are somewhere in between these two extremes. In these models the primal and dual values associated with

terminal powers are faster to converge than the voltage values. Quickly finding good voltage values appears to be key to improving the convergence of these models.

Warm Starting— We are giving the algorithm a relatively naive starting point for both the primal and dual variables, as described in section 4.8. In practice, the receding horizon control scheme will provide an excellent warm starting point, because the values from the previous horizon can be used for all but one time step.

As a sanity check, we performed warm starting experiments for the AC model. Similar to what was done in [Kraning et al., 2014], we duplicate a problem instance and then randomly resample the house background power and shiftable device power parameters according to the rule: $P \sim PN(1, \sigma^2)$. We used the solution of the original instance as a starting point for the modified instance. For $\sigma = 0.2$ the warm started run only needed 11% of the original iterations on average. In a second experiment we fully correlated the resampling step, which could represent a correlated change in solar panel output for many houses. With $\sigma = 0.2$, only 29% of the original iterations were required on average.

Communication— Communication delay will play a major part in the total solve time for an actually distributed system. The communications could be done over existing internet infrastructure, or dedicated wired or wireless communications could be built for the system to enable more direct communications. Regardless of what technology is used, for each iteration messages need to be communicated from group I components to group II components and then back again. If we assume each of these hops takes 60ms, then 1000 iterations would require up to 2 minutes of communication time.

The significant communication overhead suggests that in certain circumstances it may be beneficial to cut down the total number of iterations, even if it requires more processing time per iteration. This can be done by changing the grouping strategy which was discussed in section 4.6.1. The example on the right in figure 4.4 leaves large parts of the network coupled together. These coupled components need to be solved together in each iteration, which requires more computation per iteration in a parallel environment; however, it can reduce the total number of iterations required for convergence and therefore also the communication burden.

We expect the mechanism can be designed to be quite robust to intermittent drops in communication. For example, if a component fails to receive a message from another connected component, then they can continue working by using the last received message. If a connection is dropped for an extended period, then the system could fall back to load predictions based on historical data and some conservative pricing scheme could come into place.

4.9.3 Solution Quality

Next we show the quality of the solutions obtained for the different line models. For each model we calculate the percentage difference in objective value relative to the best known AC solution: $100 \cdot (f - f_{best})/f_{best}$. The means and standard deviations of these values for the 60 instances are shown in table 4.3.

Table 4.3: Object values for the different line models relative fo AC model. Means and standard deviations given as percentages.

QC	DF	LDC	QA
-0.03 (0.01)	0.04 (0.02)	-3.54 (0.07)	4.73 (0.09)

Because the AC equations are non-convex, we have no guarantee that the solutions they produce are globally optimal, but they do provide a feasible upper bound on the global optimal. On the other hand, the QC and DF models are convex relaxations of the AC equations, so they provide a lower bound on the global optimal. Therefore the global optimal solution resides somewhere between the values of the AC solution and the QC and DF solutions.

With this in mind, we find that the AC, QC and DF models all produce solutions which are very close to each other. The difference is within the margin of error of the objective function afforded by our stopping criteria, which we estimate to be 1% (see section 4.8). This indicates that the AC, QC and DF models produce solutions that are within 1% of the global optimal. They may in fact be closer than this, but we would need to run the experiments with tighter tolerances in order to check (on a limited number of instances we tested this was the case).

These results give us confidence that the non-convex AC model, which is the only one that guarantees Ohm's law is satisfied, produces solutions that are very close to optimal. The QC and DF models produce results with an objective that is very close to the AC model, but even with this small difference, there is the risk that the solutions violate constraints in the exact AC model. Other work has come to a similar conclusion, but in a more traditional OPF setting [Hijazi et al., 2014, Phan and Kalagnanam, 2014, Erseghe, 2014].

The approximate models tell a different story. On average the LDC model underestimates the optimal value by around 3.5% while the QA model overestimates it by around 4.7%. Part of the reason for this is that the LDC model completely ignores line losses while the QA model overestimates them. Even though the QA model has fast convergence, it is unlikely to be useful on its own in a realistic setting due to its poor accuracy.

These results show the feasibility of using the non-convex AC power flow

equations for solving a distributed OPF problem in a microgrid context. The QA model, used by Kraning et al. [2014] in a simpler distributed OPF setting, converges much faster, but it is unlikely to work with a real network, as it ignores voltages and reactive power and produces overly high costs.

4.10 Discrete Decisions

We now move onto finding solutions for the multi-period OPF problem where the binary variables in the shiftable loads are satisfied. In order to do this we extend the algorithm so that it can manage discrete decisions. The focus here is on the scheduling of shiftable loads within houses. Discrete decisions can also occur in some generator models and for automated network switching events, which will require a separate analysis.

We identify 3 different approaches to managing discrete decisions. Although quite simple, they nonetheless prove to be very effective at managing the shiftable loads within houses.

4.10.1 Methods

We investigate 3 tractable methods for dealing with integer variables which have no global optimality guarantees. Just as we did for the AC equations, we will compare our result to a lower bound in order to get an understanding of the optimality gap. We categorise these methods as:

- Relax and price (RP)
- Relax and decide (RD)
- Unrelaxed (UR)

The RP and RD approaches are broken up into 2 and 3 stages respectively. The first stage, called the negotiation stage, is common to both methods. All integer variables are relaxed and the distributed algorithm is run until convergence, just like what was done for the power flow experiments. At this point the integer variables may take on fractional values. This solution provides a lower bound on the global optimal. In the second stage each component makes a local decision in order to force any fractional values to integers. Recall from section 4.5 that shiftable devices have a binary variable u_t for each time step, only one of which can take on the value 1 to indicate the starting time.

Relax and Price— In the second stage of the RP method, each house performs a local optimisation to determine how to enforce integer feasibility of u_t . We designed a range of cost functions which penalise a component if it changes its terminal values from those that were negotiated in the first

stage. For a given cost function each house solves a MIP to obtain an integer-feasible solution.

For a particular house terminal i with final terminal values y_i , let y'_i and λ'_i be the negotiated terminal values and prices respectively. The most effective cost function that we identified was (assuming the house terminal is the first element in the connection):

$$f_{\text{RP}}(y_i, y'_i, \lambda'_i) = \lambda'^T_i y_i + \alpha \|y_i - y'_i\|_2^2 \quad (4.62)$$

where α is a penalty parameter. The function charges houses at the negotiated price for what they *actually* consume, but they are also charged a quadratic penalty for operating away from the negotiated consumption.

After this local optimisation step, we check that the solution is feasible and what the overall cost is. At this point it would be possible to restart the distributed algorithm with the integer variables fixed (as we do in the second method). Instead of this we allow the frequency regulating generators to make up the difference between the negotiated and actual power consumption.

In networks the dispatch of generators is established in advance in response to an estimated demand. This forecast is never perfect, so a certain number of generators are paid to perform frequency regulation in order to balance demand in real time (see section 2.1.1). In our RP experiments we employ both our generators for this use by allowing them to adjust their output. For these experiments we assume that if the generators have to raise their output they get paid extra, and if they have to lower their output they get paid what was originally negotiated. This means that they are incentivised to participate in frequency regulation because they make money from both raise and lower actions.

Relax and Decide— In the second stage of the RD method, the largest u_t value of each shiftable component is chosen to be fixed at 1 and the rest set to 0. In the third stage the distributed algorithm is restarted in order to converge to a new solution that is integer feasible.

Unrelaxed— The final approach, UR, consists of a single stage where it attempts to enforce the integrality requirement at each iteration of the distributed algorithm. We have already foregone theoretical convergence guarantees due to our adoption of the non-convex AC equations. Here we push the ADMM algorithm even further by allowing discrete variables into the algorithm (4.57–4.59), where Gurobi solves MIQP subproblems for houses, and Ipopt solves NLP subproblems for lines.

We ran experiments on 60 microgrid instances for each of the three approaches. We use the AC line model for each experiment and a penalty of $\alpha = 10$ for the RP approach. In the following sections we discuss the convergence of the methods and the quality of the solutions.

4.10.2 Convergence

Although none of the approaches are guaranteed to find an integer feasible solution if one exists, in practice they all converged to feasible solutions (within our stopping tolerance) except for 2 out of 60 UR experiments which ran out of time.

The RP method only marginally increases the solve time above the results in section 4.9. The RD method requires a small amount of extra time as it performs a warm restart of the distributed algorithm. The UR method takes 1.7 times longer on average, which is a result of the fact that it has to solve MIQP instead of QP subproblems for houses during each iteration.

4.10.3 Solution Quality

In order to assess the solution quality for each method, we compare the change in objective value relative to the relaxed version of the problem. The results are shown in figure 4.7, where we have separated the objective into terms for the cost of generation and the charge to houses. The charge is the sum of house objective functions, which represents the amount of money they pay for their electricity. For the RP method this is given by the cost function f_{RP} . For the RD and UR methods the charge is simply the final $\lambda_i^T y_i$ for each house.




Figure 4.7: Change in generator cost and house charge relative to relaxed solution.

For the relaxed problem itself, the true generation costs can be different from the amount houses are charged, as we are using marginal prices. In addition to this, network congestion typically produces additional revenue above the cost of generation. An increase in cost for the integer feasible

solution relative to the relaxed problem is an indication of the additional cost to the generators for balancing supply. When house charges increases relative to the relaxed solution, this indicates that houses were forced to change their consumption from the originally negotiated amount to ensure integer feasibility of the shiftable loads.

All methods produce costs that are within 1% of the relaxed problem, and hence also the global optimum. There is no significant difference between the methods as they reside within our estimated margin of error based on our stopping tolerance. What these results suggest is that we have a tight relaxation of the integer problem. A contributing factor is that each shiftable load only contributes a tiny amount to the overall power demand.

A significant gap between the relaxed and candidate solution was found for some instances that where we artificially increased the size of the shiftable loads by more than an order of magnitude and heavily congesting the network. However, the relaxation was tight for the realistically-sized residential shiftable loads utilised in our test microgrid.

The charges to houses are higher for the RP method without gaining any benefit in terms of reduced costs. We noticed that smaller α value decreased the charges without a significant increase to costs. When battery storage is introduced, we expect houses will have even more flexibility in how they reach their negotiated consumptions, therefore further reducing incurred charges. This suggests that for the sole purpose of managing shiftable loads, there is no need to have a strong penalty; however, the penalty might help manage the effects of uncertainty in the network or prevent agents from lying during the negotiation stage.

The presented methods provide efficient means for dealing with the discrete decisions in a house. The UR method should probably be avoided as it can significantly increase the runtime. Other factors such as the way they can handle uncertainty or the need for a penalty to prevent agents from attempting to game the system (as discussed in chapter 5) will influence the choice between these methods.

4.11 Pricing Uncertainty

In this section we investigate the inclusion of stochastic components into our system. Many parameters such as background house power consumption and solar PV output can only be estimated. Generators set up to regulate frequency can be used to balance this mismatch in real time, but such regulation can be expensive so we would like to minimise the error. There also exist certain circumstances where frequency regulation on its own is not enough. For example, if the output of house solar PV systems turns out lower than expected, then certain lines in the network could be overloaded if the network was already running near capacity.

Receding horizon control as we tested in chapter 3 and discussed in section 4.4 can help significantly with this by only ever acting on the most immediate time step before reoptimising. In this section we use the RP method from the previous section to encourage houses to correct for any local sources of uncertainty. To simplify the experiments we do not perform full receding horizon control. With full receding horizon control we expect the same trends, but with a further improvement to costs.

4.11.1 PV Generation Uncertainty

We perform a simple experiment on the 70 bus microgrid where we have added 2kW PV systems and 2kWh batteries independently at random to half of the houses. The battery efficiencies η were uniformly sampled from the interval $[0.85, 0.95]$. For normalised solar irradiance we use the simple relation: $I_t = \max(0, \sin(2\pi t/96 - \pi/2))$. We solve the first stage of the RP method with this irradiance, and then either lower or raise it 20% before running the RP second stage. Figure 4.8 shows the resulting costs and charges relative to the first stage result.




Figure 4.8: Performance of RP with lowering (L) and raising (R) solar output.

Lowering solar output produces increases to generation costs of 8% relative to the original solution. Nearly this full amount is recovered through extra house charges. When raising solar output the houses take advantage of the increased effectively free energy. The overall cost remains unchanged, because as described in section 4.10.1 we have to pay the frequency regulating generators at their originally negotiated price on a output lowering event. This creates a discrepancy in payments.

As we found in the previous section, generator costs are relatively independent of the RP cost function penalty parameter α ; however, α can be used to increase house charges. This means that α can be tuned over time to ensure that the market is budget balanced. For example, these extra penalties can be used to recover frequency regulation costs or to help pay for fixed network costs. Houses with batteries have more flexibility in how they can deal with local sources of uncertainty, especially in a receding horizon setting, so a system with batteries and reasonable penalties might produce the best outcome.

4.12 Related Work

Dynamic pricing (or real-time pricing) has been used extensively, including by us in the previous chapter, as a control signal for influencing household behaviour and performing DR [Ramchurn et al., 2011, Mohsenian-Rad and Leon-Garcia, 2010, Chen et al., 2012, Tischer and Verbic, 2011, Gast et al., 2014]. In these methods, participants receive a price signal and individually optimise their own behaviour so as to minimise their costs. Other approaches have utilised non-pricing control signals, which are simpler to implement, but are more limited in the types of devices that they can model [van den Briel et al., 2013, Shinwari et al., 2012].

These approaches implement a form of open loop control, because the agent that sets the control signal (dynamic price or otherwise) at best can only estimate how consumers will respond to it. In order to reduce the amount of guesswork and improve solutions, Gatsis and Giannakis [2012] developed a closed loop approach to dynamic pricing. In this scheme, the prices are iteratively updated by a central agent, with consumers communicating their best responses to the price prior to acting. Mohsenian-Rad et al. [2010] introduce an alternative iterative procedure not based on dynamic pricing, where consumers cooperate to reduce total generation costs in a distributed manner.

The approaches discussed so far do not model the electricity network, so cannot account for real power losses, reactive power, voltage limits or line thermal limits. Without these considerations, we cannot be sure that the coordinated outcome is efficient, safe or even possible. Many of the works on distributed algorithms which explicitly model the network have used ADMM as a solving technique, due to its ease in decomposition, and its convergence guarantees on a wide range of problems [Boyd et al., 2011]. However, most of these works have focused on more traditional OPF problems rather than coordinating prosumers in a microgrid context.

Some of the earliest work that used ADMM to solve power network problems was by Kim and Baldick [2000], who decomposed a convex approximation of the OPF problem into regions, and compared ADMM to

two other approaches. They found it to have a significant speed improvement over a centralised approach, and that it preserved privacy between regions. Erseghe [2014] also performed region-based decomposition of the network and found exact local solutions to the OPF problem. Instead of decomposing on the network structure, Phan and Kalagnanam [2014] decomposed across scenarios in a security-constrained OPF. The recent work by Magnússon et al. [2014] decomposes the network to a greater extent than these other methods, and they solve the underlying non-convex OPF by taking sequential convex approximations. One thing all these works have in common is that they are focused on the more traditional OPF problem, whereas in our work we consider a microgrid where distributed participants act independently.

Region-based decomposition was also used by Dall’Anese et al. [2013] to control distributed generation on radial feeders. They used ADMM to solve an unbalanced OPF problem using a semi-definite programming (SDP) relaxation. In our work we consider each customer to be independent, for privacy reasons, and we also allow for meshed microgrid topologies. Šulc et al. [2014] use the relaxed DF (SOCP) equations to perform reactive power control on radial networks. For a similar problem, Peng and Low [2014] provide closed-form solutions for ADMM subproblems, greatly reducing the computational requirements. Again these works focus on radial networks.

Distributed optimisation has been used for more specialised power networks problems, for example, managing the load on low voltage transformers [Vandael et al., 2010] and performing voltage regulation on radial networks [Zhang et al., 2012].

The work that is closest to ours is that presented by Kraning et al. [2014], and indeed we build on their approach. They decompose all network components for a multi-period OPF problem using a quadratic power flow approximation (the QA model). This procedure is effectively a principled method for settling dynamic prices for each bus. Their experiments showed that very large problems could be solved efficiently in a parallel environment.

All these works have taken different approaches to modelling the power flows on the network. There is no comparison of the relative performance between these different power flow models in a distributed algorithm for a prosumer-focused microgrid setting, which is what we achieve in this chapter for five different models. Our results in this area indicate that exact local methods can produce close to optimal solutions in a competitive number of iterations relative to other models. In addition, to the best of our knowledge, we are the first to incorporate discrete decisions into a distributed prosumer coordination method that models the network. Our work brings ADMM to the point where it can be considered a practical approach for efficiently balancing power in a prosumer focused setting.

4.13 Conclusion and Future Work

We have presented a distributed mechanism for coordinating prosumers and DETs across a network. It can coordinate a range of distributed agents with time-coupled behaviours, whilst preserving network constraints. It provides a natural way of pricing power (a distribution-level market) which enables prosumers and networks to work together to more intelligently operate the overall power system. It also allows for new flexible and efficient network designs including microgrids.

Using this mechanism, we have successfully compared the performance of a range of power flow models in a meshed microgrid, and introduced simple but effective approaches to handling the shiftable loads within households. We developed a suburb-sized test microgrid, and found that the full non-convex AC equations produce close to optimal solutions in short solve times. All three of our methods for handling household shiftable loads produce close to optimal solutions with only a moderate increases in solve times.

Our work has shown that in practice distributed algorithms are not only feasible, but also highly effective at coordinating prosumers in a microgrid context.

In future research we will investigate alternative distributed solving techniques with the aim of further improving the rate of convergence. There are opportunities for finding closed-form solutions for the exact AC equations, and to further parallelise the problem by decomposing certain components across time. It might also be possible to build a frequency regulation market into the distributed algorithm.

We need further experiments to investigate if our results carry over to larger discrete decisions, for example, those related to large industrial plant, generator start-up costs, and line switching.

Chapter 5

Receding Horizon Manipulation

5.1 Introduction

The prosumer coordination mechanism from chapter 4 converges to a socially optimal solution (from a utilitarian perspective), but only under the assumption that all prosumers (agents) act truthfully. An agent that misrepresents their true costs or constraints can lead the mechanism to a solution where they are personally better off, but which is socially suboptimal. This *manipulation* of the mechanism can harm other agents and the overall outcome, especially if multiple agents are involved in such behaviour.

Such manipulation is not uncommon in markets including current wholesale electricity markets; however, the presence of strong competition often limits such behaviour. What is different in our prosumer coordination mechanism, is the presence of overlapping receding horizons, which present unique opportunities for agents to lie about their preferences. There appears to be a strong conflict between managing uncertainty and deterring manipulation.

