Localised Demand Control for Improved Grid Integration of Renewable Energy and Electric Vehicles

Joshua Reuben Lee

A thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy in Electrical and Electronic Engineering.

The University of Auckland

2014

Abstract

Reliable electricity supply is essential to the way of life in developed countries and of extreme importance when looking to raise the standard of living in developing countries. Unfortunately electricity generation is dominated by fossil fuels which cause significant pollution and emissions – a situation that due to climate change and environmental concerns should not continue. Most power systems also tend to show significant peaks in electricity demand for only a few hours per day. Since power systems are designed to handle the peak demand, peaking plants could be replaced with renewable generation if some demand were able to be shifted in time. This would be more cost effective and emissions would be reduced.

Shifting demand in time can be achieved with Demand Side Management (DSM), and a new DSM scheme called Localised Demand Control (LDC) is introduced in this thesis. LDC is compared with a traditional DSM scheme called Dynamic Demand Control (DDC). A number of aspects that make DDC unsuitable are presented, and LDC is shown to improve on these while also providing additional benefits.

LDC is designed to regulate power flow on a single distribution transformer, and the stability and transient response of LDC is presented. Simulations of transformers with controllable load and fluctuating local generation are presented and these show good results. A method of load scheduling is also demonstrated, which reduces electricity cost without impacting the end user.

A lab-scale system has been created and its stability and transient response match those of the simulated system. Here a 3-phase power meter measures power flow and a centralised LDC controller determines a load control signal which is then distributed to controllable loads. Motor-generator sets are used to simulate fluctuating local generation sources and the system provides good power regulation with just a small number of controllable loads. A custom power-line communication system is used to distribute the load control signal, and the performance has been verified in the lab and in simulations with a large amount of signal noise.

The proposed LDC system could be used to improve integration of renewable energy and electric vehicles in the power system, increasing efficiency and reducing emissions.

Acknowledgements

I would like to thank my supervisors John Boys and Grant Covic. John you were always available, always encouraging and full of new ideas for any and every problem. Thanks for all your hard work and the huge effort you put in, especially during the last stages. Grant, I am especially grateful for the extra attention to detail you provided when reviewing papers and chapters. Thanks for your constant encouragement, direction, and advice throughout.

I would also like to thank the technical staff in the ECE department, in particular Howard Lu, Kavitha Penneru and Vic Church. Thanks for being so helpful when sorting out lab equipment, components and PCBs.

I would like to thank my colleague Aaron for his support, advice and for always being available to bounce ideas off. I would also like to thank Duncan Eason, James Scoltock, Maja Krivokuca, Nora Dittrich, Tim Roper and Vedrana Kirvokuca for their comradery, moral support and especially for the great lunchtime conversation.

I would also like to thank my colleagues in the ECE department for their help and support: Abhilash Kamineni, Abiezer Tejeda, Adeel Zaheer, Ali Abdolkhani, Baljit Riar, Craig Baguley, Daniel Robertson, David Huang, Edward van Boheemen, Frank Hao, Ganesh Nagendra, Glen Robertson, Hunter Wu, Jackman Lin, Jason James, Jerry Chen, Jimmy Peng, Jonathan Beaver, Jonathan Bradshaw, Michael Kissin, Mickel Budhia, Nick Keeling, Richard Stebbing, Robert Bowmaker, Robert Dunn, Seho Kim, Stefan Raabe and Zak Beh.

A special thanks also to Transpower New Zealand Limited for providing me with a generous PhD scholarship. This research would have been far more difficult without this assistance.

I would finally like to thank my amazing Family. Dad, thanks for your wisdom, advice and proof-reading over these many years. Mum, thanks for keeping me well fed and looked after and for encouraging me towards a university education early on. Jono, Rebekah and Sarah, thanks for friendship, support and encouragement over these many years. Thanks also for your incredible assistance in these final stages. To my grandparents Paul, Ruth, David and Glennis, thanks for always being interested in my research and constantly encouraging me on, I admire you all.

Table of Contents

List of Figures

List of Tables

Glossary

AC	Alternating Current
ADC	Analogue to Digital Convertor
AUFLS	Automated Under-Frequency Load Shedding
BPF	Band-pass Filter
DC	Direct Current
DDC	Dynamic Demand Control
DG	Distributed Generation
DSM	Demand Side Management
EPA	Environmental Protection Agency
EV	Electric Vehicle
FFT	Fast Fourier Transform
GPRS	General Packet Radio Service
HAN	Home Area Network
HPF	High-pass Filter
HV	High Voltage
HVAC	Heating, Ventilation and Air-conditioning
HWC	Hot Water Cylinder
IC	Integrated Circuit
IGBT	Insulated Gate Bipolar Transistor
IPT	Inductive Power Transfer
JFET	Junction Gate Field-Effect Transistor
LDC	Localised Demand Control
LV	Low Voltage
MG	Motor Generator
MV	Medium Voltage
NZ	New Zealand
PAR	Peak to Average Ratio
PLC	Power Line Communication
PMU	Power Management Unit
PSoC	Programmable System on Chip
PWM	Pulse Width Modulation
RC	Resistor Capacitor
RTP	Real-time Pricing
SEP	Smart Energy Profile
SiC	Silicon Carbide

