Managing Sustainable Demand-side Infrastructure
for Power System Ancillary Services

by

Simon Christopher Parkinson
B.Sc.E., University of Saskatchewan, 2008

A Thesis Submitted in Partial Fulfillment of the
Requirements for the Degree of

Master of Applied Science

in the Department of Mechanical Engineering

© Simon Christopher Parkinson, 2011
University of Victoria

All rights reserved. This thesis may not be reproduced in whole or in part, by photocopying or other means, without the permission of the author.

Managing Sustainable Demand-side Infrastructure
for Power System Ancillary Services

by

Simon Christopher Parkinson
B.Sc.E., University of Saskatchewan, 2008

Supervisory Committee

Dr. Curran Crawford, Co-supervisor
(Department of Mechanical Engineering)

Dr. Ned Djilali, Co-supervisor
(Department of Mechanical Engineering)

Supervisory Committee

Dr. Curran Crawford, Co-supervisor
(Department of Mechanical Engineering)

Dr. Ned Djilali, Co-supervisor
(Department of Mechanical Engineering)

ABSTRACT

Widespread access to renewable electricity is seen as a viable method to mitigate carbon emissions, although problematic are the issues associated with the integration of the generation systems within current power system configurations. Wind power plants are the primary large-scale renewable generation technology applied globally, but display considerable short-term supply variability that is difficult to predict. Power systems are currently not designed to operate under these conditions, and results in the need to increase operating reserve in order to guarantee stability. Often, operating conventional generation as reserve is both technically and economically inefficient, which can overshadow positive benefits associated with renewable energy exploitation. The purpose of this thesis is to introduce and assess an alternative method of enhancing power system operations through the control of electric loads. In particular, this thesis focuses on managing highly-distributed sustainable demand-side infrastructure, in the form of heat pumps, electric vehicles, and electrolyzers, as dispatchable short-term energy balancing resources. The main contribution of the thesis is an optimal control strategy capable of simultaneously balancing grid- and demand-side objectives. The viability of the load control strategy is assessed through model-based simulations that explicitly track end-use functionality of responsive devices within a power systems analysis typically implemented to observe the effects of integrated wind energy systems. Results indicate that there is great potential for the proposed method to displace the need for increased reserve capacity in systems considering a high penetration of wind energy, thereby allowing conventional generation to operate more efficiently and avoid the need for possible capacity expansions.

Contents

Supervisory Committee ii

Abstract iii

Table of Contents iv

List of Tables vi

List of Figures vii

List of Abbreviations and Symbols ix

Acknowledgements xiv

1 Introduction 1

1.1 Motivation 1

1.2 Main Contributions 3

1.3 Thesis Outline 4

2 Barriers to Low Carbon Electric Power Systems 5

2.1 Review of Electric Power Generating Sources 5

2.2 Power System Operations 8

2.2.1 Effects of Variable Renewable Generation 10

2.2.2 Distributed balancing resources 14

3 Engaging Communities in Low Carbon Power System Operations 15

3.1 Designing Load Management Networks 16

3.1.1 A Community-scale Approach 17

3.2 Deferrable Loads 18

3.2.1 Thermostatically controlled loads 19

3.2.2 Energy-constrained storage as a deferrable load 21
3.2.3 Hysteresis control of deferrable loads 24
3.3 Managing Large Populations 26
3.3.1 Controlling Population Dynamics Online 26
3.3.2 Optimal dispatch of multiple responsive load populations . . . 31
3.3.3 Target design for a self-regulating load 34

4 Computational Modelling Framework **37**
4.1 Demand-side: Load Models 38
4.1.1 Heat Pump Model 38
4.1.2 EV charging 43
4.1.3 Electrolyzers 44
4.2 Supply-side: Power System Model 46
4.2.1 Wind Energy System Model 48
4.2.2 Grid Model 49
4.2.3 Selecting the Simulation Time-step 52
4.3 Control Performance Metrics 53
4.3.1 Supply-side 53
4.3.2 Demand-side 55

5 Model Results **57**
5.1 Preliminary Model Results 57
5.1.1 Population Dynamics 60
5.1.2 Network Bus Power Dynamics 64
5.1.3 Effect on Conventional Generation Scheduling 66
5.2 Sensitivity Study 70
5.2.1 Heat Pump Integration 71
5.2.2 EV Integration 75
5.2.3 Electrolyzer Integration 78
5.3 Effects of Wind System Capacity 80
5.3.1 Application in Remote Power Systems 81
5.4 Effects of Communication Network Quality 82

6 Conclusions **87**

Bibliography **92**

List of Tables

Table 2.1 Cost and Emissions Data for Large-scale Generating Technologies. 6

Table 4.1 Parameters implemented in the heat pump model 42
Table 4.2 Parameters implemented in the EV charging model 44
Table 4.3 Parameters implemented in the electrolyzer model 46
Table 4.4 Parameters implemented in the wind energy system model . . . 50

Table 5.1 Performance of the load resource. 64

List of Figures

Figure 2.1 Power system management decision making framework 9
Figure 2.2 Part-load efficiency curves for two identical generators participating as both regulation and load following resources. 12

Figure 3.1 Closed-loop control strategy implemented. 29
Figure 3.2 Flow-chart of the multi-stage optimization problem 36

Figure 4.1 The one-dimensional coupled ETP model of a building, with thermostatically controlled heat pump currently in the inactive-state ($n = 0$). 39
Figure 4.2 Example of the ETP model under the displayed outdoor temperature. 42
Figure 4.3 Example of the EV load model. 45
Figure 4.4 Example of the electrolyzer load model. 47
Figure 4.5 Simplified DC electric power system model. 51
Figure 4.6 Unresponsive load and wind power time-series data. 52

Figure 5.1 Load trajectories associated with the preliminary simulation. . 59
Figure 5.2 Close-up view of the virtual generator responses over the period 01:00 to 02:00. 61
Figure 5.3 LC power distributions from 01:00 to 02:00. 62
Figure 5.4 Left: power gradient experienced by the generator. Right: cumulative demand for energy from the generator. 65
Figure 5.5 Regulation component that must be accommodated by the conventional generation 66
Figure 5.6 Required hourly online load-following and regulation reserves. . 67
Figure 5.7 Operating characteristics of the dispatchable capacity. 69
Figure 5.8 Effect of heat pump population size and selection of M. 72

Figure 5.9 Effect of heat pump population size and outdoor temperature regime. 74
Figure 5.10 Effect of EV population size and selection of M. 75
Figure 5.11 Effect of EV population size and number of day-charging events. 76
Figure 5.12 Effect of EV population size and width of deadband. 77
Figure 5.13 Effect of electrolyzer population size and selection of M. 78
Figure 5.14 Effect of electrolyzer population size and measurement noise. . 79
Figure 5.15 Effect of wind energy system capacity and selection of M. . . . 80
Figure 5.16 Effect of grid flexibility. 82
Figure 5.17 The effects of packet loss on the total demand and response error. 84
Figure 5.18 The effects of packet loss on the load control performance. . . . 85

List of Abbreviations and Symbols

Abbreviations

ACE	Area Control Area
AGC	Automatic Generation Control
BAN	Building Automation Network
CCGT	Combined Cycle Gas Turbine
COP	Coefficient of performance
ED	Economic Dispatch
EV	Electric Vehicle
HVAC	Heating Ventilation and Air Conditioning
LA	Load Aggregator
LC	Load Community
LP	Linear Program
OCGT	Open Cycle Gas Turbine
PCH	Programmable Communicating Hysteresis-controlller
PLR	Packet Loss Ratio
TCL	Thermostatically Controlled Load
UC	Unit Commitment
VGM	Virtual Generator Model

Symbols

α	Polynomial coefficient
β	Blade pitch angle
$\mathbf{\Gamma}$	Input transition matrix
$\mathbf{\Omega}$	State transition matrix
ΔP^*	Target deviation from uncontrolled responsive load trajectory
δ	Width of hysteresis control deadband space
$\dot{C}_{lf}$	Online load-following reserve ramp capacity
$\dot{C}_{reg}$	Online regulation reserve ramp capacity
$\dot{L}$	Load-following ramp capacity contract
$\dot{P}_G$	Generator ramp-rate
$\dot{P}_{max}$	Maximum ramp-rate
$\dot{q}_d$	Design heating rate of the heat pump
$\dot{q}_h$	Heating rate of heat pump
$\dot{q}_f$	Heating power of heat pump fan
$\dot{q}_{loss}$	Rate of heat transfer from TCL to environment
$\dot{q}_{op}$	Rate of heat transfer from heat pump to indoor air
$\dot{R}$	Regulation ramp capacity contract
ϵ	End-use state comparison
$\epsilon_{\pm}$	Operational state-transition boundary
η	Energy conversion efficiency
η_e	Efficiency of grid-interfacing power electronics
η_g	Efficiency of wind turbine gearbox

γ	Over-sizing factor
λ	Tip-speed ratio
$\mathbf{Z}$	Set of power state vectors
$\mathbf{z}$	Power state vector
$\mathcal{F}$	Flexibility of responsive load population
$\mathcal{R}$	Deadband discretization resolution
Φ	Capacity factor of responsive load population
ϕ_0	Inactive power density distribution function
ϕ_1	Active power density distribution function
ρ	Air denisty
θ_a	Indoor air temperature
θ_d	Design outdoor temperature
θ_m	Temperature of interior building thermal mass
θ_o	Outdoor temperature
θ_s	User's desired indoor air temperature
C_a	Indoor air thermal mass
C_m	Indoor building thermal mass
C_p	Turbine coefficient of performance
C_{lf}	Online load following reserve capacity
C_{reg}	Online regulation reserve capacity
D	Rotor diameter
E	Storage charging trajectory
e	Measurement or model error

E_c	Storage user's desired amount of energy
E_D	Cumulative energy dispatched from load
E_G	Cumulative demand for energy from the conventional generation
E_s	Set-point energy level of storage
g	Generator type index
i	Population index
k	Discrete time sampling index
L	Load-following capacity contract
M	Number of previous terms considered in self-regulating target trajectory
m	Deadband location index
m_s	Deadband index associated with the set-point
n	Operational state of load
N_i	Number of loads in population
N_p	Number of responsive load types
p	Responsive load type index
P^*	Target trajectory for individual responsive load population
P_C	Curtailed wind power
P_G	Conventional generation power output
P_h	Required electric power of heat pump to provide design heating rate
P_L	Total aggregate load
P_m	Mechanical power imparted to turbine blades
P_o	Uncontrolled responsive load trajectory
P_R	Self-regulating target responsive load trajectory

P_r	Rated power of individual responsive load
P_T	Target trajectory for multiple responsive load populations
P_U	Aggregate unresponsive load
P_W	Wind power
P_{base}	Base component of generator load
P_{cap}	Total installed capacity in responsive load population
P_{lf}	Load following component of generator load
P_{max}	Maximum power output
P_{min}	Minimum power output
P_{reg}	Regulation component of generator load
P_{rt}	Rated power of the wind turbine
R	Regulation capacity contract
R_{ao}	Building envelope heat transfer resistance
R_{ma}	Heat transfer resistance between building mass and air
T	Discrete sampling period
T_s	Storage user's desired charge-time
u	Set-point modulation
u_w	wind speed
u_{ci}	Cut-in wind speed
u_{co}	Cut-out wind speed
u_r	Rated wind speed
w	Generation cost function

ACKNOWLEDGEMENTS

I would like to thank Curran Crawford and Ned Djilali for providing me with the insight to pursue research in this field. Your patience and intellectual support has provided me with research skills beyond my own expectations. I would also like to thank Dan Wang for the support in developing these ideas into tangible applications, as well as the other researchers at the Institute for Integrated Energy Systems who provided their ideas and comments. Finally, I would like to thank Tom Pedersen for providing invaluable feedback on this work. The financial support from the Pacific Institute for Climate Solutions, NSERC Hydrogen Strategic Network, and the University Victoria is gratefully acknowledged.

Chapter 1

Introduction

1.1 Motivation

The carbon-intensive industrialization pursued over the past century has resulted in a rapid increase in the amount of atmospheric carbon dioxide [1]. The higher levels of atmospheric concentration in conjunction with the spectroscopic properties of this molecule result in increased atmospheric absorption of radiation. Resultant temperature differentials compromise the thermodynamic stability of earth's climate system, and foreshadow a future of extreme weather events. Furthermore, as a large portion of the carbon dioxide in the atmosphere makes its way to the ocean via the carbon cycle, producing carbonic acid, increased atmospheric concentration and ocean acidification are in concert. Acidic oceans usurp key building-block ions used by integral species in the marine food-web, reducing the ability of such species to survive in the long-term [2]. Finally, health problems linked to the inhalation of carbon emissions and concomitant pollutants associated with fossil-fuels combustion are becoming a serious health concern in many locations in which emission intensities are at their highest [3]. As the stability of these global-scale phenomena (climate, food-webs, and well-being) allowed humanity to develop societies to the current levels of sophistication, urgent action is required to reduce carbon emissions globally.

Integral to economic development is access to energy, where energy dense fossil fuels have traditionally provided a cheap, plentiful, and reliable supply framework. The widespread use of fossil fuel as an energy feedstock is the primary contributor to global anthropogenic carbon emissions, and accounts for approximately 60% of total global greenhouse gas emissions [4]. Further burdening humanity's explicit reliance

on fossil fuels are problems related to supply sustainability [5]. Fossil fuel resources are becoming harder to obtain each year, where the amount of energy being expended in the extraction processes is approaching similar levels to that achieved at end-use. In order to maintain standards of living currently commensurate with economic or social success, a rapid shift away from fossil fuel exploitation is needed, which will require a total transformation within the energy supply sector.

To meet these ambitious goals, it is expected that nearly 6 trillion dollars will be invested into restructuring the entire energy supply chain [6]. Development of replacement clean energy technologies and application of energy efficiency measures are currently pursued with varying degrees of rigour internationally, and have resulted in a plethora of individual technologies and system configurations that claim to reduce carbon emissions. On the supply-side, low-carbon generation technologies, such as wind, wave, solar, and tidal power plants are in the process of demonstrating their suitability to displace conventional fossil fuelled generation. Due to resource intermittency and unpredictability, challenges remain in the development of approaches focused on fostering the grid-integration of these technologies into the locations where they are needed, whilst ensuring supply reliability [7].

On the demand-side, electric vehicles (EV), and electric heating systems, are seen as available tools to mitigate emissions within transportation and building sectors, with the development of alternative fuels, such as hydrogen, set to further aid in the displacement of fossil fuel for both stationary and transport-based applications. The key underlying theme between these demand-side clean energy technologies is the possible use of electric power as an energy source, although widespread use will result in a large increase in electric power demand. Introduction of renewables on the supply-side is viewed as the most promising route for meeting these increasing requirements [8, 9, 10], but as mentioned, the power system is currently not designed to operate with a large portion of such generation. In fact, the clean-energy service technologies identified on the demand side have in common the requirement of electric energy and load cycles with a degree of flexibility, opening the door for possible demand side management to alleviate increased variability on the supply side. The main motivation for this thesis is then the need to develop a resilient systems integration framework that looks to tap into this flexible demand, such that optimal pathways for minimizing emissions can be achieved through the simultaneous integration of renewable generation technologies and sustainable demand-side infrastructure.

1.2 Main Contributions

The pursuit of innovative control strategies and system configurations specifically aimed at overcoming traditional operational barriers is on the horizon via coupling of power system operation with real-time networking capabilities. There is considerable ground to cover before mass deployment of this so-called *smart grid* technology occurs. Many of the issues that need to be resolved require new design tools and control models to capture a physical system-level view over relevant time-scales, as well as concurrent management strategies that display robustness and reliability during implementation [11]. To this regard, this thesis offers the following main contributions:

1. A community-scale approach to demand-side management is proposed, in which an energy-based business is to be in charge of developing a transaction between the operator of a low-carbon electric power system and community-specific electric load infrastructure. Demand-side clean energy technology, in the form of air-source heat pumps, EVs, and electrolyzers are targeted for demand response recruitment, in order to leverage the use of these devices within the low-carbon energy systems in which they are needed.

2. A novel control strategy is introduced that accurately manages the aggregate demand trajectory of a large population of electric loads to provide ancillary services to the power system. The key attribute of the proposed method is its capability to simultaneously balance grid-side objectives with those typically expected on the demand-side by the customers. The method is able to seamlessly integrate with typical power system control by aggregating responsive load population dynamics into models equivalent to conventional dispatchable generation.

3. A computational modelling framework is developed that allows for an integrated bottom-up analysis of the proposed load management strategy. The model is capable of explicitly tracking end-use functionality of responsive devices within a power systems analysis typically implemented to observe the operational effects of integrated renewable energy systems.

1.3 Thesis Outline

The thesis proceeds as follows. Chapter 2 provides an overview of power system operation. In particular, the chapter explicitly outlines problems associated with low-carbon system configurations, in which intermittent renewable generation is to be employed at large-scales. Chapter 3 introduces the prospect of involving demand-side infrastructure in the operation of low-carbon energy systems. A resilient framework for engaging communities is developed, as well as a novel design of a comfort-constrained load management strategy aiming to control large populations of heat pumps, electric vehicles, and electrolyzers for power system ancillary services. Chapter 4 provides the integrated computational modelling framework that connects the supply- and demand-side dynamics within the power systems analysis. Chapter 5 provides an overview and discussion of the results obtained from utilizing the computational model within different scenarios. Chapter 6 concludes the thesis, providing recommendations and further avenues for future research.

Chapter 2

Barriers to Low Carbon Electric Power Systems

The carbon emissions associated with an electric power system are related to the mix of resources that provide capacity in conjunction with the efficiency at which the grid's infrastructure delivers energy to end-users. Understanding tradeoffs between the two is therefore important in the development of optimal design and operational policies intended to achieve both economic and environmental objectives of system stakeholders. The purpose of this chapter is to review the current barriers to implementing low carbon power systems, and to provide insight into possible benefits associated with involving elements from the demand-side in operations.

2.1 Review of Electric Power Generating Sources

There is a myriad of generation technologies currently available to power system operators, and in turn applied within a variety of different configurations. In order to meet the demand for energy, typical utilities employ services from centralized, large-scale generating plants. Conventional large-scale generation technologies include: hydroelectric plants that capture the potential energy of falling water; nuclear plants that capture the energy liberated through nuclear fission; coal power plants that implement the energy liberated in coal combustion; and combined-cycle gas turbines (CCGT) and open-cycle gas turbines (OCGT) that use the energy liberated through combustion of gas (typically natural gas).

Alternative generating technologies that rely on naturally reoccurring sources of

energy (or renewable energy) are also under vigorous development. These include: solar, wind, tidal, and wave power generating plants. Of these alternative technologies, wind power generation has emerged as the primary scaleable technology that can be applied economically. As such, wind energy development is predicted to continue to dominate the renewable energy technology market, in particular due to its use in large-scale developments [7].

Historically, minimizing short-term project costs has prevailed in terms of generation planning and operation, and allowed fossil-fuel based generating technologies to dominate, although hydroelectric generation has played a large role within suitable geographic regions. As new fossil fuel reserves are becoming harder to locate each year, fossil fuel costs have become an increasing worry, that now combined with the long-term project costs associable to the environmental degradation resulting from high carbon feedstocks, have driven the desire to pursue alternative low carbon generation technologies. Table (2.1) displays costs and emissions data for typical large-scale generation technologies, and has been adapted from [12]. Each individual technol-

Table 2.1: Cost and Emissions Data for Large-scale Generating Technologies.

Technology	Fuel Cost [\$/MWh]	Variable O&M [\$/MWh]	Construction [$\$10^6$/MW]	Emissions [tCO_2/MWh]
Hydroelectric	1.13	0.02	1.55	0.009
Nuclear	6.20	0.07	1.70	0.012
Coal	13.70	0.70	1.10	0.980
CCGT	37.00	5.00	0.55	0.450
OCGT	41.00	4.50	0.46	0.650
Wind	0.00	0.17	1.30	0.015

ogy displays beneficial attributes that depend on whether the project's objectives are focused on short-term costs, or long-term sustainability.

In terms of emissions, both hydroelectric and nuclear developments appear superior, even when compared in relation to wind generated energy. However, simply ranking the environmental qualities of these generation technologies with respect to their CO_2 emission rate is in fact misleading. Development of large-scale hydroelectric plants can result in the need to displace conventional land-use in surrounding areas, which can mean flooding surrounding farmland or altering the surrounding aquatic / terrestrial ecosystems. Availability of undeveloped locations suitable for hydroelectric generating stations has also diminished, meaning future developments may need to

focus on locations associated with higher risk. For nuclear power projects, the resultant waste from the nuclear fission process pursued within a nuclear generating plant emits dangerous levels of radiation, and while only small amounts of waste are produced, disposal requires assumptions that the disposal location will be uninhabited until the waste decays (thousands of years). Further are added problems associated with operational security either nuclear or hydro plants introduce in the geographic regions in which they operate. Natural disasters or human errors can result in unstable operation, leading to the collapse of a hydroelectric dam, or a nuclear reactor meltdown. Either event can contaminate large geographic areas for generations. Finally, the costs associated with building a nuclear power plant or hydroelectric dam are immense, due to the complicated system required to contain and control the high energy processes. Likewise, such plants take a considerable amount of time to build, and therefore may not be able to provide the rapid decarbonization the electric power system requires to allow atmospheric CO_2 levels to stabilize and eventually diminish to non-threatening levels.

Conversely, widespread integration of renewable energy technologies is seen as a method to displace carbon emissions in the short-term [7], although many of these technologies also face social barriers to implementation. Wind for instance can be an unappealing development from the viewpoint of nearby communities, mainly due to perceived decay in landscape aesthetics and threat posed to surrounding wildlife. Nonetheless, prospects of community ownership and ecological siting considerations have been shown to ease these tensions [13]. In hopes of meeting the majority of their jurisdictions' emission reduction targets within the electric power sector, utilities in Denmark have demonstrated that wind energy systems can provide effective generation capacity, where approximately 20% of their total electricity demand is met with wind alone [14]. A problem faced by these utilities, and others hoping to pursue similar aggressive wind integration strategies, is the loss of controllability on the supply-side, as the fuel source—wind energy—is non-dispatchable, and therefore new methods of maintaining grid stability are required.

Many European nations are granted the added benefit of a meshed transmission network, providing needed support from neighbouring jurisdiction's power systems. In alternative locations in which similar connections are lacking, wind integration will require access to grid-reinforcements in order to accommodate supply variability [15]. This underlines one of the main barriers preventing widespread access to renewable electricity, which are the challenges associated with handling unpredictable short-

term supply variability within conventional approaches to power system operation and control [16].

2.2 Power System Operations

In order to understand how intermittent renewable generation effects power system operations, a brief review of power system control is now provided. Power system operators rely on specific control strategies aimed at particular planning horizons to achieve a continuous balance between supply and demand. Capacity expansion planning on the order of months to years is typically considered through forecasting demand trends in conjunction with equipment maintenance and decommissioning [17]. Actual individual generator schedules become available on the order of days to hours, as many types of generating sources take a significant amount of time to bring online (coal-fired plants are an example), and therefore must be allocated well in advance. Fittingly, this type of scheduling is typically referred to as unit commitment (UC), and relies on the use of load forecasts to develop an optimal operating schedule. Beyond UC, economic dispatch (ED) is pursued on the order of minutes to hours, again relying on updated load forecasts and generator states to directly control the output of available generating resources [17].

The common theme among these approaches is that each approach applies an optimization problem to arrive at a decision. As all methods must rely on forecasts to establish control variables, each contains inherent uncertainty, which can be somewhat accommodated through application of suitable stochastic problem formulations. Errors are inevitable though, which would result in an energy imbalance between supply and demand, leading to operational instability. Real-time control strategies commonly denoted as regulation must be applied in order to maintain a continuous balance through the adjustment of suitable generators' output.

Two separate regulation strategies are applied to achieve real-time balancing: primary and secondary control[1]. Primary control is intended to balance local frequency deviations through decentralized control of individual generating units. Speed governors are implemented on units participating in primary control (typically hydro or thermal resources), and managed in response to local frequency measurements using

[1]There is also tertiary control that focuses on system security and involves planning reserves meant to come on-line in the event of a contingency, for instance if a voltage collapse is foreseen to occur, or a generator fails.

the droop characteristics of the particular unit [18]. Secondary control, often denoted as automatic generation control (AGC), regulates frequency on the order of control areas or balancing authorities that encompass many different generating units. AGC uses measurements of tie-line flows between adjacent areas to generate a system-wide metric known as the area control error (ACE). Units participating in AGC react to the ACE signal based on their individual availability, and attempt to drive the ACE signal to zero through adjustment of their output [18].

In either case, regulation reserves must be pre-allocated, which in turn effects the other components of power system management. Figure (2.1) depicts a conceptual representation of the different decision making time-scales apparent in power system operation, where each individual control action is encompassed within the grid structure and demand profile. Feedbacks between the decision making intervals exist, and result in regulation requirements propagating through the system.

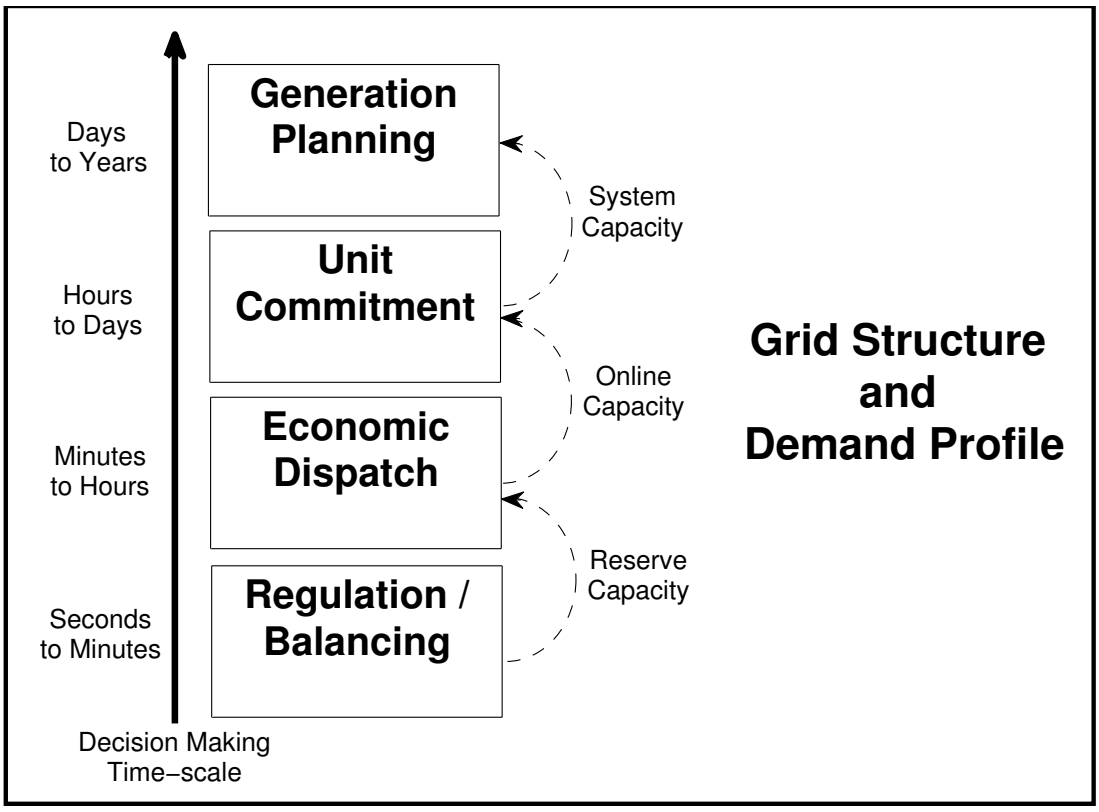

Figure 2.1: Power system management decision making framework

2.2.1 Effects of Variable Renewable Generation

The typical power system control architecture has provided a robust framework for maintaing system performance, with problems now becoming apparent in systems intended to integrate large amounts of intermittent renewable generation, in particular wind [16]. Wind displays a high level of temporal variability, and is difficult to accurately predict over short time-scales (minute-to-minute) [19]. Indication of the scales over which wind power production displays variability is given by Apt [20]. Measured and compiled performance data from a population of geographically dispersed wind turbines is used by Apt to show that wind power production follows the Kolmorgorov spectrum over a time-scale ranging from 30 seconds to over 2 days. As a result, Apt was able to conclude that although wind variability can be somewhat accommodated by the variability in demand, a considerable amount of wind power fluctuates over different time-scales.