This chapter investigates how vulnerable the receding horizon mechanism is to manipulation, and devises methods for identifying it. We simplify and abstract the prosumer coordination problem in chapter 4 to one of clearing an electrical energy exchange market. In this market agents have private preferences and constraints around their usage of power, and are paid or charged at the marginal market clearing price.

Our first contribution is the presentation of a method for calculating an optimal strategy for a manipulative agent in a complete information setting. This agent is used to establish empirically how much a single agent can gain by manipulating the mechanism, and how this harms the social objective.

Our second contribution is a privacy-preserving method for identifying when manipulation has occurred and who is responsible. Supporting theoretical and empirical results show how this approach can identify agents that

are driving change, and distinguish a manipulative agent from an uncertain agent.

In practice, these identifiers can be used in auditing the actions of agents participating in a distributed receding horizon algorithm without altering the algorithm itself. This provides the full benefits of a receding horizon mechanism whilst significantly limiting the opportunities for manipulation.

Section 5.2 provides background material on mechanism design, which we draw on when discussing the properties of our coordination mechanism. Sections 5.3 and 5.4 introduce the simplified power balancing problem and the mechanism for solving it. Section 5.5 provides a formal definition of receding horizon manipulation. Sections 5.6 and 5.7 develop and experiment with an approach for calculating optimal manipulative actions for a single agent. This is followed by the development of manipulation identifiers in sections 5.8–5.10 and experiments showing these identifiers working in section 5.11. The chapter finishes with related work and a conclusion in sections 5.12 and 5.13.

5.2 Mechanism Design

The overall objective of the prosumer coordination algorithm is to minimise the total system costs (see section 4.4.3). This can be rephrased within the field of social choice, which investigates how to combine individual agent preferences, in order to produce a collective decision [Shoham and Leyton-Brown, 2009]. Preferences are combined by defining a *social choice function* that maps agent preferences to an overall outcome.

For our power systems problem, prosumers and other participants are the agents, their preferences are over different patterns of electricity consumption and the collective outcome is the allocation of power for each agent. The social choice function for this problem returns an outcome that minimises the total system costs.

Mechanism design is a sub-field of game theory and economics that is related to social choice. A *mechanism* defines the rules around how agents interact and how decisions are made in a strategic setting where agents can lie about their preferences [Nisan et al., 2007, Shoham and Leyton-Brown, 2009]. The goal of a mechanism is to produce the same outcome as a social choice function, but in a strategic setting.

Auctions are an example of a class of mechanism, that can be used for allocating a finite resource to potential buyers. In this setting, the social choice function might be to allocate an item to the agent that values it most.

For a particular mechanism, game theory can be applied to find the outcome. When agents know the true preferences of all other agents, it is called a complete information game¹. More commonly, if agents only have beliefs

¹For such cases to remain interesting in mechanism design, the mechanism itself does

about the preferences of others, this is called an incomplete information game. A mechanism successfully implements a social choice function if it creates a game that has dominant strategies or Bayes-Nash equilibria that produce the same outcome as the social choice function applied to the true agent preferences [Shoham and Leyton-Brown, 2009, definitions 10.2.3-4].

5.2.1 Incentive Compatibility

An agent manipulates a mechanism when it gains an advantage by being untruthful (see Nisan et al. [2007, definition 9.4]), which can prevent the mechanism from achieving the social outcome. If it is not possible for agents to manipulate a mechanism (their best strategy is to be truthful) then it is *incentive compatible*. More precisely, a mechanism is:

- dominant-strategy incentive compatible if being truthful is a weakly-dominant strategy for agents; or
- Bayes-Nash incentive compatible if there is a Bayes-Nash equilibrium where all agents act truthfully.

Unfortunately, there are strong theoretical results that show incentive compatibility to be incompatible with other desirable properties.

One of the leading results is the Gibbard-Satterthwaite theorem. It states that for a social choice problem with three or more outcomes, any incentive compatible social choice function that does not preclude an outcome from winning must be a dictatorship [Nisan et al., 2007, theorem 9.8]. A dictatorship, where one agent gets to decide the outcome for the whole system, is clearly not a desirable property.

This suggests that it is futile to design any useful incentive compatible social choice function; however, this result is for the case where agents are allowed to specify arbitrary preferences. By restricting the class of preferences, it turns out that one can develop useful mechanisms that achieve incentive compatibility.

5.2.2 Vickrey-Clarke-Groves Mechanisms

One way of restricting the class of preferences is to work with *utilities* and to introduce monetary payments. An agent's utility function maps each outcome to a real value, where the larger the value the greater their preference. This utility function also takes as input any money paid to, or paid by, the agent. A common assumption is that the money appears as a linear term in the utility function, making it a quasilinear utility function:

$$u(o, m) = m + v(o) \quad (5.1)$$

not have access to the agents' true preferences, otherwise it could trivially just apply the social choice function to determine the outcome.

where o is an outcome, m is money (positive when paid to agent), u is a utility function and v is a valuation function. Here we have made the assumption that the agent is risk neutral (otherwise the money term appears as a monotonically increasing function) [Shoham and Leyton-Brown, 2009, definition 10.3.1]. We also assume that the property of transferable utility holds, allowing utility to be transferred between agents through payments.

When money is involved, mechanisms have a rule for choosing an outcome given the preferences, and a rule for allocating payments to agents. This provides mechanisms with greater flexibility, where payments can be used to induce incentive compatibility.

A commonly used social choice function in this setting is the maximisation of social welfare, where social welfare for a particular outcome is the sum of all agent valuations: $\sum_{i \in A} v_i(o)$. This is equivalent to our desire to minimise total system costs for the prosumer coordination problem.

Vickrey-Clarke-Groves (VCG) mechanisms are an important family of mechanisms, that break free of the Gibbard-Satterthwaite result within this new restricted preference setting. They maximise social welfare and are dominant-strategy incentive compatible. They achieve this because the payments are designed in such a way that each agent's utility becomes dependent on the social welfare.

Many of the mechanisms in the VCG family are impractical because they are not *budget balanced* (the net sum of payments in the mechanism is non-zero), requiring large amounts of funds to be externally sourced or sunk. One common VCG mechanism (often referred to as *the* VCG mechanism) uses the Clarke pivot rule to fully define the payments. Under some, often mild, additional conditions, this rule achieves the desirable properties of weak budget balance and individual rationality.

A mechanism is *weakly budget balanced* when the net sum of payments made to mechanism is non-negative. This means that only an external sink of funds is required, which is much easier to achieve than an external source (in the worst case the money can be burnt). A mechanism is *individually rational* if all agents receive a positive utility. The idea is that if this were not the case, then certain agents would not participate in the mechanism (assuming they have a utility of zero for not participating).

The VCG mechanism has some great properties as we have discussed, but it does come with its negatives [Shoham and Leyton-Brown, 2009, section 10.4.5]. It requires agents to fully reveal their preferences to the mechanism, which has severe privacy implications. Also, it can be computationally intractable to calculate agent payments (largely due to the Clarke pivot rule) for many real-sized problems.

Another problem is that the weak budget balance property only holds in certain domains. For example, it no longer holds in exchange settings where there are multiple buyers and sellers (as is the case in our prosumer coordination problem). Also, while VCG is incentive compatible in a fully

competitive setting, it remains susceptible to collusion. This means that when there is a group of agents that work together, it is no longer necessarily the case that their best strategy is to be truthful.

These negatives illustrate the compromises that need to be considered when adopting an incentive compatible mechanism. These compromises, especially the issues with tractability and privacy, can be too onerous in a large, complex and uncertain distributed system. Section 5.4 discusses how the mechanism we adopt foregoes incentive compatibility in order to achieve many of these other desirable properties.

5.3 Power Balancing Problem

This section introduces the power balancing problem, which is a simplification of the full OPF problem from chapter 4.

The power balancing problem has A agents and a horizon of T time steps. Each agent $i \in \{1, \dots, A\}$ has a power profile $P_i \in \mathbb{R}^T$, a cost function $f_i : \mathbb{R}^T \mapsto \mathbb{R}$ and N_i constraint functions $\forall j \in \{1, \dots, N_i\} : g_{i,j} : \mathbb{R}^T \mapsto \mathbb{R}$. For now these functions are restricted to being continuously differentiable and/or convex.

The problem now splits into three versions: the single horizon (SH), receding horizon (RH) and terminating receding horizon (TRH) problems.

5.3.1 Single Horizon Version

For the SH version, the objective is to minimise the cost functions of all agents, subject to the agent constraints and a power conservation constraint:

$$\min_{P_i} \sum_{i=1}^A f_i(P_i) \quad (5.2)$$

$$\text{s.t. } g_{i,j}(P_i) \leq 0 \quad \forall i \in \{1, \dots, A\}, j \in \{1, \dots, N_i\} \quad (5.3)$$

$$\sum_{i=1}^A P_i = 0 \quad (5.4)$$

5.3.2 Receding Horizon Version

The RH version of the problem is an infinite sequence of SH problems. Let each horizon $h \in \{1, 2, \dots\}$ be T time steps long. The functions and variables become indexed by the horizon to which they belong: $f_{h,i}$, $g_{h,i,j}$ and $P_{h,i}$.

The optimisation problem for the h -th horizon is:

$$\min_{P_{h,i}} \sum_{i=1}^A f_{h,i}(P_{h,i}) \quad (5.5)$$

$$\text{s.t. } g_{h,i,j}(P_{h,i}) \leq 0 \quad \forall i \in \{1, \dots, A\}, j \in \{1, \dots, N_i\} \quad (5.6)$$

$$\sum_{i=1}^A P_{h,i} = 0 \quad (5.7)$$

We use an index t to access individual time steps within the power vectors (and later on also the price vectors). For convenience and consistency with the final formulation, the index t is relative to the time steps in the first horizon, no matter which horizon it is used with (see figure 5.1). This means that the powers in the first and second horizons $P_{1,i} \in \mathbb{R}^T$ and $P_{2,i} \in \mathbb{R}^T$ are accessed with the indices $t \in \{1, \dots, T\}$ and $t \in \{2, \dots, 2 + T - 1\}$ respectively. For example, $t = 3$ represents the third time step no matter what the horizon is, and the values $P_{1,i,3}$ and $P_{2,i,3}$ represent the powers in time step 3 for agent i , but as calculated in two different horizons.




Figure 5.1: A visualisation showing how the time step index t is relative to the first horizon for all horizons.

5.3.3 Terminating Receding Horizon Version

We cannot directly perform experiments on the RH version of the problem because it runs over an infinite sequence of horizons. The TRH version reduces the problem to a finite number of horizons which truncate at a particular time step. It is made up of T single horizon power balancing problems $h \in \{1, \dots, T\}$, where the horizons shrink in size until the last horizon only contains a single time step.

For our formulation, instead of explicitly shrinking the horizon lengths, we model this as T perfectly overlapping horizons, each of length T , and force some of the variables in later horizons to take on the values of earlier

horizons. For the h -th horizon the optimisation problem is:

$$\min_{P_{h,i}} \sum_{i=1}^A f_{h,i}(P_{h,i}) \quad (5.8)$$

$$\text{s.t. } g_{h,i,j}(P_{h,i}) \leq 0 \quad \forall i \in \{1, \dots, A\}, j \in \{1, \dots, N_i\} \quad (5.9)$$

$$\sum_{i=1}^A P_{h,i} = 0 \quad (5.10)$$

$$P_{h,i,t} = P_{h-1,i,t} \quad \forall i \in \{1, \dots, A\}, t < h \quad (5.11)$$

5.3.4 Agents

This chapter uses a small subset of DETs to assess the opportunity for manipulation and the accuracy of our identifiers. They are: fuel-based generators, deferrable loads and fixed power loads. They were chosen as they provide the basic behaviour of network agents, without making the strategic agent that we develop too complicated to implement or solve.

Generator— We assume the fuel-based generator has the quadratic cost:

$$f(p) = \sum_{t=1}^T \psi_t P_t^2 \quad (5.12)$$

where $\psi \in \mathbb{R}_+^T$, and with bounds on power $P_t \in [P, 0]$.

Deferrable Load— A deferrable load is similar to the relaxed model of the shiftable loads that we have encountered in previous chapters. The power of the deferrable load has bounds which change over time $P_t \in [0, \bar{P}_t]$. The deferrable load has a zero cost function, but it must consume a total amount of energy E where:

$$\sum_{t=1}^T P_t = E \quad (5.13)$$

Fixed Power— The fixed power agent has a zero cost function and fixed power $P = P_F$, where $P_F \in \mathbb{R}^T$ is a constant.

5.4 Power Balancing Mechanism

The mechanism we propose to solve the power balancing problem is the same as the distributed algorithm from chapter 4. This mechanism is not incentive compatible (see section 5.5.1 for examples), but it is computationally efficient, preserves privacy, is budget balanced, and converges to a socially optimal outcome *when* agents act truthfully.

For this chapter, only the solutions (outcomes) that the mechanism achieves are important, rather than the algorithm specifics. ADMM and other distributed optimisation algorithms can be categorised as finding a Karush-Kuhn-Tucker (KKT) point for the problem. We therefore abstract our mechanism to that of iteratively finding a KKT point for each horizon in turn.

5.4.1 KKT Conditions

For a given power balancing problem, we assume that there exists at least one KKT point for each horizon (section 5.4.4 discusses what happens when a point cannot be found). Under certain regularity conditions (constraint qualifications) on the problem functions, a KKT point becomes necessary for optimality [Kuhn and Tucker, 1951, Karush, 1939]. Under other conditions such as convexity, a KKT point is sufficient for optimality [Martin, 1985].

When the functions are continuously differentiable, the KKT conditions for an agent i in horizon h are:

$$\nabla f_{h,i}(P_{h,i}) + \sum_{j=1}^{N_i} \mu_{h,i,j} \nabla g_{h,i,j}(P_{h,i}) + \lambda_h = 0 \quad (5.14)$$

$$g_{h,i,j}(P_{h,i}) \leq 0 \quad \forall j \in \{1, \dots, N_i\} \quad (5.15)$$

$$\mu_{h,i,j} g_{h,i,j}(P_{h,i}) = 0 \quad \forall j \in \{1, \dots, N_i\} \quad (5.16)$$

$$\mu_{h,i,j} \geq 0 \quad \forall j \in \{1, \dots, N_i\} \quad (5.17)$$

where $\mu_{h,i,j} \in \mathbb{R}$ are the KKT multipliers for the agent constraints and $\lambda_h \in \mathbb{R}^T$ is the KKT multipliers for the power conservation constraint. Particular powers $P_{h,i}$ and prices λ_h satisfy the KKT conditions for agent i if there exists some $\mu_{h,i,j}$ where this is true.

The KKT conditions for the overall problem in horizon h is the combination of KKT conditions for all agents along with the power conservation constraint:

$$\sum_{i=1}^A P_{h,i} = 0 \quad (5.18)$$

When dealing with convex functions, which might be non-differentiable, we allow a subgradient to be used in place of the gradient. Therefore, if f is a convex function, then we define $\nabla f(P) \in \partial f(P)$, where $\partial f(P)$ is the subdifferential (set of subgradients) of f at P . We will explicitly mention when this overloading of the differential operator leads to any ambiguities.

Using this notation the KKT stationarity condition (5.14) remains the same, except we have the implicit assumption that $\nabla f(P)$ can be used to represent a subgradient for any convex cost or constraint functions.

5.4.2 Payments

Agents are paid (charged) for their power production (consumption) at the marginal prices $\lambda_{h,t} \in \mathbb{R}$. As discussed, these are the KKT multipliers for the power conservation constraints:

$$\forall t \in \{1, \dots, T\} : \sum_{i=1}^A P_{h,i,t} = 0 \quad (\lambda_{h,t}) \quad (5.19)$$

The expected total cost $c_{h,i}$ for an agent i in horizon h is the combination of their cost function and their expected payments over the horizon:

$$c_{h,i} := f_{h,i}(P_{h,i}) + \lambda_h^\top P_{h,i} \quad (5.20)$$

This is only what is expected in horizon h , because the actual costs will change as solutions are allowed to change in later horizons. Defining the costs like this assumes that agents have quasilinear utilities and are risk neutral (as discussed in section 5.2.2). The power conservation constraint ensures that our mechanism is budget balanced, as all payments between consumers and producers sum to zero for each time step.

In the receding horizon problem we assume that agents must consume or produce the power negotiated in the first time step of each horizon, which (recalling the discussion from section 5.3.2) is where $t = h$. The powers and prices beyond the first time step in each horizon are just hypothetical at that point in time.

5.4.3 Convexity

Chapters 3 and 4 discussed how power consumption decisions can be discrete or non-convex in nature, which raises doubt about the applicability of the KKT conditions. However, there are some strong arguments for why in practice convex relaxations or approximations of prosumer device models will produce high quality results:

- Chapter 4 found that the most common source of discrete decisions for houses (shiftable loads) are well approximated by their convex relaxation when the number of participants reaches realistic levels.
- Many new household devices are becoming more continuously variable. Variable speed compressors used in devices such as air conditioners can provide a more efficient alternative to discrete on/off switched compressors.
- More intelligent control can be used to apply a duty cycle to discrete components to mimic more continuous power levels.

- Household batteries will be able to compensate for the discrete power levels of other devices.

As a result, convex functions can be used with only a small hit to quality. We require, as a condition of participation, that agents reformulate any non-convex cost and constraint functions into a convex approximations or relaxations. Any increase in cost will be more than offset by the benefits of being able to efficiently solve large systems in a distributed manner.

5.4.4 Limits

Power Limits— In order to participate in the mechanism, each agent must first negotiate a contract which restricts how much power they can consume and supply. This contract establishes limits such that $\forall t \in \{1, \dots, T\} : P_{i,t} \in [\underline{P}_i, \bar{P}_i]$. They can be set at the physical limits of the equipment connecting the agent to the rest of the network, or to a subinterval of these physical limits if the agent does not need the extra capacity.

It is likely that agents will be provided with a financial incentive to set tighter limits as it provides the market operator with more certainty. A benefit of having these limits is that it restricts the outcomes available to strategic agents. For example, without any limits, a two-bedroom house could indicate to the mechanism that they will provide the network with the equivalent output of a large nuclear power station, which would clearly have serious consequences for the mechanism.

We assume that these limits are enforced in the private constraints of each agent. In practice this would be combined with a simple check by the utility, with adequate penalties that would deter agents from choosing tighter limits than they can meet.