SNR	Signal to Noise Ratio
SoC	State of Charge
SST	Solid State Transformer
TCP/IP	Transmission Control Protocol / Internet Protocol
UART	Universal Asynchronous Receiver / Transmitter
VAR	Volt-ampere Reactive
VPN	Virtual Private Network
VSD	Variable Speed Drive

Chapter 1
Introduction

1.1 Introduction

Much is changing in our world today with respect to our energy supply. Attitudes towards the environment and climate are shifting as we begin to see the effect of decades of unchecked emissions. The general consensus in the scientific community is that we cannot continue at the current rate of emissions – but when over 80% of the global energy supply is from Oil, Coal and Gas [1], this is a very difficult problem to solve.

Renewable energy provided just 13.2% of the global energy supply in 2010, including 2.3% from hydroelectric generation [1]. While this is a very small proportion, interest in renewable generation is increasing and the installed capacity is growing by 5% annually, and the rate of this growth is also accelerating [2].

It is worth noting that only around 15% of the world energy supply is provided by electrical means, with the direct burning of fossil fuels being the most common energy source. In New Zealand, 77% of its electricity supply is generated from renewable sources, and only 39% of total energy consumption is renewable. The energy demand from the transportation sector is a significant contributor to this difference. This is because 34% of the total energy used is provided by oil, and of this 82% is used for transportation. Overall, transportation consumes 38% of the total energy used in New Zealand.

Renewable generation can be provided by many sources, including hydroelectric dams, geothermal plants, wind turbines and solar panels. These can be separated into two main classes based on whether they are dispatchable or non-dispatchable. Dispatchable generation can generate at any time of the day, and can be forecast with almost complete certainty. Hydroelectric and geothermal generation are both dispatchable. Non-dispatchable generation includes wind and solar as these are much less predictable. If the energy is not taken when it is available it is simply lost. Examples of wind, coal, and hydro plants from New Zealand are shown in Figure 1.1, Figure 1.2, and Figure 1.3 respectively.

Figure 1.1 – The 160MW Tararua wind farm is NZ's largest. Source: Wikipedia

The technology by which electricity is generated, transported and distributed has remained almost constant for many decades. This combination of technology is commonly referred to as the "power grid" and consists of the generators and multiple levels of transmission lines and power transformers. In most cases electricity is generated in large quantities at concentrated sites, and then transmitted long distances to both small and large load centers where it is consumed. Up until very recently the power grid has been a fully one way system, with the aim of constantly matching generation to whatever the instantaneous load may be – if a load is switched on anywhere in the power grid, then generation must be dispatched to cover that load.

Figure 1.2 – Huntly power station. Source: Wikipedia

Most of the new renewable generation being deployed around the world is instead in the non-dispatchable category, as most easy to access sources of hydroelectric and geothermal generation are already being exploited. The unpredictable nature of generation output from these sources can cause significant reliability and planning issues as they begin to make up a larger proportion of the electricity supply. In order to guarantee that the full electricity demand can be met, additional dispatchable generation is held in reserve such that it can be brought online quickly if there is a sudden drop in renewable output.

Figure 1.3 – 100MW Whakamaru hydroelectric power station. Source: Mighty River Power

When the proportion of non-dispatchable generation in a power system is only a few percent, the reserve requirement is of minor significance with respect to power system operation. But as proportions increase to 20% or more, the reserves become quite significant and in many cases the construction of a large new wind farm or solar installation will require the construction of additional fossil fuel generation as backup. To make matters worse, the reserve generation must be fast responding which generally implies a gas-fired power plant. These are less efficient and therefore produce even more emissions than equivalent coal-fired generation. Furthermore, in order to provide fast response, the plant must be kept running at partial load such that its output can ramp up quickly if required, further reducing the efficiency and also the profitability of the plant. It has been proposed that under some circumstances, increasing

the proportion of wind in a power system could increase the total emissions while also costing 10x as much compared to installing modern coal-fired generation [3]. It is possible that by adding additional intelligence and flexibility to the power system, these issues can be mitigated without expensive and inefficient reserve generation.

1.2 Low Voltage Network

While much of the additional renewable generation installations are large-scale and connected at high voltage levels, there continues to be an increase in small-scale generation connected at the medium and low voltage level. These are not significant in terms of the overall power system balance and the provisioning of reserves, but they can create issues of their own.

Deployment of small-scale generation can cause the system voltage to vary outside the acceptable operating range, which is of significant concern as there is no easy solution and in some cases the problem could even go completely unnoticed. This is due in part to the distribution network being originally configured with unidirectional power flow in mind. When the line or transformer is loaded, the voltage will drop by a small amount due to the impedance of the line and transformer windings. The system will be designed such that when both are unloaded or at their rated capacity, the voltage stays within a set operating range at all points in the low voltage network. Since the generation of electricity will cause the voltage to rise where it did not previously do so, it is possible for the voltage to swing outside the desired operating range, especially when the line is lightly loaded. Generally, such small-scale generation is single-phase which can also result in the creation of phase imbalances, causing one of the three phases to swing above the desired range, while the others remain within specification.