To understand what this means in terms of power system operations, integration studies usually focus on how operation of the wind plant affects each component of power system control. This process typically entails model-based power system simulations in order to quantify the incremental reserves arising from enhanced variability in the power system [21]. Using these incremental requirements in conjunction with the corresponding cost per capacity seen by these services in conventional markets, the integration-based costs of the project can be estimated, and combined to compute the levelized cost of the wind energy.

Many wind integration studies find that reserves relevant for regulation are the most affected, which can introduce problems for the power system operator, who must provide greater access to such reserve resources [22]. In particular, operating conventional generation as regulation reserve (typically hydro or thermal units) can be unappealing from the viewpoint of the system operator for the following reasons:

- In direct contrast to the other control components, no attempt is made to implement a least-cost dispatch for regulation [23].
- Both primary and secondary control require participating units to rapidly ramp output in response to corresponding frequency measurements. Excessive ramping prematurely degrades units participating in regulation, reducing performance and reliability over time [24].
- Ramping thermal units increases operating emissions [25].

- Regulation requirements (up and down) can result in the need to operate the unit partly-loaded, and therefore reduce its efficiency, though this can be minimized through the use of multiple smaller units [26].

Beyond these initially identified problems associated with regulation are the negative effects felt throughout the other components of power system planning and control. Within a particular management area, in order to ensure reliability under all operating scenarios, generators suitable for regulation are contracted to provide regulation capacity, as well as a corresponding online ability to ramp output up and down. Similarly, control actions made during ED further require online ability of suitable generators to provide load-following contracts for capacity and ramping capability. Figure (2.2) is introduced to illustrate these ideas from the viewpoint of individual generators. Part-load efficiency curves associated with two identical generating units within a hypothetical control area are depicted. These curves are meant to represent either hydro or thermal units, as both suffer a reduction in efficiency when subject to part-load, achieving a minimum efficiency η_{min} at its minimal load-level P_{min} and a maximum efficiency η_{max} near the capacity-loading P_{max}, with the efficiency between varying non-linearly [27, 26]. The regulation requirements are denoted as capacity R, with the load-following requirements denoted as capacity L. In this case, either unit will provide both load following and regulation services, with generator **(A)** contracted to provide twice the regulation reserve requirement as compared to generator **(B)**. The regulation reserve capacity covers both up- and down-ramping, and therefore, as indicated in Fig.(2.2), results in the need to buffer the maximum and minimum achievable load following levels (denoted L_{max} and L_{min}) by the same amount. The size and ramping capability of these generating units are thus constrained by:

$$\begin{aligned} P_{max}^{(g)} - P_{min}^{(g)} &\geq L^{(g)} + 2R^{(g)} \\ \dot{P}_{max}^{(g)} &\geq \mathbf{Max}\left[\dot{L}^{(g)}\,,\, \dot{R}^{(g)}\right] \end{aligned} \tag{2.1}$$

where g is the generator type index, $\dot{L}$ denoting load-following ramping requirements, and $\dot{R}$ denoting regulation ramping requirements. The first constraint in (2.1) makes sure the unit is capable of meeting the contracted reserve requirements, with the second guaranteeing that the generator can respond at a rate required to qualify as both a load-following and regulation resource.

As generator **(A)** and **(B)** are identical, with the only difference in this case being that generator **(A)** has twice the regulation reserve obligations, generator **(B)** is able

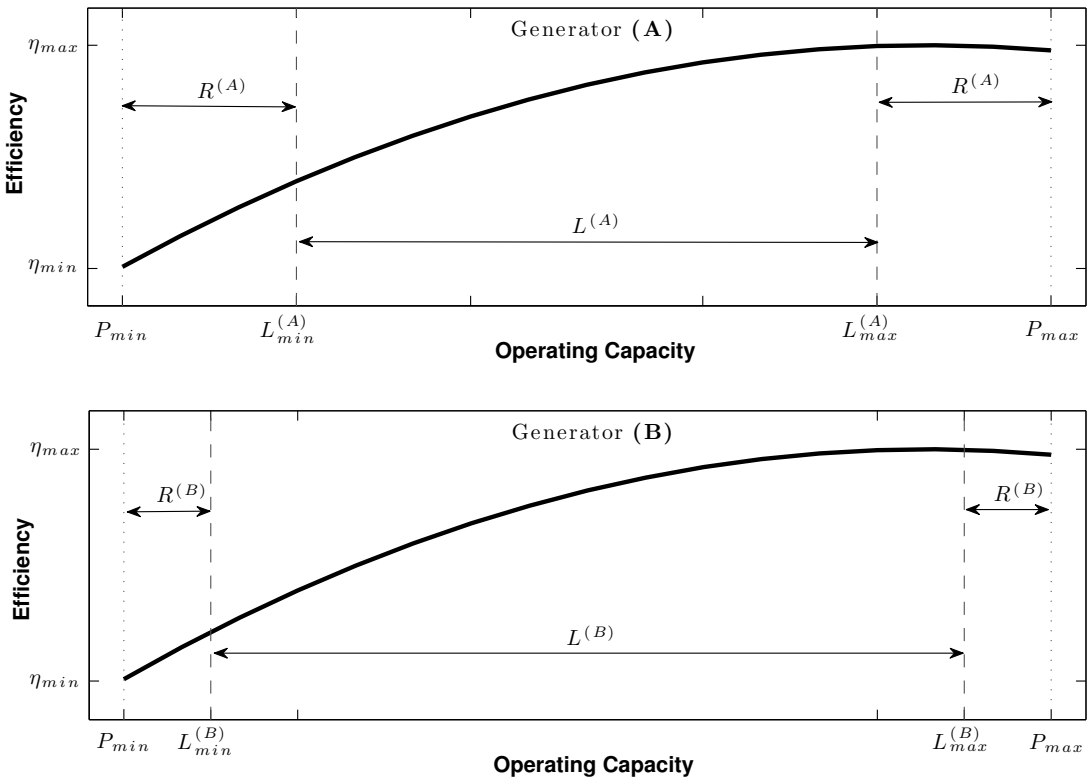

Figure 2.2: Part-load efficiency curves for two identical generators participating as both regulation and load following resources.

to commit more of its capacity to ED reserve. In fact, the first constraint in (2.1) states that for each unit of regulation reserve capacity reduced, we are able to increase the generators' load-following flexibility, or ability to participate in ED, two-fold. As there is no attempt made to implement a least-cost dispatch for regulation, it can be assumed that energy allocated for load-following, procured through ED, is allocated under more socially-optimal conditions [23]. Therefore, it can be expected that in this scenario, from the viewpoint of system operator, generator **(B)** is utilized at a greater social benefit in comparison to generator **(A)**. Finally, generator **(A)** would need to increase its total capacity by $2R^{(B)}$ in order to maintain the same operating range as generator **(B)**, which is a condition that would need to be treated within the power system's capacity planning stages.

This brief example demonstrates the unequivocal connection between regulation and the other components of power system planning and control, and can be summarized into the following main effects:

- Regulation reserve capacity results in less capacity available for ED, and therefore less capability to allocate installed capacity optimally.

- Regulation reserve capacity results in the need to allocate more online capacity during UC. This online capacity is rarely used to its full potential and is therefore under-utilized. The extra capacity that does not make it onto the grid represents missed opportunity, which would otherwise embody enhanced flexibility to participate in ED, and thereby provide additional opportunities to supply cleaner and cheaper electricity to the grid.

- Regulation reserve capacity affects the size of the required generation capacity, and therefore plays an important role in planning generation developments for a given management area. Operating generation as regulation reserve can reduce its operating efficiency, increase maintenance requirements, and overall prematurely degrade the unit. These aspects must also be considered during the capacity planning stages.

As continued large-scale wind integration will require greater access to regulation-based resources, these negative operational effects can overshadow positive benefits commonly associated with wind power generation.

2.2.2 Distributed balancing resources

Alternative methods for mitigating enhanced regulation requirements arising from integrated wind energy systems have been proposed, with the underlying thematic conclusions outlining co-located regulation reserves as an attractive option [28, 29]. This can involve capacity expansion, usually through the introduction of an alternative fast-acting energy resource classified as energy storage. While these types of devices may in fact provide a wealth of flexibility to the power system, storage accrues efficiency penalties during roundtrip energy conversion, and in most situations, current storage systems lack economic viability at required capacities. Nonetheless, similarly themed resources will be required for continued large-scale wind integration, which in terms of power system control, results in increased participation of more spatially distributed resources in the provision of the regulation ancillary service, and a need to down-scale conventional AGC dispatch to smaller localized balancing authorities [30].

As will be introduced in the following chapter, spatially-distributed regulation provides an effective foreground for the participation of communities located nearby wind project developments to aid in integration. An important aspect of regulation is that its requirements are typically zero-mean (energy requirements balance to zero), although both regulation-down and regulation-up events are considered dispatched capacity within typical ancillary service markets [31]. This is an extremely important aspect of regulation, as when it is possible to tap into flexible demand of nearby communities to offset the need for enhanced localized regulation reserves, there is little or no change to the functionality of targeted loads. The resultant effects would in turn propagate throughout the other power system control components, and help to provide a powerful platform to efficiently integrate renewable generation onto the grid.

Chapter 3

Engaging Communities in Low Carbon Power System Operations

Traditionally, energy imbalances within the power system are overcome by changing the supply to meet the demand, with operational problems arising with the integration of large amounts of intermittent renewable energy resources that lack controllability. In fact, the goal of balancing energy within the system could be met with a change in demand rather than supply. Such a system configuration requires management of the electric devices the power system services, and could be achieved through the application of suitable communication infrastructure to manipulate the operation of these loads. Based on the suggestions in the previous chapter, if these loads could be targeted to displace the need for online regulation reserve, they could provide considerable operational benefits to power systems looking to mitigate carbon emissions, by aiding integration of renewable energy. Implementation of this system configuration will incur increased capital costs, and further look to manipulate resources that are not the property of the utility. An operational strategy capable of balancing power system and demand-side objectives is therefore needed to ensure success. In response to this requirement, this chapter introduces a resilient approach based on the direct control of community-based electric loads in the efficient operation of low-carbon electric power systems through intelligent management of their aggregate demand profile.

3.1 Designing Load Management Networks

Before diving into a rigorous design framework, it is important to understand the goals and constraints of all involved stakeholders, as well as the available resources on hand. Many commercial operations have already focused on applying information technology to aid in the implementation of energy efficiency measures at the building-level. For example, a typical building automation network (BAN) is capable of linking with a web-based browser that can be accessed over the internet, allowing a building manager to benchmark energy usage and identify possible avenues to increase efficiency [32]. It would in principle be relatively straightforward for the building manager to share device-level information with an outside entity, by granting access to the BAN. The outside entity could then act as an intermediary, and develop and implement real-time operational strategies that are beneficial from the viewpoint of the power system operator, for instance to balance system energy through modulating the building's energy demand to match the available supply.

It will be the job of the intermediary to develop the transaction between the building managers and the power system operator. To be successful the intermediary should display the following four attributes:

1. **Non-intrusive**: On the customer-side, the intermediary must develop control strategies that are comfort-constrained in the sense that service levels commensurate with customer satisfaction are observed at all times. On the grid-side, the intermediary's strategy should seamlessly integrate with conventional system management.

2. **Secure**: On the customer-side, the intermediary must ensure that an acceptable level of privacy is maintained when accessing the BAN. On the grid-side, the intermediary must ensure operation does not compromise system security.

3. **Profitable**: The intermediary must develop contracts between the stakeholders so that the project provides economic benefit to both the customer and load serving entity.

4. **Ecological**: As an energy-based development, the project needs to represent a net reduction in emissions over its classical counterparts.

Both building managers and power system operators are more likely to participate in load management programs of this nature, and readily provide long-term access to

their resources.

3.1.1 A Community-scale Approach

While a power system operator could in fact encompass intermediaries within their operations, it is argued that separating an intermediary at the community-level is indeed a superior approach. This is because electric loads are inherently the property of community members, who collectively represent a primary stakeholder. Through treatment of load management as a community-based resource, the intermediary could in fact operate as a community business, offering the demand resource to the power system operator. This in turn could allow communities to diversify their business profiles, and further provide a solid foundation for future deployment of local renewable energy resources, as a community-based entity could act to provide the integration and operational support. The task of maintaining network privacy is also simplified within a community-based approach, as securing device-level information within a local network is much easier than a geographically distributed large-scale network that encompasses many thousands, perhaps even millions of customers.

This idea has further implications in terms of resource adequacy, as the issue of population size and the accompanying problems associated with communication congestion can be alleviated through operating a local system that contains fewer customers. Control actions could then focus on shorter time-frame events, at a high-level of accuracy, promoting resource exploitation to its full potential and thus providing greater overall benefits. Community members are also much more likely to trust a community-based operation working to maintain their goals, rather than an outside entity who may act, or be perceived to act, to primarily serve the interests of outside stakeholders. Localized goals could include efficient management of local energy resources that, for instance, could entail maintaining local power quality to maintain the performance / lifespan of the community's appliances, or alternatively, managing integration of a community-owned renewable generation facility, in hopes of maximizing its performance.

The goal of the community-based entity is then to recruit loads from their community, with the resultant group of recruited loads hereafter referred to as the load community (LC). The proposed community-based load management policy is therefore seen to consist of the following steps. Prospective building managers, who at the residential building-level are more specifically community residents, enable the trans-

action between power system operator and their targeted devices by joining their community-specific energy management network, with community-members retaining the right to decide whether or not to participate. By joining the network, the community member's device-level information becomes available to the intermediary, who can then change the operational-state of suitable end-use devices based on their specific strategy. In turn, the intermediary rewards the community member based on some pre-defined contractual obligations, for instance based on the amount of time a particular community member's devices participate in the program during a particular billing period. The intermediary is then rewarded by the power system operator based on the performance quality of the resource the intermediary offers their system.

The intermediary therefore represents a load aggregator (LA), where in this thesis interaction with the power system operator will be achieved through the development of models that describe LC population dynamics in a format equivalent to conventional generation, or a virtual generator model (VGM). The LA will then use these models to pursue strategies that result in loads encompassed within the LC cooperating to achieve grid-side benefits, which in this work is the efficient management of embedded renewable generation technologies. The VGM format is selected as it will integrate well with traditional deregulated ancillary service market structures [33], but it will be the focus of this work to develop a strategy that seamlessly integrates with power system operations. Control of the load management network is based solely on local inputs, and therefore communication with the power system operator, or interference with their operations, would be deemed unnecessary.

3.2 Deferrable Loads

Having determined a framework for the operation of the load management network, the next challenge becomes targeting potential demand-side infrastructure, as only specific types of loads will be suitable for recruitment. The primary objective of any load control strategy is increased demand flexibility, and therefore requires participatory load-types to display similar device-level attributes. The ability to operate under a flexible schedule results in demand that can be deferred to more opportune times, such as when an excess of renewable energy is available, and therefore devices displaying similar attributes can be defined as deferrable loads. The purpose of this section is to introduce and assess different types of deferrable load potentially avail-

able at the community-level from the viewpoint of controllability and sustainability, and then to select certain end-use appliances for further study that display the best characteristics for recruitment.

3.2.1 Thermostatically controlled loads

Upon inspection of the load-types currently dominating the market, thermostatically controlled loads (HVAC systems, water heaters, refrigerators, etc.) occupy a large portion of the residential and commercial demand profile [34]. A thermostatically controlled load (TCL) is driven by a thermostat that controls a machine capable of converting electrical energy to heat energy, which can then be used to condition a given space to a user-provided temperature set point. Other than providing the comfort settings, daily interaction with the end-user is scarce, though the device continues to operate post-interaction. If TCLs are indeed deferrable, they are therefore an ideal candidate for demand response recruitment, as changes to the operational schedule could be achieved without disrupting the end-user's experience. This is in direct contrast to other load-types, where customer interaction during load operation is frequent or continuous, and therefore any change to the operational schedule would directly impact the end-use functionality of the unit.

TCLs are typically designed to operate at a certain power rating P_r, meaning the trajectory of the load demand associated with a single TCL is given by:

$$P(t) = n(t)P_r \tag{3.1}$$

where n is the operational state of the device at time t. For thermally-based loads, operation is triggered as a result of heating / cooling requirements. If the heat transfer rate from the conditioned thermal system $\dot{q}_{loss}$ is in fact less than the rated heating power of the device, the TCL can be operated intermittently:

$$\eta|P_r| > |\dot{q}_{loss}| \tag{3.2}$$

where η is the energy conversion efficiency associated with the unit. Through communication with the thermostats controlling the operation of grid-connected TCLs conforming to (3.2), operation can be deferred for brief periods of time, coasting through prescribed intervals by relying on the thermal energy store inherent in the system. Constrained by thermal comfort bounds, the use of the device within a certain

time-frame is inevitable, but seemingly small deferral intervals could in fact provide considerable flexibility when considered in aggregate across a large population, and thereby provide grid-support under enhanced supply-side variability.

Heat pumps

In the case of water heaters and refrigerators, the costs associated with investing in the required device-level communication infrastructure may out-weigh resulting grid-side benefits, as individual units do not consume a large amount of energy, and therefore a very large population, or a large amount of communication hardware would be required. In comparison, HVAC systems consume a large amount of energy per device, and would therefore require a smaller population or less communication hardware. While this may be the case, the majority of communities in colder climates heat their buildings with systems that utilize fossil fuels as an energy source, mainly natural gas [35], and such HVAC systems would be unavailable for recruitment. This reality further places HVAC systems as one of the main sources of carbon emissions within the building sector [36]. The natural gas these furnaces commonly employ is often touted as a clean energy substitute for conventional fossil fuels, due to its lower emission rate. If there is a desire to move towards less carbon-intensive energy resources, promoting a shift towards natural gas will only provide temporary mitigation, and does not represent a long-term sustainable solution.

In temperate climates (moderate changes between summer and winter), such as those found on the Pacific Coast of Canada, air-source heat pumps represent the ideal replacement for fossil-fuel based furnaces. Operating in these geographic regions, air-source heat pumps can achieve high seasonal efficiencies, with coefficients of performance (COP) exceeding 2.5 year-round [37]. The COP of a heat pump represents how efficiently it converts the electrical energy needed to run the device compressors, relative to the heat that is transferred to the indoor air. A value of 2.5 means that for a heat pump that uses 1 kW of electrical energy, 2.5 kW of heat energy is provided. All combustion-based units have a COP less than 1. Heat pumps are also reversible, meaning that they can act as both heating and cooling units, and further can be combined into multi-function units that simultaneously provide both water heating and indoor air conditioning services [38].

If in fact heat pumps are integrated into low carbon power systems, many of the emissions associated with thermal space-conditioning can be mitigated [8]. While

such a large-scale integration would result in a shift of the heating service to the electrical load and the need for widespread retrofits, if appropriately controlled, these devices could support the implementation of a sustainable power system by effectively providing "virtual" capacity expansion. This would therefore provide possibilities for retrofit financing. Retrofits would further provide an ideal time to equip these units with the required communication hardware, which in turn should cost relatively little in comparison to the costs of installing the heat pump itself.

An effective load management strategy provides a pathway to transform conventional heat pumps into extremely efficient combined heat and power units (both thermal and electrical services are attained during operation), that further enable the mitigation of emissions within both the electric power and building sectors, simultaneously replacing fossil fuelled furnaces and aiding grid-integration of intermittent renewables.

Some previous work has recognized the potential of TCLs as a regulation resource. In [39], the potential for water heaters as a regulation resource is investigated, but the accuracy of the attained response as compared to the costs of equipping these devices with the required communication hardware would not support widespread use. In [40], Callaway proposes an open-loop control strategy aimed at controlling large populations of air-conditioning units. Callaway implements system identification techniques to develop aggregate load models capable of following the output of a large wind farm. The method relies on identification of linear models of aggregate power dynamics using large populations of first-order load models under quasi-steady-state conditions, and therefore may face problems under true dynamic conditions unless suitable adaptive identification methods are developed. The main goal of this thesis is to introduce a method, based on the implementation of higher-order load models under dynamic conditions that can adapt the aggregate load models online so that large populations can be accurately controlled.

3.2.2 Energy-constrained storage as a deferrable load

Many types of electric load convert grid energy to an alternative resource that can be utilized later for similar energy-based activities. This includes charging of device batteries, or the production of an alternative fuel. The operations of such storage devices share similar objectives of achieving a desired level of stored energy at the end of the operation cycle. This corresponds to a certain amount of charging energy

E_c to be accumulated over a user's desired charging-horizon T_s. Denoting the charging trajectory over the charging-horizon as E, and assuming the rated power and efficiency of the unit is constant, the end-use constraint on the device is then given by:

$$E(T_s) = \int_0^{T_s} \eta P_r n(t) dt \geq E_c \; ; \; E(0) = 0 \tag{3.3}$$

where η is the one-way energy conversion efficiency associated with the storage unit. Considering (3.3), it is clear that if the unit can charge in less time then the user requires, namely:

$$\eta P_r > \frac{E_c}{T_s} \tag{3.4}$$

then through changes to the operational-state trajectory, the grid-loading trajectory, given again by (3.1), can be deferred over the user-set charging-horizon, without disrupting end-use functionality of the unit.

Electric vehicles

The transportation sector is responsible for approximately 13% of global carbon emissions, of which personal transportation vehicles are the single largest component [4]. In light of this, many automobile manufacturers have recognized this statistic and have pursued vehicle designs that are capable of all-electric operation, with some modern designs capable of achieving distances of up to 160 km on a single charge [41]. As vehicles that travel less than this distance per day are responsible for a considerable portion of total travel requirements [42], EVs may be able to power a substantial portion of daily travel with electricity, and could thus displace a large fraction of gasoline use.

Nonetheless, if these units are not integrated into low carbon power systems, any mitigatory benefits decrease rapidly [9]. However, if these units can be controlled to provide support to the integration of renewable energy sources, their use can be leveraged within these system types, thereby maintaining the desired demand-side environmental qualities. As some of the EVs available on the market come equipped with the technology to communicate important state attributes [41], the prospect of communicating with these loads to modulate demand could be achieved with little added investment in end-use hardware.

The idea of using EVs to provide grid-side benefits is not new. Introduced in

[43] and further in [44], the effectiveness of EVs in providing ancillary services is discussed, where vehicle-to-grid energy services are demonstrated. The problem with a vehicle-to-grid approach is that it relies on both the storage and extraction of energy from the battery to attain the grid-side service, and therefore energy dispatched is subject to the roundtrip efficiency losses associated with the battery. Furthermore, enhanced cycling of the batteries degrades performance over time, and therefore may be unappealing from the viewpoint of the end-user, who owns and operates the unit for alternative reasons. It is therefore the focus of this work to obtain grid-side services from EV fleets without extracting energy from the battery.

Other methods for involving EVs in the provision of ancillary services have been proposed. In [45], a sequential algorithm aimed at engaging energy constrained loads (EVs or TCLs) is investigated, while in [46] a decentralized method for utilizing large populations of EVs to fill diurnal valleys in demand (a load-following service) is proposed. These initial studies have demonstrated that in order to attain ancillary services from EVs, vehicle-to-grid is not a necessity, and that simply deferring on/off status of the charging process can in fact be quite beneficial.

Electrolytic hydrogen production infrastructure

Hydrogen fuel cells can be employed to extract energy from an electrochemical reaction involving hydrogen and oxygen. As the product of this reaction is mainly water, operational carbon emissions are negligible. Fuel cells thus provide an ideal pathway to fully replacing the traditional internal combustion engine for both stationary and mobile applications. While fuel cells are enticing from an environmental perspective, how the hydrogen feedstock is produced is of utmost importance [10]. It is expected that locations that will first use fuel cell technologies will coincide with those in which electric power and water supply represent the most reliable demand-side services. As electricity can be used to split water molecules into its constituent gases via the process of electrolysis, hydrogen can be produced at the location of end-use (hydrogen fuelling stations) through utilization of distributed grid-connected electrolyzers.

Hydrogen lacks density, which causes issues related to storage and transport. By producing this fuel locally using electrolyzers distributed throughout communities, and operating these based on an on-demand production schedule (producing daily quotas to be used immediately), problems associated with the transportation and storage costs associated with centralized production schemes and seasonal storage

can be avoided.

Again, use of electrolyzers results in traditionally fossil fuel-based services shifting to the electrical load. Expected environmental benefits are only to be observed if increased capacity requirements are met with low carbon generating sources. A similar argument to that made for EVs can be made for leveraging electrolyzers to support the integration of renewable energy resources in order to attain the desired low carbon electricity supply. Numerous projects have focused on the use of regenerative fuel cell systems that encompass both the storage and regeneration aspects of operating side-by-side fuel cells and electrolyzers [47]. The roundtrip efficiency losses achieved in regenerative systems are problematic, which are indeed much greater than those associated with conventional energy storage. Also, the rapid-cycling of electrolyzers that accompany their use as a short term energy buffer has been shown to effect their long-term performance [47]. Therefore, it is the focus of this work to obtain the desired grid-side services from the electrolyzers without ever extracting energy from units, whilst further ensuring that device operation maintains the long-term performance of the participating units.