Price Limits— A second set of limits is on the allowed prices λ , which are commonly referred to as market caps in wholesale electricity markets. This can be set at a point where participants become indifferent towards being disconnected (shed) from the network. If the solver finds a KKT point with costs outside the price caps or if no KKT point is found at all, then the market operator can shed loads in an attempt to find a feasible solution.

Agents have the ability to respond to prices, so they will tend to reduce their loads on their own as the prices increase. The price caps and load shedding by the operator will likely only be needed as a last resort, for example, if agents do not have enough flexibility or if something goes seriously wrong.

The market operator could choose to shed those agents that are operating outside their normal behaviour combined with some degree of randomisation. Although we do not go into specific approaches to shedding [Concordia and Fink, 1995], it is worth noting that the actions of strategic agents could cause them to be shed if they push the price outside the market caps. This

is how the price limits further restrict the possible outcomes for strategic agents.

5.5 Manipulation

This section defines truthfulness, consistency and receding horizon manipulation. It then provides some simple examples of manipulation with and without a receding horizon and discusses agent strategies. First some notation is introduced.

When we focus on an arbitrary pair of consecutive horizons h and $h + 1$, we drop the horizon subscript from functions, powers and prices, and instead indicate the later horizon with a prime. For example, the cost functions for agent i in two consecutive horizons are $f_i := f_{h,i}$ and $f'_i := f_{h+1,i}$. To make comparisons easier, we time shift powers and prices from the later horizon so that they line up with the values from the earlier horizon (see figure 5.2). These vectors are marked with an asterisk and are defined as:

$$P_{i,t}^* := \begin{cases} P_{h,i,t} & \text{if } t = h \\ P_{h+1,i,t} & \text{if } h < t < h + T \end{cases} \quad (5.21)$$

$$\lambda_t^* := \begin{cases} \lambda_{h,t} & \text{if } t = h \\ \lambda_{h+1,t} & \text{if } h < t < h + T \end{cases} \quad (5.22)$$

We define vectors for the change in the power and price as $\Delta P_i := P_i^* - P_i$ and $\Delta \lambda := \lambda^* - \lambda$. We also define the change in cost with respect to the earlier horizon as $\Delta c_i := f_i(P_i^*) + \lambda^{*\top} P_i^* - f_i(P_i) - \lambda^\top P_i$.




Figure 5.2: For consecutive horizons, the shaded time steps represent the values used for the time shifted price and power vectors.

As discussed in chapter 3, agents, especially household prosumers, are exposed to a lot of uncertainty that influences their electricity consumption. Instead of modelling this through random parameters, we instead more abstractly allow agents to update their cost function and constraints for each horizon. We use a hat to indicate an agent's best estimate of these functions for a particular horizon $\hat{f}_{h,i}$ and $\hat{g}_{h,i,j}$.

The functions that the agent uses to interact with the mechanism are $f_{h,i}$ and $g_{h,i,j}$. When these are the same as their predicted functions, the agent is truthful. Formally:

Definition 1 (Truthfulness). For a horizon h , let $f_{h,i}$ and $g_{h,i,j}$ be the functions that agent i uses with the mechanism and $\hat{f}_{h,i}$ and $\hat{g}_{h,i,j}$ be the functions the agent expects are most likely to be true at that moment in time. Agent i is *truthful* in a horizon h if $f_{h,i} = \hat{f}_{h,i}$ and for all j : $g_{h,i,j} = \hat{g}_{h,i,j}$. Otherwise agent i is *untruthful*.

An agent is consistent between consecutive horizons when the preferences they provide in the earlier horizon could have produced the result from the later horizon (for those time steps that overlap). This is more complicated than saying the functions are equivalent in both horizons, because in the RH problem the functions accept powers from different time steps. Formally:

Definition 2 (Consistency). Let the KKT points for the power balancing problem in two consecutive horizons be (P_i, λ) and (P'_i, λ') for agent i . If the later horizon time shifted point (P_i^*, λ^*) satisfies the KKT conditions for agent i in the earlier horizon, then agent i is *consistent* between the horizons. That is, the following conditions must hold:

$$\nabla f_i(P_i^*) + \sum_{j=1}^{N_i} \mu_{i,j}^* \nabla g_{i,j}(P_i^*) + \lambda^* = 0 \quad (5.23)$$

$$g_{i,j}(P_i^*) \leq 0 \quad \forall j \in \{1, \dots, N_i\} \quad (5.24)$$

$$\mu_{i,j}^* g_{i,j}(P_i^*) = 0 \quad \forall j \in \{1, \dots, N_i\} \quad (5.25)$$

$$\mu_{i,j}^* \geq 0 \quad \forall j \in \{1, \dots, N_i\} \quad (5.26)$$

for some multipliers $\mu_{i,j}^*$. Otherwise agent i is *inconsistent* between the horizons.

An agent manipulates the receding horizon mechanism when it is untruthful in order to create an inconsistent result. Formally:

Definition 3 (Receding Horizon Manipulation). If agent i is inconsistent between consecutive horizons and untruthful in the earlier horizon, then agent i is *manipulating* the receding horizon problem between these horizons.

5.5.1 Examples

This section provides two examples of how untruthful agents can manipulate the power balancing problem. The first example demonstrates how agents can manipulate the SH problem and the second example demonstrates receding horizon manipulation (which is the focus of this chapter) in the TRH problem.

Consider a SH problem instance with a single time step on a network with a fixed load of $\hat{P}_L = 5$ and two generators (A and B) with true generation prices $\hat{\psi}_A = 2$ and $\hat{\psi}_B = 4$ (equation (5.12) provides their cost functions).

Figure 5.3 shows how the outcome of the mechanism changes as generator A misreports its cost function. The optimal total cost (social optimal) is achieved when it reports truthfully, i.e. $\psi_A = \hat{\psi}_A = 2$; however, the figure shows that generator A can reduce its individual costs if it falsely reports a higher cost function. This action from A also reduces the costs of generator B , which is reporting truthfully in this example, with the savings coming at the expense of the load.




Figure 5.3: Change in outcome as generator A changes its reported cost function. The social optimal occurs at $\psi_A = 2$, when the generator truthfully reports its cost.

The lack of incentive compatibility in the general case can be insignificant in a particular problem instance if there is sufficiently strong competition, as in other markets [Makowski et al., 1999]. We expect there to be a reasonable level of competition in future prosumer driven networks, especially in comparison to existing electricity markets.

For the second example that demonstrates receding horizon manipulation, consider a TRH problem with two time steps on a network with two deferrable loads (A and B) with $\hat{E}_A = \hat{E}_B = 2$ (to simplify assume no upper power bounds) and a single generator with prices in the two time steps of $\hat{\psi}_1 = 4$ and $\hat{\psi}_2 = 2$. Load A can take advantage of the cheaper prices in the second time step by discouraging B from consuming at this time.

The results in figure 5.4 show how the overall outcome changes as load A lies during the first horizon about how much it needs to consume in the second time step $P_{1,A,2}$. In each instance A reports zero consumption in the first time step $P_{1,A,1} = 0$, and it corrects the lie in the second horizon

by stating truthfully how much power it still needs in the second time step (after which it is too late for B to make a correction).

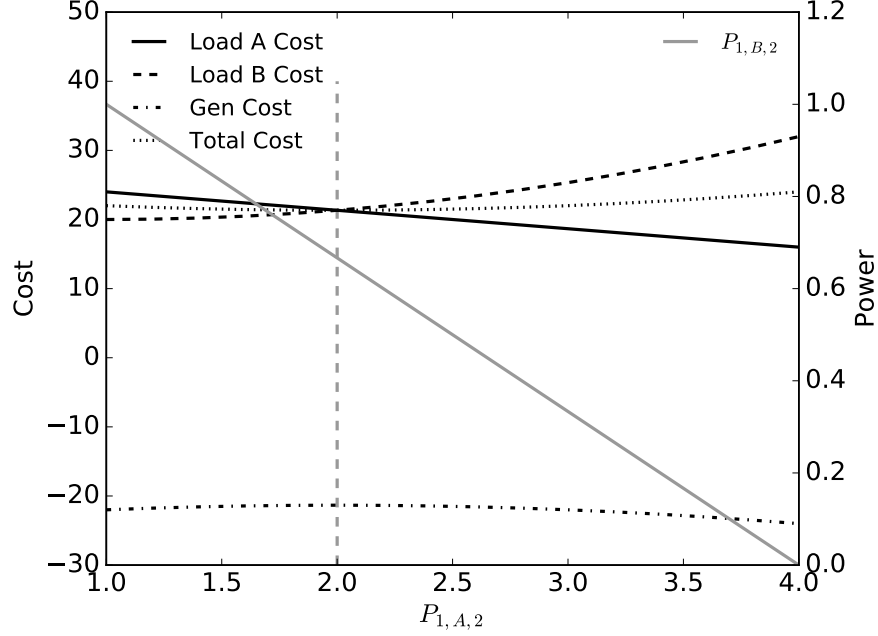


Figure 5.4: Change in outcome as deferrable load A changes its reported power consumption for the second time step $P_{1,A,2}$. The total costs are equivalent to the social outcome when $P_{1,A,2} = 2$.

As load A increases $P_{1,A,2}$, load B responds by decreasing its second time step consumption $P_{1,B,2}$, which in turn reduces the costs for A . The outcome where $P_{1,A,2} = 2$ is socially optimal, and is equivalent to outcome of the mechanism with A acting truthfully.

5.5.2 Strategies

There are many reasons why it might be difficult for an agent to manipulate the mechanism in a practical setting, including the:

- computational intractability of computing a beneficial strategy in realistically sized settings;
- limited information about the private preferences and constraints of others;
- strong competition between agents;
- need for collusion with other agents with complementary needs; and
- rarity of circumstances where a worthwhile benefit can be achieved.

However, we cannot rule out the existence of very simple (sub-optimal) strategies that, both in theory and in practice, can cause significant problems when adopted by many agents. In this section we discuss some possible strategies.

In general, the power balancing mechanism results in an incomplete information game. This means that each agent knows how the mechanism works and the type of functions that it allows, but not necessarily the number of other agents or their types. Some information, for example, a distribution over the types of agents that might be present, will be needed to make sensible strategic decisions. In general, only knowing how the mechanism works is not enough.

For example, consider the task of developing a pure strategy for an agent that has a fixed power that is non-zero at the very first time step of the horizon. The agent must be honest about the first time step, but they can lie in future ones. We find that for any strategy that involves lying about what the agent can achieve in a future time step, we can always construct another agent that will actually increase the costs for the fixed power agent relative to their truthful strategy (see lemma 3 in appendix B).

As agents gain information about the other agents in the market, they may start to form mixed or pure strategies that can reliably manipulate the mechanism. They may also decide to form coalitions of agents that work together to have a larger impact. The exact formation of such coalitions and any equilibria in the mechanism will be strongly dependent on what agents are present, and the nature and amount of information that is known. We leave it to future work to investigate such Bayes-Nash equilibria in incomplete information cases.

Instead, in this chapter we focus on a simple concrete case where a single “greedy” agent has complete information about all other agents.

5.6 Greedy Agent Strategy

The “greedy” agent we develop has complete information about all other agents (which act truthfully) and optimises its costs by manipulating the mechanism. This provides it with the best possible chance of manipulating the mechanism.

This section describes a general algorithm for calculating an optimal strategy for such a greedy agent. This is an artificial scenario designed to establish the worst case harm that a single agent (or coalition of such agents working together) can cause. Later on in the chapter we will demonstrate techniques for identifying untruthful agents such as this greedy one.

The approach relies on solving a bilevel program, where the greedy agent at the top level chooses its optimal preferences given that it knows how the mechanism will respond at the lower level [Migdalas et al., 1998]. In general,

even linear bilevel programming is NP-hard [Migdalas et al., 1998, chapter 6]. This will dramatically limit the size and type of problems that can be solved optimally. We will use bounds and solve simple instances in order to gain insights despite the complexity.

The algorithm is formulated for the TRH problem. The greedy agent has index $i = 1$ and it knows the preferences of all other agents who act truthfully. The greedy agent can lie about their preferences by fixing their power consumptions to fixed values during each horizon. There is no uncertainty: for all h let $f_{h,i} = f_i$ and $g_{h,i,j} = g_{i,j}$. We assume that there are market cap prices $\underline{\lambda}$, $\bar{\lambda}$ and that the greedy agent will avoid a solution which violates these because of the risk of being shed from the network.

We make the *optimistic* bilevel assumption, which allows the greedy agent to choose between lower level solutions if more than one exists for a given upper level decision. This is the same as allowing the greedy agent to choose between KKT points of the lower level mechanism. This allows the problem to be immediately flattened into a single level problem:

$$\min_{P_{h,i}, \lambda_h} f_1(P_{T,1}) + \lambda_T^T P_{T,1} \quad (5.27)$$

$$\text{s.t. } g_{1,j}(P_{T,1}) \leq 0 \quad \forall j \in \{1, \dots, N_1\} \quad (5.28)$$

$$P_{h,i,t} = P_{h-1,i,t} \quad \forall i \in \{1, \dots, A\}, h \in \{2, \dots, T\}, t < h \quad (5.29)$$

$$\lambda_{h,t} = \lambda_{h-1,t} \quad \forall h \in \{2, \dots, T\}, t < h \quad (5.30)$$

$$P_{h,1} \in [\underline{P}_1, \bar{P}_1]^T \quad \forall h \in \{1, \dots, T\} \quad (5.31)$$

$$\lambda_h \in [\underline{\lambda}, \bar{\lambda}]^T \quad \forall h \in \{1, \dots, T\} \quad (5.32)$$

$$(P_{h,1}, P_{h,2}, \dots, P_{h,A}, \lambda_h) \in \text{KKT}_h \quad \forall h \in \{1, \dots, T\} \quad (5.33)$$

The optimisation problem is over the powers for all agents and the prices in all horizons. The objective is to minimise the actual costs of the greedy agent ($i = 1$) at the end of the TRH problem. This is given by the agent's true cost function f_1 and monetary payments. The powers and prices in the last horizon ($h = T$) are used for these calculations, because at this point they are fully finalised (in earlier horizons some of these are still hypothetical).

The constraints for the agent are enforced on the final true powers for the greedy agent with inequality (5.28). Equations (5.29) and (5.30) tie together powers and prices between horizons (those that have been finalised), in accordance with the TRH formulation. Constraints (5.31) and (5.32) limit the greedy agent's power values and the prices according to the limits discussed in section 5.4.4.

KKT_h is defined to be the set of feasible KKT points for the problem when for all time steps where $t < h$ the prices and powers are fixed to some arbitrary values, and for all time steps the powers of agent 1 are fixed to some arbitrary values. For a given h , constraint (5.33) can be expanded into

the constraints:

$$\sum_{i=1}^A P_{h,i,t} = 0 \quad \forall t \geq h \quad (5.34)$$

and for all $i \in \{2, \dots, A\}$:

$$\nabla_t f_i(P_{h,i}) + \sum_{j=1}^{N_i} \mu_{h,i,j} \nabla_t g_{i,j}(P_{h,i}) + \lambda_{h,t} = 0 \quad \forall t \geq h \quad (5.35)$$

$$g_{i,j}(P_{h,i}) \leq 0 \quad \forall j \in \{1, \dots, N_i\} \quad (5.36)$$

$$\mu_{h,i,j} g_{i,j}(P_{h,i}) = 0 \quad \forall j \in \{1, \dots, N_i\} \quad (5.37)$$

$$\mu_{h,i,j} \geq 0 \quad \forall j \in \{1, \dots, N_i\} \quad (5.38)$$

The stationarity (5.35) and complementary slackness (5.37) equalities and the multiplication of λ_T in the objective (5.27) are all possible sources of non-convexity. We will now further develop our three devices so that they can be used in this formulation as one of the truthful agents.

5.6.1 Generator

The KKT conditions for a truthful generator i can be reduced to $\forall t \geq h$:

$$2\psi_{i,t} P_{h,i,t} + \mu_{h,i,t}^u - \mu_{h,i,t}^l + \lambda_{h,t} = 0 \quad (5.39)$$

$$P_{h,i,t} \leq 0 \quad (5.40)$$

$$\underline{P}_i - P_{h,i,t} \leq 0 \quad (5.41)$$

$$\mu_{h,i,t}^u P_{h,i,t} = 0 \quad (5.42)$$

$$\mu_{h,i,t}^l (\underline{P}_i - P_{h,i,t}) = 0 \quad (5.43)$$

$$\mu_{h,i,t}^u \geq 0 \quad (5.44)$$

$$\mu_{h,i,t}^l \geq 0 \quad (5.45)$$

The generator has upper and lower power bounds at each time step, for which we associate the dual variables $\mu_{h,i,t}^u$ and $\mu_{h,i,t}^l$. The time index is used to indicate which time step the dual variables are for. This is reformulated as a series of mixed-integer linear constraints by introducing binary variables $z_{h,i,t}^u$ and $z_{h,i,t}^l$, and combining the KKT multipliers. The complementary slackness requires that $\mu_{h,i,t}^u > 0 \implies \mu_{h,i,t}^l = 0$ and $\mu_{h,i,t}^l > 0 \implies \mu_{h,i,t}^u = 0$. This means that we can replace each pair of upper and lower bound KKT multipliers by a single multiplier, $\nu_{h,i,t} \in \mathbb{R}$, where $\nu_{h,i,t} = \mu_{h,i,t}^u - \mu_{h,i,t}^l$. Big-M style constant bounds $\underline{\nu}_{i,t}$ and $\bar{\nu}_{i,t}$ are used for these new multipliers. The reformulation is $\forall t \geq h$:

$$2\psi_{i,t} P_{h,i,t} + \nu_{h,i,t} + \lambda_{h,t} = 0 \quad (5.46)$$

$$\nu_{h,i,t} \leq z_{h,i,t}^u \bar{\nu}_{i,t}, \quad P_{h,i,t} \geq (1 - z_{h,i,t}^u) \underline{P}_i \quad (5.47)$$

$$\nu_{h,i,t} \geq z_{h,i,t}^l \underline{\nu}_{i,t}, \quad P_{h,i,t} \leq z_{h,i,t}^l \underline{P}_i \quad (5.48)$$

The market price caps and the stationarity condition provide bounds for the multiplier $\nu_{h,i,t}$:

$$\underline{\nu}_{i,t} = -\bar{\lambda} \quad (5.49)$$

$$\bar{\nu}_{i,t} = -\underline{\lambda} - 2\psi_{i,t}\underline{P}_i \quad (5.50)$$

5.6.2 Deferrable Load

For a truthful deferrable load, the KKT conditions can be transformed into mixed-integer linear constraints in the same way as the generator:

$$\gamma_{h,i} + \nu_{h,i,t} + \lambda_{h,t} = 0 \quad \forall t \geq h \quad (5.51)$$

$$\sum_{t=1}^T P_{h,i,t} = E_i \quad (5.52)$$

$$\nu_{h,i,t} \leq z_{h,i,t}^u \bar{\nu}_{i,t}, \quad P_{h,i,t} \geq z_{h,i,t}^u \bar{P}_{i,t} \quad \forall t \geq h \quad (5.53)$$

$$\nu_{h,i,t} \geq z_{h,i,t}^l \underline{\nu}_{i,t}, \quad P_{h,i,t} \leq (1 - z_{h,i,t}^l) \bar{P}_{i,t} \quad \forall t \geq h \quad (5.54)$$

A difference here is the presence of a KKT multiplier $\gamma_{h,i}$ associated with the energy consumption constraint, which makes $\nu_{h,i,t}$ unbounded in general. However, there exist finite bounds which do not cut off any feasible KKT point if we are only interested in the values of $P_{h,i}$ and λ_h (which are all that matter for the overall problem).