Another significant issue is that of safety in the presence of line faults. If there is a fault that needs repairing, such as when a tree brings down a line, small-scale generation can liven lines where they are not expected to be live which can endanger line workers. As the output may be intermittent, lines could even be checked and found safe one minute, and then become live the next. Renewable units that can sense faults and shut off generation are available but these are more costly than those without such protection.

In many parts of the world, government mandated feed-in tariffs are available which incentivise the deployment of distributed generation. This is not the case in New Zealand, and as such there is much less interest than is found in other locations such as parts of Europe and

the USA. These tariffs can cause headaches for the utility, as they require that power be bought at a certain rate, even if that power is not needed - and due to the previously mentioned issues the power could even be highly unwanted. In Germany in particular, there has been significant uptake of solar generation due to tariffs, and in many households the rating of the solar array is much larger than the original rating of the line supplying the house. This results in further cost to the utility as the lines and transformers must then be upgraded to integrate this power source that may be unwanted in the first place. As a consequence, the end user receives a smaller power bill while still making the same or even greater use of the distribution network, and a smaller pool of users end up shouldering the cost of a more expensive network.

Overall the issues caused by the deployment of small-scale renewable energy are significant and likely to get worse as this generation becomes more common. It is possible that intelligent use of controllable loads could help mitigate many of these issues, as discussed further in this thesis.

1.3 Electric Vehicles

As previously mentioned, the transportation sector accounts for a significant proportion of energy and fossil fuel consumption. Transportation electrification would therefore decrease fossil fuel consumption and also total energy consumption, as electric drives are far more efficient than gasoline equivalents. Given current trends, it looks quite likely that electric vehicles (EVs) will indeed become a large part of the transportation landscape.

The performance of electric vehicles has always been tied to the performance of the batteries. In the early days of the automobile, electric vehicles actually outsold gasoline powered vehicles, but as the prevalence of gas stations increased and people started driving longer distances, EVs could no longer compete due to their limited range and higher purchase cost.

Due to a recent rapid increase in consumer demand for portable electronics over the last two decades, battery technology has received significant research and development. In 2008 lithium-ion batteries began approaching a cost and performance threshold where they could again compete with gasoline vehicles in a number of usage scenarios.

A major step in EV development came with the Tesla Roadster in 2008. The Roadster (shown in Figure 1.4) was an attempt at a no compromise electric sports car, and with its lithium-ion battery it had an EPA rated range of 344 miles. Around 2500 roadsters were built

between 2008 and 2012. With a 53kWh battery, best charging results were obtained with a dedicated outlet and charging unit. The highest capacity unit available was a 16.8kW (240V, 70A) charger, and it could recharge from empty in just under 4 hours [4]. A charger of this size is a significant load, and is certainly of concern to utility companies.

Figure 1.4 – 2008 Tesla Roadster. Source: Tesla Motors

Although the Roadster was a successful vehicle for a start-up company to produce, 2500 vehicles is insignificant in terms of global auto production. In late 2010 Nissan introduced the Leaf, which was targeted more as a mass market electric vehicle. With a 24kWh battery, a range of 135km and a price of around $35,000USD without subsidies, the Leaf is an economically viable alternative to a gas powered car if the range is acceptable [5]. The standard Leaf charger is rated at 3.6kW, with DC fast charging available on some models at around 40kW. As the fast chargers are quite expensive, the 3.6kW charger is likely to be the most common method of recharging, and while smaller than the Tesla charger, it is still of concern for electrical utilities especially if deployed widely. The Nissan Leaf is currently the best-selling electric vehicle ever, with around 100,000 sales globally since its introduction in 2010 [6].

Also released in 2010 was the Mitsubishi i-MiEV, with slightly lower specifications and price compared to the Nissan Leaf. The i-MiEV has not sold quite as well as the Leaf, but it was the first ever EV to reach 10,000 total sales.

There are now a number of full electric or plug-in hybrid vehicles available on the market, but they still represent a very small slice of the overall vehicle market. It is likely that sales will continue to climb, with initial growth in fact higher than that of hybrids, which now represent a small but significant amount of vehicles sold.

EVs have the capability to add around 1/3 to the average household energy demand. While electric vehicles are very efficient with their energy consumption, this extra demand could mean that network and generation upgrades are required. Electric vehicles could also be particularly problematic due to the timing of the charge cycle. Many electricity networks (including NZ) experience a peak in consumption in the evening due to people arriving home from work and performing energy intensive activities such as cooking, space heating and consuming hot water. If an electric vehicle is also plugged in at this time, the peak will be further increased.