3.2.3 Hysteresis control of deferrable loads

The focus now shifts to a common device-level control strategy that will look to maintain levels of end-use functionality commensurate with user satisfaction, while further providing a foreground for an effective system-level control strategy. For TCLs, device-level control is maintained by thermostats, which control the operational-state of the TCL through comparing the current level of end-use function periodically, and comparing it in relation to the desired end-use function (the set-point) in order to decide whether or not to operate the device over the next operational period. Temperature measurements are usually accompanied by a considerable amount of volatility, and therefore thermostats typically employ hysteresis-based control logic. Hysteresis control involves defining a range of end-use measurements that can occur, or a deadband space, over which end-use state measurements contained within result in no change in the operational state. Any measurement outside this region results in a state-transition to either the active or inactive machine state, depending on the current state of the device. This type of strategy prevents rapid cycling of the TCL that can occur under a noisy input and a definite set-point objective. As many types of loads are subject to design constraints that result in the need for minimum run- and

down-times, hysteresis control maintains these operational qualities autonomously through the choice of the deadband width.

As thermostats sample the end-use state and temperature periodically to control the TCL, they thereby discretize the system, where mathematically, hysteresis control logic can be written for a heating condition as [48]:

$$n[k+1] = \begin{cases} 1 & \epsilon[k] \leq \epsilon_- \\ 0 & \epsilon[k] \geq \epsilon_+ \\ n[k] & otherwise \end{cases} \tag{3.5}$$

where ϵ is the sampled end-use state comparison, k the thermostat's sampling index associated with the current operational interval, ϵ_+ representing the upper deadband boundary, and ϵ_- denoting the lower deadband boundary. When the device- level controller measures a charging-trajectory outside the deadband space, or beyond the deadband boundary, it will switch the charging status of the TCL. Thus, ϵ_- and ϵ_+ represent state-transition boundaries, where measurements below ϵ_- cause TCLs in the inactive-state to transition to the active-state, while end-use measurements above ϵ_+ cause active TCLs to transition into the inactive-state.The transition boundaries are assumed to be centred around the set-point datum:

$$\epsilon_\pm = \pm\frac{\delta}{2} \tag{3.6}$$

where δ is the deadband width. In the case of heat pumps, the sampled end-use state comparison is given by:

$$\epsilon[k] = \frac{\theta_a[k] - \theta_s[k]}{\delta} \tag{3.7}$$

where θ_a is the indoor air temperature measurement made by the device-level controller, and θ_s the user's current set-point temperature. In the case of heat pumps, the end-use state comparison is non-dimensionalized by dividing it by the deadband width.

As suggested by Callaway and Hiskens [11], the operation of deferrable energy-constrained storage can further be controlled at the device-level by hysteresis controllers. For EVs and electrolyzers, the primary objective of the device-level controller is to maintain a charging trajectory from which (3.3) can always be attained. The set-point energy-level E_s therefore represents a charging-schedule that must be maintained to ensure the device can meet the end-user's desired energy-level by the

end of the charging horizon, and is given by:

$$E_s[k+1] = E_s[k] + \frac{E_c}{T_s}T \; ; \; E_s[0] = 0 \tag{3.8}$$

By periodically querying the end-use state of the storage unit (state-of-charge), and updating the charging-schedule, the device-level controller can then apply the hysteresis control logic given by (3.5) to control the operational state by converting to the following non-dimensionalized state-of-charge:

$$\epsilon[k] = \frac{E[k] - E_s[k]}{E_c} \tag{3.9}$$

3.3 Managing Large Populations

With an operational strategy common to both TCLs and energy-constrained storage defined at the component-level capable of assuring normal device operation is maintained, the focus now shifts to the development of a system-level strategy in order to manage large populations effectively.

3.3.1 Controlling Population Dynamics Online

Considering a population of responsive loads, denoting the population index as i, the aggregate demand P from a particular population of N_i individual units is:

$$P(t) = \sum_{i=1}^{N_i(t)} n_i(t) P_{r,i} \tag{3.10}$$

The operational state of the device at time t is determined by the individual digital controllers:

$$n_i(t) = n_i[k_i] \; ; \; k_i T_i \leq t \leq [k_i + 1]T_i \tag{3.11}$$

where each sample index and period has been explicitly written as k_i and T_i, as each individual unit is operating based on its own clock. Clearly, through controlling the individual operational states of the units, control of the aggregate load can be achieved, but will require not only solving for each individual machine-state, but also tracking of the individual end-use states of the devices to make sure device-level constraints are maintained (temperature/charging-levels and minimum run- and shut-

down times). Management of this type may prove unsuccessful in controlling loads in near-real-time, such as that required for regulation-based ancillary services, as the computational complexity resulting from the multi-period optimization required to handle these inputs will prove inefficient at the system-level except in trivial cases considering very small load populations.

The alternative approach taken in this thesis is to rely on the device-level hysteresis controllers to maintain normal operating limits. An aggregate-level management strategy based on the characteristics of hysteresis control is then formulated. This relieves the system-level task of tracking each individual unit's constraints, providing a more efficient method of managing larger populations.

Aggregate population dynamics can be described by the power density distribution function for both the active (ϕ_1) and inactive (ϕ_0) machine-states. These functions describe the amount of power at a given air temperature or state-of-charge relative to the total installed power that exists in the participating population. As the total power existing in the active state defines the aggregate load, the current level can be expressed in terms of the total power in the current responsive population P_{cap}, and a capacity-factor Φ as:

$$P(t) = \sum_{i=1}^{N_i(t)} P_{r,i} \int_{-\infty}^{\infty} \phi_1(\epsilon, t)\, d\epsilon = P_{cap}(t)\Phi(t) \tag{3.12}$$

If each load operates based on the same device-level clock, we can use the fact that the hysteresis control logic given by (3.5) introduces a discontinuity between sampling intervals, wherein the end-use state measurement is utilized to determine the individual machine-states thereafter, thus re-distributing the power density distribution in the aggregate system accordingly. Any distribution in a given state that has traversed past the corresponding state-transition boundary (ϵ_+ for $n = 1$, and ϵ_- for $n = 0$), will be transferred to the opposite distribution. Boundary conditions on either side of the discontinuity event occurring at time t^* require that:

$$\lim_{t_+^* \to t^*} \Phi(t_+^*) = \lim_{t_-^* \to t^*} \left\{ \int_{-\infty}^{\epsilon_-} \phi_0(\epsilon, t_-^*)\, d\epsilon + \int_{-\infty}^{\epsilon_+} \phi_1(\epsilon, t_-^*)\, d\epsilon \right\} \tag{3.13}$$

where t_+^* is the time just after the discontinuity event, and t_-^* the time just before.

Traditionally, the capacity factor can not be controlled, but considering (3.13)

it is clear that changes to the state-transition boundary locations could in fact be used to control the capacity-factor of the aggregate load over the next sampling interval. As was first suggested by Callaway in [40], if each element within a targeted population is equipped with communication hardware capable of enabling a network through which operational data from a LA and LC can be shared, or programmable communicating hysteresis controllers (PCH), the capacity factor could indeed be systemically controlled through perturbation u to the end-use state comparison. At the aggregate-level, again assuming the loads are synchronized, the capacity-factor is then:

$$\lim_{t_+^* \to t^*} \Phi(u, t_+^*) = \lim_{t_-^* \to t^*} \left\{ \int_{-\infty}^{u(t_-^*)+\epsilon_-} \phi_0(\epsilon, t_-^*)\, d\epsilon \quad + \int_{-\infty}^{u(t_-^*)+\epsilon_+} \phi_1(\epsilon, t_-^*)\, d\epsilon \right\} \tag{3.14}$$

From the viewpoint of the device-level controllers, the hysteresis control logic is now given by:

$$n_i[k+1] = \begin{cases} 1 & \epsilon_i[k] \leq \epsilon_- + u[k] \\ 0 & \epsilon_i[k] \leq \epsilon_+ + u[k] \\ n_i[k] & otherwise \end{cases} \tag{3.15}$$

The control signal synchronizes loads near state-transition boundaries to attain the response, and therefore in the case of energy storage, never involves extracting energy from the unit. This mitigates all possible conflicts associated with enhanced performance degradation and roundtrip conversion losses apparent in vehicle-to-grid and regenerative fuel cell approaches.

To measure the distributions and synchronize their response, each load willing to participate in the current responsive population provides the LA with its current power-state vector $\mathbf{z}_i$, then waiting for the system-level set-point decision before reacting to (3.15). By waiting for the system-level response, the individual sampling intervals are synchronized to the same central clock ($k_i = k$), and will therefore commit any managed aggregate load over the next measurement cycle, allowing us to schedule this load in real-time. The power-state vectors convey all information needed to classify the loads into the distributions, and as each load in the population

is a hybrid-state system, observability requires power-state vectors to be given by:

$$\mathbf{z}_i[k] = \begin{bmatrix} n_i[k] \\ \epsilon_i[k] \\ P_{r,i} \end{bmatrix} \tag{3.16}$$

The LA therefore obtains a certain set of power-state vectors $\mathbf{Z}$ from the participating loads where:

$$\mathbf{Z}[k] = \{\, \mathbf{z}_1[k], \mathbf{z}_2[k], \ldots, \mathbf{z}_{N_i}[k] \,\} \tag{3.17}$$

The proposed closed loop control strategy is depicted in Fig.(3.1). The resultant

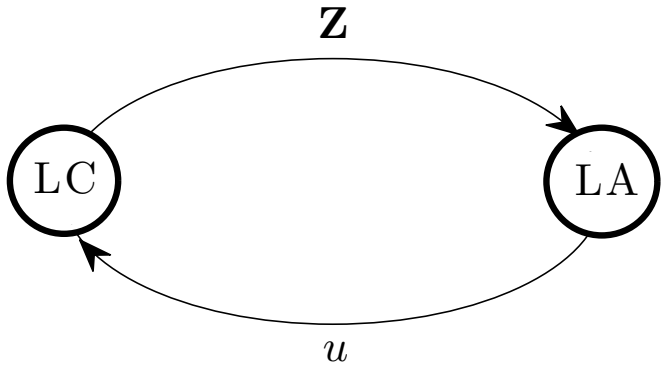

Figure 3.1: Closed-loop control strategy implemented.

delay-time between LA signal broadcast and LC response must be considerably less than the PCH update time to ensure compatibility, and will be directly related to the number of participating customers, network quality, and efficiency of system-level decision making.

As the device-level controllers are constrained by measurement resolution, the integral in (3.14) can be simplified over a region encompassing all possible end-use state measurements, which is taken to be twice the deadband width[1]. Each end-use

[1]In the case of heat pumps, if the customer was to change the temperature set-point, temperature measurements could be made that exist well outside the 2δ-region. Therefore, these customers would be unfit, or ineligible to participate in this aggregate load model formulation. As these customers choose to maintain control authority, they would in turn forfeit the benefits associated with participation. Therefore, this customer's thermostat would not send a power-state vector to the LA until the end-use state was again within this interval.

measurement therefore corresponds to an index m by:

$$\begin{gathered}\epsilon[m] = \delta\left(\frac{2m}{\mathcal{R}} - 1\right) \;\; ; \;\; \mathcal{R} \in \mathbb{N} \\ m = 0, 1, 2, ..., \mathcal{R}\end{gathered} \tag{3.18}$$

where $\mathcal{R}$ is the discretization resolution. If all device-level controllers operate at the same measurement resolution, through ordering $\mathbf{Z}$ into the distributions, an exact formulation of the forthcoming capacity-factor can be obtained as a function of the set-point index location:

$$\Phi[m_s, k+1] = \sum_{m=0}^{m_s[k]+m_-} \phi_0[m,k]\Delta\epsilon \; + \sum_{m=0}^{m_s[k]+m_+} \phi_1[m,k]\Delta\epsilon \tag{3.19}$$

where $\Delta\epsilon$ is the distance between possible state-measurements, m_s the index corresponding to the set-point, and $m_\pm$ denoting the indices associated with the transition boundaries. As the resultant function is convex, once obtained it can be reliably implemented by the LA within the following deviation minimization problem:

$$\textbf{Min} \quad \left(P^*[k+1] - P_{cap}[k+1]\Phi\left[\, m_s^*, k+1 \,\right] \right)^2 \tag{3.20}$$

where optimal solutions m_s^* cause a $k+1$ aggregate responsive load demand corresponding to the target P^*. This target will be pre-defined by the LA to achieve some grid-side objective. The optimal set-point index then converts to a corresponding set-point change by:

$$u[k] = \delta\left(\frac{2m_s^*[k]}{\mathcal{R}} - 1\right) \tag{3.21}$$

Without ramp-constraints, set-point modulations can access any load within the deadband space, resulting in possible rapid cycling of units that can degrade devices prematurely. This can be avoided through constraining LA decisions to set-point modulations existing within the quarter-deadband width:

$$|\, u[k] \,| \leq \frac{\delta}{4} \; \forall \; k \tag{3.22}$$

The above constraint prevents access to loads for over half of the deadband each operating cycle, guaranteeing that each system displays a somewhat natural duty-cycle, passing through the customer's desired set-point. All of the targeted devices

are designed to operate over a range in duty-cycle, and therefore should not face degradation issues under selection of (3.22). This constraint converts to corresponding maximum and minimum set-point indices $m_{s,max}$ and $m_{s,min}$:

$$\begin{aligned} m_{s,min} &= \frac{3}{8}\mathcal{R} \\ m_{s,max} &= \frac{5}{8}\mathcal{R} \end{aligned} \tag{3.23}$$

Therefore, the LA must define targets that exist within the feasible region:

$$\begin{aligned} P_{min}[k+1] \leq P^*[k+1] \leq P_{max}[k+1] \\ P_{min}[k+1] = P_{cap}[k+1]\Phi[m_{s,min}, k+1] \\ P_{max}[k+1] = P_{cap}[k+1]\Phi[m_{s,max}, k+1] \end{aligned} \tag{3.24}$$

The feasible region given by (3.24) limits ability to ramp the aggregate responsive load. This interval directly relates the end-use constraints (reliability-constrained charging-schedule) to power system-level planning (aggregate demand), and thus represents the VGM. As the responsive population-basis consistently evolves (especially in the case of EVs), the closed loop control strategy is critical, and allows the LA to accurately adapt the VGM online.

3.3.2 Optimal dispatch of multiple responsive load populations

Due to device-level dynamics, it can be expected that a LA that manages different types of loads will most likely need to generate individual VGMs for each, and will therefore need to determine a dispatch process that optimally allocates power within the management area. Of particular interest is minimizing both the variability and amplitude of the end-use function perturbation u. Reducing the variability of the control signal accelerates dispersion of end-use states within the system-level equivalent deadband space, widening the feasible region given in (3.24). This will result in a greater capability of the VGM to ramp up and down the next operational period, and thus enhance its long-term performance as a balancing resource. The benefit to also selecting an objective that minimizes amplitude is of course that the deflection from the customer's settings will be minimized. As it can be expected that this amplitude

will be proportional to forfeited end-use benefits, this objective also minimizes the cost of attaining the service from the LC.

Integer programming would need to be applied to explicitly solve for the different set-point indices, which may in turn present a bottleneck in terms of system-level decision making efficiency. Alternatively, a linear programming (LP) approach can be pursued by assuming that the amplitude of the control signal will be proportional to the amplitude of the deflection from the uncontrolled trajectory. The control signal that minimizes the variability and amplitude occurs at $u[k] = 0 \; \forall \; k$, or the trajectory corresponding to the uncontrolled capacity-factor P_o, which can be determined using (3.19) at the uncontrolled set-point index ($m_s = \mathcal{R}/2$):

$$P_o[k+1] = P_{cap}[k+1]\Phi\left[\mathcal{R}/2, k+1\right] \tag{3.25}$$

Denoting the controlled deviation from P_o as ΔP^*, the responsive load trajectory, as defined in (3.20), is then given by:

$$P^*[k+1] = P_o[k+1] + \Delta P^*[k+1] \tag{3.26}$$

Dropping the time-dependence, the controlled deflection is decomposed into positive and negative deflections from the uncontrolled datum:

$$\begin{gathered} \Delta P^* = \Delta P^*_+ - \Delta P^*_- \\ \Delta P^*_+, \Delta P^*_- \geq 0 \end{gathered} \tag{3.27}$$

A similar process can be pursued for the control signal:

$$\begin{gathered} u = u_+ - u_- \\ u_+, u_- \geq 0 \end{gathered} \tag{3.28}$$

Based on the maximum allowable set-point change, as defined in (3.22), in combination with the uncontrolled setting, we know the value of the controlled deflections in terms of the decomposed control signal at two points:

$$\begin{gathered} \Delta P^*_+(u_+ = 0) = 0 \\ \Delta P^*_+(u_+ = \delta/2) = P_{max} - P_o \end{gathered} \tag{3.29}$$

$$\begin{aligned} \Delta P_{-}^{*}(u_{-}=0) &= 0 \\ \Delta P_{-}^{*}(u_{-}=-\delta/2) &= P_o - P_{min} \end{aligned} \tag{3.30}$$

where the maximum and minimum are defined by the VGM feasible region in (3.24). While the relationship between these points may not be linear, to first order we can assume the deflections in either direction are proportional to the ratio of the controlled deflections and their maximums, namely:

$$\begin{aligned} u_{+} &= \frac{\delta}{4}\frac{\Delta P_{+}^{*}}{P_{max}-P_o} \\ u_{-} &= \frac{\delta}{4}\frac{\Delta P_{-}^{*}}{P_o - P_{min}} \end{aligned} \tag{3.31}$$

If there are now N_p responsive load types within the particular control area, the LA will attempt to minimize the control signal deflections, while ensuring that the grid-side objective target P_T and individual VGM ramp constraints are maintained. An optimization problem considering the grid-side inputs achieving these objectives can then be formulated as an LP by:

$$\begin{aligned} \textbf{Min} \sum_{p=1}^{N_p}\left(\frac{\Delta P_{+}^{(p)*}}{P_{max}^{(p)}-P_o^{(p)}}+\frac{\Delta P_{-}^{(p)*}}{P_o^{(p)}-P_{min}^{(p)}}\right) \\ \textbf{s.t.} \sum_{p=1}^{N_p}\left(\Delta P_{+}^{(p)*}-\Delta P_{-}^{(p)*}\right) = P_T - P_U + P_W - \sum_{p=1}^{N_p} P_o^{(p)} \\ \Delta P_{+}^{(p)*} \le P_{max}^{(p)} - P_o^{(p)} \\ \Delta P_{-}^{(p)*} \le P_o^{(p)} - P_{min}^{(p)} \\ \Delta P_{+}^{(p)*}, \Delta P_{-}^{(p)*} \ge 0 \end{aligned} \tag{3.32}$$

where P_U is the area's unresponsive load, P_W the wind power generation, and p representing the VGM number associated with a particular load-type. Solving the LP results in controlled deflections from the uncontrolled trajectories capable of meeting the grid-side target, that are in turn optimal in the sense that the amplitude of the control signal u is minimized across each load-type. This formulation assumes the defined grid-side target is feasible, and is a condition that must be checked before the LP can be solved.

3.3.3 Target design for a self-regulating load

As the LA develops models that are equivalent to conventional online generation, this virtual generation can be dispatched to mitigate variability in place of traditional regulation reserve. In this case, the losses accrued at end-service (converting to battery state-of-charge or air temperature) is only a problem for the end-user, not the grid, as there is no need to convert this stored energy back to electricity (changing the state-trajectory of the deferrable loads provides the desired response). Therefore, from the grid's perspective, the service is obtained at near 100% efficiency, with the inefficiencies in this case accruing due to transmission and distribution losses. To obtain this service, the target response P_T must be defined to provide regulation-based services to the particular area, or *self-regulate* the area load. As wind power is non-dispatchable, it is often treated as a component of the area demand (negative load) and therefore a self-regulating objective attempts to simultaneously smooth the total demand and variable wind energy supply.

In order to achieve this goal, the grid-side objective is set to attain an aggregate demand target that does not require dispatch of regulation P_R. This target represents an aggregate demand trajectory that causes demand fluctuations that occur outside of the regulation time regime. A weighted moving average of the sampled area load can be used in an attempt to smooth the demand over some defined range:

$$P_R[k+1] = \sum_{j=1}^{M} a_j P_{tot}[k-j+1] \; ; \; \sum_{j=1}^{M} a_j = 1 \tag{3.33}$$

where P_{tot} is the total area load, M is the number of past load samples to be considered in the moving average, and a_j the weighting associated with the particular load sample. In this thesis, these weightings are taken to be equivalent, and therefore (3.33) represents a moving average. It will be the focus of future work to develop an online procedure to set these coefficients to ensure the long-term performance of the controller through integration of possible wind / load forecasting techniques. As this regulation target is defined at the level of the LA, there is no need to interact with the power system operator if the LA can obtain the area load measurement.

Even in unison, the responsive load populations may be unable to meet such a trajectory, due to the end-use constraints contained in the individual VGMs defined in (3.24). Therefore, the target must be refined to reflect these limits through con-

sideration of each VGM's ramp-constraints when defining the target:

$$P_T[k+1] = \begin{cases} \sum_{p=1}^{N_p} P_{max}^{(p)}[k+1] & P_R[k+1] \geq \sum_{p=1}^{N_p} P_{max}^{(p)}[k+1] \\ \sum_{p=1}^{N_p} P_{min}^{(p)}[k+1] & P_R[k+1] \leq \sum_{p=1}^{N_p} P_{min}^{(p)}[k+1] \\ P_R[k+1] & otherwise \end{cases} \tag{3.34}$$

This corrected target is what is sent to the optimization problem in (3.32), and guarantees a feasible LP solution.

Figure (3.2) depicts a flow-chart of the multi-stage optimization problem. While in this thesis the self-regulating objective has been defined as the target, it would be relatively straightforward to implement other types of objectives (i.e. frequency and voltage regulation, relays from conventional power system management, etc.), if these objectives can be formulated in terms of an equivalent change in demand.

The following section will introduce the computational modelling study implemented to observe the performance of the proposed load control strategy.

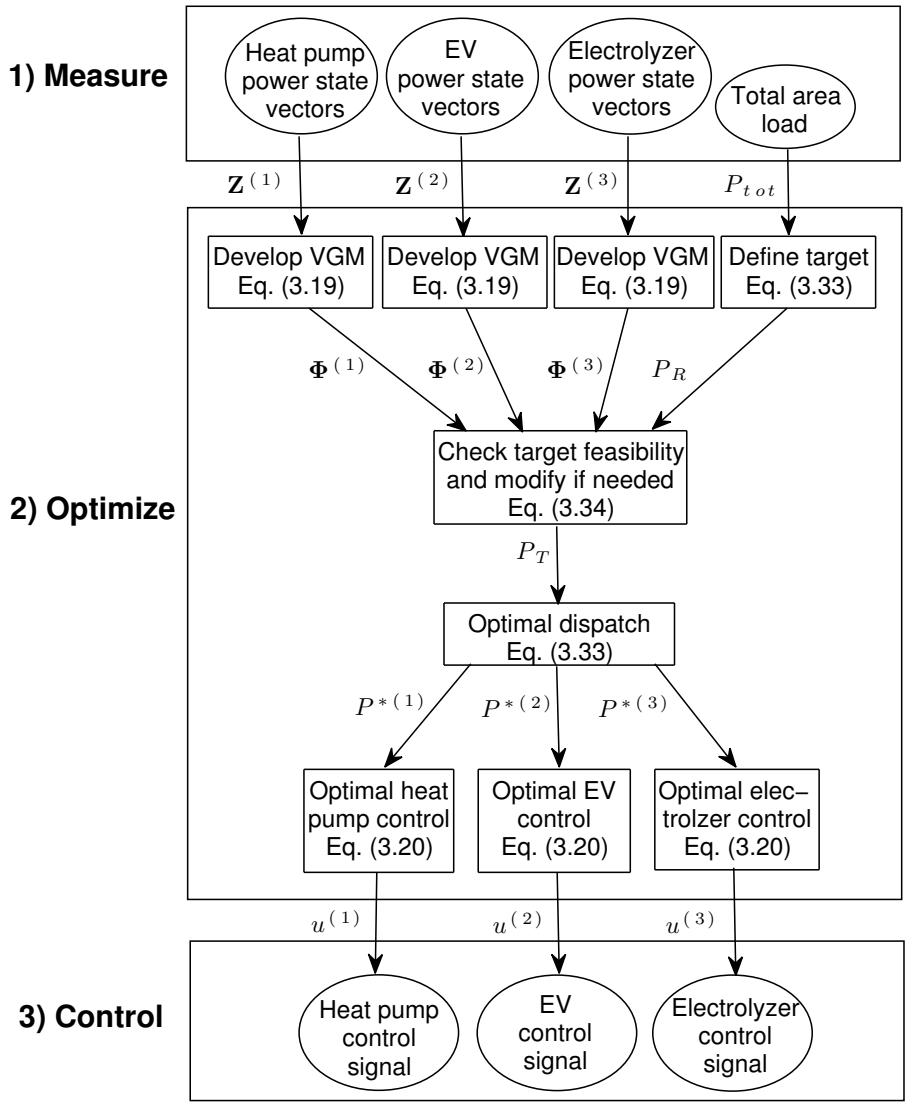

Figure 3.2: Flow-chart of the multi-stage optimization problem

Chapter 4

Computational Modelling Framework

The previous chapter presented a mathematical framework for engaging sustainable demand-side infrastructure in the provision of power system ancillary services. Obtaining experimental evidence regarding the viability of this system configuration will prove extremely costly. Thus, a computational model was developed to test the performance of the proposed load control policy. The strategy is inherently multi-scalar, as it looks to encompass dynamics occurring within the supply- and demand-side, and therefore requires a similarly themed modelling framework to understand its potential.

Of particular interest is simulation of the device-level hysteresis controllers. These govern both the resultant load as well as providing the power-state vectors to the LA, and will require sufficient treatment of the end-use states of the individual loads (air-temperature for heat pumps and charging trajectory for energy-constrained storage) that are input to the hysteresis control logic. This approach of modelling the individual device-states, that are in turn used by a local control authority that takes part in the aggregate response of interest, is typically referred to as an agent-based model. In this case, the hysteresis controllers are the autonomous agents, as they look to fulfill the local control objectives, which is to maintain the end-use state comparison within a certain deadband width, as given by (3.15). The interaction between the agents emerges as an aggregate load on the power system, as given by (3.10). The purpose of this chapter is to introduce the individual dynamic models of these agents (the loads), providing a framework for rating the performance of the management

strategy. The latter looks to manipulate the interaction between these agents (the response), for system-level benefits (ancillary services).

4.1 Demand-side: Load Models

Each targeted end-use device is a complicated systems unto itself. Large-scale simulations that look to capture the behaviour of an individual load will therefore require simplification. The following sections present the dynamic load models for each targeted end-use device (heat pumps, EVs, and electrolyzers).