If $(\tilde{P}_{h,i}, \tilde{\lambda}_h, \tilde{z}_{h,i}^u, \tilde{z}_{h,i}^l, \tilde{\gamma}_{h,i}, \tilde{\nu}_{h,i})$ is a KKT point, then so is $(\tilde{P}_{h,i}, \tilde{\lambda}_h, \tilde{z}_{h,i}^u, \tilde{z}_{h,i}^l, \tilde{\gamma}_{h,i} + \epsilon, \tilde{\nu}_{h,i} - \vec{1}\epsilon)$, where $\epsilon = \tilde{\nu}_{h,i,\tau}$ given $\tau = \arg \min_{t \in \{h, \dots, T\}} |\tilde{\nu}_{h,i,t}|$, and $\vec{1}$ is the appropriate all-ones vector. This means that for any values $(P_{h,i}, \lambda_h)$ that satisfy the KKT conditions, we can choose the other multipliers in such a way that there exists some τ where $\nu_{h,i,\tau} = 0$. By cancelling $\gamma_{h,i}$, the stationarity conditions require that for all $t \geq h$:

$$\nu_{h,i,t} + \lambda_{h,t} = \nu_{h,i,\tau} + \lambda_{h,\tau} \quad (5.55)$$

$$\implies \nu_{h,i,t} + \lambda_{h,t} = \lambda_{h,\tau} \quad (5.56)$$

Therefore the market price caps provide the bounds:

$$\underline{\nu}_{i,t} = \underline{\lambda} - \bar{\lambda} \quad (5.57)$$

$$\bar{\nu}_{i,t} = -\underline{\lambda} + \bar{\lambda} \quad (5.58)$$

5.6.3 Fixed Power Device

The conditions for a truthful fixed power device are trivial $\forall t \geq h$: $P_{h,i,t} = P_{F,t}$. When the fixed power device is greedy the problem is also nicely simplified, as the objective becomes linear.

5.7 Greedy Agent Experiments

We perform experiments which show how much the greedy agent has to gain, and how much this hurts the efficiency of the overall system. We have two sets of experiments, one where the greedy agent is a deferrable load and the other where it is a fixed power load.

5.7.1 Problem Instances

Each experiment has 4 truthful deferrable loads and 1 truthful generator ($A = 6$ including the greedy agent). We randomly sampled TRH instances with 4 time steps ($T = 4$), where the prices, energy and power requirements vary.

The required energy of each deferrable load is sampled from a uniform distribution $\hat{E}_i \sim \mathcal{U}(1, 5)$. The power bound for each time step for each load is also sampled from a uniform distribution $\hat{P}_{i,t} \sim \mathcal{U}(0, \hat{E}_i)$. If this sampling results in a $\sum_t \hat{P}_{i,t} < \hat{E}_i$, then the power bounds are resampled until a feasible result is obtained. The generator has no limit on how much power it can supply and its prices are sampled from a uniform distribution $\hat{\psi}_{i,t} \sim \mathcal{U}(1, 5)$.

The greedy agent is sampled differently so that we can vary its size. We introduce a scale parameter s , which represents this agent's size relative to the other agents. A value of $s = 1$ means that the greedy agent has the same energy requirements (on average) as the truthful agents. Because there are in total 5 loads, this means that the agent requires 1/5th of the energy in our problem. We run experiments with the greedy agent scale set at 0.1, 0.2, 0.5 and 1, which ranges from 1/40th to 1/5th of the total system load.

The greedy agent's energy requirement is $\hat{E}_i \sim s \times \mathcal{U}(1, 5)$. The power bounds are sampled using this energy requirement as for the other deferrable agents. For the experiments with the deferrable greedy agent we stop here, for the fixed greedy agent we take these values as a basis for building a fixed load: $\hat{P}_{F,i,t} = \hat{E}_i \hat{P}_{i,t} / \sum_{t=1}^T \hat{P}_{i,t}$. The greedy agent is given negotiating power limits (see section 5.4.4) of $P_i = 0$ and $\bar{P}_i = s \times 5$.

Ideally we want to find globally optimal solutions for the greedy agent, so that we have a bound on the worst case manipulation. This, along with the complexity of bilevel programming, severely limits the size of the problems in our experiments. Our experiments only have a total of 6 agents whereas in chapter 4 our microgrid had over 3000. The experiments have been designed to capture the main behaviour of a more realistic setting, so that we can still draw useful insights into receding horizon manipulation.

The first design choice is the use of truthful deferrable loads which have time coupled behaviour and a total energy consumption requirement. Many DETs including hot water heating, space heating and cooling, electric vehicle

charging and smart appliances are well represented by deferrable loads as a first approximation.

In most scenarios small agents will have very little influence on the rest of the network. In order for small agents to gain a meaningful benefit from manipulation, they will have to collude with other agents. The single greedy agent in our experiments can represent such a coalition. By varying the scale parameter we investigate the cases where such a coalition (or single agent) controls between 1/40th and 1/5th of the total system load.

5.7.2 Solving

When the greedy agent is a fixed load the problem can be reduced to a MILP. These problems were solved to optimality using Gurobi 6.5. The problem is a MINLP when a deferrable greedy agent is used. An attempt was made to use the global solver Couenne (versions 0.4 and 0.5), which in many cases appeared to return an optimal solution. Unfortunately the powers and prices that were returned from the solver were not feasible for the problem when independently checked. Instead Bonmin 1.8 was used to find local solutions and was given a 5 hour time limit for each instance.

Bonmin is only expected to return local solutions for the deferrable greedy agents (due to the bilinear term in the objective); however, there are a few signs to indicate that it found close to optimal solutions. Firstly, the objectives returned by Bonmin and Couenne were almost the same. We do not know what went wrong with the Couenne returned solution, but if this is a simple bug in the interface rather than the solver itself then this would indicate that the Bonmin results are near optimal. Secondly, the deferrable greedy agent results followed the same trend as the fixed greedy agent, which we do know to be optimal.

5.7.3 Greedy Agent Results

For both types of greedy agent (fixed and deferrable) and each scale parameter value, 50 random problem instances were generated. For comparison, we calculated an optimal solution when all agents are truthful for each instance, which is also known as the social outcome.

Figures 5.5 and 5.6 show the results for the fixed and deferrable greedy agents. The results are provided as changes in costs for the greedy agent and the total system costs, relative to the optimal social outcome. These are provided as an average, and 10th and 90th percentile regions over all instances for each scale parameter.

As expected, greedy agents can reduce their costs, which also ends up increasing the total system costs. When the fixed and deferrable greedy agents have 1/5th of the total load ($s = 1$), they can reduce their average costs by 10% and 25% respectively. With 1/40th of the total load ($s = 0.1$),




Figure 5.5: Percent change in fixed power greedy agent and total costs relative to social solution. Shaded areas give 10th and 90th percentiles for all 50 problem instances.

these values are 2% and 5% respectively. The fixed and deferrable greedy agents consume the same amount of power on average for a given scale, but the greater flexibility of the deferrable greedy agent allows it to form more successful strategies.

At small scale values the total costs do not deviate far from the optimal social solution. This means that the cost increases (for the sum of the truthful agents) are almost equal to the cost reductions (for the greedy agent). At large scales, the cost increases are significantly more than the cost decreases, so the greedy agent does much more harm relative to what they gain.

We have shown that agents have the potential to gain significant benefits through receding horizon manipulation, but the question remains as to whether or not it will be plausible in practice. Take the suburb-sized microgrid from chapter 4 as an example. To achieve the scale value of $s = 0.1$ would require around 90 of the 3674 houses to work together. Assuming they are made up of a combination of fixed and deferrable loads, in the best case they could achieve a cost reduction of a few percent.

There are two main reasons why achieving even this small reduction of a few percent will be challenging. Firstly, the greedy agents will not have complete information, so at times they will make mistakes that actually cost themselves more than the social outcome. Secondly, the computational burden of calculating a strategy in such a setting is high. Approximate methods would almost certainly have to be used, which in turn will lead a

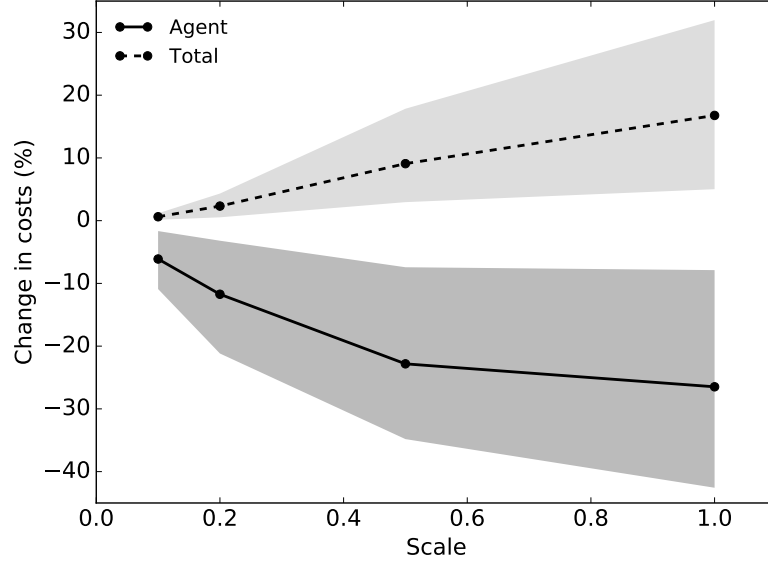


Figure 5.6: Percent change in deferrable greedy agent and total costs relative to social solution. Shaded areas give 10th and 90th percentiles for all 50 problem instances.

further degradation of their expected outcome. Overall it will not be a very attractive proposition.

However, the formation of larger coalitions may remain attractive (assuming the computational challenge can be overcome), and also smaller numbers of agents can have a larger impact in parts of the network where network constraints (e.g., thermal or voltage) are commonly active. When active, the constraints can effectively isolate one part of the network from another. For example, if the line feeding an area is at its maximum capacity, then any additional generation needs to come locally from within that area. The number of active agents is now reduced to the smaller area, and so individual agents can now have a much larger impact.

In practice these situations are likely to be relatively rare, but they are plausible enough to make it worth investigating what we can do to prevent or reduce the impact of manipulation. In the section that follows we investigate ways of identifying receding horizon manipulation.

5.8 Identifying Manipulative Agents

Instead of attempting to adopt an incentive compatible mechanism and forgoing the other desirable properties discussed in section 5.4, we propose a lightweight monitor which can detect receding horizon manipulations. Importantly, it does not affect the other desirable properties of the mechanism when agents act truthfully.

It may not identify all manipulation, but it will be able to deter the worst offenders, and limit the social harm. The major challenge is to distinguish an agent that is behaving strategically from one that is uncertain about their consumption for legitimate reasons (e.g., uncertain occupant behaviour or weather patterns). For the most part we do not want to penalise uncertainty because it will discourage agents from providing their latest predictions, which is the main advantage of receding horizon control.

Another complicating factor is that truthful agents will change their behaviour across successive horizons in response to the actions of others, even if they have no uncertainty. Indeed, they might even gain a benefit by the unscrupulous behaviour of others. We do not want to penalise these agents because they are not the cause of the change.

In the next section we will show how agents can be correctly identified as the cause of change. The section after investigates identifiers that can be applied to an agent that has caused change, to determine whether they are manipulative or just uncertain.

5.9 Inconsistency Identification

In this section we develop an identifier that in most cases can identify whether or not an agent is being inconsistent between horizons. This identifier is made up of two parts: the revealed preference activity rule and the cost change proxy.

5.9.1 Revealed Preference Activity Rule

Activity rules are used in iterative mechanisms, for example, clock auctions and simultaneous ascending auctions, to limit the bidding strategies of agents [Shoham and Leyton-Brown, 2009, chapter 11]. They help ensure that an agent's bids are "consistent" with earlier bids. The revealed preference activity rule (RPAR) proposed by Ausubel et al. [2006] has been of particular interest. Chapman and Verbic [2015] used a RPAR within a single horizon to prevent agents from manipulating their bids in a electricity demand allocating clock auction.

We derive a rule analogous to the RPAR that works with the notion of consistency between horizons to identify those agents that have changed their convex preferences. Unlike a conventional activity rule, we do not use this to enforce particular behaviour as this would prevent agents from reacting as they gain more information about uncertain events. Instead we use it to monitor the behaviour of agents to determine who is driving change.

Assume that agent i is consistent between consecutive horizons. From definition 2 this means that the time shifted solution from the later horizon satisfies the earlier KKT conditions as given in equations (5.23–5.26). The

KKT conditions for agent i in the earlier horizon must also hold:

$$\nabla f_i(P_i) + \sum_{j=1}^{N_i} \mu_{i,j} \nabla g_{i,j}(P_i) + \lambda = 0 \quad (5.59)$$

$$g_{i,j}(P_i) \leq 0 \quad \forall j \in \{1, \dots, N_i\} \quad (5.60)$$

$$\mu_{i,j} g_{i,j}(P_i) = 0 \quad \forall j \in \{1, \dots, N_i\} \quad (5.61)$$

$$\mu_{i,j} \geq 0 \quad \forall j \in \{1, \dots, N_i\} \quad (5.62)$$

for some multipliers $\mu_{i,j}$. Subtracting (5.59) from (5.23) and multiplying by ΔP_i :

$$\begin{aligned} 0 &= [\nabla f_i(P_i^*) - \nabla f_i(P_i)]^\top \Delta P_i \\ &\quad + \sum_{j=1}^{N_i} [\mu_{i,j}^* \nabla g_{i,j}(P_i^*) - \mu_{i,j} \nabla g_{i,j}(P_i)]^\top \Delta P_i \\ &\quad + \Delta \lambda^\top \Delta P_i \end{aligned} \quad (5.63)$$

The inequality $f(y) - f(x) \geq \nabla f(x)^\top (y - x)$ holds for any convex function f (recall $\nabla f(x)$ is a subgradient when f is non-differentiable). Assuming that f_i and $g_{i,j}$ are convex, applying the above rule and using (5.62) and (5.26):

$$\begin{aligned} 0 &\geq f_i(P_i^*) - f_i(P_i) + f_i(P_i) - f_i(P_i^*) \\ &\quad + \sum_{j=1}^{N_i} [\mu_{i,j}^* g_{i,j}(P_i^*) - \mu_{i,j}^* g_{i,j}(P_i) + \mu_{i,j} g_{i,j}(P_i) - \mu_{i,j} g_{i,j}(P_i^*)] \\ &\quad + \Delta \lambda^\top \Delta P_i \end{aligned} \quad (5.64)$$

By cancelling and using (5.61) and (5.25) this simplifies to:

$$\begin{aligned} 0 &\geq - \sum_{j=1}^{N_i} [\mu_{i,j}^* g_{i,j}(P_i) + \mu_{i,j} g_{i,j}(P_i^*)] \\ &\quad + \Delta \lambda^\top \Delta P_i \end{aligned} \quad (5.65)$$

The inequalities (5.60), (5.62), (5.24) and (5.26) require that the sum term must be non-negative. Finally this produces an inequality which is called the revealed preference activity rule (RPAR):

$$0 \geq \Delta \lambda^\top \Delta P_i \quad (5.66)$$

Any convex agent that is consistent between consecutive horizons will satisfy the RPAR. The converse is not necessarily true: a convex agent that is inconsistent between consecutive horizons can also possibly satisfy the RPAR.

5.9.2 Cost Change Proxy

The cost change proxy is used to calculate how the costs of a consistent convex agent changes between consecutive horizons when the RHS of (5.66) is zero.