It has been proposed that the low voltage distribution transformer could be the most severely affected piece of grid hardware in the event of EV deployment [7]. This is of concern as most transformers are not monitored directly in any way. In addition, it is likely that EVs will be bought by people in close geographic proximity due to both the higher affluence and word of mouth effects of the average EV buyer. It has been noted by Tesla Motors that the sale of a Model S EV usually spurs 2-3 additional sales in the same area, likely due to word of mouth. It is sensible to assume that many distribution transformers could not handle the addition of multiple electric vehicles, especially if they are charged at peak times.

1.4 Demand Side Management

One way of reducing peaks is by intentionally time-shifting loads away from periods of peak loading. Demand Side Management (DSM) is a common term for the technique whereby electrical loads become involved in the general operation of the power system. In the same way that the system operator is able to dispatch various generation sources to meet demand, controllable loads can be delayed, curtailed or even switched on to achieve the required load balance. There are a number of other terms for DSM that effectively describe the same

procedure including load control, interruptible load, demand response, demand dispatch, controllable load etc.

A number of household loads are suitable for DSM including hot water cylinders, fridge/freezers, heat pumps and air-conditioning. EVs are also a perfect load for DSM as generally they are a large load that can be charged at any point during the night as long as they are ready for the following day's commute. Large loads with high energy storage and good time flexibility are the best loads for DSM. For this reason EVs and hot-water cylinders are ideal. Also loads that only switch on for a short time periodically are suitable, as this switching cycle can easily be adjusted without the end user noticing a difference in operation. Loads in this category include heating, ventilation and air-conditioning (HVAC) as well as fridges and freezers. All together these make up a significant portion of residential and commercial load.

Power system peaks are an issue for a number of reasons. Firstly, the entire power system must be designed to handle the peak load expected, even if this peak is only reached for a small portion of the day or year. The peak to average ratio (PAR) is the ratio between the peak load in a power system, and the average system load. The lower the PAR, the more electricity is transported for a given maximum capacity, and therefore the more efficient and cost effective the system is. A lower PAR is also desirable with regards to generation, as peaking plants that are only online for a few hours per day are highly inefficient when compared to generation designed for steady output. Firstly, they represent a larger capital cost, but operate infrequently so the investment per unit output is very high. They are also less efficient in terms of fuel use when operating at partial load, and due to being cycled frequently they are also likely to fatigue faster.

Without DSM a power system is strictly one way, with generation always directly matched to supply regardless of how inefficient this may be. By adding DSM the power system effectively becomes two-way, with both the supply and demand side acting together to maintain system balance. This can reduce peaks throughout the power system while also more closely matching demand to supply, especially in the presence of fluctuating renewable generation.

It is apparent that in order to achieve this new functionality, additional control and communication hardware needs to be added to the power system and linked together with smart loads. The current state of DSM technology is not necessarily capable of providing the features

needed for a power system with increasing levels of renewable energy and electric vehicles, and this will be discussed more in the next chapter.

1.5 Research Contribution

A new DSM system, proposed as part of this research is called ‘Localised Demand Control’ (LDC). The aim of this research was to look at how electric vehicle chargers can be better integrated into the power system, especially in the presence of increasing renewable energy. LDC is the result of this aim, and was designed with many goals in mind, including protecting the local distribution transformer, monitoring the low voltage network, integrating local renewable generation and integrating EVs safely.

The system works by placing a device at the local distribution transformer to perform monitoring, determine appropriate load control actions and signal controllable loads. The system showed promise in simulated results and a laboratory system has been built along with a load signalling system. The laboratory system performs well and successfully manages controllable loads in response to simulated changes in local generation and operating constraints.

1.6 Thesis Outline

This thesis will be structured as follows. More detailed background information on renewable energy and power system operation, as well as an overview of current DSM technology and research will be presented in Chapter 2. In Chapter 3, Dynamic Demand Control will be discussed as the first technology considered for better integrating EVs and renewable generation. Localised Demand Control will be presented in Chapter 4 with a focus on the design, modelling and simulation of the system. In Chapter 5 the experimental system will be introduced, along with results from the system running under various typical operating conditions. The design of the control signal transmission system will be covered in Chapter 6. In Chapter 7 the work on practical smart loads, smart EV chargers and future work will be covered. A summary of thesis conclusions and research outputs will be presented in Chapter 8.

Chapter 2

Demand Side Management and the Electricity Grid

2.1 Introduction

In this chapter some of the present and future issues that power systems are facing will be discussed in more detail, with a focus on the New Zealand power system. These relate mainly to the integration of further renewable energy generation, and the prospective introduction of a large electric vehicle fleet. Potential solutions to these issues that involve demand side management (DSM) will also be investigated both in terms of the commercial systems presently available, and future systems detailed in the literature. General comments about the likely performance of these systems and their suitability as a tool for future power systems will also be made. The goal of this chapter is to establish that a state of the art DSM system could provide significant improvements to the power system, especially when combined with electric vehicles and further renewable generation, and that at the present time no suitable DSM system exists.

2.2 New Zealand's Electricity Grid

The power system in New Zealand (NZ) is an ideal place for testing new technologies, due to the system's small size, islanded nature and modern design. With an already high proportion of renewable generation, NZ could also be one of the first nations to demonstrate that 100% renewable energy is possible on a large scale.