4.1.1 Heat Pump Model

Accurately modelling an individual air-source heat pump requires treatment of the building physics that govern its operation. Building heat transfer dynamics is complex, where multi-physics approaches are commonly employed at the building-level. Such tactics are computationally intense, and therefore unsuitable within a modelling framework that is primarily interested in large populations of buildings. An alternative is a simplified model structure, that typically takes the form of a one-dimensional equivalent thermal parameter (ETP) model. An ETP model decomposes the building into different thermal masses that are connected to each other through coefficients of heat transfer. One of the simplest ETP-type models considers a small- to medium-sized building as a coupled thermal system, where the indoor air thermal mass C_a is connected to a single building-interior thermal mass C_m, through a heat transfer resistance R_{ma}. Outdoor air temperature conditions θ_o interact with the coupled masses, through the building envelope resistance R_{ao}. The complete coupled ETP model is displayed in Fig.(4.1) [49].

The temperature trajectories of the individual masses (θ_a for the air temperature, and θ_m for the interior contents) are then coupled through the following linear dynamic system representation:

$$\dot{\boldsymbol{\theta}} = \mathbf{A}\boldsymbol{\theta} + \mathbf{B}\mathbf{U} \tag{4.1}$$

The state-equation defines the temperature trajectories of the coupled thermal masses under the current system inputs (outdoor air temperature θ_o, and heat input $\dot{q}$). The

state- and input-matrix are determined by the structure of Fig.(4.1):

$$a_{11} = -\left(\frac{1}{R_{ma}C_a} + \frac{1}{R_{ao}C_a}\right) \quad a_{12} = \frac{1}{R_{ma}C_a} \quad a_{21} = \frac{1}{R_{ma}C_m} \quad a_{22} = -\frac{1}{R_{ma}C_m}$$

$$b_{11} = \frac{1}{R_{ao}C_a} \qquad b_{12} = \frac{1}{C_a} \qquad b_{21} = 0 \qquad b_{22} = 0$$

$$\mathbf{A} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \quad \mathbf{B} = \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} \quad \boldsymbol{\theta} = \begin{bmatrix} \theta_a \\ \theta_m \end{bmatrix} \quad \mathbf{U} = \begin{bmatrix} \theta_o \\ \dot{q} \end{bmatrix}$$

$$\dot{\boldsymbol{\theta}} = \mathbf{A}\boldsymbol{\theta} + \mathbf{B}\mathbf{U} \tag{4.2}$$

In this model, the imposed heat input is determined by

$$\dot{q}(t) = n(t)\,\dot{q}_{op}(t) \tag{4.3}$$

where $\dot{q}_{op}$ is the operational heating power of the device, and the operational state of the device is determined by the hysteresis control logic given by (3.15).

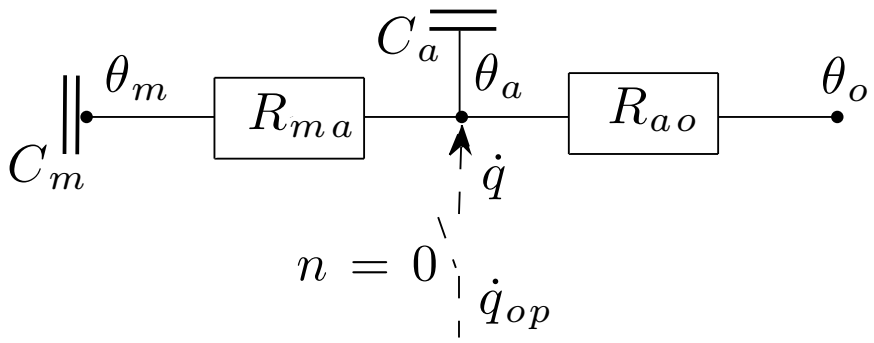

Figure 4.1: The one-dimensional coupled ETP model of a building, with thermostatically controlled heat pump currently in the inactive-state ($n = 0$).

A solution to the resultant discrete-time version of (4.1) under steady-inputs is [50]:

$$\begin{aligned} \boldsymbol{\theta}[k+1] &= \boldsymbol{\Omega}\boldsymbol{\theta}[k] + \boldsymbol{\Gamma}\mathbf{U}[k] \\ \boldsymbol{\Omega} &= exp\,(\mathbf{A}T) \\ \boldsymbol{\Gamma} &= \int_0^T \boldsymbol{\Omega}\mathbf{B}d\tau \end{aligned} \tag{4.4}$$

where $\mathbf{\Omega}$ is the state-transition matrix, and $\mathbf{\Gamma}$ the input-transition matrix. The dynamics over which outdoor temperature gradients effect device-level duty-cycle (diurnal) are much longer than that of the thermostat (typically 1 minute), making the assumption of a constant uniform outdoor temperature regime over this period reasonable. For added accuracy, the outdoor temperature input to the building models should reflect the sol-air temperature in hopes of representing any solar-induced heat flux in the system. All other unmodelled system internal gains / losses (operation of other electronic devices, opening doors, body-heat, etc.) are considered through inclusion of a stochastic white noise term $\mathbf{e}$ displaying independent components of standard deviation σ_n.

$$\boldsymbol{\theta}[k+1] = \mathbf{\Omega}\boldsymbol{\theta}[k] + \mathbf{\Gamma}\mathbf{U}[k] + \mathbf{e}[k] \tag{4.5}$$

Heat pump unit sizing

The thermal-integrity level of each building, contained within its corresponding ETPs, can further be utilized to size an individual heat pump. To ensure reliability, each device must be sized under the design-level conditions that the building will be subject to. These include worst-case scenario estimates of all terms in (4.5), under the climatic temperature extremes expected in the particular geographic region. By balancing the system heat flux that occurs under these conditions, so that the desired indoor air temperature represents the equilibrium level of the linear system, and applying an appropriate oversizing factor, the design heating power of the heat pump $\dot{q}_d$ can be determined by:

$$\dot{q}_d = \gamma \left(\frac{\theta_s - \theta_d}{R_{ao}} \right) \tag{4.6}$$

where γ is the over-sizing factor, and θ_d is the heating design temperature. The operational capacity-rating is time dependent, due to the fact that the efficiency η of an air-source heat pump depends on outdoor temperature conditions. This characteristic is treated through fitting to a manufacturer's data-sheet third order polynomials with constant coefficients α_x.

$$\eta[k] = \sum_{x=0}^{2} \alpha_x \theta_o[k]^x \tag{4.7}$$

The individual unit can then be sized under the design heating power of the heat pump, with the efficiency corresponding to the design efficiency at this temperature extreme η_d:

$$P_h = \frac{\dot{q}_d}{\eta_d} \tag{4.8}$$

where P_h is the required electrical power to produce the design heating power. Each rated power is then rounded up to the nearest 0.5 kW to account for discretization between unit capacities. Assuming the device will always be operated at this level, $\dot{q}_h$ varies according to:

$$\dot{q}_h[k] = \eta[k] P_h \tag{4.9}$$

The heat generated by the heat pump is circulated through the building by a fan. The operational capacity of the device is then given by the heating power of the unit $\dot{q}_h$, increased by the heating power attained from fan operation $\dot{q}_f$:

$$\dot{q}_{op}[k] = \dot{q}_h[k] + \dot{q}_f \tag{4.10}$$

The total electric power associated with the operational heat pump is then modelled by:

$$P_r = P_h + P_f \tag{4.11}$$

where in this work, the fan is assumed to operate at an ideal efficiency ($\dot{q}_f = P_f$).

Example

To demonstrate the device-level dynamics, a simple example considering a single building is presented. The parameters implemented in the model can be found in [51] and [40], and are given in Tab.(4.1). The resulting ETP model is meant to represent a typical small- to medium-scale building or single-detached residential home.

The device is sized under a design heating temperature of -5^oC, and results in a heat pump power rating of $P_r = 5.5$ kW. The model is ran under the outdoor temperature profile given in the top panel of Fig.(4.2), with the results given in the bottom two panels. The indoor air temperature is seen to rebound between the thermostat deadband, at a rate directly related to the outdoor temperature regime, while the indoor building thermal mass remains centred around the set-point. The warmer

Table 4.1: Parameters implemented in the heat pump model

Parameter	Description	Value
R	Envelope thermal resistance	2 ^{o}C/kW
C	Bulk thermal mass	10 kWh/^{o}C
C_m	Indoor contents thermal mass	75% of C
C_a	Indoor air thermal mass	25% of C
R_{ma}	Indoor heat transfer resistance	0.5 ^{o}C/kW
$\dot{q}_{fan}$	Heating power of the fan	500 W
θ_s	Temperature setpoint	20 ^{o}C
δ	Thermostat deadband width	1 ^{o}C
T	Sampling time interval	1 minute

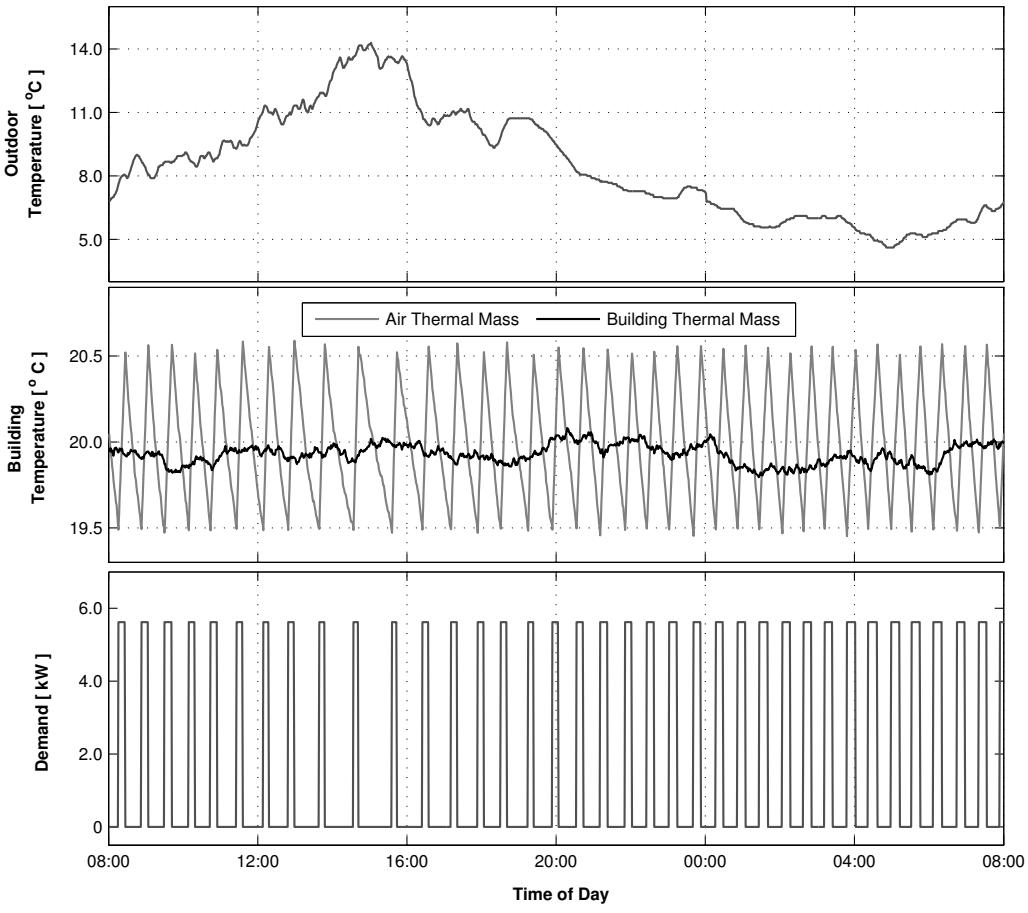

Figure 4.2: Example of the ETP model under the displayed outdoor temperature.

temperatures experienced by the building in the late afternoon results in less of a temperature gradient to exist across the building envelope, and therefore the building loses less energy over these periods. This results in less need to operate the unit, as the indoor temperature trajectory decays at a slower rate in comparison to that observed at colder temepratures, while the operational intervals of the device diminish due to less of a need to add heat energy into the system. It is therefore extremely important to accurately capture these dynamics within the load model, as they directly effect the device's duty-cycle (percentage of time the unit is active relative to the time it is inactive), or the loading the heat pump places on the grid. The coupled ETP model allows for these characteristics to carry through, while further providing the computational efficiency for application in population-based simulations.

4.1.2 EV charging

The charging of an EV is strongly tied to the travel history associated with the unit, which is complicated to capture without full treatment of the social factors governing charging schedules. We can in fact represent the charging trajectory using a simplified linear model of the charging process alone, and neglect the travel history through stipulation of a suitable charging capacity in (3.3). The charging process is an electrochemical reaction, and is non-spontaneous, meaning that work equivalent to the Gibbs free energy change must be provided for the reaction to proceed in the proposed direction. The power system will provide the energy needed by the reaction, as well as the encompassing intrinsic heat loss to the environment, due to a change in entropy. As the characteristic of interest is the behaviour of the hysteresis controller governing the charging status, the approach pursued in this work is to simply model the local controller using a quasi-steady state model of the charging reaction through application of a defined efficiency. This controller will need to sample the charging trajectory periodically, with the state-equation for charging-trajectory measurements occurring at a sample period T modelled by:

$$E[k+1] = E[k] + n[k]\eta P_r T + e[k] \tag{4.12}$$

An error term e has been added, that will be taken to display constant variance, and is meant to represent all unmodelled physical phenomena associated with the charging process (ambient temperature changes, transient electrochemical effects, leaks, sensor noise, etc.). Future work should focus on extending this load model to account for

the transient phenomena accompanying the charging process.

Example

An example of the EV charging model is given in Fig.(4.3), with Table (4.2) overviewing the model parameters implemented. The unit modelled requires 14 kWh of

Table 4.2: Parameters implemented in the EV charging model

Parameter	Description	Value
P_r	Rated power of the charger	4 kW
η	Charging efficiency	0.95
E_c	Desired charging capacity	14 kWh
T_s	Desired time to accumulate E_c	7 hours
δ	Hysteresis control deadband width	0.025
T	Sampling time interval	1 minute

charging energy over a charging horizon of approximately 7 hours, with the rated power of the charger being 4 kW and a charging efficiency of 95%. This could represent a typical EV owner, who depletes the majority of their battery's state of charge over the course of the day, and then charges the vehicle overnight. The charging trajectory is mapped to the hysteresis controller's deadband space by (3.9), using the set-point defined by (3.8), and results in the duty-cycle-based charging rate, and the staggered charging trajectory. The choice of deadband width somewhat sets the cycling frequency. A larger deadband space results in a greater allowable drift from the user's set-point or distance for loads to traverse and therefore longer times between operational switching events, with the other extremely important input driving device-level duty-cycle being the user's set-point. For this example it has been set at $\delta = 0.025$ so that the operational periods last between about 20 and 30 minutes.

4.1.3 Electrolyzers

Electrolyzers utilize the process of electrolysis to decompose water into its constituent gases oxygen, and hydrogen. This reaction is non-spontaneous, and therefore (4.12) can also be applied for electrolyzers as well through stipulation of a suitable "charging" capacity. It is assumed that this capacity corresponds to a daily production quota, and therefore there is no need to account for units passing into and out of the available responsive population (units participate continuously).

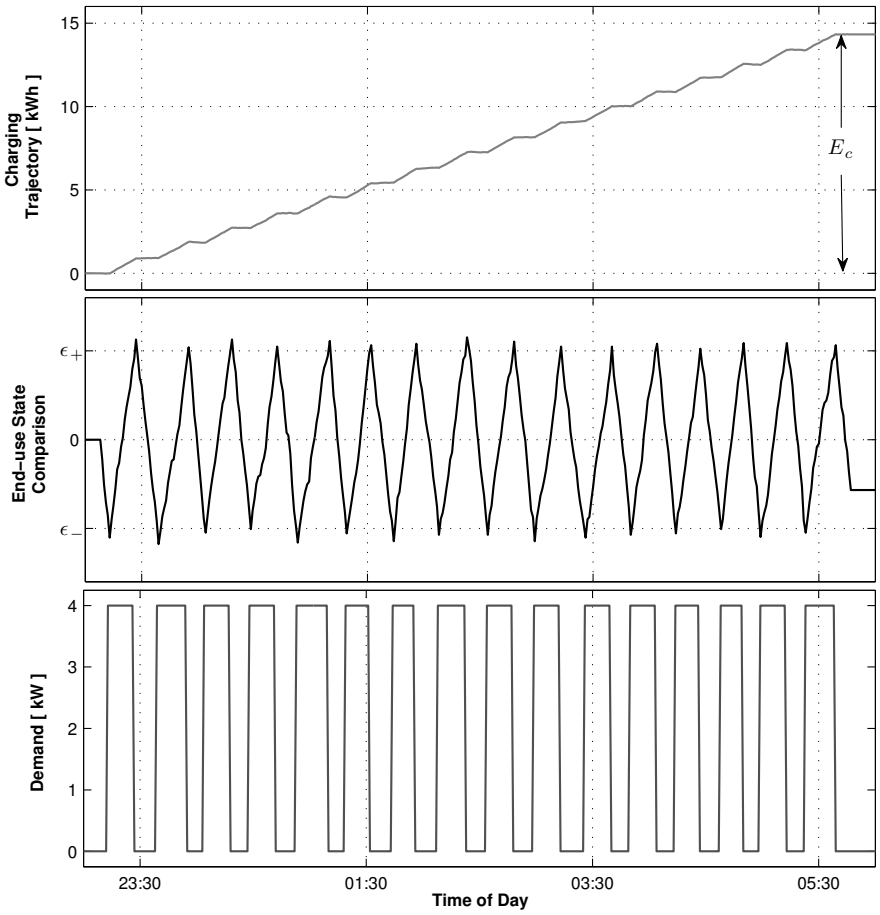

Figure 4.3: Example of the EV load model.

Example

An example of the electrolytic hydrogen production model is given in Fig.(4.4), where a 20 kW unit is commissioned to produce 170 kWh of hydrogen (based on lower-heating value) over the production period (1 day). The efficiency of the production process is sensitive to the voltage applied across the electrolytic cells, where it is assumed that in each case the rated power that the unit will be operated at corresponds to the maximum efficiency associated with the number of electrolytic cells encompassed in the unit ($\eta = 0.75$). A summary of the parameters implemented in the electrolyzer example are give in Table (4.3). Again, the charging trajectory is

Table 4.3: Parameters implemented in the electrolyzer model

Parameter	Description	Value
P_r	Rated power of electrolyzer	20 kW
η	Electrolyzer energy conversion efficiency	0.75
E_c	Desired amount of hydrogen (energy content)	170 kWh
T_s	Desired time to accumulate E_c	24 hours
δ	Hysteresis control deadband width	0.035

mapped to the hysteresis controller's deadband space by (3.9). In the case of electrolyzers, the choice of the deadband width is important, as rapid cycling of these units has been shown to degrade performance [47]. A deadband width of $\delta = 0.035$ is selected, as this produces minimum run- and down-times exceeding 45 minutes, which should allow transient phenomena to sufficiently diminish during operation.

4.2 Supply-side: Power System Model

As the primary interest of this work is the prospect of managing demand-side infrastructure for power system ancillary services, a model for the power system in which the load models interact is needed. Furthermore, these interests lie with the operation of power systems that employ a low carbon mixture containing large amounts of wind generation. Treatment of wind generation systems within the model is therefore needed. The purpose of this section is to introduce the modelling framework for the power system, as well as the model of the wind energy system that will promote low carbon operation.

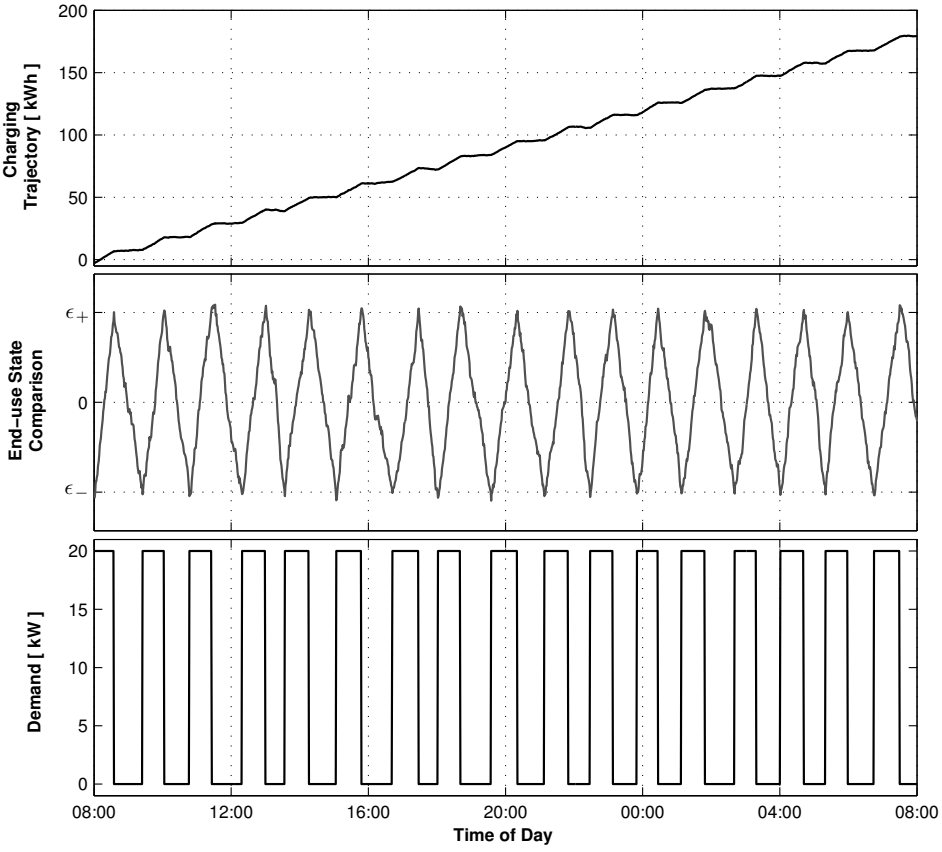

Figure 4.4: Example of the electrolyzer load model.

4.2.1 Wind Energy System Model

In modern turbines, wind that passes over the turbine's blades generate lift based on the structural properties of the blade (material and shape). This lifting force spins the turbine blades, which are connected to an induction generator. The generator produces an electrical current due to the time-varying magnetic flux brought on by the spinning generator, which when passed through power electronics, can be used as power supply on the grid. Due to the large number of components that make up a system of this type, development of rigorous analytical models capable of accurately quantifying the systems performance is an extremely arduous undertaking. Typically, simplification is pursued, and a general transfer function between the wind speed input and power output is defined.

The quasi-steady total mechanical power extracted by a wind turbine P_m is related to the wind speed u_w through [52]:

$$P_m[t] = \frac{1}{8} C_p\left(\lambda[k], \beta[k]\right) \rho \pi D^2 u_w[k]^3 \tag{4.13}$$

where D is the rotor diameter, ρ the air density, and C_p the turbine's coefficient of performance, which in turn is a function of the tip-speed ratio λ and blade pitch angle β. Equation (4.13) relates the mechanical power imparted to the turbine, to the total power in the wind stream passing over its blades through considering the wind turbine aerodynamic interaction as an inefficiency, contained in C_p. Often the full $C_p(\lambda, \beta)$ data are not available so that (4.13) cannot be used directly in power performance evaluations. According to [53], a good approximation is found by using the following relation:

$$C_p[k] = \begin{cases} 0 & u_w[k] < u_{ci} \\ C_{p,m}\left[1 - F_1\left(\frac{u_m}{u_w[k]} - 1\right)^2 - F_2\left(\frac{u_m}{u_w[k]} - 1\right)^3\right] & u_{ci} \le u_w[k] \le u_r \\ C_{p,r}\left(\frac{u_r}{u_w[k]}\right)^3 & u_r \le u_w[k] \le u_{co} \\ 0 & u_w[k] > u_{co} \end{cases} \tag{4.14}$$

where u_{ci} is the wind speed at which the turbine begins to generate power (cut-in), $C_{p,m}$ is the maximum coefficient of performance, u_m the wind-speed at $C_{p,m}$, u_r the

wind speed at which the turbine is rated, $C_{p,r}$ coefficient of performance at u_r, and u_{co} the speed at which the turbine furls (cut-out). The coefficients F_1 and F_2 are determined using boundary conditions, namely that $C_p = 0$ at $u_w = u_{ci}$ and $C_p = C_r$ at $u_w = u_r$.

In this thesis, the mechanical power imparted to the turbine blades is taken to be passed through a gear-box before the generator. Power electronics then interface the system to the grid. These conversions further reduce the available work the wind turbine can produce on the grid-side. Based on the method described in [54], the gearbox is modelled as a number of stages S, where the efficiency at which the gearbox operates is related to the current mechanical input by:

$$\eta_g[k] = \begin{cases} 0.1 & P_m[k] \leq \dfrac{SP_{rt}}{90} \\ 1 - \dfrac{0.01SP_{rt}}{P_m[k]} & otherwise \end{cases} \tag{4.15}$$

where P_{rt} is the rated power of the turbine, which occurs at the rated wind speed. The piecewise approach is applied to ensure the gearbox efficiency remains positive. The model for the generator simply assumes that all electrical inefficiencies are described as a constant η_e, or that the wind system provides a constant power input unto the grid and does not effect voltage. The total wind power injected onto the grid P_W can then be obtained using the mechanical power imparted to the blades less the inefficiencies that accompany the transferral of this power through the mechanical and electrical components:

$$P_W[k] = \eta_e \eta_g[k] P_m[k] \tag{4.16}$$

Tab. (4.4) displays the inputs used throughout this work to describe the wind turbine, and are meant to represent a generic 1.5 MW system.

4.2.2 Grid Model

The wind energy system and load models interact through the power system model. A DC formulation of a single network bus is implemented to simulate the power system, as it provides a computationally efficient method of modelling both the load and renewable generation, while also providing a platform to test both grid-connected and isolated systems. Future work should focus on extending the model to AC simulations and accounting for distribution-level power-flow, but for the purpose of this work, the

Table 4.4: Parameters implemented in the wind energy system model

Parameter	Description	Value
P_{rt}	Rated power of the unit	1.5 MW
D	Rotor diameter	82.5 m
u_{ci}	Cut-in wind speed	4 m/s
u_{co}	Cut-out wind speed	26 m/s
u_r	Rated wind speed	12 m/s
$C_{p,r}$	Rated coefficient of performance	0.250
$C_{p,m}$	Maximum coefficient of performance	0.300
u_m	Wind speed at $C_{p,m}$	7 m/s
S	Number of stages in the gearbox	3
η_e	Efficiency of the electrical components	0.95

DC formulation will allow us to sufficiently capture the primary characteristics of the proposed load managment policy, which is balancing energy within the system.