Theorem 1 (Cost Change Proxy). *For consecutive horizons, if $\Delta\lambda^\top \Delta P_i = 0$, then the change in costs with respect to the earlier horizon for a consistent convex agent i is $\Delta c_i = \Delta\lambda^\top P_i = \Delta\lambda^\top P_i^*$.*

Proof. As agent i is consistent and $\Delta\lambda^\top \Delta P_i = 0$, equation (5.63) provides:

$$\begin{aligned} 0 &= [\nabla f_i(P_i^*) - \nabla f_i(P_i)]^\top \Delta P_i \\ &\quad + \sum_{j=1}^{N_i} [\mu_{i,j}^* \nabla g_{i,j}(P_i^*) - \mu_{i,j} \nabla g_{i,j}(P_i)]^\top \Delta P_i \end{aligned} \quad (5.67)$$

This equation along with the other consistent KKT conditions (5.23–5.26) and (5.59–5.62) are sufficient to invoke lemma 1 (see appendix B). Because the agent is convex we know that:

$$\nabla f_i(P_i^*)^\top \Delta P_i \geq f_i(P_i^*) - f_i(P_i) \quad (5.68)$$

$$\nabla f_i(P_i)^\top \Delta P_i \leq f_i(P_i^*) - f_i(P_i) \quad (5.69)$$

Lemma 1 has the result that $[\nabla f_i(P_i^*) - \nabla f_i(P_i)]^\top \Delta P_i = 0$. Adding this to the second inequality above, cancelling and comparing it to the first inequality requires:

$$\nabla f_i(P_i^*)^\top \Delta P_i = f_i(P_i^*) - f_i(P_i) \quad (5.70)$$

From the convexity of the constraints we have for all $j \in \{1, \dots, N_i\}$:

$$\nabla g_{i,j}(P_i^*)^\top \Delta P_i \geq g_{i,j}(P_i^*) - g_{i,j}(P_i) \quad (5.71)$$

$$\nabla g_{i,j}(P_i)^\top \Delta P_i \leq g_{i,j}(P_i^*) - g_{i,j}(P_i) \quad (5.72)$$

Lemma 1 gives the result that $[\mu_{i,j}^* \nabla g_{i,j}(P_i^*) - \mu_{i,j} \nabla g_{i,j}(P_i)]^\top \Delta P_i = 0$ and $\mu_{i,j} g_{i,j}(P_i) = \mu_{i,j}^* g_{i,j}(P_i^*) = \mu_{i,j}^* g_{i,j}(P_i) = \mu_{i,j} g_{i,j}(P_i^*) = 0$. Multiplying the first inequality above by $\mu_{i,j}^*$, the second by $\mu_{i,j}$ and applying these equalities results in:

$$\mu_{i,j}^* \nabla g_{i,j}(P_i^*)^\top \Delta P_i = 0 \quad (5.73)$$

Multiplying (5.23) by ΔP_i and substituting results (5.70) and (5.73) produces:

$$f_i(P_i^*) - f_i(P_i) + \lambda^{*\top} \Delta P_i = 0 \quad (5.74)$$

$$\implies f_i(P_i^*) - f_i(P_i) = -\lambda^{*\top} \Delta P_i \quad (5.75)$$

By adding $\lambda^{*\top} P_i^* - \lambda^\top P_i$ to both sides we get:

$$\Delta c_i = \lambda^{*\top} P_i^* - \lambda^\top P_i - \lambda^{*\top} \Delta P_i \quad (5.76)$$

$$\implies \Delta c_i = \Delta \lambda^\top P_i = \Delta \lambda^\top P_i^* \quad (5.77)$$

□

If all agents are convex and consistent then lemma 2 (see appendix B) ensures that $\Delta \lambda^\top \Delta P_i = 0$, and hence theorem 1 applies to all agents. Multiplying the power balancing constraint by the change in price produces $\sum_{i=1}^A \Delta \lambda^\top P_i = 0$. Applying theorem 1 to this gives $\sum_{i=1}^A \Delta c_i = 0$. Therefore, if the expected cost for an agent increases, one or more other agents must have a decrease in expected cost. Those that have a decrease in expected cost stand to gain an advantage by the change.

5.9.3 Inconsistency Identifier

We now bring these results together to form an identifier for inconsistency. Given the powers P_i and P_i^* and prices λ and λ^* for two consecutive horizons, an agent i is labelled as:

- *inconsistent* if $\Delta \lambda^\top \Delta P_i > 0$;
- *notionally inconsistent* if $\Delta \lambda^\top \Delta P_i = 0$ and $\Delta \lambda^\top P_i < 0$; and
- *notionally consistent* otherwise.

The *notional* adjective is used to indicate labels that are not exact; there may exist false positives or false negatives.

Agents labelled as notionally inconsistent may actually be consistent. What is important is that they are not further falsely identified as being manipulative in the next step.

Although agents labelled as notionally consistent may actually be inconsistent, the impact of their inconsistency is not as significant. To understand this consider the two possible cases.

The first case is if $\Delta \lambda^\top \Delta P_i < 0$. The conservation of power requires that $\sum_{i=1}^A \Delta \lambda^\top \Delta P_i = 0$ always holds (see lemma 2 proof). This sum ensures that there must be some other agent j that definitely is inconsistent $\Delta \lambda^\top \Delta P_j > 0$, and therefore has a stronger impact on the system.

The second case is if $\Delta \lambda^\top \Delta P_i = 0$ and $\Delta \lambda^\top P_i \geq 0$. Here either no one is any worse off or the inconsistent agent itself is expected to be worse off. This is not the action of a rational manipulative agent, so this inconsistency is most likely caused by agent uncertainty.

Some inconsistent agents will be missed, but the agents that have the most significant impact on the system will be correctly identified.

5.10 Manipulation Identification

The inconsistency identifier can identify who caused a solution to change, the next step is distinguish whether this change was due to uncertainty or whether it was manipulative. As agents' true preferences are kept private, the only way to distinguish between these two is to apply statistical tests to the data that is public. No test will perform perfectly in every case, and they will likely have to be developed over time to match the type of participants and the form of their uncertainty. Here we propose a simple test, which we will demonstrate to be effective on our greedy agent.

5.10.1 Cost Anticipation Test

The cost anticipation test analyses how agents anticipate their future monetary costs. In many instances an agent manipulates the actions of others by pretending to have high electricity requirements in a future time step. As the horizons progress forwards, the strategic agent eventually changes their estimate to their true consumption requirements. They appear to anticipate higher costs in earlier horizons for particular time steps.

For a given agent i , horizon h and time step $t > h$, the anticipated cost error $\alpha_{h,i,t}$ is defined as the difference between the anticipated monetary cost and the actual monetary cost that eventuates:

$$\alpha_{h,i,t} := \lambda_{h,t}P_{h,i,t} - \lambda_{t,t}P_{t,i,t} \quad (5.78)$$

When this value is positive the agent has overestimated what their costs will be, and when it is negative they have underestimated them. The index $d := t - h$ measures the distance of a time step t from the beginning of a horizon h . It represents how many horizons must advance before the value in the time step is finalised. In our experiments the anticipated cost errors are grouped by d to help identify patterns in the strategies of agents. When analysing up to H horizons the set of anticipated cost errors for a given d is defined as:

$$\mathcal{A}_{i,d} := \{\Delta\alpha_{h,i,t} | h \in \{1, \dots, H\}, d = t - h\} \quad (5.79)$$

The mean of $\mathcal{A}_{i,d}$ should tend towards zero for a particular d if the agent is subject to simple forms of uncertainty. If the agent is being manipulative, then in many cases patterns will show up as non-zero mean values. These statistical results can be used to further distinguish between the inconsistent agents, and label them as either being uncertain or manipulative.

5.11 Identification Experiments

In this section we test the consistency and cost anticipation identifiers on the greedy agent experiments from section 5.7, and new problem instances

where the agent is uncertain instead of being strategic.

5.11.1 Uncertain Problem Instances

The uncertain deferrable agent resamples its power bounds $\hat{P}_{h,i,t}$ after each horizon according to the rule $\hat{P}_{h,i,t} \sim \hat{P}_{1,i,t} \max(0, \mathcal{N}(1, 0.5^2))$. These values are then scaled to ensure that the sum of these new power bounds is equal to $\max(\sum_{t=h}^T \hat{P}_{h-1,i,t}, \check{E}_{h,i})$, where $\check{E}_{h,i}$ is the power remaining to be consumed for the agent in horizon h . This ensures that the problem remains feasible for the uncertain agent. These values are used as the basis for calculating fixed powers for the fixed uncertain agent (using the same approach as in section 5.7).

The uncertain agent experiments are based off the instances generated in section 5.7, with 50 instances per scale parameter value. They were solved using Gurobi.

5.11.2 Inconsistency Results

The first experiments look at how the inconsistency identifier performs when applied to the strategic, truthful uncertain and truthful certain agents. As we are using a numerical method to solve the problem, we introduce tolerances for the inconsistency identifier. These tolerances are biased so that we avoid misreporting inconsistency.

Any consecutive horizons where $\Delta\lambda^\top \Delta P_i > 10^{-4}$ are labelled as inconsistent. They are notionally inconsistent when $0 \leq \Delta\lambda^\top \Delta P_i \leq 10^{-4}$ and $\Delta\lambda^\top P_i < -10^{-4}$. Otherwise the horizons are notionally consistent.

In tables 5.1 and 5.2 we show the percentage of horizons that are labelled in each of these categories for the fixed and deferrable agent experiments respectively. Within each table the results are shown for the greedy and uncertain agent experiments. Further, the results are then split up between either the greedy (Greedy) or uncertain (Uncer) agent, and one of the certain truthful (Truth) deferrable agents within each type of experiment. The results are provided as a percentage identified out of all consecutive horizons in all 50 instances generated for each experiment.

The greedy agents are found to be inconsistent for almost all horizons. The uncertain agents are identified as being inconsistent much less often, 10–19% of the time. Even though they are uncertain in each horizon, this does not always lead to an inconsistent result. The deferrable uncertain agents are marginally less inconsistent than the fixed ones because of their greater flexibility.

The truthful certain agents are often identified as notionally inconsistent, when we know that they are actually consistent. They are identified in this way when they appear to gain a benefit from the actions of the greedy or

Table 5.1: Percent of horizons identified as inconsistent and notionally inconsistent (in brackets) for different agents in fixed power experiments.

Scale	Greedy Exp					Uncertain Exp				
	Greedy		Truth			Uncer		Truth		
0.1	99	(0)	0	(33)		14	(0)	0	(5)	
0.2	99	(0)	0	(36)		14	(0)	0	(3)	
0.5	99	(0)	0	(33)		15	(1)	0	(5)	
1.0	99	(0)	0	(37)		19	(1)	0	(5)	

Table 5.2: Percent of horizons identified as inconsistent and notionally inconsistent (in brackets) for different agents in deferrable load experiments.

Scale	Greedy Exp					Uncertain Exp				
	Greedy		Truth			Uncer		Truth		
0.1	99	(0)	0	(47)		10	(1)	0	(3)	
0.2	99	(0)	0	(43)		10	(1)	0	(2)	
0.5	100	(0)	0	(40)		13	(1)	0	(3)	
1.0	100	(0)	0	(35)		13	(1)	0	(3)	

uncertain agents. There is no strong trend that is observed as the scale of the problem varies.

5.11.3 Anticipated Cost Results

The next step is to calculate the anticipated cost errors for the inconsistent agents. We only calculate anticipated costs for an agent in horizon h if it has been identified that it was inconsistent (not notionally inconsistent or notionally consistent) between horizons h and $h + 1$. The results for the greedy and uncertain agents are shown in figures 5.7 and 5.8 for the two load types respectively.

These results show the anticipated cost errors obtained for each scale value. These are split up into separate means depending on how far in advance the anticipated cost error occurred (see section 5.10.1). Also provided are the 10th and 90th percentiles over all anticipated costs aggregated together for all instances. The mean of the uncertain agent remains around zero, which indicates that on average the agent underestimates its costs in advance just as often as it overestimates them.

The overestimation of costs in earlier horizons is exactly how the greedy agents manipulate the other agents in this setting to gain an advantage. This is clearly seen as a non-zero mean for the greedy agents. This is more pronounced for the deferrable greedy agent because it gains a greater ad-

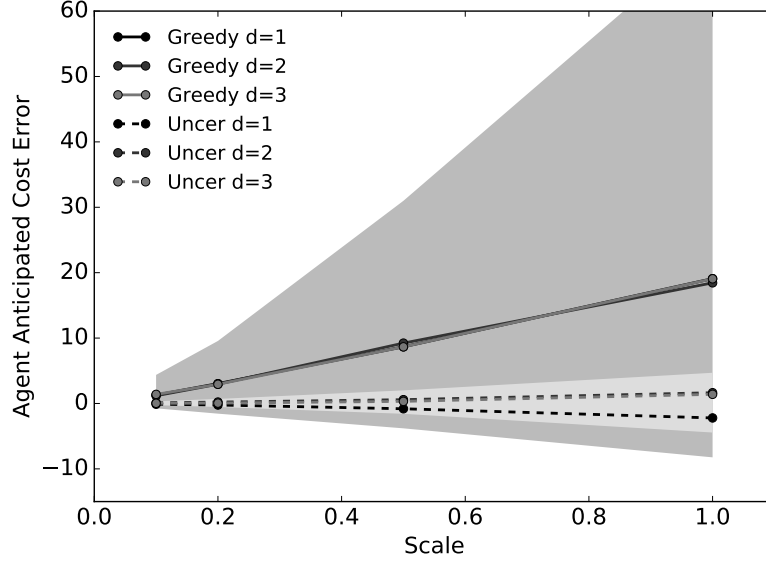


Figure 5.7: Anticipated cost errors for fixed power greedy and uncertain agents. Means for each instance are taken separately based on how far in advance the cost was anticipated (given by d). The shaded sections provide the 10th and 90th percentiles over all instances for the greedy and uncertain agents.

vantage.

A threshold can be used to separate anticipated cost error means of the truthful uncertain agents from the greedy agents. This threshold can increase as the contracted agent power limits (scale parameter) varies. After a number of measurements, if the agent crosses over the threshold then they can either be targeted for further auditing or fined for manipulative behaviour. This will provide a strong deterrent.

It might still be possible for agents to gain some benefit by staying within the threshold we select, but the benefit is reduced. The computational burden is increased for the calculation of strategies that attempt to remain within the threshold. For example, it introduces bilinear terms for the fixed greedy agent, converting it from a MILP to a MINLP.

Using the anticipated cost error identifier with a threshold will limit the type of uncertainty that agents can exhibit. Those with more complex forms of uncertainty may be penalised if they produce results that look similar to manipulative actions. If too many false positives are made in a particular setting then new identifiers can be developed. The great advantage of having privacy-preserving identifiers is that they can always be modified or more can always be added without affecting the operation of the underlying distributed algorithm.




Figure 5.8: Anticipated cost errors for deferrable greedy and uncertain agents. Means for each instance are taken separately based on how far in advance the cost was anticipated (given by d). The shaded sections provide the 10th and 90th percentiles over all instances for the greedy and uncertain agents.

5.12 Related Work

Our development and calculation of greedy agent strategies is related to other work that has looked at equilibria in electricity markets. Hu and Ralph [2006] study equilibria in electricity markets with locational marginal prices where each agent solves a bilevel problem to obtain their strategy. They find sufficient conditions for the existence of pure Nash equilibria. Weber and Overbye [1999] similarly develop a method for finding Nash equilibria for producers and consumers that have linear price functions. Li and Shahidehpour [2005] develop a method for the case where agents have incomplete information about other agents.

Instead of searching for Nash equilibria, Kozanidis et al. [2013] develop optimal bidding strategies for a strategic producer in a single time period market with no network constraints. This is closer to what we developed in section 5.6, but we look at multiple time steps which overlap between horizons, and with different types of agents.

Compared to these works our problem is more complicated in certain areas but simpler in others. For example, instead of a single time horizon, we consider strategies over multiple overlapping horizons given by the receding horizon structure. We also focus more on a prosumer oriented setting where each agent can have a diverse set of preferences and constraints instead of a market dominated by large generator units. However, as a first step we

ignore network constraints.

Mechanism design and game theory have been used in demand response [Mohsenian-Rad et al., 2010, Chen et al., 2011, Tanaka et al., 2012, Akasidis and Chalkiadakis, 2013, Samadi et al., 2012], as well as other network problems including electricity markets and storage adoption [Andrianesis et al., 2010, Tanaka et al., 2012, Vytelingum et al., 2011].

VCG mechanisms have been utilised by Samadi et al. [2012] and Tanaka et al. [2012]. As discussed in section 5.2, VCG quickly becomes intractable for realistically sized problems, and requires agents to fully disclose their preferences. Tanaka et al. [2012] acknowledges these problems and proposes that future work looks at the development of approximate methods.

Chapman et al. [2013] provide a detailed review into the practicality of these existing game-theoretic approaches. They develop four important considerations which any realistic mechanism should take into account, which paraphrased are:

1. The power consumption in houses can take on both discrete and continuous values.
2. Each house has a private state that represents the state of the occupants' goals.
3. Houses have private preferences which are state-based, combinatorial and non-convex.
4. Household behaviour is strategic.

In section 4.10 we showed how the most common source of household discrete decisions can be effectively dealt with, and furthered this argument in section 5.4.3 by discussing other practical methods for dealing with non-convexities. To this list we add a fifth consideration which acknowledges the uncertain nature of the problem:

5. The system is inherently uncertain as a result of household occupant behaviour and weather patterns.

A realistic mechanism should at least work in the presence of uncertainty and ideally reduce its impact on the system performance.

Mhanna et al. [2015] use a scoring rule to charge consumer agents based on both their actual consumption and their deviation from day-ahead allocations. By requiring agents to provide information on their uncertainty, they can reduce the incentive for agents to lie about their requirements over the day ahead. They find it to be “asymptotically” incentive compatible as system size or reported precision increases; however, the approach does not enable agents to share information they gain throughout the day and re-optimize their allocation in an online fashion.

Other works have adopted or suggested a receding horizon approach [Venkat et al., 2008, Chen et al., 2012, Kraning et al., 2014], but Chapman and Verbic [2015] is the only other work to consider the impact of agents manipulating the receding horizon mechanism. They use clock auctions [Ausubel et al., 2006] for allocating loads to agents within horizons. The prices are discounted between horizons to give some flexibility for uncertain agents. This discount factor provides a trade off between allowing agents to recover from uncertain events and preventing manipulation.

Our approach avoids this trade off in the cases where an identifier can be developed to distinguish uncertain actions. We found that this is possible with the anticipated cost error identifier, at least when agents have simple forms of uncertainty. These identifiers can also be applied to more general settings, such as those where agents are both producers and consumers. To the best of our knowledge, our work is the first to formalise and focus on manipulation of receding horizon power balancing, and to provide a practical solution for the identification of this form of manipulation.

5.13 Conclusion and Future work

We formally introduced the notion of receding horizon manipulation in a power systems setting. By developing a strategic agent, we empirically identified in the worst-case setting how much advantage an agent can gain by manipulating the power balancing mechanism. We also showed how this harms the overall social outcome.

We then developed an inconsistency identifier and an anticipated cost error identifier which can be used to monitor the interactions of agents in a non-invasive manner. We successfully used these identifiers to distinguish between an uncertain agent and one that is intentionally manipulating the receding horizon mechanism. Such privacy-preserving identifiers can be used to deter such actions by initiating further auditing or applying fines.

Future work will expand these results to the full network algorithm developed in chapter 4. This requires the consideration of multiple flow variables per time step (real and reactive power) and the deployment of identifiers at each bus in the network.

This is expected to be a simple extension of the existing identifiers. For example, the agent power vectors P used in this chapter can be doubled in length to account for the reactive power and most results should follow with minor modification. The calculations for an agent will be done with the prices for the bus that they connect to. This could enable the development of more powerful identifiers which uses spatial information about the location of inconsistent agents and the changing power flows on the network.

Experiments with more diverse agents and diverse forms of uncertainty will also be valuable, and with them the development of additional identifiers

if the anticipated cost error identifier is insufficient. Experiments with a greedy agent that has varying amounts of information about other agents will also help to further narrow down the bounds on how much advantage manipulation can offer.

Chapter 6

Conclusion

The objective of this thesis was to develop techniques for operating and coordinating prosumer-owned DETs, so as to maximise the benefits to both prosumers and the network. Three key parts of the problem were investigated in chapters 3, 4 and 5, with the techniques developed in these chapters combining to form an overall solution to the problem. This solution increases the value of DETs for both prosumers and networks, and facilitates cooperation between the two.