NZ is a small islanded nation in the South Pacific with a population of around 4.4 million. In 2011 the annual electricity consumption was 43TWh or 155PJ [8]. Compared to the 2011 US consumption of 3,856TWh, this number makes the NZ power consumption and power system about 1.1% the size of the US. This shows a slightly lower per capita usage as the NZ population is about 1.4% that of the US.

Much of the data presented in this section has been obtained from the New Zealand Energy Data File 2012 [8], which is an in-depth resource of information on the NZ energy supply and power system. The focus of this section is NZ's power system with respect to the way in which

demand side management (DSM) schemes could be beneficial. DSM itself is discussed in further detail in the next section.

New Zealand has a very high proportion of renewable energy generation, generating some 77% of its electricity from renewable sources in 2011. A figure this high is especially notable given the islanded nature of NZ. Some smaller European countries have similar or higher proportions of renewable generation but these also have the advantage of being able to export electricity to the rest of Europe when there is a surplus and import electricity when the domestic supply is insufficient.

New Zealand does not have this luxury, and with the goal of further increasing the proportion of renewable generation, faces a significant technical challenge in relation to grid integration of further unpredictable generation sources.

2.2.1 Energy Sources

The NZ power system had an installed capacity of 9,751MW at the end of 2011 and experienced a peak demand of 6654MW in August of that year [8]. The total electrical energy generated was ~43TWh, and 7.3% was lost in transmission and distribution. The generation and capacity from various sources is shown in Table 2.1.

Table 2.1 – Electricity Generation in New Zealand, 2011

Source	Total Generation (%)	Energy (GWh)	Installed Capacity (MW)
Hydro	57.6	24831	5252
Gas	18.4	7955	1942
Geothermal	13.4	5770	731
Coal	4.7	2026	920
Wind	4.5	1931	614
Bioenergy	1.3	215	108
Other Thermal	0.1	108	184
Total:	100	42836	9751

In terms of total energy supply (not just electricity), renewable sources account for just under 40% of the total, as shown in Table 2.2. This is because oil is a huge source of energy, particularly for transportation. Electrification of the transportation network in New Zealand would significantly increase the amount of energy that could be sourced from renewable generation.

Table 2.2 – Energy use in New Zealand, 2011

Energy Source	Consumption (PJ)	Percentage (%)
Renewables	320.6	39.2
Oil	276.21	33.8
Natural Gas	158.68	19.4
Coal	60.74	7.4
Waste Heat	1.52	0.2

2.2.2 Renewable Energy

2.2.2.1 Hydroelectricity

Hydroelectricity accounted for 57.6% of electricity generation in NZ in 2011. For the last 20 years, between 75 and 100 PJ of energy has been generated each year from 39 hydro installations around NZ. There are limited locations that are suitable for hydro, and no new major installations have been commissioned since 1992.

2.2.2.2 Wind

Wind power has seen the largest increase of any renewable generation source in the last decade. Due to the low initial installed capacity, wind power has so far been quite economic to install. The first significant wind turbine was installed in NZ in 1993 and had a capacity of 225kW. As of 2011, there is 614MW of installed capacity. 1931GWh of energy was produced from wind power in 2011, and given the capacity figure, this results in a capacity factor of 35.9%. This is above the typical figure of around 30% for land based wind installations. Given NZ's situation it would appear that wind power is the most viable source for NZ to increase significantly the proportion of renewable energy generated, but as the proportion of wind gets higher, reliability issues could start being created.

A study was conducted by Garrad Hassan in 2007 on wind power variability and forecasting errors [9]. The impacts of various penetration levels of wind on the NZ power system and electricity market were investigated. They concluded that for most scenarios, the unpredictability of wind generation output would have a significant effect on the combined forecast error. It was also mentioned that a significant improvement in wind forecasting would be required to preserve the present overall forecast accuracy.

An interesting result was also reached with respect to ramp rates. The rate of change of a wind farm's output is of considerable concern to grid operators, since if it changes too quickly, other generation may not be able to keep up, and the system metrics could fall outside

acceptable limits. Assuming that most wind farms are on the order of 100MW and have reasonable regional diversity, it was found that load changes would still dominate the ramp rates on the 5 minute dispatch interval. Note that this is not necessarily the case for small scale wind installations as discussed later in relation to Low Voltage Network health. Forecasting difficulties and fast ramping of wind and demand make a reliable DSM scheme a valuable tool for the power system.

2.2.2.3 Geothermal

Geothermal generation is a stable and renewable energy source and NZ has a good amount with an installed capacity of 731MW. Generation in 2011 totalled 5770GWh resulting in a capacity factor of 90.1%. As it is such a stable generation source, the only benefit from the addition of DSM would be in the situation where demand exceeds supply, and the energy will be wasted if not consumed. While this may be unlikely given the low capacity, when both wind and geothermal are taken into account, non-storable renewable generation during the night could exceed demand, and if employed, DSM could make sure this energy is utilized instead of being wasted. This is similar to the situation in many other power systems with large amounts of base nuclear generation and increasing proportions of wind.