The selected energy system is given in Fig.(4.5), where the total aggregate community load is given by P_L, the wind power by P_W, and the conventional generation by P_G. The energy system is modelled through consideration of the grid management problem, which is to minimize operating costs subject to the operating constraints. Generators are modelled as capacity- and ramp-constrained units, which results in the need to consider power curtailment P_C at the network bus in the case the power-gradient exceeds the ramping duty of the generators $\dot{P}_{G,max}$ or capacity constraints $P_{G,min}/P_{G,max}$. The best possible outcome can be obtained by assuming all inputs are known from the outset (perfect knowledge assumption), and solving the grid-management problem as a dynamic optimization, namely:

$$\begin{aligned} \textbf{Min} \quad & OC = \sum_{k=1}^{N_k} \left(w_G[k] P_G[k] + w_C[k] P_C[k] \right) \\ \textbf{s.t.} \quad & P_G[k] + P_C[k] = P_L[k] - P_W[k] \; ; \; \forall \; k \\ & P_G[k] - P_G[k-1] \leq \dot{P}_{G,max} T \; ; \; \forall \; k \\ & P_G[k-1] - P_G[k] \leq \dot{P}_{G,max} T \; ; \; \forall \; k \\ & P_{G,min} \leq P_G[k] \leq P_{G,max} \; ; \; \forall \; k \\ & P_C[k] \geq 0 \; ; \; \forall \; k \end{aligned} \tag{4.17}$$

where OC is the total operating cost associated with the planning horizon consisting of N_k decision intervals, w_G the cost associated with the generator's output, and w_C the costs associated with the curtailed power. No attempt is made at this point to

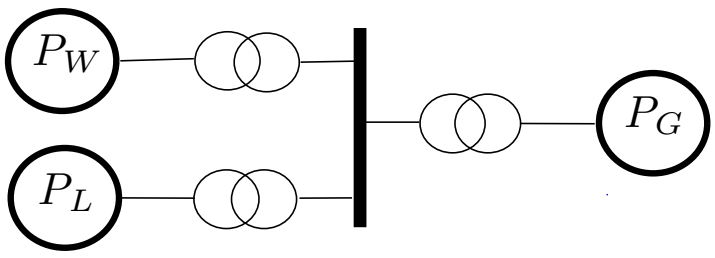

Figure 4.5: Simplified DC electric power system model.

distinguish the control components, as the main interest is obtaining the aggregate load placed on the conventional generation units. Under this main objective, the

costs are set at $w_G = 0$ and $w_C = 1$ for further simplification, so that the management system focuses on the single objective of attempting to minimize the amount of curtailed energy from the bus.

The total aggregate load is made up not only of the responsive load populations, but further the unresponsive loads existing within the community as well. The unresponsive load profile used throughout this work is given in Fig.(4.6), along with the wind power data that has been generated using the wind plant model and then scaled up to represent a 10 MW plant. The unresponsive load is generated to represent a typical load profile that displays diurnal variations driven by the synchronization of community resident schedules. The total load on the network bus at each time-step

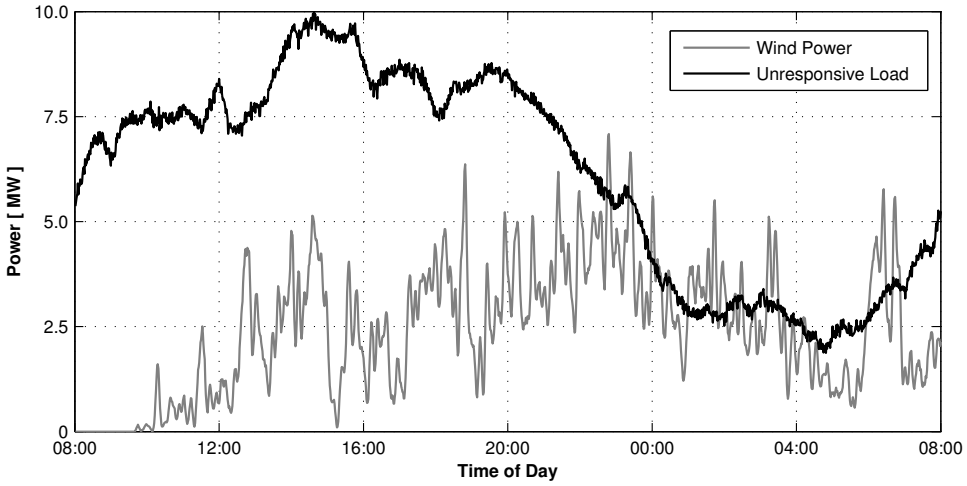

Figure 4.6: Unresponsive load and wind power time-series data.

is then given by the aggregation of the individual load profiles in (3.10), added to the uncontrolled inputs to the bus.

$$P_L[k] = \sum_{p=1}^{N_p} \sum_{i=1}^{N_i^{(p)}[k]} \left(n_i^{(p)}[k] P_{r,i}^{(p)} \right) + P_U[k] \tag{4.18}$$

4.2.3 Selecting the Simulation Time-step

The complete model integrating each component description requires the ability to capture both the dynamics of the load models and generation while further preserving

the accuracy of the simplifications pursued to accelerate computational times. A one-minute sampling time ($T = 1$ min) provides these qualities, as it will allow transient dynamics associated with the component models to sufficiently dissipate, allowing the quasi-steady-state descriptions to remain viable, while further coinciding with the device and grid-level control strategies of interest. Below one-minute, generation has been shown to not follow load exactly [55], while ill-considered turbine structural dynamics are also observed to effect the performance of the unit below a one-minute sampling frequency [20]. Wind displays considerable variability at this scale, which similarly themed models that focus on longer time-scales may neglect.

From the viewpoint of the control strategy, individual loads will be entering into and out of the LC, especially in the case of EVs. End-users wishing to interact with the loads can not be expected to remain passive until the control interval is complete. Selection of one-minute provides some room for user's accommodation, in which they may choose to wait in order to reap the benefits associated with participation. Furthermore, the device-level controllers typically operate at this sampling frequency, meaning that longer simulation steps may in fact be incompatible. Finally, grid regulation is the primary ancillary service of interest, which takes place between economic dispatch intervals. These intervals typically take place every 5 to 15 minutes, meaning a one-minute sampling time should provide a sufficient framework to capture the regulation dynamics.

4.3 Control Performance Metrics

To understand the viability of the proposed load management strategy, the following section introduces the metrics used to rate performance, which should look to directly quantify the merits for success outlined at the beginning of Chapter (3).

4.3.1 Supply-side

As introduced in Chapter (2), increased integration of intermittent renewable generation will require greater access to real-time balancing services (regulation) in order to maintain system security. As suggested in [23], different ancillary service types can be estimated through separating the loading profile on the grid's generation into base-load, load-following, and regulation components. The base-load and load-following components are meant to represent energy allocated during UC and ED, and can be

estimated using the 15-minute moving average of the total area load:

$$P_{base} + P_{lf}[k] = \frac{1}{15} \sum_{m=-7}^{7} P_L[k+m] \tag{4.19}$$

The base load P_{base} represents the minimum constant load that can be achieved:

$$P_{base} = \mathbf{Min} \left(\frac{1}{15} \sum_{m=-7}^{7} P_L[k+m] \right) \tag{4.20}$$

while the load-following component P_{lf} is given by:

$$P_{lf}[k] = \frac{1}{15} \sum_{m=-7}^{7} P_L[k+m] - P_{base} \tag{4.21}$$

The regulation component P_{reg} therefore encompasses the remaining load:

$$P_{reg}[k] = P_L[k] - P_{lf}[k] - P_{base} \tag{4.22}$$

Generators contracted to provide either load-following or regulation requirements must ensure reliability under all operating scenarios. These contracts were introduced in Chapter (2) to understand their effect on conventional online generation, and now using the separated components, the required characteristics of these resources can be quantified. In the case of load-following, the required capacity C_{lf} corresponds to the maximum of the load-following sequence less the base-load requirements:

$$C_{lf} = \mathbf{Max} \left(P_{lf} \right) - P_{base} \tag{4.23}$$

As the regulation capacity C_{reg} must cover both up- and down-ramping events, it is defined using three standard deviations of the absolute value of the regulation sequence:

$$C_{reg} = 3 \cdot \mathbf{stddev} \left(|P_{reg}| \right) \tag{4.24}$$

This approach assumes the regulation sequence is normally distributed so that three standard deviations encompass over 99% of occurrences.

Load-following and regulation also requires a certain level of responsiveness, or online ability to ramp participating resources up and down. Finite-differencing the

absolute value of the load-following signal provides the load-following power gradient $\dot{P}_{lf}$, with the regulation power gradient $\dot{P}_{reg}$ obtained from pursuing the same procedure on the regulation signal. Assuming these power gradients are again normally distributed, the required load-following reserve ramp capacity $\dot{C}_{lf}$ and regulation reserve ramp capacity $\dot{C}_{reg}$ are given by:

$$\dot{C}_{lf} = 3 \cdot \mathbf{stddev}\left(|\dot{P}_{lf}|\right) \tag{4.25}$$

$$\dot{C}_{reg} = 3 \cdot \mathbf{stddev}\left(|\dot{P}_{reg}|\right) \tag{4.26}$$

The process of determining the required load-following and regulation characteristics resources providing these services must exhibit is usually pursued on an hourly basis. In order for the load control strategy to be deemed suitable, it should result in a decrease in the required regulation reserve capacity, as this would result in an increase in the operating efficiency of the effected generation and displace possible needs to expand capacity under wind integration.

Another supply-side attribute of interest with respect to load control performance is the amount of wind power curtailed during the simulation. Along with obtaining these data through simulating the grid management strategy given by (4.17), the ramp-rate of the conventional generation $\dot{P}_G$ further provides a method of assessing how well the control strategy is able to mitigate wind-induced variability across the network bus, and can be obtained through finite-differencing the generator's trajectory:

$$\dot{P}_G[k] = \frac{P_G[k] - P_G[k-1]}{T} \tag{4.27}$$

4.3.2 Demand-side

On the demand-side, the characteristics of interest in terms of control performance relate to both the benefits achieved through equipping the targeted infrastructure with the proposed smart grid technology and the end-use services forfeited as a result. As each load is explicitly modelled, the cumulative energy demand from the generator E_G over the simulated time-horizon provides a method of observing the

end-use functionality of individual loads at the system-level, and is given by:

$$E_G[k] = \sum_{\hat{k}=0}^{k} P_G[\hat{k}]T \tag{4.28}$$

By comparing the E_G trajectory achieved subject to control to that seen uncontrolled, discrepancies should remain minimal throughout the time-series in order to draw conclusions that the control strategy was in fact non-invasive. Furthermore, the statistics of the control signal u sent to the individual populations can provide insight into both the cost and end-use disruption that accompanies the control process, and will be used to draw conclusions with respect to control performance.

As the different targeted load-types display different device-level qualities that will surely effect population dynamics, the performance of the each in comparison to the others can be achieved through observing the total energy dispatched over a given period in the optimization problem given by (4.17). Using the deviations from the uncontrolled trajectories, the total energy dispatched E_D from load-type p is taken to be equivalent to regulation in the sense that both up and down requirements are to be considered as dispatched capacity:

$$E_D^{(p)} = \sum_{k=1}^{N_k} \left|\Delta P^{(p)}[k]\right| T \tag{4.29}$$

The ramping capability of the load resource is also of keen interest in terms of the reserve requirements, and can be quantified through obtaining the width of the VGM feasible region given in (3.24). Hereafter this quality will be referred to as the load flexibility $\mathcal{F}$, and is given by:

$$\mathcal{F}^{(p)}[k] = P_{max}^{(p)}[k] - P_{min}^{(p)}[k] \tag{4.30}$$

Chapter 5

Model Results

The main objective of this section is to analyze the operating characteristics of the energy system model, and draw conclusions with respect to the observed performance of the proposed load management strategy.

5.1 Preliminary Model Results

In order to test the load control strategy, populations of the targeted demand-side infrastructure are required. The unresponsive load given in Fig.(4.6) will be used throughout the simulations. It is assumed that due to its peak demand (10 MW), that it could represent a suburban or rural community connected to a single distribution system (approximately 2000 customers) [56]. Assuming somewhat arbitrarily that three quarters of the customers have both a heat pump and EV, initially populations of 1500 heat pumps and 1500 EVs are generated.

To ensure realistic conditions, these populations must be heterogenous so that the device diversity expected in typical communities is observed. For the heat pump population, this is achieved through drawing normally distributed ETPs using the example given in Table (4.1) as the average values. For the EV population, a similar process is pursued, where the user's initial input, the charging capacity to be accumulated over their desired charging horizon, is drawn again from a normal distribution. Charging rates of 3, 4 or 5 kW are set by generating a uniform distribution over these three values. Complicating the EV model is the definition of the exact time the units should initialize charging. It can be expected that most EV drivers will deplete their battery during the day, and charge overnight, which would prove beneficial under

time-of-use electricity pricing scenarios (overnight would correspond to the cheapest customer electricity rates). Desired times to be fully charged are initially drawn from a uniform distribution spanning 5 - 9 AM, with the charge initialization times then correlated to these instances through drawing charging horizons uniformly between a range of 6 - 10 hours, so that the majority of charge initializations begin between 6 - 12 PM. As it can be expected that some EV drivers would also charge during the day, 50% of the vehicles (750 in this case) were also selected to partially charge over short periods during the day, at smaller required charge capacities (selected as approximately 50% of their nightly charging capacity). All other parameters in the EV population are left as is given in Table (4.2). In total, the EV population charges 22.6 MWh of energy from the grid, which can then be used to displace gasoline for transportation. Assuming an average charge depletion mode efficiency of 0.25 kWh/km [56], and a discharge efficiency of 95%, this amount of EV charging energy corresponds to an average of approximately 60 km per customer, more range than an average driver would require.

As the electrolyzers consume more power per unit than the heat pumps and EVs, a smaller population of 350 electrolyzers is initially implemented. Their rated power was set uniformly over the range 10, 15, or 20 kW. The time requiring full charge was set through defining a normal distribution centred about 24 hours. This choice avoids having each individual unit traversing a very similar deadband trajectory, or more particularly, the population synchronizing within the deadband space. The required hydrogen capacity was then defined so that each system would display a 50% duty cycle, which corresponds to setting $E_c = \eta P_r T_s/2$. All other parameters in the electrolyzer population are left as is given in Table (4.2). In total the electrolyzers produce a daily quota of 48 MWh of distributed local hydrogen feedstock (based on the lower heating value), which could be used in various fuel cell applications. Fuel cells achieve a wide arrange in operating efficiency, especially in applications such as a vehicle engine in which a variety of loading conditions are expected. In these scenarios the system-level fuel cell efficiency typically attains approximately 50% [57], which means that the locally produced hydrogen could provide about 24 MWh of energy for similar stationary or transport-based requirements.

Once the model inputs were defined, two separate scenarios were considered, with the first simulating the system without access to the load management resource, and the second with the load management strategy enabled. The temperature profile input to the building models is given in the top plot of Fig.(4.2). The scenarios

result in the uncontrolled and controlled load trajectories given in Fig.(5.1). The top

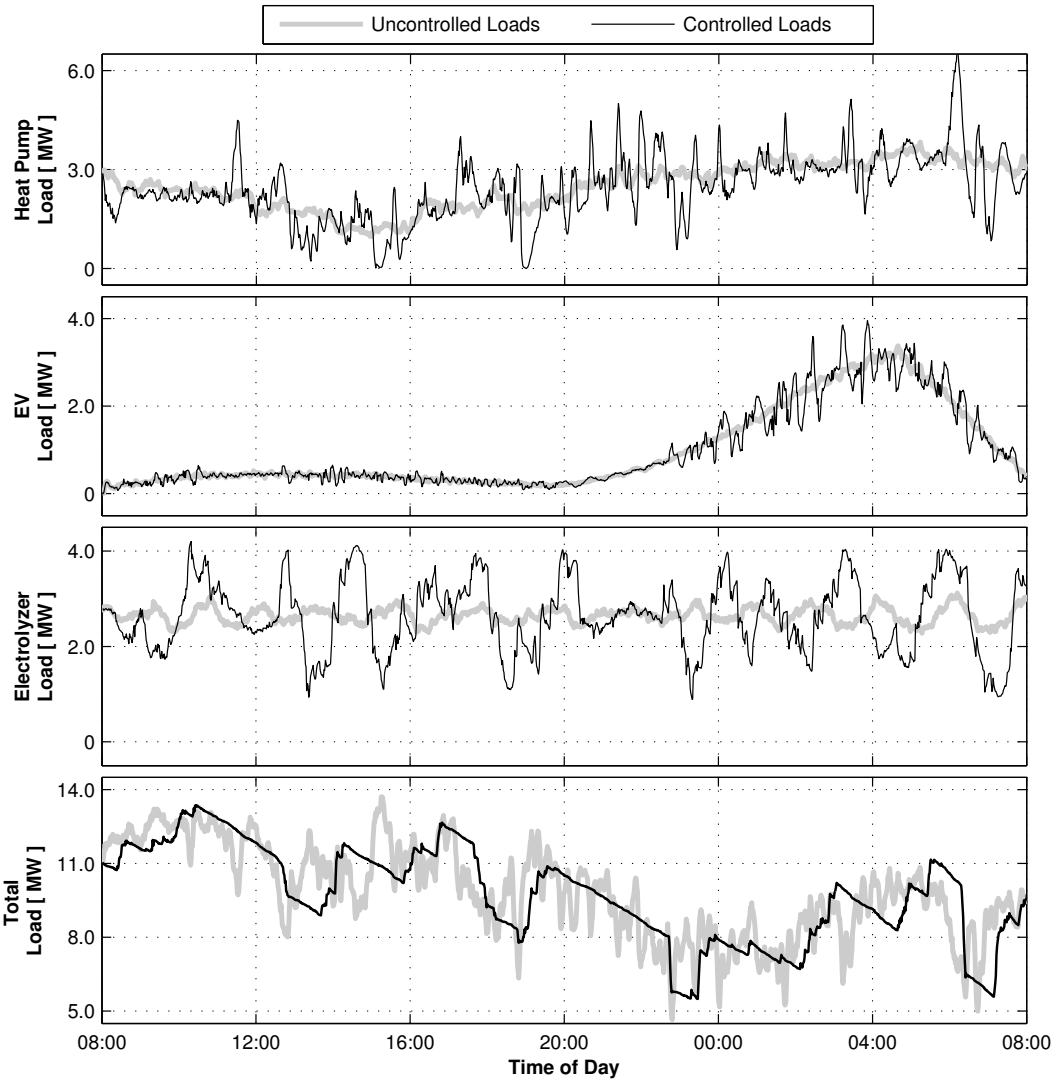

Figure 5.1: Load trajectories associated with the preliminary simulation.

plot displays the aggregate heat pump load, the second plot the aggregate EV load, the third plot the aggregate electrolyzer load, with the bottom plot indicating the total load on the network bus (including wind input). The total load is input to the grid management problem in (4.17), which in this preliminary case, does not result in any power curtailed from the network bus, meaning the bottom plot of Fig.(5.1) also represents the load on the conventional dispatchable generators. As can be seen, the controlled case displays considerably less variability. The heat pump trajectory is observed to display the greatest ability to ramp up and down over short-periods, while

the electrolyzer population displays oscillatory behaviour occurring over longer time-scales. These characteristics are directly attributable to the device-level duty-cycles that accompany the choice of deadband width, as the shorter duty-cycles observed in the heat pump and EV populations result in accelerated replenishment of regions near state-transition boundaries, while the oscillations observed in the electrolyzer population, even in the uncontrolled case, seem to indicate that the population's states may be synchronizing (many individual electrolyzers existing within a similar region in the deadband space) near the end of the simulated horizon. These aspects will be further investigated in the following sections.

5.1.1 Population Dynamics

To further investigate population dynamics, the individual responsive load trajectories are examined in greater detail over the hour 01:00 to 02:00, and are displayed in Fig.(5.2). The top plot depicts the trajectory associated with the VGM, which encompasses the aggregate demand from each responsive load type. Also depicted is the target trajectory defined in (3.34), as well as the time-varying feasible region attained through online assessment of the demand resource, determined through combining the device-level boundaries given by (3.24). Beneath the total responsive load trajectory are the individual responsive load-trajectories, computed by (3.10) for each type, as well as the associated time-varying feasible region boundaries, and target trajectories computed through solution to the optimization problem in (3.32). The targets are observed to be followed at an accuracy exceeding 98.5% over the entire simulation, where these minute discrepancies can be attributed to the discretization of the deadband width and quantized device-level power-ratings. Over certain intervals the trajectory associated with the total VGM is observed to bottom-out against the minimum feasible region boundary. These occurrences indicate that the regulation target defined in (3.33) is unattainable in the sense that the loads are unable to meet the target without compromising the feasible region, and thereby reduces the capability of the responsive loads to regulate the total demand, and results in the need to dispatch conventional regulation resources. Therefore, while the electric load populations may reduce the regulation requirements, it can be expected that they will be unable to fully replace generating units traditionally participating.

Delving deeper into the population dynamics, the power density distributions are further tracked over the same interval, and are depicted graphically as contour plots

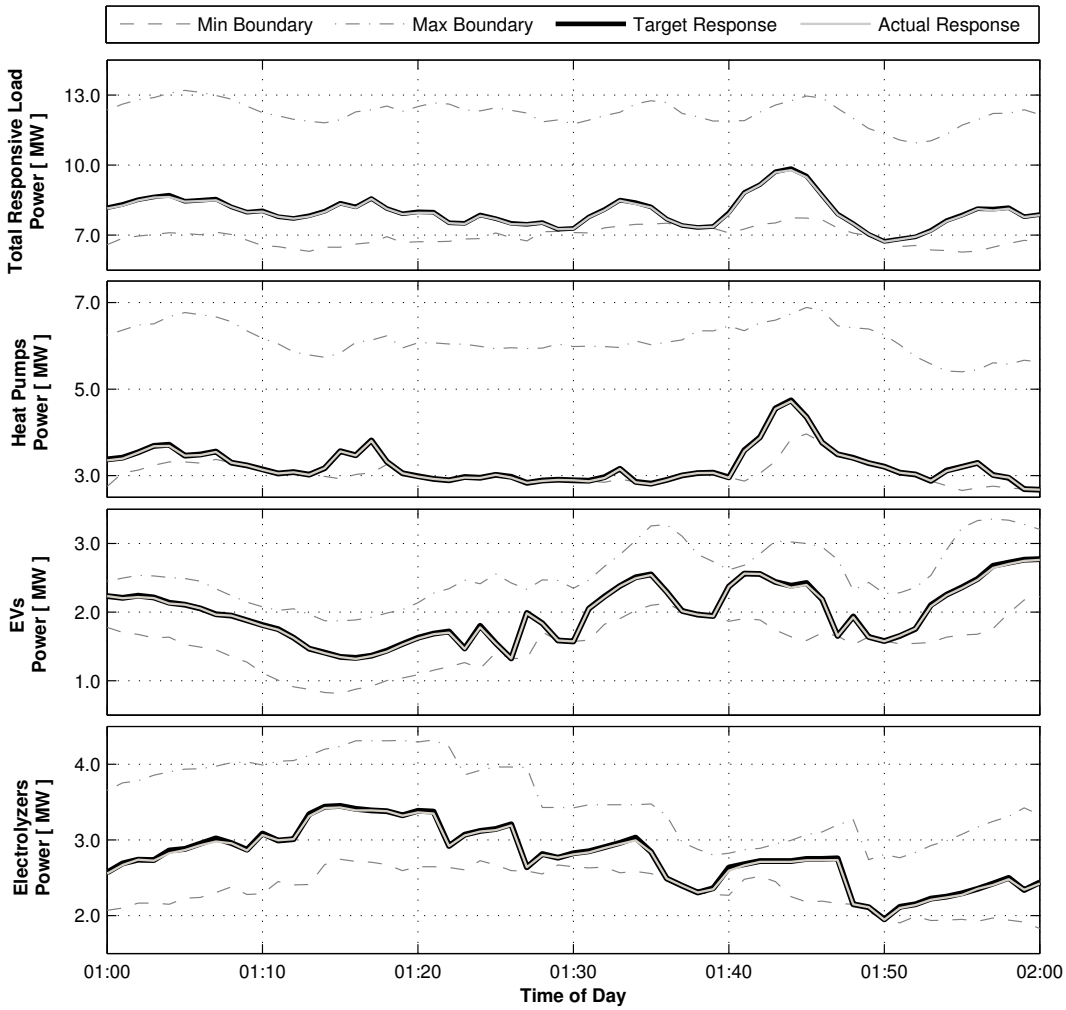

Figure 5.2: Close-up view of the virtual generator responses over the period 01:00 to 02:00.

in Fig.(5.3). The amount of power at each individual end-use state is shown over the discretized 2δ-region encompassing all end-use state measurements. Also depicted are the controlled state-transition boundaries computed upon solution to (3.20) under the target power allocated to the particular load type. Apparent state-transition

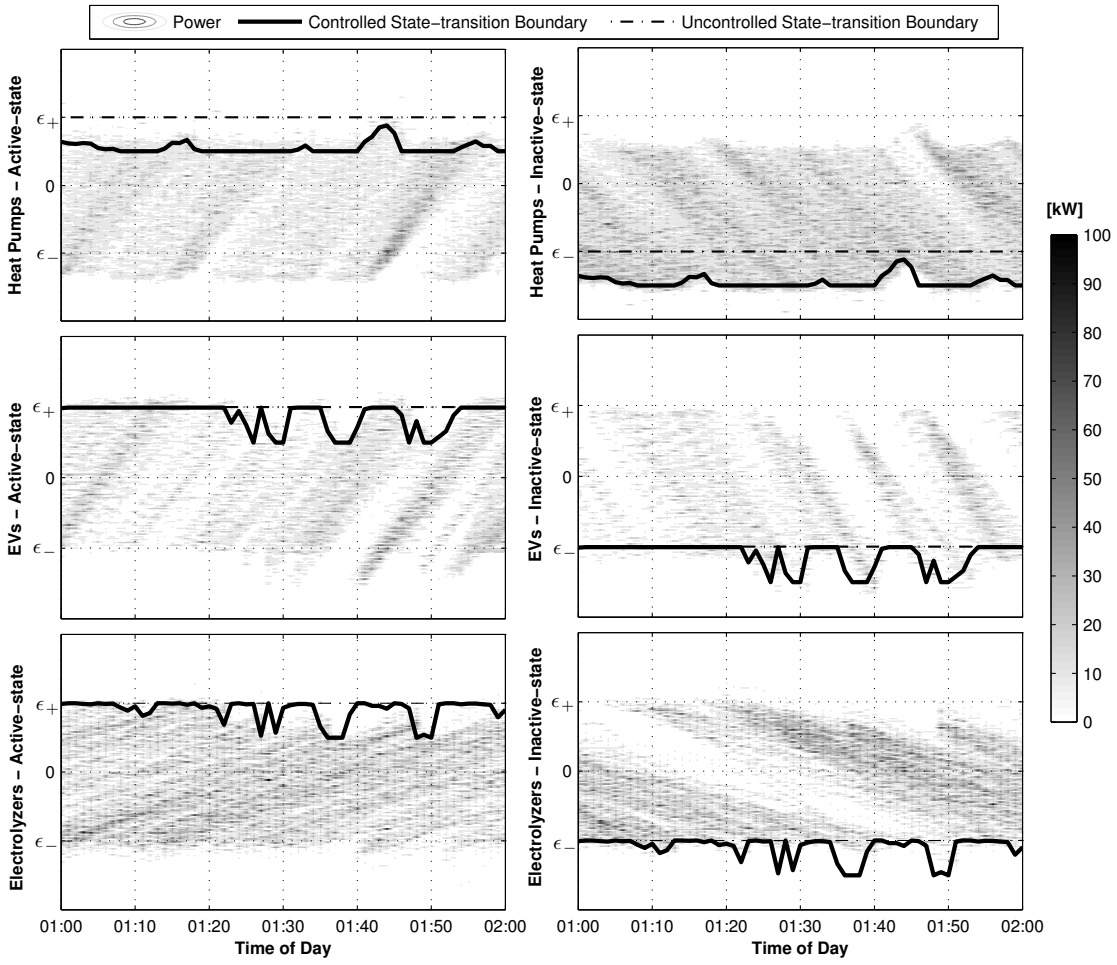

Figure 5.3: LC power distributions from 01:00 to 02:00.

boundary overshoots are a result of devices in the preceding interval diffusing into the state-transitioning regions, which are then the focus of current time-step management. It can be seen that the intervals over which the total responsive load bottomed-out in Fig.(5.2) correspond to events in which the control signal also lies along its minimal level in Fig.(5.3), as set in (3.22) so as to ensure the control remains non-disruptive at the device-level. Further noteworthy are the relative differences observed in the evolution of the distribution shapes. The heat pump population's distribution remains relatively uniform. Device-level state synchronization is observed to occur

between 01:40 and 01:50, but in turn dissipates quickly, and indicates the heat pump population is capable of maintaining diversity.