The results can be viewed as the development of an electricity market that extends down to the distribution level, along with automated methods for prosumer participation and market clearing. The method has a large number of advantages, as discussed throughout this thesis, the most significant of which are:

- A distributed architecture that relies on the computational resources of prosumers, enabling the problem to scale to thousands of participants.
- Prosumer privacy preservation by only requiring expected power profile responses to be exchanged with the mechanism.
- The ability to model and effectively account for the uncertainty that prosumers are exposed to.
- Full automation, with only high-level preferences to be communicated by people.
- Accurate power flow models to ensure network losses and constraints are accounted for.

If deployed on scale, this approach will help the electricity systems to transition to one that is supplied by high penetrations of distributed renewables, and for the further electrification of our societies.

Key Learnings

The three parts of the problem that were investigated in this thesis produced their own key learnings.

Chapter 3 developed a residential EMS to automate the operation of prosumers DETs, and investigated several approaches to managing uncertainty in weather, occupant behaviour and prices. Online receding horizon optimisation using stochastic models of external sources of uncertainty was found to provide clear benefits over reactive control. The technique was found to have short solve times, and to produce solutions with quality close to a clairvoyant solution.

Chapter 4 investigated the use of distributed optimisation for coordinating EMSs over a network. Distributed optimisation based on ADMM was found to be an efficient, privacy-preserving technique for coordinating prosumers in a microgrid or on conventional distribution networks. The ADMM algorithm can work directly with the non-convex AC power flow equations, producing comparable results and performance to convex relaxations. Common sources of prosumer discrete decisions, smart appliance shiftable loads, can be efficiently handled by this distributed algorithm with little degradation to performance.

Chapter 5 formalised receding horizon manipulation, explored the harm it can cause, and developed identifiers for it. It was found that prosumers can manipulate the receding horizon mechanism, but the circumstances where they can gain a significant benefit will be rare. For these rare cases, privacy-preserving identifiers can successfully identify receding horizon manipulation by differentiating between the actions of manipulative and uncertain agents.

Future Research

This thesis has raised a number of important questions that can form the basis of future research.

In chapter 3 we utilised semi-Markov models to generate human-like behaviour, but it remains an open question whether it is possible to fit such models in an online environment. Other stochastic models might provide a better match to human behaviour, which would need to be investigated. The level of abstraction at which the models are developed is an important consideration, as this will impact accuracy and the ability to learn the model. For example, instead of modelling individual occupants, which will be difficult to learn automatically, it might be sufficient to just model the resulting DET interactions.

There are three areas worth pursuing for the distributed algorithms used in chapter 4. The first is whether convergence guarantees can be established for the non-convex power flow equations. It might be possible to obtain a theoretical result that explains why the ADMM algorithm converged in all

our experiments, or if not ADMM then another algorithm.

The second area relates to the rate of convergence. More sophisticated optimisation algorithms can speed up convergence, but distributing such algorithms tends to be more challenging. Future research could develop distributed algorithms specific to power flows that strike the right balance between rate of convergence, subproblem complexity and communications overhead.

The third area relates to the development of anytime distributed algorithms. The ADMM algorithm works on a dual decomposition of the problem, which becomes feasible at the point of convergence. An anytime algorithm keeps the problem feasible as it searches for better solutions, which could be beneficial in a power systems context where a feasible solution might be required at short notice. Future work would develop and investigate the performance of distributed anytime algorithms in this power systems setting.

The approach developed in this thesis improves the efficiency of the electricity system, but it also makes it more complicated. While automation prevents people from being directly exposed to much of this increased complexity, dynamic prices also apply to household loads that are not directly controlled by the EMS. This might be unacceptable for some prosumers, even if their costs average out to something less than what is offered by current fixed electricity tariffs.

This might be just a problem of perception that can be overcome with education, information and framing. But if not, then more consistent and predictable electricity tariffs can be used to charge prosumers, with the dynamic prices just working behind the scenes as an EMS control signal. Future research would investigate such tariffs, and whether or not they can be designed to encourage the right behaviour without exposing prosumers to volatility.

Industry Application

Further developments and testing of the approach in this thesis will be pursued in the upcoming CONSORT project. CONSORT is a collaboration between universities, industry and a network utility, with the goal of demonstrating the use of consumer-owned batteries for solving distribution network problems.

The project will deploy PV-battery systems in homes in order to solve an existing network congestion problem. This will provide a testbed for demonstrating and improving the distributed coordination algorithms in this thesis. The project will also experiment with different prosumer reward structures and analyse the response of residents to these new technologies.

For the distributed algorithms, this project will focus on further optimisations and making them robust enough for deployment, including:

- Reducing the communication burden by developing smart communication strategies and improving convergence.
- Improving the responsiveness of the EMS solvers.
- Handling rare cases where discrete loads cause a problem, or the algorithm fails to converge in time.
- Identification of faults, lost or intermittent communications and the development of robust offline behaviour.
- Development of efficient methods for dealing with unbalanced 3-phase settings.
- Enabling seamless integration into a network with non-participants.
- Developing tariffs and rewards that buffer agents from the dynamic prices.

Summary

To summarise, this thesis has addressed the question of how to utilise and coordinate prosumer-owned DETs for balancing and supporting the network. It has expanded the knowledge in the areas of optimisation under uncertainty, distributed optimisation and mechanism manipulation, as they relate to this topic. It has discovered a range of interesting future research topics, and will continue to be developed and demonstrated in a research collaboration with industry.

Transitioning away from fossil fuels still presents a formidable challenge, but one that is becoming ever more achievable due to advances such as those presented in this thesis. This challenge, along with the rapid pace of technological developments, makes it an exciting time to work in power systems.

Appendix A

AC Power

When dealing with power flows on the network, this thesis makes a quasi-steady-state approximation. That is, we discretise time and assume a steady state in each discrete time step. We do not look at what happens between time steps but just assume there is a stable way of transitioning between steady states. In practice what we produce is target set points for controllers.

A.1 Steady-State AC

In steady-state AC, the voltages and currents in an electrical circuit are assumed to be sinusoidal, with time-independent amplitudes and phase angles, and sharing a common time-independent frequency. An example steady-state voltage is:

$$V(t) = V_p \cos(\omega t + \theta) \quad (\text{A.1})$$

where the angular frequency ω , peak amplitude V_p and phase angle θ are all time independent. These sinusoidal voltages and currents are commonly transformed into a complex number representation, as it makes it easier to apply Ohm's and Kirchhoff's Laws. Representing powers in the complex domain is also useful, but care needs to be taken in calculations and their interpretation, as these tend to be less intuitive. The complex representation of the above voltage in polar and rectangular coordinates is:

$$\mathbf{V}(t) = V_p e^{i(\omega t + \theta)} = V_p \cos(\omega t + \theta) + i V_p \sin(\omega t + \theta) \quad (\text{A.2})$$

The actual voltage is just the real part of this complex number (the same for currents):

$$V(t) = \Re(\mathbf{V}(t)) \quad (\text{A.3})$$

This complex representation is often further simplified by treating the time dependence implicitly for voltages and currents. The resulting quantity is called a phasor. In phasors the root mean square (RMS) amplitude V is typically used in place of the peak amplitude, where $V = \frac{V_p}{\sqrt{2}}$ for a

sinusoid. As we will, see this makes the power calculations simpler. The phasor representation for the voltage in polar coordinates, polar using angle notation and rectangular coordinates is:

$$\mathbf{V} = V e^{i\theta} = V \angle \theta = U + iW \quad (\text{A.4})$$

Operations such as the sum of two phasors and the multiplication of a phasor with a complex constant produces another phasor; however, the multiplication of two phasors in general does not produce another phasor with the same frequency.

Passive two-terminal components (e.g., resistors, capacitors, inductors or any circuit combination of these with two terminals exposed to the outside world) can be represented by a complex constant (not a phasor) called the impedance:

$$\mathbf{Z} = R + iX \quad (\text{A.5})$$

where R is the resistance and X is the reactance. The inverse of the impedance is the admittance \mathbf{Y} , which has components of conductance G and susceptance B :

$$\frac{1}{\mathbf{Z}} = \mathbf{Y} = G + iB \quad (\text{A.6})$$

Ohm's law for impedances in the complex/phasor representation is intuitive. A voltage $\mathbf{V}_{ab} = \mathbf{V}_a - \mathbf{V}_b$ across a passive two-terminal component (see figure A.1) with impedance \mathbf{Z} produces a current $\mathbf{I}_{ab} = -\mathbf{I}_{ba}$ where:

$$\mathbf{V}_{ab} = \mathbf{Z} \mathbf{I}_{ab} \quad (\text{A.7})$$




Figure A.1: Passive two-terminal component with impedance Z .

Kirchhoff's current law (KCL) is also intuitive. If \mathcal{T} represents the set of terminals connected to a node (bus) and \mathbf{I}_j is the current entering (negative if leaving) the node from terminal j , then the sum for currents must be zero:

$$\sum_{j \in \mathcal{T}} \mathbf{I}_j = 0 \quad (\text{A.8})$$

Going back to the time domain for a moment, the instantaneous power is given by the multiplication of voltage and current. Assuming a voltage of $V_p \cos(\omega t + \theta)$ and current of $I_p \cos(\omega t + \phi)$ the instantaneous power is:

$$P(t) = V_p I_p \cos(\omega t + \theta) \cos(\omega t + \phi) \quad (\text{A.9})$$

The instantaneous power can be reformulated as the sum of 90°-out-of-phase sinusoids, one with a DC offset applied to it:

$$P(t) = \frac{V_p I_p}{2} \cos(\theta - \phi)(1 + \cos(2\omega t + 2\phi)) - \frac{V_p I_p}{2} \sin(\theta - \phi) \sin(2\omega t + 2\phi) \quad (\text{A.10})$$

The first term is known as the *instantaneous* real/active power and the second term the *instantaneous* reactive power. The constants real/active power P and reactive power Q characterise these sinusoidal terms:

$$P = \frac{V_p I_p}{2} \cos(\theta - \phi) = VI \cos(\theta - \phi) \quad (\text{A.11})$$

$$Q = \frac{V_p I_p}{2} \sin(\theta - \phi) = VI \sin(\theta - \phi) \quad (\text{A.12})$$

Physically, real power P represents the average power consumed/produced. It also represents the peak amplitude of the element of power that does work. Reactive power Q is the peak amplitude of the element of power that oscillates back and forth without doing any net work. It is the result of energy being periodically stored and released in reactive components like capacitors and inductors. The sign of reactive power modifies the phase of the element it represents by 180°.

Together real and reactive power characterise instantaneous power, and they form a complex constant called the complex power \mathbf{S} :

$$\mathbf{S} = P + iQ \quad (\text{A.13})$$

The magnitude of the complex power $S = |\mathbf{S}|$ is called the apparent power. The complex power consumed/produced by a two-terminal component can be calculated in phasor notation by multiplying the complex conjugate of the current by the voltage across the component:

$$\mathbf{S} = \bar{\mathbf{I}}_{ab} \mathbf{V}_{ab} \quad (\text{A.14})$$

Intuitively, one might expect this operation to produce a complex form of the instantaneous power, but this is not the case. The time dependence of the phasors cancel out when multiplied in this way, producing a complex constant. Note that if peak amplitudes for the current and voltage are used instead of RMS values, then a factor of $\frac{1}{2}$ would need to be applied to the RHS of this equation.

A.1.1 Balanced 3-Phase

A balanced 3-phase system is one where the load and line impedances in each phase are identical. The voltages (currents) are the same in each phase, apart from being 120° out of phase from each other. If the phases are labelled

a , b and c , and the neutral labelled n , the relations between the different line-to-neutral voltages (same for line currents) is given by:

$$\mathbf{V}_{an} = \mathbf{V}_{bn} \angle 120^\circ = \mathbf{V}_{cn} \angle 240^\circ \quad (\text{A.15})$$

The same relation holds for the line-to-line voltages:

$$\mathbf{V}_{ab} = \mathbf{V}_{bc} \angle 120^\circ = \mathbf{V}_{ca} \angle 240^\circ \quad (\text{A.16})$$

which are related to the line-to-neutral voltages by:

$$\mathbf{V}_{ab} = \sqrt{3} \mathbf{V}_{an} \angle 30^\circ \quad (\text{A.17})$$

The total complex power is the sum of power on each phase or just three times the power on one of those phases in a balanced system:

$$\mathbf{S} = 3\mathbf{S}_a = 3\bar{\mathbf{I}}_a \mathbf{V}_{an} \quad (\text{A.18})$$

where \mathbf{V}_{an} is the phase a to neutral voltage and \mathbf{I}_a is the phase a current.

The balanced 3-phase system can be treated as a single phase system by converting Δ connected loads to their Y equivalent:

$$\mathbf{Z}_Y = \frac{\mathbf{Z}_\Delta}{3} \quad (\text{A.19})$$

There are a few options when it comes to dealing with the factor of three that appears in the total complex power equation. Either it can be used in all calculations, the power in just one phase can be used instead, or the factor of three can be absorbed by other terms. It does not really matter what approach is taken as long as it is applied consistently and the physical significance as it ties back to the 3-phase system is remembered. We adopt the latter approach, and have the voltage and current terms soak up the factor:

$$\mathbf{Z} = \mathbf{Z}_Y \quad (\text{A.20})$$

$$\mathbf{V} = \sqrt{3} \mathbf{V}_{an} = \mathbf{V}_{ab} \angle -30^\circ \quad (\text{A.21})$$

$$\mathbf{I} = \sqrt{3} \mathbf{I}_a \quad (\text{A.22})$$

After applying this transformation, Ohm's law, Kirchhoff's laws and complex power are calculated as for any single phase system, and the calculated powers are the total powers for the full 3-phase system.

Appendix B

Proofs

Lemma 1. *Given the convex functions $f : \mathbb{R}^M \mapsto \mathbb{R}$ and $g_j : \mathbb{R}^M \mapsto \mathbb{R}$, if $(\nabla f(x) - \nabla f(y))^\top(x - y) = \sum_{j=1}^N (b_j \nabla g_j(y) - a_j \nabla g_j(x))^\top(x - y)$, $a_j g_j(x) = b_j g_j(y) = 0$ and $g_j(x), g_j(y) \leq 0$ for some $x, y \in \mathbb{R}^M$, $a_j, b_j \geq 0$ and subgradients $\nabla f(\cdot) \in \partial f(\cdot)$, $\nabla g_j(\cdot) \in \partial g_j(\cdot)$, then $b_j g_j(x) = a_j g_j(y) = 0$, $(\nabla f(x) - \nabla f(y))^\top(x - y) = 0$ and $(a_j \nabla g_j(x) - b_j \nabla g_j(y))^\top(x - y) = 0$.*

Proof. We define Γ to be equal to the LHS of the first equation:

$$\Gamma := (\nabla f(x) - \nabla f(y))^\top(x - y) = \sum_{j=1}^N (b_j \nabla g_j(y) - a_j \nabla g_j(x))^\top(x - y) \quad (\text{B.1})$$

Using the convexity of f and g_j we get the following bounds on Γ :

$$\Gamma \geq (f(x) - f(y)) + (f(y) - f(x)) = 0 \quad (\text{B.2})$$

$$\Gamma \leq \sum_{j=1}^N b_j (g_j(x) - g_j(y)) + a_j (g_j(y) - g_j(x)) \quad (\text{B.3})$$

The condition $a_j g_j(x) = b_j g_j(y) = 0$ simplifies the second inequality to:

$$\Gamma \leq \sum_{j=1}^N b_j g_j(x) + a_j g_j(y) \quad (\text{B.4})$$

The conditions $g_j(x), g_j(y) \leq 0$ and $a_j, b_j \geq 0$ imply that $b_j g_j(x) \leq 0$ and $a_j g_j(y) \leq 0$, but (B.2) and (B.4) requires that they sum to a non-negative value. This is only possible if the first result holds:

$$b_j g_j(x) = a_j g_j(y) = 0 \quad (\text{B.5})$$

This means that $\Gamma \leq 0$, which combined with (B.2) implies $\Gamma = 0$, which is the second result. Using the convexity of g_j and applying the first result we get:

$$\begin{aligned} & (b_j \nabla g_j(y) - a_j \nabla g_j(x))^\top (x - y) \\ & \leq b_j(g_j(x) - g_j(y)) + a_j(g_j(y) - g_j(x)) = 0 \end{aligned} \quad (\text{B.6})$$

The definition of Γ and the result $\Gamma = 0$ requires that these non-positive terms sum to zero, which is only possible if the final result holds:

$$(b_j \nabla g_j(y) - a_j \nabla g_j(x))^\top (x - y) = 0 \quad (\text{B.7})$$

□

Lemma 2 (No Violation). *If all agents satisfy the RPAR, then for each agent i : $\Delta \lambda^\top \Delta P_i = 0$.*

Proof. The power conservation rule requires:

$$\sum_{i=1}^A P_i = 0 \quad \sum_{i=1}^A P_i^* = 0 \quad (\text{B.8})$$

Taking the difference and multiplying through by $\Delta \lambda$:

$$\sum_{i=1}^A \Delta \lambda^\top \Delta P_i = 0 \quad (\text{B.9})$$

This is a sum of non-positive values because all agents satisfy the RPAR: $\Delta \lambda^\top \Delta P_i \leq 0$. Therefore for each agent i :

$$\Delta \lambda^\top \Delta P_i = 0 \quad (\text{B.10})$$

□

Lemma 3 (Feasible Powers Adversary). *Assume a convex agent i with fixed power $\hat{P}_i \in \mathbb{R}^T$ that has non-zero power requirements in the first time step of the horizon $\hat{P}_{i,1} \neq 0$. For any untruthful pure strategy that changes the feasible set of powers for the agent at a future time step $t \in \{2, \dots, T\}$ (i.e. from $P_{i,t} \in \{\hat{P}_{i,t}\}$ to $P_{i,t} \in [\underline{P}_i, \bar{P}_i]$ for some $\underline{P}_i \leq \bar{P}_i$), there exists a convex adversary which will increase the costs for the agent relative to their truthful strategy.*