2.2.2.4 Other Sources

There are a number of other renewable generation sources, the most significant being bioenergy with 108MW of installed capacity. Many of these are wood fired generators using waste from the timber industry.

Solar power is not currently a significant source in NZ, but as the price of solar panels continues to decline, the economics are steadily improving. The supply curve of solar is also quite a good fit for the daytime load curve of most power systems (including NZ). As NZ does not have feed-in tariffs, there has not been the solar boom seen in other parts of the world such as Europe and parts of the US where tariffs are available. Without tariffs, a solar installation must be funded purely on its own merits, and here DSM has the capability to improve the load factor of intermittent sources and therefore improve solar economics.

2.2.3 Frequency Keeping and Load Balancing

In New Zealand, Transpower New Zealand Limited owns and operates the national power grid. Transpower is a state owned enterprise and in its role as system operator schedules and dispatches generation in the New Zealand electricity market and procures ancillary services

such as frequency keeping, instantaneous reserves, over frequency reserves, voltage support and black start capability.

For frequency keeping, a total of ±50MW and ±25MW are procured for the North and South island respectively from 4 companies operating 6 hydro stations and 6 thermal units. Multiple units can operate based on what combination gives the lowest cost for each trading period. Between 2010 and 2014 the cost of frequency keeping ranged from $39,831,913 to $55,940,913 [10] [11] annually.

Additionally, instantaneous reserves, over frequency reserves and black start services are also contracted out, with instantaneous representing the next highest cost, ranging from $29,161,768 to $38,766,777 [10] [11] for the same period. New Zealand does regularly experience large changes in frequency due to loss of load or generation. Examples include the Tiwai Point Smelter (630MW), the 220kV North Auckland and Northland line (900MW) or the HVDC bi-pole link (1200MW). The loss of the HVDC link at full transfer can cause South Island frequency to rise above 55Hz.

Around 20 companies provide these reserves, with some already utilising interruptible load. It is of interest to note that EnerNOC, a demand response aggregator is listed as under-provisioning capacity for 3 under-frequency events in 2012. This highlights the need for demand response services to accurately know what load can be dropped over the contracted period in order to provide accurate reserve estimates.

Total annual ancillary service costs between 2010 and 2014 ranged from $72,331,503 to $104,458,872 [10] [11].

2.2.4 Electricity Market

New Zealand has a highly competitive electricity market, with significant restructuring occurring in the late 80s and early 90s with this aim in mind. Power companies were forced to separate the parts of their business that could be monopolized from the parts where there could be competition. All of the lines companies are publically or trust owned, and multiple retail companies including those owned by generators, compete to sell electricity to the end user that has been purchased on the spot electricity market.

New Zealand has a location marginal pricing framework. Each generator offers generation at a certain price at specific locations (grid injections points). If dispatched by the System Operator, the generator is paid at the marginal price for their injection location. The System

Operator dispatches generation every half hour according to offers that minimise the overall cost of generation (including instantaneous reserves) to meet demand.

The NZ electricity industry participation code allows for certain generators (i.e. those that wish to generate regardless of price) to offer in at zero or $1c/MWh into the market. Wind generation is required to do so (i.e. price taker rather than price setter). This is done through the must run dispatch auction. Generation wishing to offer at zero or $1c/MWh must be successful in the Must Run Dispatch Auction, otherwise they can be curtailed. Wind farms are required to make a forecast of what they will generate, but are not held to this forecast. The System Operator relies to some extent on that forecast but reverts to a persistence forecast close to real-time.

Having "price taker" generation offers has the effect of lowering overall wholesale electricity prices, but this also makes the expected return of a wind farm harder to forecast. Again, if a significant amount of DSM were to be deployed, such that interruptible load reduces in cost and increases in availability, the consequences of forecast inaccuracies would reduce and wind farms could be deployed at a wider scale and with increased profitability.

In summary, a DSM scheme that can monitor and respond to sudden voltage and frequency fluctuations and consume power closer to the source of generation could help mitigate some of the issues with renewable generation.

2.3 Integrating Wind into the Power System

The difficulty of integrating wind power into a modern power system in meaningful proportions has been covered by a number of authors. As the authors of [12] discuss, the fundamental causes are the high inter-temporal variation and forecasting difficulty of wind, and the need for better primary stabilisation and frequency regulation. How difficult the wind is to integrate is largely dependent on the generation and load mix within the power system. Integration is easiest if there is flexibility in generation, and low variability or high flexibility in the load. Given the unpredictable nature of wind, problems can occur when there is either more or less power than forecast. If there is less power, then there must always be sufficient reserves to cover any shortfall in wind generation. If there is more power, then there must be generation that is available to be curtailed. Both of these requirements become significant as the proportion of wind power in the system increases.

Having generation that can easily be curtailed sounds simple, but there are many situations where reducing output is not economically viable. The author of [13] investigates the rise of negative power prices in Germany during high wind and low load conditions. In particular it is noted that nuclear power and thermal units that have reserve commitments cannot be brought below a certain level, and there are also limits on how fast thermal units can ramp up and down, and how much wear will be caused by thermal cycling.