On the other hand, the EV and electrolyzer distributions display considerable non-uniformity, as well as intervals over which considerable state synchronization is observed. Even though the EV population is operating on a deadband width that should promote accelerated replenishment of the deadband transitioning regions, the dynamic population size that accompanies the connection / disconnection of units, as well as the relative distribution sparsity, results in less flexibility in the aggregate population or a need to implement large amplitude control signals to achieve any sort of response.

While the electrolyzer population size remains static, the synchronization observed dissipates slowly, partly due to slower device trajectories accompanying a larger deadband width, but also further indicates the population in fact displays less of the heterogeneity that is needed to promote state diversity. The combined ability to quickly damp disturbances in state diversity and maintain distribution uniformity allows the heat pump VGM to display the greatest flexibility, or more particularly, that smaller set-point changes result in greater overall ability to move power. As can be seen in Fig.(5.2), the heat pump VGM consistently displays the largest feasible region, with boundaries that in turn remain relatively stable. These characteristics promote selection of the heat pump population over its counterparts within the optimal dispatch process in (3.32), as it looks to minimize the cost of attaining the desired demand response across each load-type through minimization of the control signals. Overall, through consistently dispatching the heat pump population, the heat pump distributions have been completely shifted downwards towards the lower state-transition boundary (ϵ_{-}). Nonetheless, the device-level mean temperature has only shifted about -0.25^{o}C, and hence should go unnoticed by the end-user.

The total dispatched capacities over the simulated time-horizon, mean VGM flexibilities, mean control signal amplitudes, and control signal standard deviations are given in Tab.(5.1). The heat pump population is dispatched for the most energy, as computed by (4.29), which can be directly attributed to the flexibility the population consistently displays, as computed by (4.30). The superior flexibility results in a lower cost associated with energy dispatched from the heat pumps, as the aggregate heat pump demand is more sensitive to set-point modulations. Therefore the heat pump VGM is selected ahead of its counterparts due to a reduced cost within the viewpoint of the optimization in (3.32).

While these results may indicate that the heat pump population is indeed the superior load-type, they must be accepted with caution. The first-order linear load models pursued in the case of EVs and electrolyzers determines the rate at which the devices traverse the deadband region, which in turn governs the evolution of the power distributions. These models do not capture load dynamics sufficiently. In contrast, models that preserve the transient phenomena associated with the electrochemical energy conversion processes should enhance heterogeneity within the load populations, thereby increasing device-level flexibility in the case of EVs and electrolyzers. Furthermore, the EV charging strategy has been defined somewhat arbitrarily, and it may in fact be expected that more units are available throughout the day, which would in turn increase aggregate flexibility. Also, in many locations, heat pumps will only be operated seasonally, and therefore may become unavailable for control. For these reasons, while the EVs and electrolyzers may in fact be outperformed by the heat pumps in this preliminary simulation, under realistic conditions EVs and electrolyzers should perform somewhat better, and the heat pumps somewhat worse. It will be the focus of Sec.(5.2) to investigate results sensitivity to certain initial conditions.

Table 5.1: Performance of the load resource.

Load Type	Dispatched Capacity [MWh]	Mean VGM Flexibility [MW]	Mean Control Signal Amplitude [% of u_{max}]	Control Signal Standard Deviation [% of u_{max}]
Heat Pumps	35.6	3.2	79.9	30.3
EV Charging	2.8	0.5	31.5	42.7
Electrolyzers	5.2	1.3	39.4	43.5

5.1.2 Network Bus Power Dynamics

The power dynamics across the network bus can be examined through computing the power gradient and cumulative energy demand given by (4.27) and (4.28), using the total load given in Fig.(5.1). The power gradient data can be ordered by occurrence frequency to observe the data's spectrum, and is given for each scenario (controlled or uncontrolled loads) in the left plot of Fig.(5.4). The controlled case is observed to act as a low pass filter, as the majority of the larger fluctuations seen in the uncontrolled

case are eliminated when the loads are subject to control. This result clearly indicates that through synchronization of deferrable device-level schedules, the proposed control strategy is capable of buffering variability across the network bus. Considering that the emission rate of thermal generation increases with power gradient, if these types of generating units are servicing the network bus, the controlled case will result in a reduction in emissions. Furthermore, ramping of any generation unit can degrade performance reliability over time. The controlled case mitigates this variability, and therefore the units servicing the network bus are able to operate steadily and smoothly, prolonging these unit's lifecycle and reducing maintenance costs.

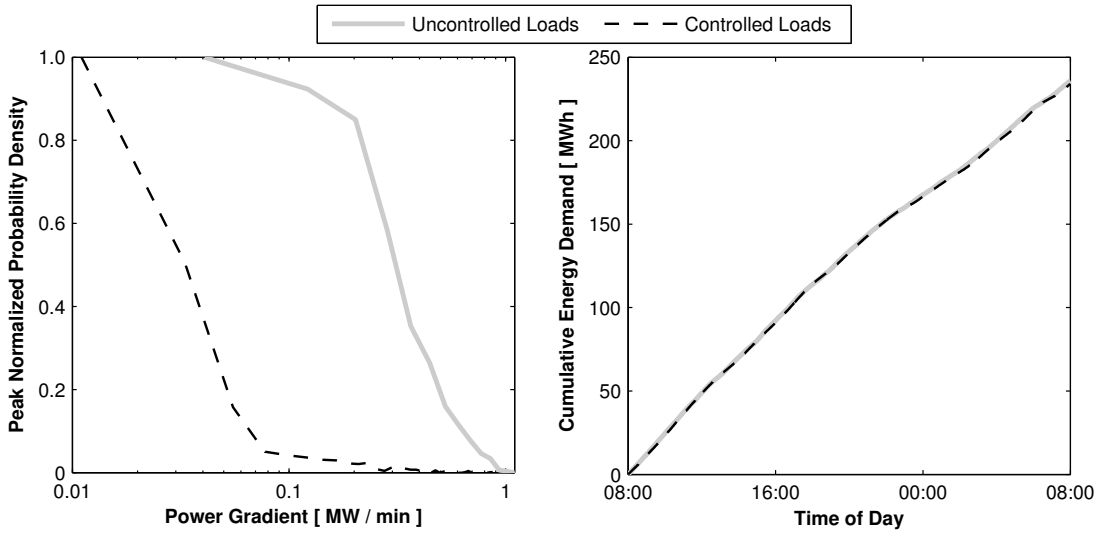

Figure 5.4: Left: power gradient experienced by the generator. Right: cumulative demand for energy from the generator.

The cumulative demand for energy is given in the right plot of (5.4). As can be seen the trajectories overlap over the entire simulated time-horizon, meaning that while we obtain the variability buffering seen in the left panel of Fig.(5.4), the demand for energy in the controlled and uncontrolled case remains equivalent. This result underlines that each modelled end-use device, controlled or not, consistently demands the same amount of energy throughout the day, and as this energy demand appears as an equivalent level of end-service, it therefore guarantees that end-use functionality of the participating units is maintained when subject to the proposed control strategy.

5.1.3 Effect on Conventional Generation Scheduling

While a reduction in power gradient has been observed across the network bus, which in turn should allow conventional generation to operate more efficiently, further analysis is needed to understand effects on power system operations. Initially, the regulation component is removed from the total load on the conventional generation by using (4.22) with the trajectories given in the bottom panel of Fig.(5.1). The result is given in Fig.(5.5), and as can be seen, the case in which the load control straegy has been enabled results in considerably less need to dispatch regulation. In total,

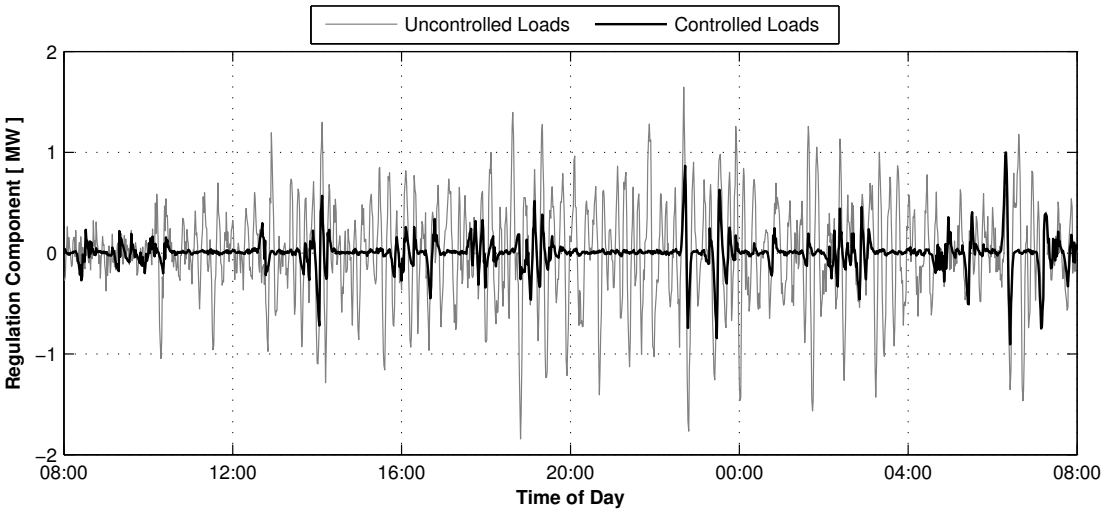

Figure 5.5: Regulation component that must be accommodated by the conventional generation

the uncontrolled case requires 8.8 MWh of regulation to be dispatched, while the controlled case requires only 1.8 MWh (80% decrease).

The required hourly online load-following and regulation reserve characteristics can be computed through application of Fig.(4.23) to (4.26), with the results displayed graphically in Fig.(5.6). Over each hour, the controlled case is observed to reduce both the required regulation capacity and ramp-rate. The load-following reserves are seen to in fact increase over certain hours of the day, although considering the cumulative energy demand given in Fig.(5.4), these increases can be met with the equivalent capacity liberated from the regulation reserves, with the major benefit now being an increase in the conventional generation's energy utilization efficiency. The reduction in dispatched regulation capacity also results in a greater capability of the affected units to provide actual capacity to the grid, as the units operating as online reserve

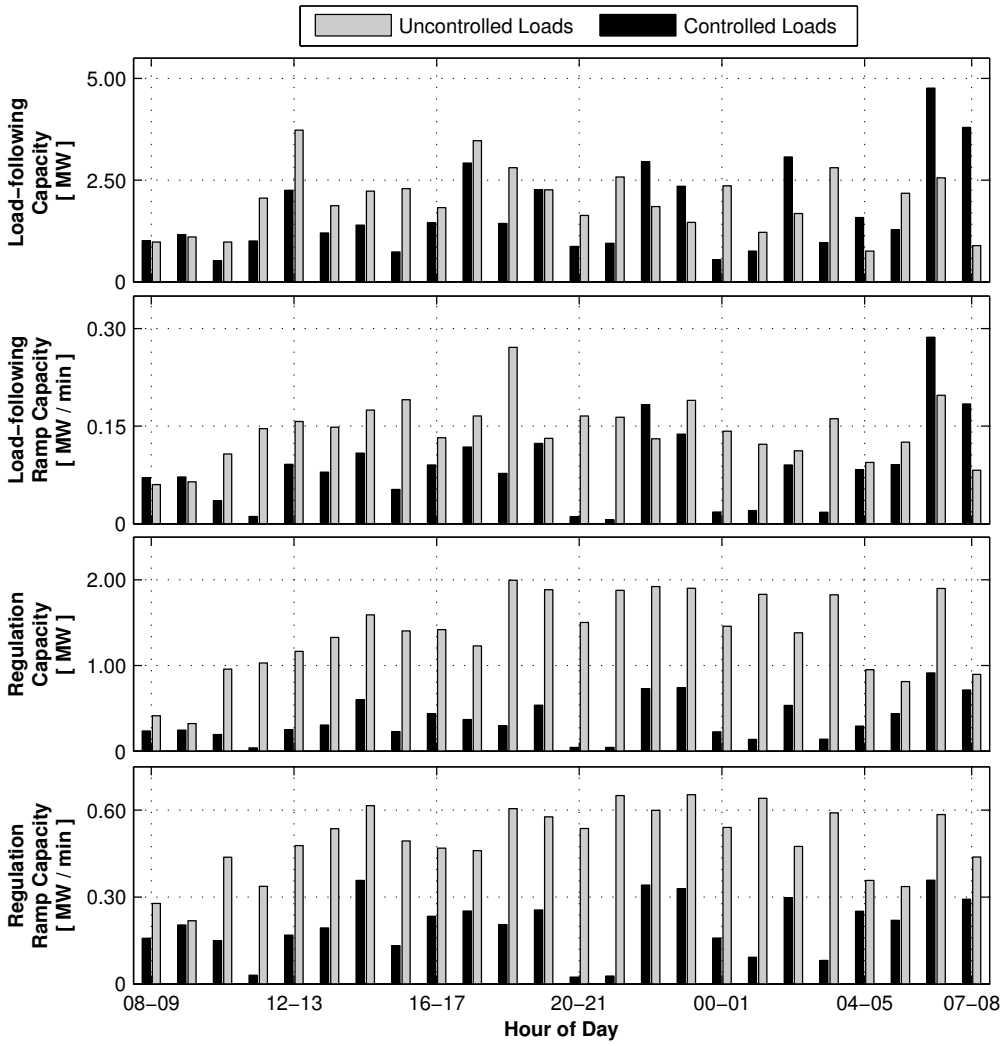

Figure 5.6: Required hourly online load-following and regulation reserves.

have associated missed opportunity capacity that must remain online but does not actually make it onto the grid. Through subtracting the hourly regulation reserve capacity from the hourly dispatched capacity, it was found that in the controlled case 24.2 MWh of missed opportunity capacity occurred, with only 6.9 MWh occurring in the controlled case. Therefore, the proposed load control strategy has essentially freed 17.3 MWh of energy that could take part in ED.

To understand what the preliminary results mean in terms of the operating characteristics of the dispatchable generation, the required load-following and regulation reserves can be combined using (2.1) to understand both the required operating range and ramp-capacity the participating units must display in order to remain viable choices for reserve. The results of this process are displayed in Fig.(5.7) for both the controlled and uncontrolled cases, along with the hourly dispatched capacity from each load type and corresponding flexibility achieved when subject to the proposed control strategy. The required ramping duty of the dispatchable generation is seen to decrease over each hour when loads are subject to control, which means that slower moving resources can be implemented. Over the majority of the day, the operating range of the conventional dispatchable generation in the controlled case is reduced, even over hours in which the required load-following characteristics were in fact observed to increase. This can be attributed to the structure of (2.1), as the operating range is twice as sensitive to the regulation reserve reductions achieved in the controlled case as compared to changes in load following reserves. A smaller operating range means that participating generators can be continuously operated closer to ideal conditions, or near loading conditions that correspond to maximum output efficiency, thereby reducing the need to operate multiple smaller units in order to conserve operating efficiency. Over the final two hours of the day, the operating range is in fact seen to increase in the controlled case, which in turn corresponds to the interval over which the electrolyzer population is observed to display oscillatory behaviour. As this attribute is brought about by synchronization of device-level states, it underlines the importance of ensuring control stability through maintaining load diversity.

The bottom two panels of Fig.(5.7) display the hourly dispatched capacity and flexibility of each responsive load population. The loads are seen to be dispatched for more energy each hour than the observed reductions in regulation achieved. Therefore, these results indicate that it would be unsuitable to set the cost of ancillary services generated by responsive loads at the same rate as power from from more traditional reserve capacity generators. Controlling loads intrinsically changes the

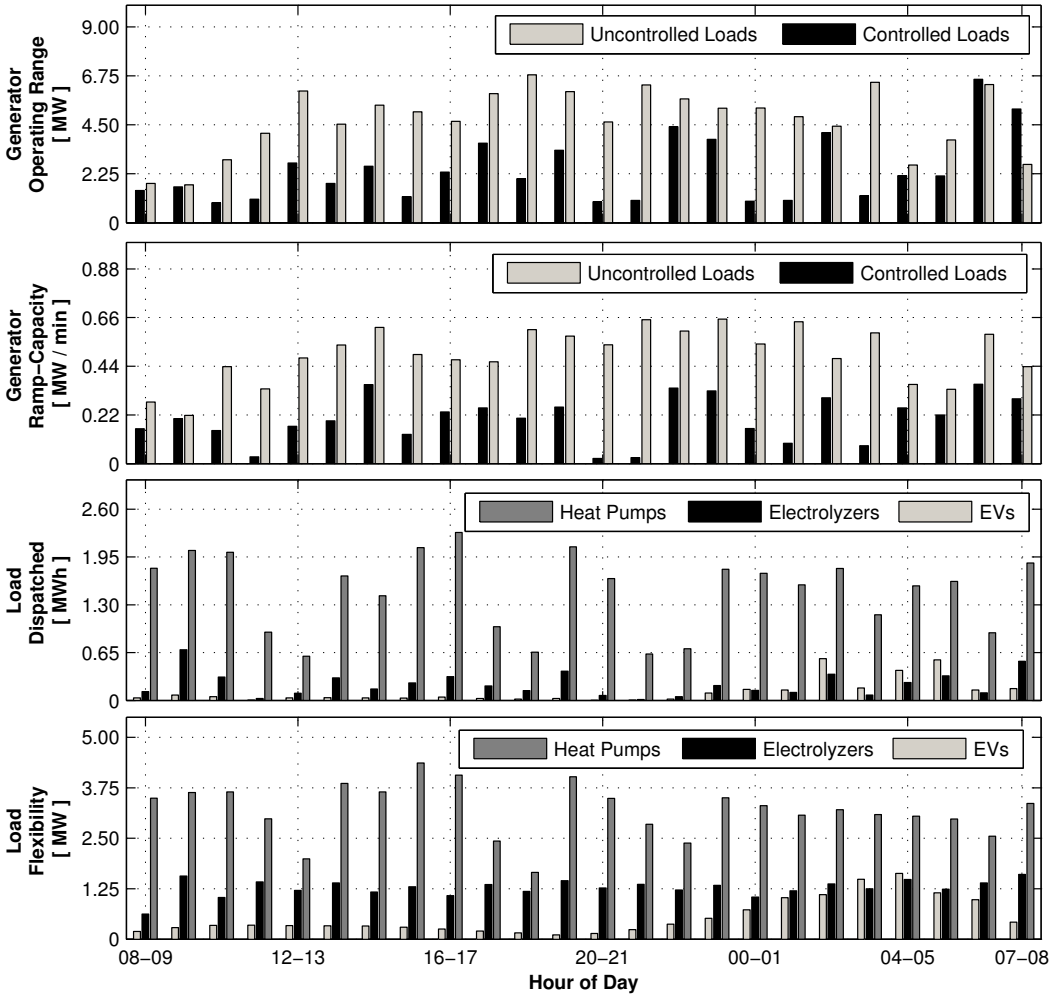

Figure 5.7: Operating characteristics of the dispatchable capacity.

fundamental demand profile, which is used as the datum from which deflections are considered dispatched capacity from the LC, as given by (3.26). Pricing services attained from the loads based on this deflection under the idea that the uncontrolled trajectory represents the ideal uncontrolled case from the viewpoint of the LC is therefore unsuitable. Once load control commences the uncontrolled trajectory may in fact result in sub-optimal conditions as direct result of a previous control action. Under these scenarios it would make most sense from the viewpoint of the load management network's stakeholders to pursue control strategies so as to ensure the long-term performance of the resource. A single stakeholder (the LA) should then not be faced with the prospect of paying another stakeholder (the LC) during pursuit of a mutually beneficial action. It is therefore argued here that it does not make sense to attempt to price direct load control actions in real-time. Rather the LA should pay participating load groups based on the positive effects they add to the system post-occurrence. These attributes could be obtained by quantifying the decreases in online reserve requirements and emissions, as well as any efficiency gains.

5.2 Sensitivity Study

The computational model is modular and extensible, and can further be used to understand the sensitivity of preliminary results to initial conditions. Here, each load-type is investigated under varying model initializations in order to gain further insight into the performance of the load management strategy. The effects of the recruited population-size on load control performance are of particular interest . Recruitment of each load-type can be modelled through generating populations of varying sizes, and then simulating these groups both controlled and uncontrolled, under model inputs that are expected to vary considerably across diverse populations.

The hourly regulation reserve requirements (capacity and ramp-rate) obtained in each case (controlled and uncontrolled) can be averaged over the simulated time-horizon, and then subtracted to quantify the reserves liberated when deferrable loads are subject to self-regulating control. Obtaining these values across different population-sizes and load-types therefore provides a method of determining the marginal benefit associated with recruiting these targeted loads through the application of the smart grid technology, and provides valuable insight into the least-cost population-basis that should be sought by the LA in order to provide the greatest overall benefits.

5.2.1 Heat Pump Integration

Different heat pump population sizes are generated using the parameters given in Tab.(4.1). In order to reflect the large decrease in demand accompanying removing the other load-types from the system, the wind power data are scaled down to represent a 2.5 MW unit. The initial input parameter investigated is the number of terms in the aggregate controller's moving average horizon M, as defined in (3.33). Selection of M sets how intensely the controller attempts to smooth the load profile, or more importantly, to operate in the load-following regime. A population performing well at a large value of M therefore reduces load-following reserves, increasing base-load requirements proportionally, which in turn would further provide considerable benefits to the power system by decreasing the gap between demand peaks and valleys. This allows the power system to operate more frequently at the loading levels it is designed to handle. The problem is that more flexibility from the load populations would be required, which may be incompatible with the comfort-constraints embedded in the optimal dispatch procedure.

Figure (5.8) graphically depicts the marginal benefit of recruiting heat pumps (decremental regulation reserve capacity and ramping capacity) over the choice of M = 2,3, and 4. It is initially observed that the marginal benefits associated with recruiting heat pumps diminishes after approximately 1000 units exist in the responsive population. As can be seen for the decremental regulation reserve capacity in the left panel, the initial $M = 2$ case performs somewhat better for population sizes less than approximately 1000 units, with the $M = 4$ case overtaking at larger levels. As these larger population sizes contain more thermal mass, more energy is able to accumulate within the population without significantly affecting the temperature. This allows the controller to operate more effectively in the load-following regime, and is why the larger values for M perform better at larger population sizes. A large amount of thermal mass is therefore seen as a positive building-level characteristic for heat pump loads, and buildings possessing this quality should be sought for recruitment over those that may lack thermal mass in comparison. Nonetheless, it is more beneficial to recruit a population of smaller buildings with a total thermal mass equivalent to a single larger building, as a larger population enhances the power density distribution diversity.

Building envelope thermal integrity is also an important attribute to consider when recruiting customers. Buildings with lower envelope resistance traverse the deadband

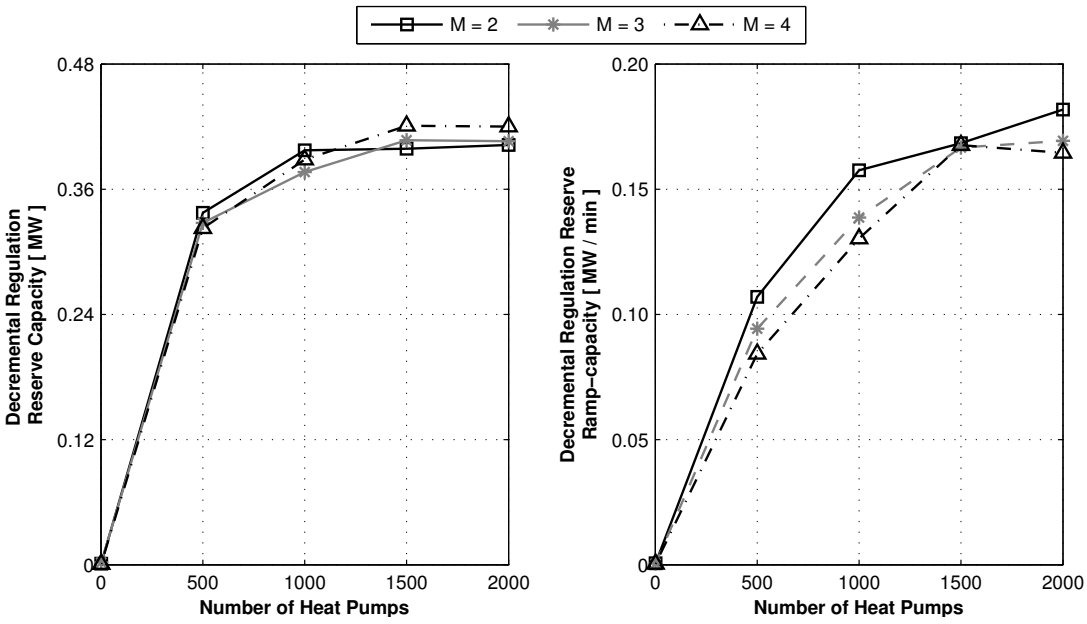

Figure 5.8: Effect of heat pump population size and selection of M.

at a greater rate, and therefore, it can be expected that buildings that have poor insulation in fact represent ideal candidates for recruitment, as these populations will in turn result in a VGM with the greatest flexibility (deadband transitioning regions replenish quicker). While intriguing in terms of the proposed load management strategy, it should not be an issue of promotion amongst community residents, as this would in turn result in a greater steady-state demand for energy from the buildings. This may offset the benefits associated with the load management strategy and further promote inefficient use of energy at the building-level.