Proof. We build an environment where there is only one other agent j , the adversary, with power $P_j \in \mathbb{R}^T$ and convex cost function:

$$f_j(P_j) = (P_{j,1} + aP_{j,t} + b)^2 \quad (\text{B.11})$$

for some constants a and b . Because there are only two agents, the power conservation constraint requires $P_j = -P_i$. We set the power bounds on the adversary so that they are looser than those for the fixed power agent (i.e. they are never active). The KKT conditions for the adversary at the two time steps of interest 1 and t are (the rest are trivially satisfied):

$$\nabla_1 f_j(P_j) + \lambda_1 = 0 \quad (\text{B.12})$$

$$\nabla_t f_j(P_j) + \lambda_t = 0 \quad (\text{B.13})$$

Evaluating and substituting $-P_i$ for P_j :

$$2(-\hat{P}_{i,1} - aP_{i,t} + b) + \lambda_1 = 0 \quad (\text{B.14})$$

$$2a(-\hat{P}_{i,1} - aP_{i,t} + b) + \lambda_t = 0 \quad (\text{B.15})$$

The value of $P_{i,t}$ can take on one of three values. First assume $\lambda_t = 0$ (bounds from fixed power agent not active). This gives the value $P'_{i,t}$:

$$2a(-\hat{P}_{i,1} - aP'_{i,t} + b) = 0 \quad (\text{B.16})$$

$$\implies P'_{i,t} = (-\hat{P}_{i,1} + b)/a \quad (\text{B.17})$$

Note that for any $\hat{P}_{i,1}$ and $a \neq 0$ we can choose a b such that $P'_{i,t} \in \mathbb{R}$. Since f_j is convex, the value of $P_{i,t}$ is:

$$P_{i,t} = \min(\max(P'_{i,t}, \underline{P}_i), \bar{P}_i) \quad (\text{B.18})$$

The cost for the agent in the first time step is:

$$c_{i,1} = \lambda_1 \hat{P}_1 \quad (\text{B.19})$$

Eliminating λ_1 and $P_{i,t}$:

$$c_{i,1} = 2\hat{P}_{i,1}(\hat{P}_{i,1} + a \min(\max(P'_{i,t}, \underline{P}_i), \bar{P}_i) - b) \quad (\text{B.20})$$

For the case that the agent tells the truth $\bar{P}_i = \underline{P}_i = \hat{P}_{i,t}$:

$$\hat{c}_{i,1} = 2\hat{P}_{i,1}(\hat{P}_{i,1} + a\hat{P}_{i,t} - b) \quad (\text{B.21})$$

The cost change relative to this truthful case is:

$$c_{i,1} - \hat{c}_{i,1} = 2a\hat{P}_{i,1}(\min(\max(P'_{i,t}, \underline{P}_i), \bar{P}_i) - \hat{P}_{i,t}) \quad (\text{B.22})$$

For given $\underline{P}_i \leq \bar{P}_i$ and $\hat{P}_{i,1}$ we can always select an a and $P'_{i,t}$ where this change is greater than zero, unless $\hat{P}_{i,1} = 0$ (which the lemma prohibits) or $\underline{P}_i = \bar{P}_i = \hat{P}_{i,t}$ (the agent is truthful). \square

Nomenclature

Acronyms

AC	Alternating current
AC	Exact AC (line model)
ACT	Australian Capital Territory
ADMM	Alternating direction method of multipliers
AEMO	Australian energy market operator
AM	Additive model
BOM	Bureau of meteorology
DC	Direct current
DC	Linear DC (line model)
DET	Distributed energy technology
DF	Dist-flow (line model)
DLC	Direct load control
DLMP	Dynamic locational marginal prices
DR	Demand response
DSM	Demand-side management
EV	Electric vehicle
EMS	Energy management system
FIT	Feed-in tariff
GAM	Generalised additive model
GLM	Generalised linear model

HVAC	Heating, ventilation and air conditioning
K	Quadratic approximation (line model)
KCL	Kirchhoff's current law
KKT	Karush-Kuhn-Tucker
LDC	Linear DC (line model)
MDP	Markov decision process
MILP	Mixed-integer linear programming
MIP	Mixed-integer programming
MIQP	Mixed-integer quadratic programming
MPC	Model predictive control
NEM	National energy market
NSW	New South Wales
OPF	Optimal power flow
PV	Photovoltaics
QA	Quadratic approximation (line model)
QC	Quadratic constraint (line model)
QP	Quadratic programming
SH	Single horizon (problem)
RD	Relax and decide (method)
RH	Receding horizon
RP	Relax and price (method)
RTP	Real-time pricing
TOU	Time-of-use (pricing)
TRH	Terminating receding horizon (problem)
UR	Unrelaxed (method)
VPP	Virtual power plant

Symbols

A	Number of agents
A	Area
\mathcal{A}	Connection constraint matrix
\mathcal{A}	Anticipated cost error set
α	Anticipated cost error
α	Relax and price penalty parameter
B	Susceptance
c	Cost
d	Device
d	Component
\mathcal{D}	Device set
\mathcal{D}	Component set
E	Energy
\mathcal{E}	Connection set
η	Efficiency
f	Cost function
g	Constraint function
G	Conductance
γ	KKT multiplier
h	Power function
h	Connection constraint function
h	Horizon
I	Global irradiance
I	Current
k	Iteration
k	Random parameter index

K	Known random parameter index set
κ	Specific heat capacity
L	Number of parameters
\mathcal{L}	Lagrangian
λ	KKT multiplier
λ	Dynamic price
Λ	Relax and price penalty matrix
m	Number of sampled scenarios
m	Mass
M	Number of variables
μ	KKT multiplier
n	Bus/node
n	Number of inter-phase terminal constraints
N	Number of constraints
ν	Big-M style values
ω	Angular frequency
p	Probability
P	Real power
ϕ	Current phase angle
ψ	Price/cost coefficient
Ψ	Price/cost coefficient
q	Stochastic programming subproblem function
Q	Reactive power
r	Device parameters
r	Residual
R	Thermal resistance
R	Electrical resistance

ρ	Step/penalty parameter
s	Scenario
s	Scale parameter
S	Scenario set
S	Complex power
t	Time step
T	Number of time steps
T	Horizon length
T	Temperature
\mathcal{T}	Terminal set
τ	Time
θ	Voltage phase angle
V	Voltage
w	Random parameter value
W	Random parameter
X	Reactance
Y	Admittance

List of Publications

Publications

Paul Scott and Sylvie Thiébaux. Distributed Multi-Period optimal power flow for demand response in microgrids. In *ACM e-Energy*, Bangalore, India, jul 2015. “Best Paper Award”.

Paul Scott, Sylvie Thiébaux, Menkes van den Briel, and Pascal Van Hentenryck. Residential demand response under uncertainty. In *International Conference on Principles and Practice of Constraint Programming (CP)*, pages 645–660, Uppsala Sweden, sep. 2013.

Menkes van den Briel, Paul Scott, and Sylvie Thiebaux. Randomized load control: A simple distributed approach for scheduling smart appliances. In *International Joint Conference on Artificial Intelligence (IJCAI)*, Beijing, China, August 2013.

Other Papers

Carleton Coffrin, Dan Gordon, and Paul Scott. NESTA, the NICTA energy system test case archive. In *arXiv, CoRR*, USA, nov 2014.

Paul Scott and Sylvie Thiébaux. Distributed Multi-Period optimal power flow for demand response in microgrids. In *OptMAS*, Istanbul, Turkey, may 2015.

Paul Scott, Sylvie Thiébaux, Menkes van den Briel, and Pascal Van Hentenryck. Residential demand response under uncertainty. In *Green CO-PLAS*, Beijing, China, aug 2013.

Bibliography

- AEMO. Rooftop PV information paper. Technical report, AEMO, 2012.
- AEMO. National Electricity Market (NEM), October 2015a. URL <http://www.aemo.com.au/About-the-Industry/Energy-Markets/National-Electricity-Market>.
- AEMO. Current registration and exemption lists, October 2015b. URL <http://www.aemo.com.au/About-the-Industry/Registration/Current-Registration-and-Exemption-lists>.
- AER. State of the energy market 2014. Technical report, Australian Energy Regulator, 2014.
- AER. Generation capacity and output by fuel source, October 2015. URL <https://www.aer.gov.au/wholesale-markets/wholesale-statistics/generation-capacity-and-output-by-fuel-source>.
AER reference: D11/2285335[V2].
- R. Aguiar and P. Cuervo. Capacitor placement in radial distribution networks through a linear deterministic optimization model. In *Proceedings of the 15th Power Systems Computation Conference (PSCC)*, Liège Belgium, 2005.
- Charilaos Akasiadis and Georgios Chalkiadakis. Agent cooperatives for effective power consumption shifting. In *AAAI Conference on Artificial Intelligence*, 2013. URL <http://www.aaai.org/ocs/index.php/AAAI/AAAI13/paper/view/6193>.
- Joel Andersson. *A General-Purpose Software Framework for Dynamic Optimization*. PhD thesis, Arenberg Doctoral School, KU Leuven, Department of Electrical Engineering (ESAT/SCD) and Optimization in Engineering Center, Kasteelpark Arenberg 10, 3001-Heverlee, Belgium, October 2013.
- Panagiotis Andrianesis, George Liberopoulos, George Kozanidis, and Alex D Papalexopoulos. Recovery mechanisms in a joint energy/reserve day-ahead electricity market with non-convexities. In *Energy Market (EEM)*,

- 2010 7th International Conference on the European, pages 1–6. IEEE, 2010.
- D. Angeli and P.-A. Kountouriotis. A stochastic approach to “dynamic-demand” refrigerator control. *Control Systems Technology, IEEE Transactions on*, 20(3):581–592, May 2012. ISSN 1063-6536. doi: 10.1109/TCST.2011.2141994.
- ASHRAE. *2013 ASHRAE Handbook – Fundamentals*. ASHRAE, 2013. ISBN 1936504456.
- Australian Energy Market Operator (AEMO). URL <http://www.aemo.com.au>.
- Lawrence M. Ausubel, Peter Cramton, and Paul Milgrom. The clock-proxy auction: A practical combinatorial auction design. In *Combinatorial Auctions*, chapter 5, pages 115–138. MIT Press, 2006.
- Xiaoqing Bai, Hua Wei, Katsuki Fujisawa, and Yong Wang. Semidefinite programming for optimal power flow problems. *International Journal of Electrical Power & Energy Systems*, 30(67):383 – 392, 2008. ISSN 0142-0615. doi: <http://dx.doi.org/10.1016/j.ijepes.2007.12.003>. URL <http://www.sciencedirect.com/science/article/pii/S0142061507001378>.
- M.E. Baran and F.F. Wu. Optimal sizing of capacitors placed on a radial distribution system. *Power Delivery, IEEE Transactions on*, 4(1):735–743, 1989. ISSN 0885-8977. doi: 10.1109/61.19266.
- J.F. Benders. Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik*, 4(1):238–252, 1962. ISSN 0029-599X. doi: 10.1007/BF01386316. URL <http://dx.doi.org/10.1007/BF01386316>.
- D.P. Bertsekas and J. N. Tsitsiklis. *Parallel and distributed computation: numerical methods*. Prentice-Hall, 1989.
- Bloomberg New Energy Finance. Renewable energy now cheaper than new fossil fuels in australia, February 2013.
- Subhonmesh Bose, Steven H. Low, Thanchanok Teeraratkul, and Babak Hassibi. Equivalent relaxations of optimal power flow. *CoRR*, abs/1401.1876, 2014. URL <http://arxiv.org/abs/1401.1876>.
- Stephen Boyd, Lin Xiao, Almir Mutapcic, and Jacob Mattingley. Notes on decomposition methods. *Notes for EE364B, Stanford University*, 2007.

- Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, Jan 2011. ISSN 1935-8237. doi: 10.1561/22000000016. URL <http://dx.doi.org/10.1561/22000000016>.
- Bureau of Meteorology (BOM). URL <http://www.bom.gov.au>.
- Anna Cain, Iain MacGill, and Anna Bruce. Assessing the potential impacts of electric vehicles on the electricity distribution network'. In *Solar2010, the 48th AuSES Annual Conference, Australian Solar Energy Society, Canberra, Australia, December*, pages 1–3, 2010.
- S. Caron and G. Kesidis. Incentive-based energy consumption scheduling algorithms for the smart grid. In *Smart Grid Communications (Smart-GridComm), 2010 First IEEE International Conference on*, pages 391–396, oct. 2010. doi: 10.1109/SMARTGRID.2010.5622073.
- Marco Carvalho, Carlos Perez, and Adrian Granados. An adaptive multi-agent-based approach to smart grids control and optimization. *Energy Systems*, 3:61–76, 2012. ISSN 1868-3967. URL <http://dx.doi.org/10.1007/s12667-012-0054-0>. 10.1007/s12667-012-0054-0.
- CAT Projects and ARENA. Investigation the impact of solar variability on grid stability. Technical report, CAT Projects and ARENA, 2015.
- G. Chalkiadakis, V. Robu, R. Kota, A. Rogers, and N. R. Jennings. Cooperatives of distributed energy resources for efficient virtual power plants. In Tumer, Yolum, Sonenberg, and Stone, editors, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent System - Innovative Applications Track (AAMAS 2011)*, Taipei, Taiwan, May 2011.
- A. C. Chapman and G. Verbic. An Iterative On-Line Mechanism for Demand-Side Aggregation. *ArXiv e-prints*, June 2015.
- A. C. Chapman, G. Verbic, and D.J. Hill. A healthy dose of reality for game-theoretic approaches to residential demand response. In *Bulk Power System Dynamics and Control - IX Optimization, Security and Control of the Emerging Power Grid (IREP), 2013 IREP Symposium*, pages 1–13, Aug 2013. doi: 10.1109/IREP.2013.6629395.
- Chen Chen, Shaline Kishore, and Lawrence V. Snyder. An innovative rtp-based residential power scheduling scheme for smart grids. In *ICASSP*, pages 5956–5959, 2011.
- Zhi Chen, Lei Wu, and Yong Fu. Real-time price-based demand response management for residential appliances via stochastic optimization and

- robust optimization. *Smart Grid, IEEE Transactions on*, 3(4):1822–1831, dec. 2012. ISSN 1949-3053. doi: 10.1109/TSG.2012.2212729.
- C. Coffrin, H.L. Hijazi, K. Lehmann, and P. Van Hentenryck. Primal and dual bounds for optimal transmission switching. In *Power Systems Computation Conference (PSCC), 2014*, pages 1–8, Aug 2014. doi: 10.1109/PSCC.2014.7038446.
- C. Coffrin, H.L. Hijazi, and P. Van Hentenryck. The qc relaxation: A theoretical and computational study on optimal power flow. *Power Systems, IEEE Transactions on*, PP(99):1–11, 2015. ISSN 0885-8950. doi: 10.1109/TPWRS.2015.2463111.
- Carleton Coffrin and Pascal Van Hentenryck. A linear-programming approximation of ac power flows. *Inform Journal on Computing*, May 2014. ISSN 1833-9646-6662.
- Yann Collette and Patrick Siarry. *Multiobjective optimization: principles and case studies*. Springer Science & Business Media, 2013.
- C. Concordia and L.H. Fink. Load shedding on an isolated system. *Power Systems, IEEE Transactions on*, 10(3):1467–1472, Aug 1995. ISSN 0885-8950. doi: 10.1109/59.466502.
- Emiliano Dall’Anese, Hao Zhu, and Georgios B. Giannakis. Distributed optimal power flow for smart microgrids. *Smart Grid, IEEE Transactions on*, 4(3):1464–1475, 2013.
- D. Das. Reconfiguration of distribution system using fuzzy multi-objective approach. *International Journal of Electrical Power & Energy Systems*, 28(5):331 – 338, 2006. ISSN 0142-0615. doi: <http://dx.doi.org/10.1016/j.ijepes.2005.08.018>. URL <http://www.sciencedirect.com/science/article/pii/S0142061506000184>.
- J. de Hoog, D.A. Thomas, V. Muenzel, D.C. Jayasuriya, T. Alpcan, M. Brazil, and I. Mareels. Electric vehicle charging and grid constraints: Comparing distributed and centralized approaches. In *Power and Energy Society General Meeting (PES), 2013 IEEE*, pages 1–5, July 2013. doi: 10.1109/PESMG.2013.6672222.
- Jr. Douglas, Jim and Jr. Rachford, H. H. On the numerical solution of heat conduction problems in two and three space variables. *Transactions of the American Mathematical Society*, 82(2):pp. 421–439, 1956. ISSN 00029947. URL <http://www.jstor.org/stable/1993056>.
- Matthias Ehrgott and Xavier Gandibleux. A survey and annotated bibliography of multiobjective combinatorial optimization. *OR-Spektrum*, 22(4):425–460, 2000.

- T. Erseghe. Distributed optimal power flow using admm. *Power Systems, IEEE Transactions on*, 29(5):2370–2380, Sept 2014. ISSN 0885-8950. doi: 10.1109/TPWRS.2014.2306495.
- Boi Faltings. Chapter 20 - distributed constraint programming. In Peter van Beek Francesca Rossi and Toby Walsh, editors, *Handbook of Constraint Programming*, volume 2 of *Foundations of Artificial Intelligence*, pages 699 – 729. Elsevier, 2006. doi: [http://dx.doi.org/10.1016/S1574-6526\(06\)80024-6](http://dx.doi.org/10.1016/S1574-6526(06)80024-6). URL <http://www.sciencedirect.com/science/article/pii/S1574652606800246>.
- Masoud Farhoodnea, Azah Mohamed, Hussain Shareef, and Hadi Zayandehroodi. Power quality impacts of high-penetration electric vehicle stations and renewable energy-based generators on power distribution systems. *Measurement*, 46(8):2423 – 2434, 2013. ISSN 0263-2241. doi: <http://dx.doi.org/10.1016/j.measurement.2013.04.032>. URL <http://www.sciencedirect.com/science/article/pii/S0263224113001437>.
- Masoud Farivar, Christopher R. Clarke, Steven H. Low, and K. Mani Chandy. Inverter var control for distribution systems with renewables. In *SmartGridComm*, pages 457–462. IEEE, 2011. ISBN 978-1-4577-1704-8. URL <http://dblp.uni-trier.de/db/conf/smartgridcomm/smartgridcomm2011.html#FarivarCLC11>.
- E.B. Fisher, R.P. O’Neill, and M.C. Ferris. Optimal transmission switching. *Power Systems, IEEE Transactions on*, 23(3):1346–1355, Aug 2008. ISSN 0885-8950. doi: 10.1109/TPWRS.2008.922256.
- Daniel Gabay and Bertrand Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & Mathematics with Applications*, 2(1):17 – 40, 1976. ISSN 0898-1221. doi: [http://dx.doi.org/10.1016/0898-1221\(76\)90003-1](http://dx.doi.org/10.1016/0898-1221(76)90003-1). URL <http://www.sciencedirect.com/science/article/pii/0898122176900031>.
- Nicolas Gast, Jean-Yves Le Boudec, and Dan-Cristian Tomozei. Impact of demand-response on the efficiency and prices in real-time electricity markets. In *Proceedings of the 5th International Conference on Future Energy Systems*, e-Energy ’14, pages 171–182, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2819-7. doi: 10.1145/2602044.2602052. URL <http://doi.acm.org/10.1145/2602044.2602052>.
- N. Gatsis and G.B. Giannakis. Residential load control: Distributed scheduling and convergence with lost ami messages. *Smart Grid, IEEE Transactions on*, 3(2):770–786, June 2012. ISSN 1949-3053. doi: 10.1109/TSG.2011.2176518.