The author of [3] discusses both the curtailment and provision of backup generation if the UK were to move to high levels of wind generation. It is noted that new wind farms would need additional thermal units as backup generation, and these units would have to be fast responding and therefore less efficient than other available thermal units. In the long run this could potentially result in more emissions than simply building efficient thermal units.

In the United States a study of the New York State area found that moving to 10% wind would be cost effective even without any forecasting, and that state of the art forecasting would improve profits by 55% [14]. A study by the National Renewable Energy Labs [15] found that in five western states, it is technically feasible to move to 30% wind and 5% solar, though a number of operational changes at unknown cost would need to be made to accommodate the variability.

2.4 The Low Voltage Network

Renewable installations are found not only at the grid level, but also at the residential level, and these present a different set of issues. Studies have been done on the effects of distributed renewable generation on the low voltage network, and in particular voltage stability issues. These generally arise from high generation when the network is lightly loaded as might be the case during the night with wind power, or in the mid afternoon with solar. One study on the effects of wind generation showed that voltage rises of up to 15% could occur, but that these events were very uncommon [16]. Another looked into the effects of solar generation and found voltage rises of 3-7% were likely, and that the most common length for voltage swings was on the 5-12 minute range, but some could be as low as 5 seconds [17]. The low voltage network is also of particular interest with respect to power system operation as roughly 90% of all power outages and disturbances have their roots there [18].

A DSM system that takes metrics related to LV network health into account when performing control could help stabilise the voltage here. The ability to respond to local

conditions is one of the factors taken into account when evaluating available and proposed DSM systems later in this chapter.

2.5 Transformer Overloading & EVs

It is not just renewable generation that has the potential to impact the power system. As electric vehicles (EVs) become more common, they too may cause issues. Boston consulting group has predicted that EVs could be 8% of the European market by 2020 [19]. A number of authors have investigated what impact EVs may have, particularly on low voltage transformers, since putting a higher load on the transformer will cause it to age faster. The authors of [20] predict that EVs would only lightly effect the local grid, while the authors of [21] note that a lightly loaded grid will have no issues at 100% EV penetration, but already heavily loaded systems will have problems. The authors of [22] investigated the relationship between temperature and transformer aging, concluding that in warmer climates, EVs could cause significant degrading of the transformer. The authors of [23] found that the aging of transformers is proportional to the THD of the charger current. The authors of [24] also looked at phase loadings, and concluded that limits on phase and temperature would be exceeded at ~25% EV penetration levels. The authors of [7] investigated the reduced lifespan of multiple network components and found EVs could have a significant impact on HV/MV transformers, and lesser but still noticeable impact on MV cables and MV/LV transformers.

In summary, for best results when integrating EVs into the power system, the current harmonics should be minimised, the transformer phases should be kept balanced, and there should be a means of handling transformer overloading. Again, if a DSM system is used that takes transformer loading into account, and EV chargers are used that integrate with the DSM system and also have a low THD, then transformer longevity can be maximised.

The potential issues created by EVs aren't just limited to the distribution network, but also overall generation. As noted by [25] EVs are usually parked up at times when the power usage is already maximum. In many power systems including New Zealand, there is an evening peak when people arrive home from work and an EV plugging in at this time would further exacerbate the issue. A study by Pacific Northwest National Laboratory in the United States found that there is enough spare generation capacity now to power 84% of the light vehicle fleet, but only if charging were managed in some way [26].

2.6 The Future Smart Grid & Demand Side Management

There is much commercial and academic work being conducted towards creating the "smart grid". There isn't an exact consensus on what is actually meant by this term, but in general it entails adding extra monitoring, communications and intelligence to the existing power network to make it more reliable and efficient. The authors of [27] and [28] give good overviews of the potential benefits afforded, with [28] noting that smart grids enable sending 50-300% more electricity down conventional corridors. Most of the prospective benefits of a smarter grid involve distributed resources (including demand) being included in the operation of the power grid, and this will be expanded on here.

In 2002 the authors of [29] noted that the technology was not quite present for a truly smart grid at the time, as current grid hardware had polling periods of multiple seconds and the computation needed for real-time control was very complex and poorly understood. It is only recently that the communications and computational technology have improved such that real-time control is now very achievable.

2.6.1 Demand Side Management

While the issues created by increased renewable generation and electric vehicle deployment could be solved by simply reinforcing network infrastructure and provisioning more backup generation, this is an expensive solution. A good alternative is DSM, where portions of the demand side are controlled to support grid operation.

As discussed in Chapter 1 DSM is possible because there are many residential and commercial loads that can be shifted in time without any impact on their overall operation. Any device that is storing energy can be delayed when the power system is stressed, or pre-emptively engaged when there is spare capacity. Examples of such loads, in addition to EVs, include space heating, hot water cylinders and fridge/freezers. These make up 34%, 29% and 10% respectively of present day household energy use in New Zealand [30] and a New Zealand Treasury report from 2005 also stated that the residential load is highly suitable for DSM [31]. The authors of [32] estimated that up to 33% of all loads could be used for DSM without significantly affected the end user. A study on loads in Belgium found a total of 358MW of load could be available for DSM. [33], and a study in Germany found that controlling EVs, whiteware and heat pumps could be used to plane the load profile, reducing peaks by up to 26% [34].