In the right panel of Fig.(5.8), the decremental regulation reserve ramp-capacity is given, where the $M = 2$ case performs better over all population-sizes. This result is to be expected, as smaller values for M result in the primary controller objective shifting towards the minimization of the power gradient across the network bus, while larger values focus on longer time-scale smoothing.

Next, the effect of outdoor temperature regime on control performance is investigated. The outdoor temperature regime drives the apparent operational duty-cycle of the heat pump, as the gradient that it places across the building envelope governs the rate at which heat is lost from the air. In terms of load control performance, a 50% duty cycle should be ideal, as it promotes a steady-state aggregate demand that provides an equal opportunity to ramp-up or down. To check the effects of outdoor temperature regime on the online regulation reserve requirements, the initial temperature input given in Fig.(4.2) is shifted by $\pm 5^o$C, resulting in three temperature profiles (mean of 3.5^oC, 8.5^oC and 13.5^oC) that display similar trends in temperature gradient. The wind plant is kept at a 2.5 MW capacity, and $M = 2$ is again implemented. Each temperature profile is input to the varying population sizes, and the regulation reserves computed, and are given in Fig.(5.9). There is no clear trend observed in the reserve requirements displayed in the left panel, though the colder temperature regime is observed to result in a better ramping performance in the right panel. Considering that the heat pumps are sized at a design temperature of -5^oC, and that the efficiency of these units have been modelled as a non-linear function of temperature, it can be expected that the lowest regime indeed results in an average duty-cycle lying closest to 50%, and therefore provides a greater ability to ramp the resource up and down. A 50% duty-cycle corresponds to air temperature trajectories that dissipate equally in either the active or inactive-state, and results in similar attributes emerging at the system-level as well. Namely, the power density distributions dissipate equally in either direction, replenishing the deadband transitioning regions

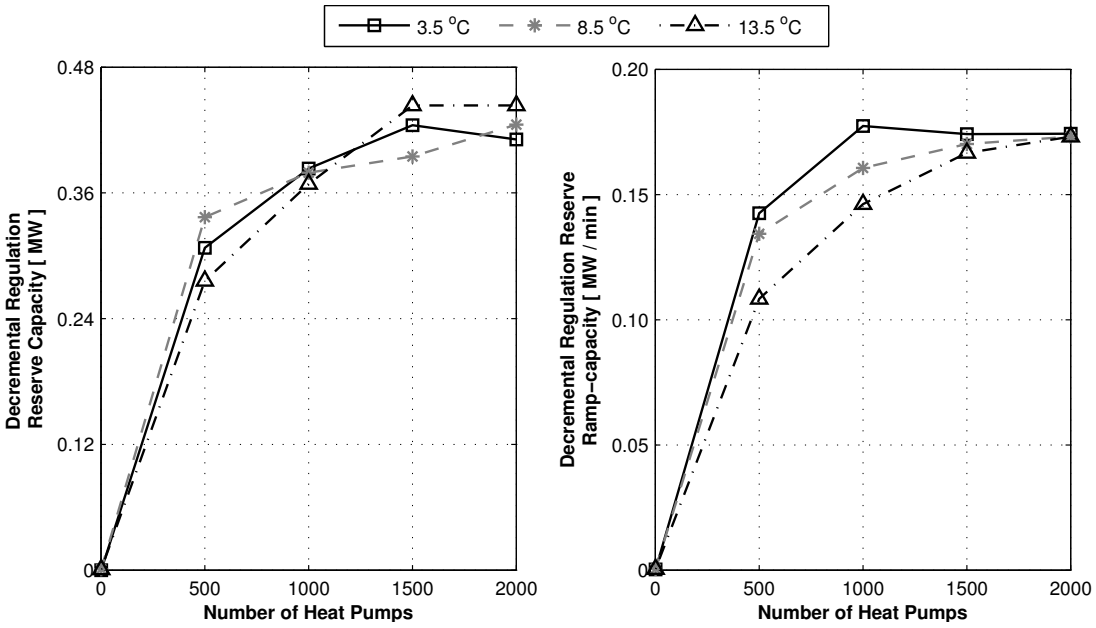

Figure 5.9: Effect of heat pump population size and outdoor temperature regime.

similarly and promoting overall diversity. This enhances the population's flexibility, or its effective ramping ability and not its capability to store energy, and is the main reason why the colder regime is observed to perform better in the case of online ramping capacity and that no clear trend is observed in the reserve capacity.

5.2.2 EV Integration

The effects of EV integration are now checked over different initial conditions. Firstly, selection of M is again checked, with results given in Fig.(5.10). It is initially observed that the EVs do not provide the same level of benefits as that observed in the case of heat pumps. As the predefined charging schedule results in EVs being disconnected from the grid for the majority of the day, they are likewise unavailable for recruitment. This results in a greater need to implement conventional regulation resources, although the extra benefit now being that a reduced capacity is consumed overall by the EV population in comparison to the other load-types. The selection of

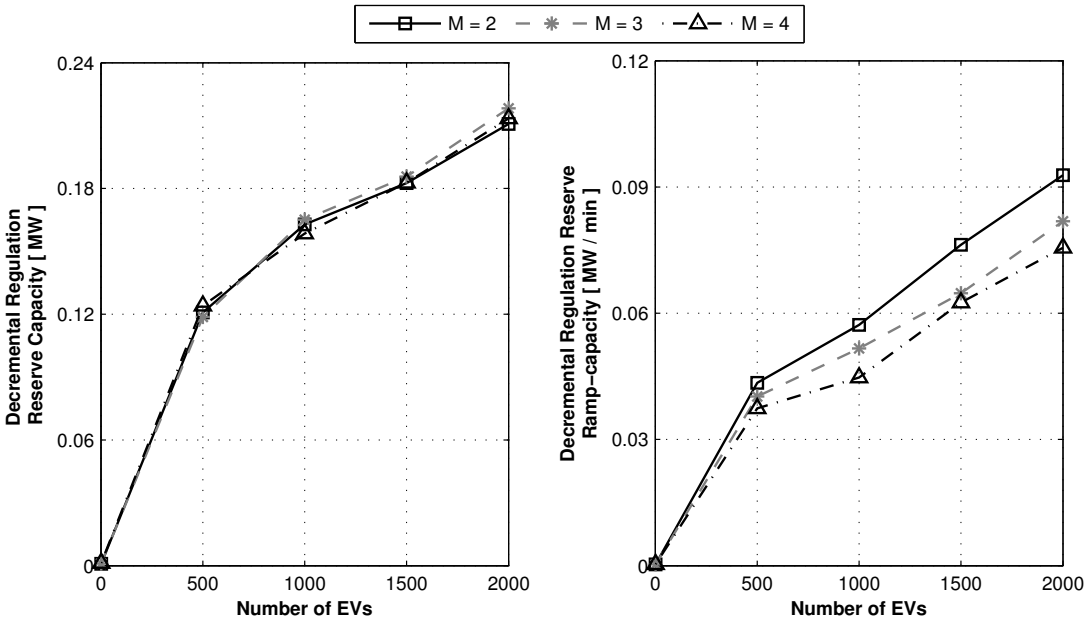

Figure 5.10: Effect of EV population size and selection of M.

M is observed to have little effect on the regulation capacity, while the ramp-capacity is derceased most in the $M = 2$ case, again due to the fact that the controller is attempting to minimize power gradients most under this selection of M.

As the charging schedule has been defined somewhat arbitrarily, resulting in low availability of grid-connected vehicles during the day, the sensitivity of the results

with respect to the percentage of vehicles that also charge during the day is now investigated. The initial 50% level is checked, along with levels of 75% and 100%, with the results given in Fig.(5.11). The results indicate that increased access to EVs during the day can provide moderate increases in control performance. While availability may increase with more EVs charging during the day, the extra variability resulting from an increased number of units connecting and disconnecting from the grid, as well as the induced duty-cycle based charging, in fact offsets some of the expected positive benefits.

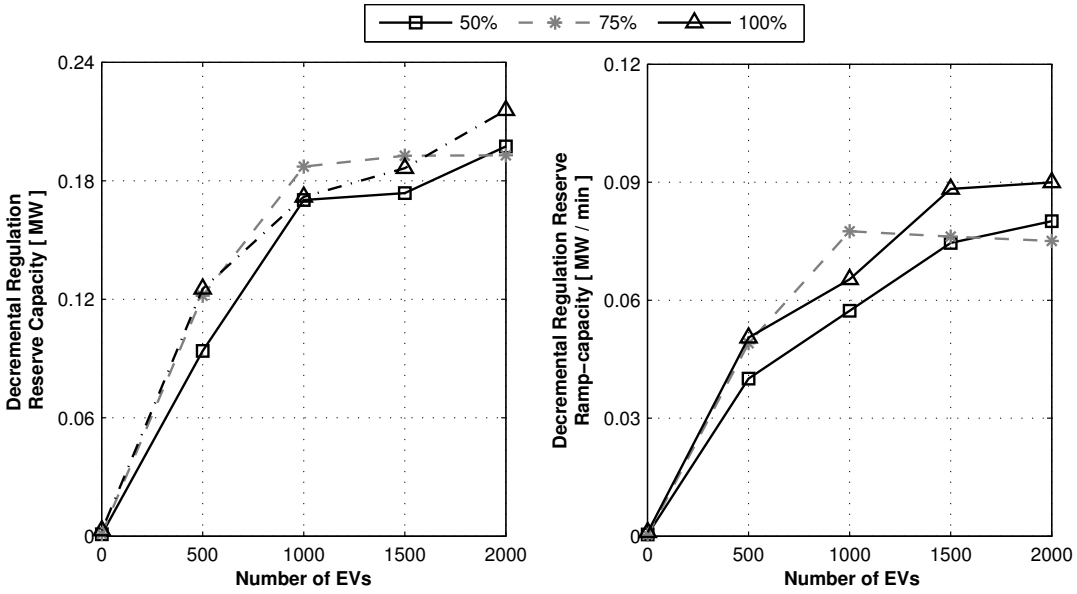

Figure 5.11: Effect of EV population size and number of day-charging events.

Finally, sensitivity with respect to deadband width δ is investigated. The choice of δ drives the duty-cycle frequency. Depending on the operational constraints of the charger itself, which would bear the brunt of any cycling induced degradation, EVs may in fact embody the greatest flexibility in terms of the selection of δ. There may be room to manipulate this parameter considerably, whereas the other load-types of interest contain definite constraints in terms of cycling frequency (electrolyzers and heat pumps have minimum run/down-times). Values of $\delta = 0.025, 0.050$ and 0.075 are checked (unitless due to the non-dimensionalization of the end-use state performed in equation (3.9)), with results displayed in Fig.(5.12). Little effect is observed on the regulation capacity, while the smallest choice of δ results in the greatest reduction in ramping rate. A smaller deadband space effectively results in EV loads traversing

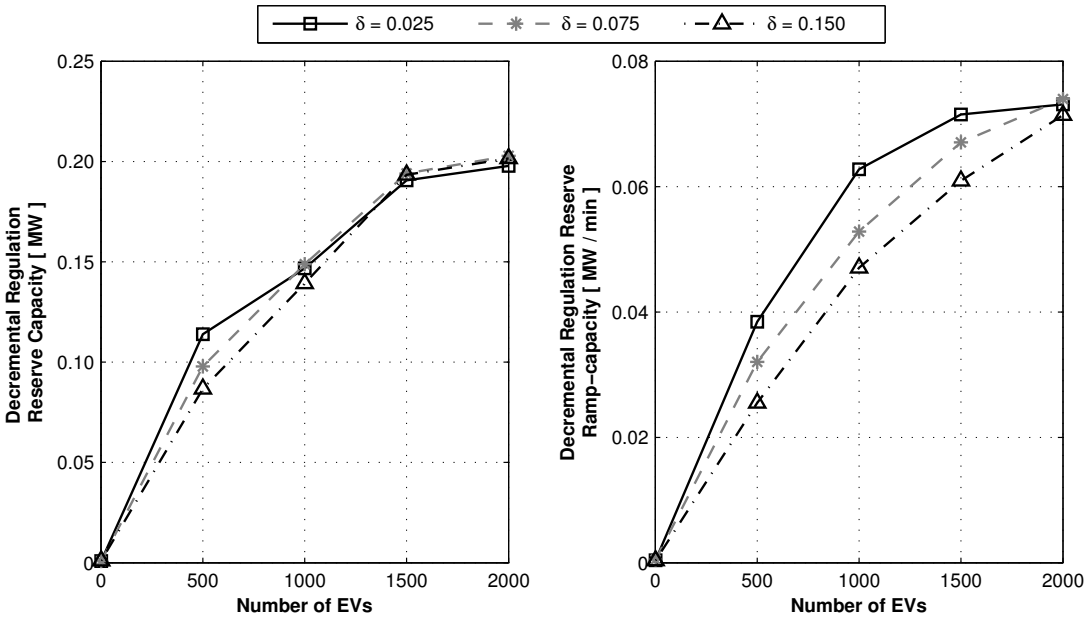

Figure 5.12: Effect of EV population size and width of deadband.

the space at a greater rate, or more particularly, accelerates replenishment of the deadband transitioning regions, and therefore increases the ramping capability of the VGM.

5.2.3 Electrolyzer Integration

Sensitivity to initial conditions for the electrolyzer population are now investigated. Three values of M are assessed over a range in population size in Fig.(5.13). The electrolyzers consume more energy per device, and therefore a smaller range in population size is considered. The electrolyzer population provides similar overall benefits as the heat pumps, with marginal benefits increasingly rapidly up to a population size of approximately 500 units, whereupon these benefits saturate. Similar conclu-

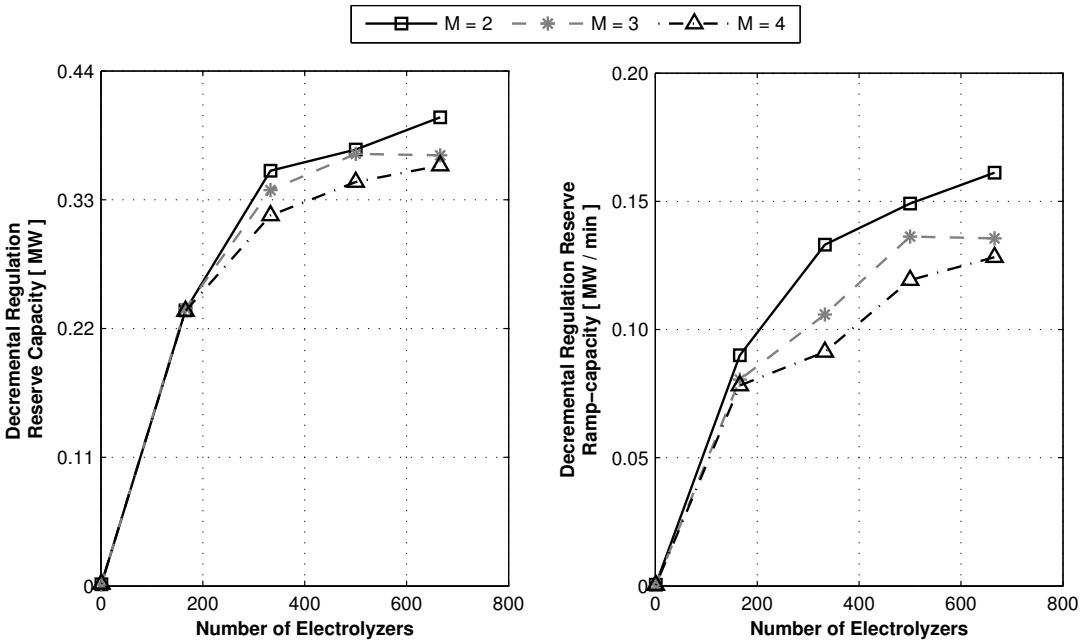

Figure 5.13: Effect of electrolyzer population size and selection of M.

sions apply to M: the selection of $M = 2$ also performs best in the regulation reserve capacity case.

The electrolyzer population yielded the greatest problems in terms of maintaining diversity within the deadband space; thus the standard deviation of the device-level measurement noise σ_n was modulated in order to model enhanced diversity. Results are given in Fig.(5.14), where the largest selection of σ_n performs best. Increasing the measurement noise pushes the device-level electrolyzer model to deviate from its

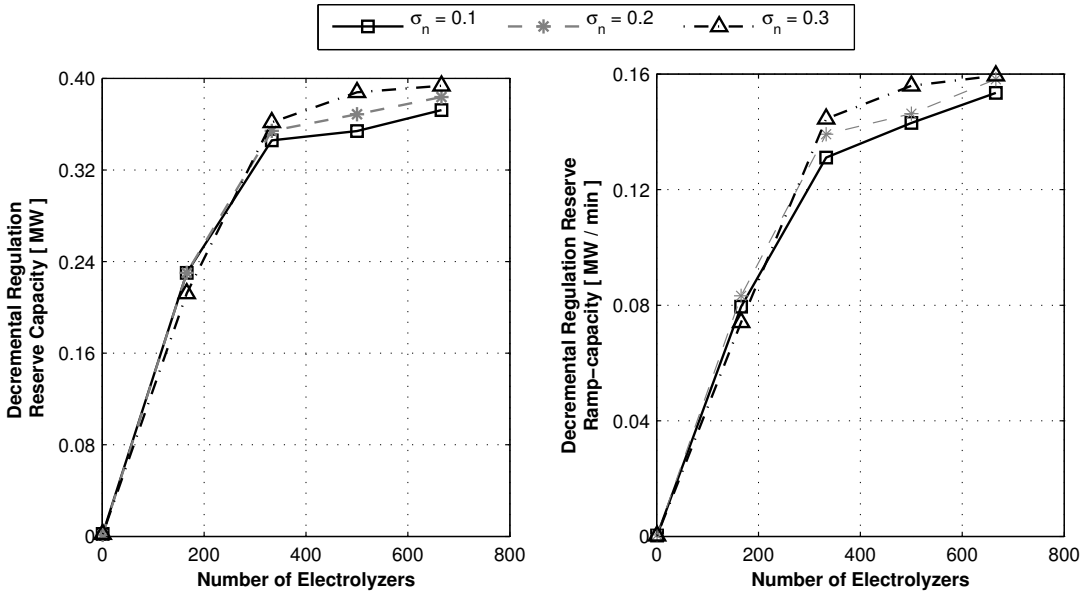

Figure 5.14: Effect of electrolyzer population size and measurement noise.

linear characteristics, and this result indicates that this increased variability in the measurement noise enhances the performance of the load management strategy. As under actual conditions it can be expected that the electrolyzer will in fact operate under non-linear conditions, dependent on transient phenomena, a real population should yield a greater ability to maintain diversity. This result underlines the need to accurately model loads at the device-level. Moreover, transient models, such as that used in the case of heat pumps, should be pursued in order to obtain higher quality results.

5.3 Effects of Wind System Capacity

By scaling the wind power time-series data given in Fig.(4.6), the effects of wind plant size on load control performance can be investigated. The smoothing effects that do occur with an increased number of turbines are neglected, and is sufficient for this study, as we are primarily concerned with how larger fluctuations can be accommodated by the load control strategy, and possible limitations or tradeoffs that may exist. Each responsive load type is reinitialized to the properties implemented in the preliminary case, with the simulation results provided in Fig.(5.15). The decre-

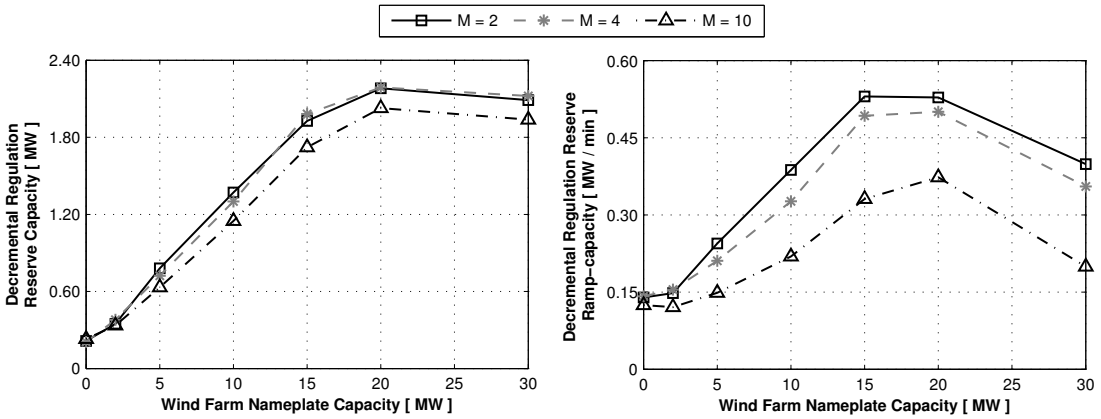

Figure 5.15: Effect of wind energy system capacity and selection of M.

mental regulation reserve capacity and ramp rates have been plotted versus the wind plant capacity, where it can be seen that the benefits associated with controlling the responsive loads increase until the wind plant capacity reaches 20 MW, whereupon benefits diminish. The responsive load populations are unable to provide the flexibility required to follow the power gradients present in the large-scale cases, and

therefore the performance of the control strategy becomes less effective. Also plotted are the results for different M, where again, the $M = 2$ provides the largest decreases in required regulation reserve capacity and ramping capability. More responsive loads would be needed at larger wind capacities, as this would provide greater flexibility to follow the larger power gradients.

As incremental regulation reserve requirements are typically used to apply integration costs to wind development projects [22], these results also indicate that grid-integration into communities enabled with the proposed load control strategy would in turn face lower integration-based costs. Therefore, from the standpoint of project cost, this type of configuration leverages community-owned wind projects, as attaining the regulatory services from wholly owned infrastructure (customers own the responsive loads) comes at reduced costs in comparison to an independent project developer.

5.3.1 Application in Remote Power Systems

Remote grids typically need to transport fossil fuels into the location in order to use diesel generators. This requirement results in an extremely expensive and vulnerable supply. Although many remote locations are ideal candidates for wind energy exploitation, in many cases they lack the ability to accommodate grid integration. These isolated grids are usually associated with low mechanical inertia that is negatively effected by an integrated wind turbine, due to rapid changes resulting from wind variability. This variability also prematurely degrades the diesel gensets. These are major drawbacks for isolated grids, that would benefit considerably from a localized renewable energy resource.

Access to the proposed load control should then provide dual benefits. It can act to protect these diesel generators from premature degradation, by reducing their ramping duty, and allowing wind energy to partially displace the need for their operation. Furthermore, the load control strategy could prove beneficial in reducing wasted wind energy. If the weak grid cannot handle large wind fluctuations, wind energy needs to be curtailed, reducing the amount of diesel fuel displaced. Within the model used in this thesis, the grid flexibility, or the maximum allowable power gradient across the network bus, describes the grid's ability to accommodate variability, and is now manipulated in order to understand any positive effects the load control strategy brings to weaker systems.

The responsive load populations and wind plant capacity are reinitialized to those implemented in the preliminary case, and simulations are then ran under a varying maximum generator ramp-rate $\dot{P}_{G,max}$, as defined in (4.17). Results are provided in Fig.(5.16), as the range of grid flexibilities tested versus the decremental wind energy that is curtailed when the loads are subject to control. As can be seen, as the grid

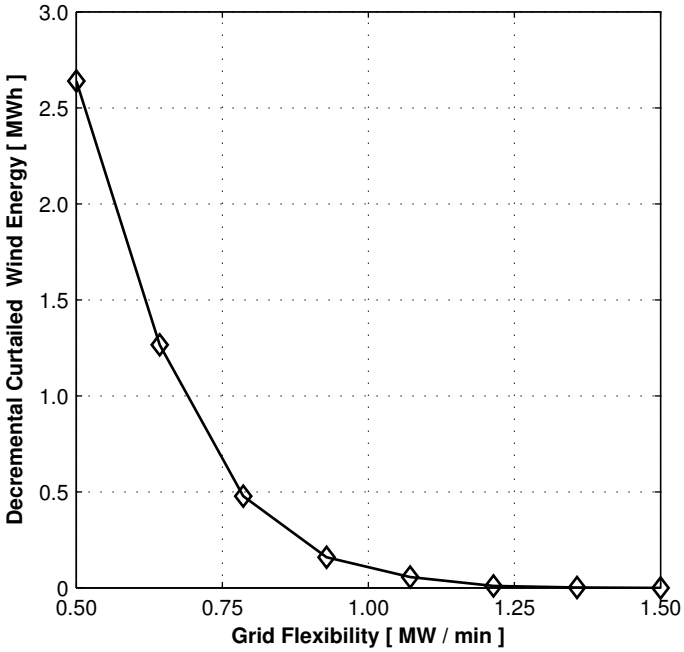

Figure 5.16: Effect of grid flexibility.

flexibility decreases, the benefits associated with pursuing the proposed load control strategy increase rapidly, as less wind energy is curtailed from the network bus. The responsive loads reinforce the grids inertia through adapting demand to explicitly follow wind fluctuations. Therefore, it can be expected that implementing similar load control strategies within low inertia grids should provide considerable benefits in terms of wind integration by allowing more wind energy to penetrate into the demand profile. Auxiliary benefits would be a reduction in the ramping duty of the gensets, thus preventing wind-induced degradation.

5.4 Effects of Communication Network Quality

The proposed control strategy is closed loop and therefore presents problems, as a considerable amount of data needs to travel among the participants in order for

the control to remain effective. If a very large number of spatially-distributed loads are to be controlled, the probability of losing information vital for accurate VGM development increases. Typically, the packet loss ratio (PLR) is used to describe the quality of a network with respect to its ability to retain information. While packet loss can occur for a number of different reasons, it is expected that packet loss will most readily occur during the most heavily congested communication intervals, which in this case represents the LA's process of gathering the power-state vectors from the LC. The PLR in this case then represents the percentage of power-state vectors (packets) that do not make it to the LA, relative to the total number of packets sent from the participating load group. It can be expected that this results in an aggregate load model that underestimates the resource, as packet loss results in more loads currently participating in the LC than the LA expects, potentially leading to the dispatch of more demand than actually desired. Packet loss therefore results in uncertainty in the aggregate load model, and should be dealt with appropriately through modification of the dispatch process.