- Philip Wolfe George B. Dantzig. Decomposition principle for linear programs. *Operations Research*, 8(1):101–111, 1960. ISSN 0030364X, 15265463. URL <http://www.jstor.org/stable/167547>.
- J.D. Glover, M. Sarma, and T. Overbye. *Power System Analysis and Design*. Cengage Learning, 2011. ISBN 9781111425777. URL <https://books.google.com.au/books?id=U77A2C37QesC>.
- S. Grillo, M. Marinelli, S. Massucco, and F. Silvestro. Optimal management strategy of a battery-based storage system to improve renewable energy integration in distribution networks. *Smart Grid, IEEE Transactions on*, PP(99):1, 2012. ISSN 1949-3053. doi: 10.1109/TSG.2012.2189984.
- Ying Guo, Rongxin Li, G. Poulton, and A. Zeman. A simulator for self-adaptive energy demand management. In *Self-Adaptive and Self-Organizing Systems, 2008. SASO '08. Second IEEE International Conference on*, pages 64–73, oct. 2008. doi: 10.1109/SASO.2008.71.
- Gurobi Optimization, Inc. Gurobi optimizer reference manual, 2014. URL <http://www.gurobi.com>.
- T.J. Hastie and R.J. Tibshirani. *Generalized Additive Models*. Monographs on Statistics and Applied Probability Series. Chapman & Hall, CRC Press, 1990. ISBN 9780412343902. URL <http://books.google.com.au/books?id=qa29r1Ze1coC>.
- Hassan L. Hijazi, Carleton Coffrin, and Pascal Van Hentenryck. Convex quadratic relaxations for mixed-integer nonlinear programs in power systems. *NICTA Technical Report*, March 2014. URL http://www.optimization-online.org/DB_HTML/2013/09/4057.html.
- HSL Archive. A collection of fortran codes for large scale scientific computation, 2014. URL <http://www.hsl.rl.ac.uk>.
- Xinmin Hu and Daniel Ralph. Using epecs to model bilevel games in restructured electricity markets with locational prices. Cambridge working papers in economics, Faculty of Economics, University of Cambridge, 2006. URL <http://EconPapers.repec.org/RePEc:cam:camdae:0619>.
- Yann-Chang Huang, Hong-Tzer Yang, and Ching-Lien Huang. Solving the capacitor placement problem in a radial distribution system using tabu search approach. *Power Systems, IEEE Transactions on*, 11(4):1868–1873, Nov 1996. ISSN 0885-8950. doi: 10.1109/59.544656.
- R.A. Jabr. Radial distribution load flow using conic programming. *IEEE Transactions on Power Systems*, 21(3):1458–1459, Aug 2006. ISSN 0885-8950.

- Mireille Jacomino and MinhHoang Le. Robust energy planning in buildings with energy and comfort costs. *4OR*, 10:81–103, 2012. ISSN 1619-4500. doi: 10.1007/s10288-011-0192-6. URL <http://dx.doi.org/10.1007/s10288-011-0192-6>.
- Ali Jadbabaie, Asuman Ozdaglar, and Michael Zargham. A distributed newton method for network optimization. In *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, pages 2736–2741. IEEE, 2009.
- D. Jakovetic, J.M.F. Moura, and J. Xavier. Distributed nesterov-like gradient algorithms. In *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, pages 5459–5464, Dec 2012. doi: 10.1109/CDC.2012.6425938.
- Warwick Johnston, Chris Taeni, and Renate Egan. National survey report of pv power applications in australia. Technical report, SunWiz and APVI, 2015. URL <http://apvi.org.au/wp-content/uploads/2015/09/PV-in-Australia-2014.pdf>.
- W. Karush. Minima of functions of several variables with inequalities as side conditions. Masters Thesis, Department of Mathematics, University of Chicago, 1939.
- B.H. Kim and R. Baldick. A comparison of distributed optimal power flow algorithms. *Power Systems, IEEE Transactions on*, 15(2):599–604, 2000. ISSN 0885-8950. doi: 10.1109/59.867147.
- T.T. Kim and H.V. Poor. Scheduling power consumption with price uncertainty. *Smart Grid, IEEE Transactions on*, 2(3):519–527, sept. 2011. ISSN 1949-3053. doi: 10.1109/TSG.2011.2159279.
- Koen Kok, Martin Scheepers, and René Kamphuis. Intelligence in electricity networks for embedding renewables and distributed generation. In R. R. Negenborn, Z. Lukszo, and J. Hellendoorn, editors, *Intelligent Infrastructures*. Springer, 2009.
- J. Zico Kolter and Matthew J. Johnson. Redd: A public data set for energy disaggregation research. In *Proceedings of the SustKDD workshop on Data Mining Applications in Sustainability*, 2011.
- J. Zico Kolter, Siddharth Batra, and Andrew Y. Ng. Energy disaggregation via discriminative sparse coding. In *24th Annual Conference on Neural Information Processing Systems (NIPS)*, pages 1153–1161, 2010.

- George Kozanidis, Eftychia Kostarelou, Panagiotis Andrianesis, and George Liberopoulos. Mixed integer parametric bilevel programming for optimal strategic bidding of energy producers in day-ahead electricity markets with indivisibilities. *Optimization*, 62(8):1045–1068, 2013. doi: 10.1080/02331934.2013.801473. URL <http://www.tandfonline.com/doi/abs/10.1080/02331934.2013.801473>.
- M. Kraning, E. Chu, J. Lavaei, and S. Boyd. Dynamic network energy management via proximal message passing. *Foundations and Trends in Optimization*, 1(2), 2014.
- 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. URL <http://projecteuclid.org/euclid.bsmsp/1200500249>.
- Lazard. Lazard’s levelised cost of energy analysis – version 9.0. Technical report, Lazard, Nov 2015. URL <https://www.lazard.com/perspective/levelized-cost-of-energy-analysis-90/>.
- Marco Levorato, Andrea Goldsmith, and Urbashi Mitra. Residential demand response using reinforcement learning. In *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, pages 409–414, 2010.
- Tao Li and M. Shahidehpour. Strategic bidding of transmission-constrained gencos with incomplete information. *Power Systems, IEEE Transactions on*, 20(1):437–447, Feb 2005. ISSN 0885-8950. doi: 10.1109/TPWRS.2004.840378.
- S. Magnússon, P. C. Weeraddana, and C. Fischione. A Distributed Approach for the Optimal Power Flow Problem Based on ADMM and Sequential Convex Approximations. *ArXiv e-prints*, January 2014.
- Louis Makowski, Joseph M. Ostroy, and Uzi Segal. Efficient incentive compatible economies are perfectly competitive. *Journal of Economic Theory*, 85(2):169 – 225, 1999. ISSN 0022-0531. doi: <http://dx.doi.org/10.1006/jeth.1998.2494>. URL <http://www.sciencedirect.com/science/article/pii/S0022053198924942>.
- D. H. Martin. The essence of invexity. *Journal of Optimization Theory and Applications*, 47(1):65–76, 1985. ISSN 1573-2878. doi: 10.1007/BF00941316. URL <http://dx.doi.org/10.1007/BF00941316>.
- S. Mhanna, G. Verbic, and A.C. Chapman. A faithful distributed mechanism for demand response aggregation. *Smart Grid, IEEE Transactions on*, PP (99):1–1, 2015. ISSN 1949-3053. doi: 10.1109/TSG.2015.2429152.

- A. Migdalas, P.M. Pardalos, and P. Värbrand. *Multilevel Optimization: Algorithms and Applications*. Nonconvex Optimization and Its Applications. Springer, 1998. ISBN 9780792346937. URL <https://books.google.com.au/books?id=zNAwFXU0dgsC>.
- Pragnesh Jay Modi, Wei-Min Shen, Milind Tambe, and Makoto Yokoo. Ad-opt: asynchronous distributed constraint optimization with quality guarantees. *Artificial Intelligence*, 161(12):149 – 180, 2005. ISSN 0004-3702. doi: <http://dx.doi.org/10.1016/j.artint.2004.09.003>. URL <http://www.sciencedirect.com/science/article/pii/S0004370204001511>. Distributed Constraint Satisfaction.
- A. 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. *Smart Grid, IEEE Transactions on*, 1(3):320 – 331, dec. 2010. ISSN 1949-3053. doi: 10.1109/TSG.2010.2089069.
- A.-H. Mohsenian-Rad and A. Leon-Garcia. Optimal residential load control with price prediction in real-time electricity pricing environments. *Smart Grid, IEEE Transactions on*, 1(2):120 – 133, sept. 2010. ISSN 1949-3053. doi: 10.1109/TSG.2010.2055903.
- J.A. Momoh, R. Adapa, and M.E. El-Hawary. A review of selected optimal power flow literature to 1993. i. nonlinear and quadratic programming approaches. *Power Systems, IEEE Transactions on*, 14(1):96–104, Feb 1999. ISSN 0885-8950. doi: 10.1109/59.744492.
- Manfred Morari and Jay H. Lee. Model predictive control: past, present and future. *Computers & Chemical Engineering*, 23(45): 667 – 682, 1999. ISSN 0098-1354. doi: [http://dx.doi.org/10.1016/S0098-1354\(98\)00301-9](http://dx.doi.org/10.1016/S0098-1354(98)00301-9). URL <http://www.sciencedirect.com/science/article/pii/S0098135498003019>.
- Nest. Nest learning thermostat. URL <https://nest.com/>.
- Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007. ISBN 0521872820.
- Sina Khoshfetrat Pakazad, Anders Hansson, and Martin S Andersen. Distributed interior-point method for loosely coupled problems. *arXiv pre-print arXiv:1312.5440*, 2013.
- Oliver Parson, Siddhartha Ghosh, Mark Weal, and Alex Rogers. Non-intrusive load monitoring using prior models of general appliance types. In *Proceedings of Twenty-Sixth Conference on Artificial Intelligence (AAAI-12)*, Toronto, CA, July 2012.

- M. Pedrasa, T.D. Spooner, and I.F. MacGill. Scheduling of demand side resources using binary particle swarm optimization. *Power Systems, IEEE Transactions on*, 24(3):1173–1181, aug. 2009. ISSN 0885-8950. doi: 10.1109/TPWRS.2009.2021219.
- Q. Peng and S. H. Low. Distributed Algorithm for Optimal Power Flow on a Radial Network. *ArXiv e-prints*, April 2014.
- Dzung Phan and J. Kalagnanam. Some efficient optimization methods for solving the security-constrained optimal power flow problem. *Power Systems, IEEE Transactions on*, 29(2):863–872, March 2014. ISSN 0885-8950. doi: 10.1109/TPWRS.2013.2283175.
- WarrenB. Powell, HugoP. Simao, and Belgacem Bouzaiene-Ayari. Approximate dynamic programming in transportation and logistics: a unified framework. *EURO Journal on Transportation and Logistics*, 1(3):237–284, 2012. ISSN 2192-4376. doi: 10.1007/s13676-012-0015-8. URL <http://dx.doi.org/10.1007/s13676-012-0015-8>.
- S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings. Agent-based control for decentralised demand side management in the smart grid. In Tumer, Yolum, Sonenberg, and Stone, editors, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems - Innovative Applications Track (AAMAS 2011)*, pages 330–331, Taipei, Taiwan, May 2011.
- Sarvapali D Ramchurn, Perukrishnen Vytelingum, Alex Rogers, and Nicholas R Jennings. Putting the ‘smarts’ into the smart grid: a grand challenge for artificial intelligence. *Communications of the ACM*, 55(4):86–97, 2012.
- Prashant P. Reddy and Manuela M. Veloso. Factored models for multiscale decision making in smart grid customers. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI-12)*, 2012.
- Reposit. Reposit power. URL <http://www.repositpower.com/>.
- M. Roozbehani, M.A. Dahleh, and S.K. Mitter. Volatility of power grids under real-time pricing. *Power Systems, IEEE Transactions on*, 27(4):1926–1940, 2012. ISSN 0885-8950. doi: 10.1109/TPWRS.2012.2195037.
- P. Samadi, H. Mohsenian-Rad, R. Schober, and V.W.S. Wong. Advanced demand side management for the future smart grid using mechanism design. *Smart Grid, IEEE Transactions on*, 3(3):1170–1180, Sept 2012. ISSN 1949-3053. doi: 10.1109/TSG.2012.2203341.

- F.C. Schweppe and D.B. Rom. Power system static-state estimation, part ii: Approximate model. *Power apparatus and systems, IEEE transactions on*, PAS-89(1):125–130, 1970.
- A. Shapiro. On Complexity of Multistage Stochastic Programs. *Operations Research Letters*, 34(1):1–8, January 2006.
- A. Shapiro, D. Dentcheva, and A.P. Ruszczyński. *Lectures on Stochastic Programming: Modeling and Theory*. MPS-SIAM series on optimization. Society for Industrial and Applied Mathematics (SIAM, 3600 Market Street, Floor 6, Philadelphia, PA 19104), 2009. ISBN 9780898718751. URL <http://books.google.com.au/books?id=7h7v97qUwPYC>.
- M. Shinwari, A Youssef, and W. Hamouda. A water-filling based scheduling algorithm for the smart grid. *Smart Grid, IEEE Transactions on*, 3(2): 710–719, June 2012. ISSN 1949-3053. doi: 10.1109/TSG.2011.2177103.
- Yoav Shoham and Kevin Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, New York, 2009.
- B. Stott, J. Jardim, and O. Alsac. Dc power flow revisited. *Power Systems, IEEE Transactions on*, 24(3):1290–1300, Aug 2009. ISSN 0885-8950. doi: 10.1109/TPWRS.2009.2021235.
- Arif Syed. The australian energy assessment (aeta) 2013 model update. Technical report, BREE, Canberra, December 2013.
- T. Tanaka, A.Z.W. Cheng, and C. Langbort. A dynamic pivot mechanism with application to real time pricing in power systems. In *American Control Conference (ACC), 2012*, pages 3705–3711, June 2012. doi: 10.1109/ACC.2012.6315337.
- J.A. Taylor and F.S. Hover. Linear relaxations for transmission system planning. *Power Systems, IEEE Transactions on*, 26(4):2533–2538, Nov 2011. ISSN 0885-8950. doi: 10.1109/TPWRS.2011.2145395.
- H. Tischer and G. Verbic. Towards a smart home energy management system - a dynamic programming approach. In *Innovative Smart Grid Technologies Asia (ISGT), 2011 IEEE PES*, pages 1 –7, nov. 2011. doi: 10.1109/ISGT-Asia.2011.6167090.
- UOW. Ripple injection load control systems. Technical Report Technical Note 14, University of Wollongong, 2014.
- Menkes van den Briel, Paul Scott, and Sylvie Thiebaux. Randomized load control: A simple distributed approach for scheduling smart appliances. In *International Joint Conference on Artificial Intelligence (IJCAI)*, Beijing, China, August 2013.

- P. Van Hentenryck and R. Bent. *Online Stochastic Combinatorial Optimization*. The MIT Press, Cambridge, Mass., 2006.
- Stijn Vandael, Nelis Boucké, Tom Holvoet, and Geert Deconinck. Decentralized demand side management of plug-in hybrid vehicles in a smart grid. In *Proceedings of the First International Workshop on Agent Technologies for Energy Systems (ATES 2010)*, pages 67–74, May 2010. URL <https://lirias.kuleuven.be/handle/123456789/272354>.
- A.N. Venkat, I.A. Hiskens, J.B. Rawlings, and S.J. Wright. Distributed MPC strategies with application to power system automatic generation control. *Control Systems Technology, IEEE Transactions on*, 16(6):1192–1206, 2008. ISSN 1063-6536. doi: 10.1109/TCST.2008.919414.
- Vivid Economics. Analysis of electricity consumption, electricity generation emissions intensity and economy-wide emissions. Technical report, Vivid Economics prepared for Climate Change Authority, 2013.
- P. Šulc, S. Backhaus, and M. Chertkov. Optimal Distributed Control of Reactive Power via the Alternating Direction Method of Multipliers. *ArXiv e-prints*, 2014.
- Perukrishnen Vytelingum, Sarvapali D. Ramchurn, Thomas Voice, Alex Rogers, and Nicholas R. Jennings. Trading agents for the smart grid. In van der Hoek, Kaminka, Lesprance, Luck, and Sen, editors, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, Toronto, Canada, May 2010.
- Perukrishnen Vytelingum, Thomas Voice, Sarvapali D. Ramchurn, Alex Rogers, and Nicholas R. Jennings. Theoretical and practical foundations of large-scale agent-based micro-storage in the smart grid. *J. Artif. Intell. Res. (JAIR)*, 42:765–813, 2011.
- Andreas Wächter and Lorenz T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106(1):25–57, 2006. ISSN 0025-5610. doi: 10.1007/s10107-004-0559-y. URL <http://dx.doi.org/10.1007/s10107-004-0559-y>.
- J.D. Weber and T.J. Overbye. A two-level optimization problem for analysis of market bidding strategies. In *Power Engineering Society Summer Meeting, 1999. IEEE*, volume 2, pages 682–687 vol.2, 1999. doi: 10.1109/PESS.1999.787399.
- S. N. Wood. Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society (B)*, 73(1):3–36, 2011.

- Fredrik Ygge and Hans Akkermans. Power load management as a computational market. In *In Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS)*, pages 393–400. AAAI Press, 1996.
- Zhe Yu, Linda McLaughlin, Liyan Jia, Mary C. Murphy-Hoye, AnnaBelle Pratt, and Lang Tong. Modeling and stochastic control for home energy management. In *2012 Power and Energy Society general meeting*, July 2012.
- Michael Zargham, Alejandro Ribeiro, Asuman Ozdaglar, and Ali Jadbabaie. Accelerated dual descent for network optimization. In *American Control Conference (ACC), 2011*, pages 2663–2668. IEEE, 2011.
- B. Zhang, A. Y. S. Lam, A. Dominguez-Garcia, and D. Tse. Optimal Distributed Voltage Regulation in Power Distribution Networks. *ArXiv e-prints*, April 2012.