To understand the effects of packet loss on the control performance, and gain insight into possibilities of dealing with uncertainty within the proposed dispatch framework, a model comprising a heat pump population of 1000 units and a 2.5 MW wind plant is generated. Packet loss is then simulated through deletion of individual state-vectors from the set $\mathbf{Z}$ each sampling interval, and is achieved through drawing three random variables from a uniform distribution spanning the total number of loads in the population. These variables are then used to directly select and delete individual power-state vectors. The number of vectors to be deleted are determined by setting the PLR at the start of simulation, which in this preliminary case was checked over a range the 0 to 4%. As as example, a PLR of 1% therefore represents a loss of on average 10 power-state vectors each sampling interval from the population of 1000 heat pumps.

Model results are given in Fig.(5.17) for three of the tested PLRs, as the total demand trajectories over the simulated time horizon. The resultant response error is given for each scenario in the bottom panel of this figure. As can be seen, packet loss reduces the controller accuracy. As the observed response error is positively biased, our initial hypothesis that packet loss causes an underestimate of the demand potential was indeed correct. Increasing the packet loss results in a controller that consistently over dispatches the load, as the aggregate load model does not consider all power-state vectors that are actually participating in the current model.

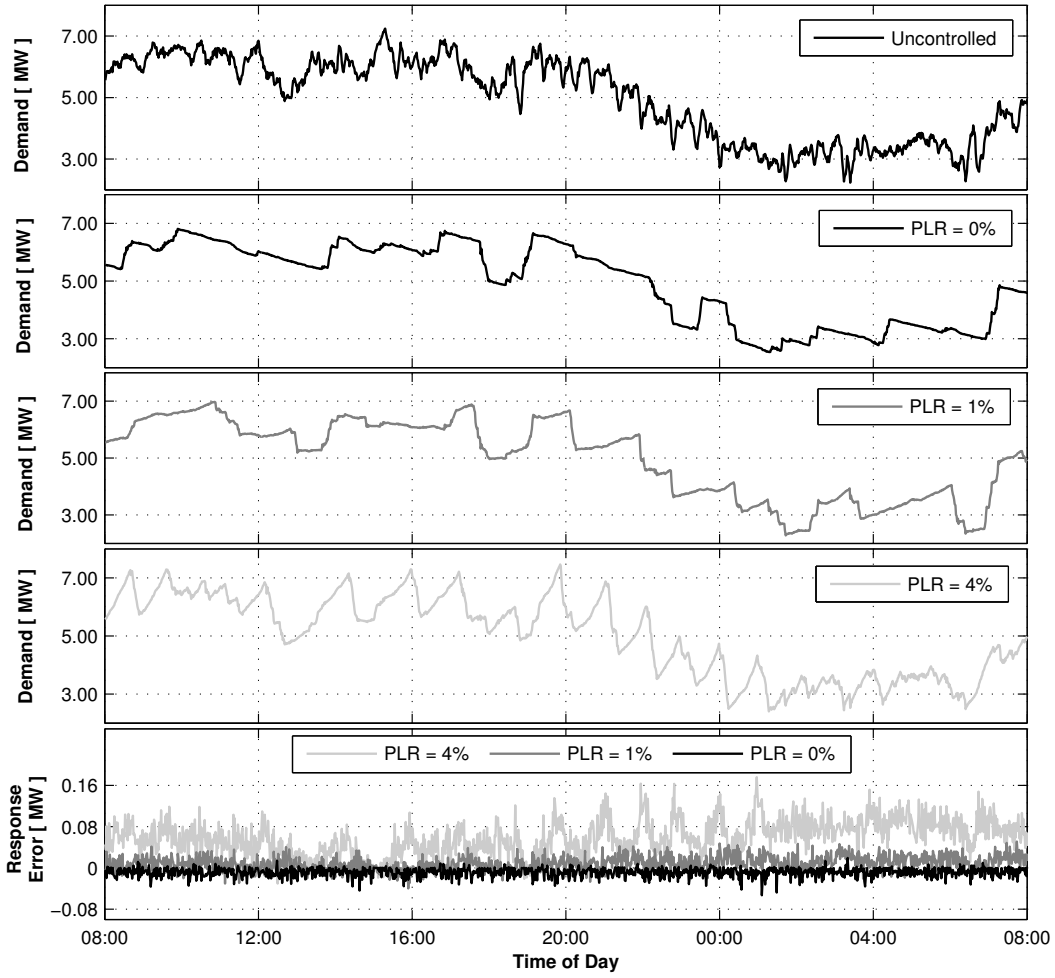

Figure 5.17: The effects of packet loss on the total demand and response error.

While these results indicate that packet loss is in fact a major problem for the control accuracy, it is important to consider the true objective associated with employing such a system configuration. The goal of the controller is to reduce demand variability or the need for regulation reserve, and therefore these attributes should be checked to understand the true impacts of packet loss. Both the regulation reserve capacity and power gradient are computed for each of the simulated levels of packet loss, with the results given with the control accuracy in Fig.(5.18). Each of these

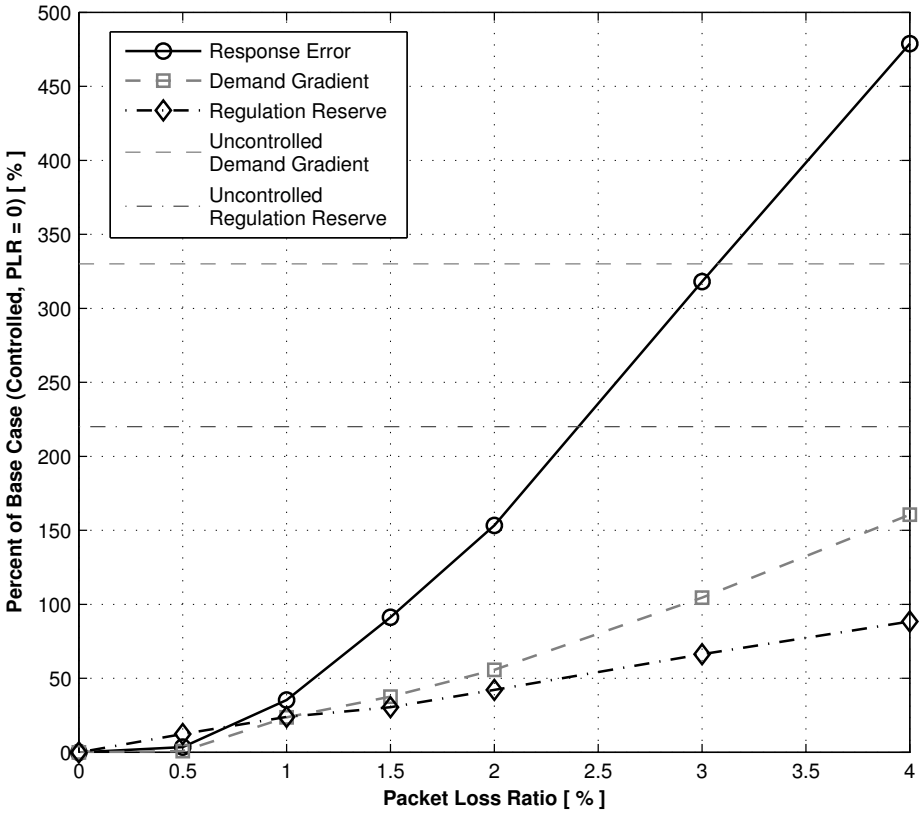

Figure 5.18: The effects of packet loss on the load control performance.

metrics is plotted as a percent of the base case (PLR = 0, and controlled) over the range of PLRs simulated. As can be seen, while the accuracy quickly deteriorates with increased PLR, the power gradient and regulation reserve capacity do not display the same level of sensitivity. In fact, the required regulation reserve and average power gradient consistently remain well below the uncontrolled case for both the gradient and regulation reserve, even at a higher PLR. This result can be attributed to the fact that while packet loss causes state-vectors to not make it to the LA, these vectors may not take part in the aggregate load model formulation. Considering (3.19), it is

clear that if power-state vectors do not fall within the regions encompassed by the summation, they do not play a role in defining the VGM. Therefore, if these types of vectors are missed due to packet loss, there is no effect on the aggregate load model. Furthermore, PLR may increase the response error considerably, but when the amplitude of these inaccuracies are considered in tandem with the required response amplitude based on the whole population, they remain minimal.

These results indicate that control accuracy is sensitive to packet loss, although these impacts diminish within the viewpoint of the controller's true objective. In fact, at a PLR less than 0.5% packet-loss has no effect at all on the control performance metrics. As packet loss will be directly related to the size and density of the population, these results promote the proposed community-scale approach to load management, as it can be expected that operating the smaller network containing fewer customers should in turn be accompanied by less packet-loss.

Chapter 6

Conclusions

This thesis has provided a framework for engaging community-based sustainable demand-side infrastructure in the provision of power system ancillary services. Clean energy technologies in the form of heat pumps, electric vehicles, and electrolyzers were targeted to act as a responsive load basis, as the inherent operational flexibility and autonomous device-level control strategy associated with these load-types provides an ideal platform to mitigate enhanced variability within a power system. By adapting their demand to reinforce the short-term operating capabilities of the grid, these end-use devices can provide dual benefits in terms of carbon emissions mitigation. Their adoption can displace fossil fuel-based demand-side devices, while their grid-side control increases the ability of the electric power system to accomodate intermittent renewable energy technologies.

Problems associated with electric power system management under increased uncertainty were discussed in Chapter 2. Real time short-term energy balancing, denoted as regulation, was singled out as a major shortcoming in terms of renewable energy integration. In particular, current methods of providing regulation require allocation of specific reserves, which in turn result in both technical and economic inefficiencies within current grid-management configurations. As renewable energy generation, in particular wind power production, displays variability over time-scales relevant to regulation, integration of these sources can result in a need to increase regulation reserves, which can thereby overshadow many of the positive economic and environmental benefits associated with wind energy.

In response to these problems, Chapter 3 presented a framework for engaging communities in the provision of regulation-based ancillary services. The idea was introduced of a community-based business acting as an intermediary between the power

system operator and community residents in possession of electric loads suitable for demand response recruitment. It was argued that such an intermediary can provide economic benefits to the community through offering the community's flexible demand resource to the power system operator at a cost equivalent to the market-based benefits achieved due to resultant increases in grid operating efficiency. Furthermore, security, quality, and reliability of the demand resource were deemed superior within the bounds of a community-based approach. Chapter 3 also provided the mathematical framework for the novel load management strategy. The strategy is framed around the characteristics of hysteresis control, and by doing so, is able to accelerate decision-making efficiency. End-use digital communication hardware would be used to track and control the population of loads through the development of online aggregate load models equivalent to conventional dispatchable generation.

In Chapter 4, the computational model used to test the proposed load control policy was put forward. A bottom-up approach to energy systems modelling was used to explicitly connect end-use functionality of the participating responsive load groups to a power systems analysis typically pursued to understand the effects of integrated renewable energy systems. Numerical metrics for rating the performance of the load control strategy within both the viewpoint of the end-user and the power system operator were further developed.

Chapter 5 presented applications of the computational model within various scenarios, with the main conclusions from these results summarized in the following:

- Large populations of hysteresis controlled loads can be managed for regulation-based ancillary services, without sacrificing the end-use functionality of participating units. Direct, closed-loop control of the demand response network provides control accuracy and reliability by allowing the LA to adapt the VGM online. Different load-types can be controlled similarly by mapping their operation to a common device-level control space.

- Model results indicate that populations of heat pumps provide an effective basis for regulation-based ancillary services, but are ill-suited for load following. Increasing the thermal mass of a building population would provide greater opportunities to operate within the load-following regime, as more energy could accumulate within the buildings without significantly affecting the indoor air temperature. The populations of EVs and electrolyzers were outperformed by the heat pumps, but EVs may in fact provide a more versatile resource to the

power system, due to greater device-level flexibility and year-round utilization.

- If a population exists, there is no need to implement vehicle-to-grid or regenerative fuel cell approaches in order to provide variability buffering, as simply controlling the on / off status of an EV or electrolyzer can provide significant short-term flexibility.

- The proposed load control strategy can significantly reduce the need to operate online regulation reserves. As a result, integration of renewable generation that lacks controllability in locations near communities enabled with the proposed load control strategy will face lower integration-based costs. Modelling results further indicated the benefits associated with pursuing the load control strategy in weak grids, as access allowed larger wind plants to be included, with the available energy able to penetrate deeper into the demand profile.

- Although the proposed load control strategy is computationally efficient at the level of the LA, it is limited due to the need for explicit closed-loop communication with each responsive load. Communication network error in the form of packet loss is a major problem for the load control accuracy, but can be alleviated through operating an optimally designed network with fewer customers.

The overarching conclusion of this thesis is the identified benefits associated with applying hierarchal systems integration to the development of low carbon energy systems. This thesis has demonstrated that by aggregating large groups of spatially-distributed loads into comfort-constrained virtual embedded generation, we can seamlessly integrate the control of these loads into conventional power system operational strategies. The insight then is to realize the applicability of these hierarchal control architectures to further development of alternative small-scale distributed energy resources, such as micro-turbines, biodiesel generators, and fuel cell systems. These devices are typically implemented in locations close to the end-user, and thereby can provide opportunities for cogeneration, as well reductions in transmission losses. As communities often have access to local renewable feedstocks that can be reformed to produce low carbon fuels, such as hydrogen or biodiesel, distributed generation that exploits these fuel types would be extremely well-suited to provide the small amounts of capacity needed for many typical ancillary services. However, the inclusion of these devices within typical ancillary service markets is difficult. Large numbers of small-scale units complicates power system management, and would reduce the ability of

the power system to perform ED efficiently (too many control variables). Alternatively, networking these devices within a given area could allow aggregation of both small-scale distributed generation and responsive loads into virtual generation that would display capacity comparable to larger-scale generation. This would reduce the computational complexity associated with optimally dispatching a large population of distributed energy resources.

The main recommendation of this thesis is that community-based energy management authorities should be developed. These entities would oversee and operate localized energy management networks. If such authorities can balance the interests of all parties within a strategy that is non-intrusive, secure, profitable and ecological from the perspective of both the community and the power system operator, they can provide a bridge that simultaneously provides economic benefit to the community and to the electric power system.

Aside from these conclusions is the observed superiority of direct load control as compared to approaches that rely on real-time pricing to achieve control actions, such as those seen in [49, 58]. Not only do these methods inherently contain uncertainties brought on by customer-behaviour, and expect participating customers to take part in the somewhat daunting task of dealing with the energy market that governs power system operation, real-time price-based control of loads requires an associated market, or more particularly, combining load control with ED. Market clearing occurs over periods of approximately 15 minutes, which means that price-based control actions would need to focus on longer-time scales services (load-following). The results in this thesis have indicated that the targeted loads are incapable of providing such services, unless end-users were to agree to forfeit typical load functionality, or a very large load population was available. The comfort-constrained load control pursued in this work is capable of acting over shorter time-frames, providing fast ramp rates and control accuracy that allows it to displace the need for regulation. Price-based control would be unable to tap into this ancillary service market, which can be provided by the loads through direct control, without effecting end-use function.

Regulation with direct load control can reduce reserve capacities associated with conventional generation, effectively increasing the active power capacity of the power system. This extra capacity could be used to meet a growing demand and perhaps relieve the need to pursue load control actions aimed at peak load reduction. This could, in fact, be classified as a load-following service (shifting generation instead of shifting demand).

A considerable number of open questions remain and provide opportunities for future research:

- Incorporation of detailed physically-based models into the systems analysis, in particular for EVs and electrolyzers. This would involve explicitly connecting the model to a full transportation model capable of accurately capturing the demand for both electric vehicle and hydrogen energy. It would be beneficial to add hydrogen storage into the model, in terms of compression requirements versus the length / amount of storage needed. The distribution system should be treated through full AC power simulations, as many load types are voltage / frequency sensitive. To this regard, there may be further opportunities to achieve innovative control objectives within the distribution system, such as minimizing power losses and voltage regulation [59].

- Further consideration of the communication network design is required. The aspect of response delay as well the energy required to in fact power the network itself should be determined [60].

- Market and environmental factors should be fully modelled to understand the viability of pursuing this system configuration. As the targeted demand-side infrastructure represents relatively new technology, tradeoffs between possible increases in capacity brought about by widespread use, and the potential ability to control these units as a resource to liberate conventional generation must be quantified.

In any case, the opportunities to make fundamental contributions within this relatively new research field are numerous, and will provide considerable benefits to the development of sustainable low carbon energy systems.

Bibliography

[1] IPCC. *Climate Change 2007: The Physical Science Basis*. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 2007.

[2] K. Caldeira and M.E. Wicket. Anthropogenic carbon and ocean pH. *Nature*, 425:365, 2003.

[3] H. Kan and B. Chen. Particulate air pollution in urban areas of Shanghai, China: health-based economic assessment. *Science of the Total Environment*, 322(1):71–79, 2004.

[4] IPCC. *Climate Change 2007: Synthesis Report*. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 2007.

[5] M. Asif and T. Muneer. Energy supply, its demand and security issues for developed and emerging economies. *Renewable and Sustainable Energy Reviews*, 11(7):1388–1413, 2010.

[6] IEA. World Energy Outlook 2010. International Energy Agency, 2010.

[7] O. Edenhofer, R. Pichs-Madruga, Y. Sokona, and et al. Special Report on Renewable Energy Sources and Climate Change Mitigation. Technical report, IPCC Working Group III - Mitigation of Climate Change, 2011.

[8] E.P. Johnson. Air-source heat pump carbon footprints: HFC impacts and comparison to other heat sources. *Energy Policy*, doi:10.1016/j.enpol.2010.12.009, 2011.

[9] C. Samaras and K. Meisterling. Life cycle assessment of greenhouse gas emissions from plug-in hybrid vehicles: Implications for policy. *Env. Sci. Tech.*, 42(2):3170–3176, 2009.

[10] J. Barton and R. Gammon. The production of hydrogen fuel from renewable sources and its role in grid operations. *J. of Power Sources*, 195(24):8222–8235, 2010.

[11] D.S. Callaway and I. Hiskens. Achieving controllability of electric loads. *Proc. of IEEE*, 99(1):184–199, 2011.

[12] G. Cornelis van Kooten. Wind power: the economic impact of intermittency. *Letters in Spatial and Resource Sciences*, 3(1):1–17, 1990.

[13] C.R. Warren and M. McFadyen. Does community ownership affect public attitudes to wind energy? A case study from south-west scotland. *Land Use Policy*, 27(2):204–213, 2010.

[14] Danish Energy Authority. Energy statistics, 2005.

[15] R. Vilafafila, A. Sumper, A. Suwannarat, B. Bak-Jensen, R. Ramirez, and A. Sudria. On wind power integration into electrical power system: Spain vs. Denmark. In *Proc. of Int. Conf. Renewable Energies and Power Quality*, 2007.

[16] M. Ilic, Y.V. Makarov, and D.L. Hawkins. Operations of electric power systems with high penetration of wind power: Risks and possible solutions. *Proc. of IEEE PES General Meeting*, pages 1–4, 2007.

[17] A.J. Allen and B.F. Wollenberg. *Power Generation Operation and Control (2nd Edition)*. John Wiley and Sons, 1996.

[18] M. Ilic and J. Zabrorszky. *Dynamics and Control of Large Electric Power Systems*. New York: Wiley-Interscience, 2000.

[19] G. Giebel and et al. The state-of-the-art in short-term prediction of wind power: A literature overview, 2nd edition. Technical report, Ris DTU, Wind Energy Division, 2011.

[20] J. Apt. The spectrum of power from wind turbines. *J. of Power Sources*, 169(2):369–374, 2007.

[21] B. Parsons, M. Milligan, B. Zavadil, D. Brooks, B. Kirby, K. Dragoon, and J. Caldwell. Grid impacts of wind power: A summary of recent studies in the United States. *Wind Energ.*, 7:87–108, 2004.

[22] BCHydro. 2008 Long Term Acquisition Plan: Appendix F3 Wind Integration Cost Assessment. Technical report, BC Hydro, 2008.

[23] E. Hirst and B. Kirby. Separating and measuring the regulation and load-following ancillary services. *Util. Policy*, 8:75–81, 1999.

[24] N.V. Zorchenko et al. Evaluating the effect of frequency regulation on the reliability and economic efficiency of thermal power generation units. *Power Tech. and Eng.*, 45(2):132–136, 2011.

[25] W. Katzenstein and J. Apt. Air emissions due to wind and solar power. *Env. Sci. Tech.*, 43(2):253–258, 2009.

[26] F. Katiraei and C. Abbey. Diesel plant sizing and performance analysis of a remote wind-diesel microgrid. In *Proc. of IEEE PES General Meeting*, pages 1–8, 2007.

[27] RETScreen. *Clean Energy Project Analysis: RETScreen Engineering and Cases Textbook*. Small Hydro Project Analysis. Natural Resources Canada, 2001.

[28] J. Kennedy, B. Fox, and D.J. Morrow. Distributed generation as a balancing resource for wind generation. *IET Renew. Power Gener.*, 1(3):167–174, 2007.

[29] E.J. Hittinger, F. Whitacre, and J. Apt. Compensating for wind variability using co-located natural gas generation and energy storage. *Energy Syst.*, 1:417–439, 2010.

[30] L. Xie, P.M.S. Carvalho, L.A.F.M. Ferreira, J. Liu, B.H. Krogh, N. Popli, and M.D. Ilic. Wind integration in power systems: Operational challenges and possible solutions. *Proc. of the IEEE*, 99(1):214–232, 2011.

[31] CAISO. 2009 annual report: Market issues and performance. Technical report, California Independent System Operator, 2009.

[32] W. Kastner, G. Neugschwandtner, S. Soucek, and H.M. Newmann. Communication systems for building automation and control. *Proc. of IEEE*, 93(6):1178–1203, 2005.

[33] N. Ruiz, I. Cobelo, and J. Oyarzabal. A direct load control model for virtual power plant management. *IEEE Trans. on Power Syst.*, 24(2):959–966, 2009.

[34] Marbek Resource Consultants Ltd. BC Hydro 2007 Conservation Potential Review. Technical report, Prepared for BC Hydro, 2007.

[35] NRC OEE. 2007 Survey of Household Energy Use: Detailed Statistical Report. Technical report, Natural Resources Canada Office of Energy Efficiency, 2007.

[36] Environment Canada. National Inventory Report 1990 - 2008: Greenhouse Gas Sources and Sinks in Canada. Technical report, Greenhouse Gas Division: Environment Canada, 2010.

[37] NRC OEE. Heating and Cooling with a Heat Pump. Technical report, Natural Resources Canada Office of Energy Efficiency, 2004.

[38] J. Stene. Residential CO2 heat pump system for combined space heating and hot water heating. *Int. J. of Refrig.*, 28(8):1259–1265, 2005.

[39] J. Kondoh, N. Lu, and D.J. Hammerstrom. An evaluation of the water heater load potential for providing regulation service. *IEEE Trans. on Power Syst.*, 26(3):1309–1316, 2010.

[40] D.S. Callaway. Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy. *Energ. Convers. and Manag.*, 50(9):1389–1400, 2009.

[41] Nissan. 2012 Nissan Leaf: Features and Specifications. http://www.nissanusa.com/leaf-electric-car/, 2011.

[42] Statistics Canada. Census: Catalogue number 97-561-xcb2006010, 2006.

[43] Kempton et al. A test of vehicle-to-grid (V2G) for energy storage and frequency regulation in the PJM system. Technical report, Mid-Atlantic Grid Interactive Cars Consortium, 2009.

[44] C. Quinn, D. Zimmerle, and T.H. Bradley. The effect of communication architecture on the availability, reliability, and economics of plug-in hybrid electric vehicle-to-grid ancillary services. *J. of Power Sources*, 195(5):1500–1509, 2010.

[45] A. Kashyap and D. Callaway. Controlling distributed energy constrained resources for power system ancillary services. In *Proc. of IEEE PMAPS*, pages 407–412, 2010.

[46] Z. Ma, D. Callaway, and I. Hiskens. Decentralized charging control for large populations of plug-in electric vehicles: Application of the Nash certainty equivalence principle. In *Proc. of IEEE ICCA*, pages 191–195, 2010.

[47] A. Bergen, L. Pitt, A. Rowe, P. Wild, and N. Djilali. Transient electrolyser response in a renewable-regenerative energy system. *Int. J. of Hydrogen Energy*, 34(1):64–70, 2009.

[48] R. Mortensen and K. Haggerty. A stochastic computer model for heating and cooling loads. *IEEE Trans. Power Syst.*, 3(3):854–860, 1990.

[49] S. Katipamula, D.P. Chassin, D.D. Hatley, R.G. Pratt, and D.J. Hammerstrom. Transactive Controls: Market-based GridWise Controls for Building Systems. Technical report, Prepared for the U.S. Department of Energy under Contract DE-AC05-76RL0183 and PNNL-15921, 2006.

[50] H. Madsen and J. Holst. Estimation of continuous-time models for the heat dynamics of a building. *Energy and Buildings*, 22(1):67–79, 1995.

[51] PNNL. GridLabD - Power System Simulation Software. U.S. Department of Energy: Pacific Northwest National Laboratory http://www.gridlabd.org/, 2011.

[52] J.F. Manwell, J.G. McGowan, and A.L. Rogers. *Wind Energy Explained: Theory, Design and Application, 2nd Edition.* John Wiley and Sons, 2010.

[53] C.G. Justus. *Winds and Wind System Performance.* Franklin Institute Press, Philadelphia PA, 1978.

[54] J. Cotrell. A preliminary evaluation of a multiple-generator drivetrain configuration for wind turbines. In *Proc. of 21st ASME Wind Energy Symposium*, pages 345–352, 2002.

[55] E. Hirst and B. Kirby. Ancillary-service details: Regulation, load following, and generator response. Technical report, Oak Ridge National Laboratory, 1999.

[56] L. Kelly. Probabilistic modelling of plug-in hybrid electric vehicle: Impacts on distribution networks in British Columbia. Masters Thesis: University of Victoria Department of Mechanical Engineering, 2010.

[57] A. Burke. Batteries and ultracapacitors for electric hybrid and fuel cell vehicles. *Proc. of IEEE*, 95(4):806–820, 2006.

[58] J. Joo and M.D. Ilic. A multi-layered adaptive load management (ALM) system: Information exchange between market participants for efficient and reliable energy use. In *Proc. of IEEE PES TDC*, pages 1–7, 2010.

[59] D. Wang, B. de Wit, S. Parkinson, J. Fuller, D. Chassin, C. Crawford, and N. Djilali. A test bed for self-regulating distribution systems: Modelling integrated renewable energy and demand response in the GridLAB-D/MatLab environment. In *Proc. of IEEE ISGT*, 2012.

[60] L. Zheng, S. Parkinson, D. Wang, L. Cai, and C. Crawford. Energy efficient communication network design for demand response in smart grid. In *Proc. of IEEE WCSP*, 2011.