2.6.2 Deferrable Load

The entire concept of DSM is predicated on the fact that many high consumption household and commercial loads can easily be deferred without significantly impacting their operation or inconveniencing the end user. This is especially true of any device that is storing energy, as it can be delayed when the power system is stressed, or pre-emptively engaged when there is spare capacity.

Many proponents of renewable energy insist that batteries are the solution to the fluctuations in supply, but batteries are expensive and not necessarily very efficient. It is potentially far more cost effective to use controllable loads, and the energy that can be stored is significant as shown following.

2.6.2.1 Hot Water Cylinders

According to a fact sheet from BRANZ [35], the heating of water is done electrically in 88% of homes in New Zealand. This is the highest of any country in the world where statistical data is available. The most common hot water cylinder (HWC) sizes are 135L and 180L with 42% and 40% shares respectively. 180L is far more common in new installations though, with 46% share vs. 24% for 135L.

A 180L HWC represents a very large energy storage device. If the temperature could be varied over a range of 25°C, 5.225kWh (4.18J/gK x 25°C x 180L x 1kg/L = 18.81MJ) of energy storage is available with a common daily consumption of 5.5-11kWh per day.

25°C is a wide temperature band and represents a poor level of consumer comfort, but with temperature mixing valves on the HWC output this is achievable. A mixing valve may add ~$5 to the cost of an HWC installation but it has benefits aside from allowing a wider temperature band. By reducing the temperature of water flowing in the pipes, less heat is lost in transit so the system becomes more efficient. A mixing valve also adds safety, as if there is a failure of the HWC controller, temperatures that reach burning levels could occur at the tap without a mixing valve.

Most HWC's store water just above 60°C so that Legionella bacteria cannot live in the tank. The most common consumer uses of hot water are showers and hand washing, and these require water at 40°C. If a HWC is intelligent, the control mechanism could make sure that the temperature was brought above 60°C regularly to kill bacteria. Assuming a common heat loss of 5°C in the pipes, this means the cylinder can go as low as 45°C without impacting user

comfort. For the maximum temperature, the main constraints are efficiency and boiling point. Obviously the absolute maximum temperature before the cylinder boils is 100°C, but the standing losses of the cylinder are proportional to the cylinder temperature. If the average temperature was kept at 60°C the temperature could range from 47.5°C to 72.5°C (25°C range) with no added losses, but if the HWC is used for DSM, the electricity at certain times of the day could be very cheap and in turn make up for a less efficient operating point.

The ability of HWC's to participate in DSM and use renewable energy sources is causing other countries to encourage HWC use. This includes Denmark where many coal-fired combined heat and power plants are running only because they are the source of heating for the local area, and the electricity they produce is considered a waste product.

Most HWC's in New Zealand are already utilised for DSM with ripple control as discussed later, but if a more powerful DSM scheme were to be implemented, it would almost certainly use HWC control.

A HWC with a more variable consumption would potentially be better suited to the DSM schemes that are discussed later. Some techniques that could achieve this are:

- Integral cycle control of the heating element
- Splitting of the element into individually controlled smaller elements
- PWM control of the heating element
- A PFC front-end for a DC PWM controlled element

These all add extra cost in order to lower the impact on the power grid, and for any design a cost-benefit analysis would need to be undertaken.

2.6.2.2 Space Heating

With space heating representing 34% of household energy use [30], any reduction or deferral could have a significant impact on peak electricity demand. Heat pumps are becoming increasingly common for providing both heating and cooling, and according to a 2009 survey [36], are in 21% of homes in New Zealand – up from 4% in 2005. Not only are heat pumps more efficient than standard resistance heaters, but they are also more suitable for DSM due to electronically controlled pumps that could be varied in speed. A DSM controlled heat pump could reduce consumption during peak hours, or pre-heat in the time leading up to a demand peak to provide peak reduction without any reduced comfort for the end user.

2.6.2.3 Electric Vehicles

In a future where electric vehicles (EVs) are commonplace, a significant amount of controllable load could be provided if EV charging were managed. A typical EV such as a Nissan Leaf has an EPA rated fuel economy of 212Wh/km. If driven for the daily New Zealand average distance of 27.7km [37], it would use around 6kWh/day, representing 20% of the daily power consumption of a typical household [30].

Most users will require a minimum level of charge in their EV for emergencies, regardless of whether it is doing a smart charge or not. If 27.7km is the most common driving distance, and the EV has range of over 100km (120 for the Leaf) then the vehicle really only needs charging once per day, and can most likely be charged during the night even if minimum charge levels are considered. This coincides well with wind power, which commonly approaches zero cost at night due to such low demand. In fact one utility in Texas offers EV owners free night time charging as there is so much excess wind available [38].

An additional factor that makes EVs desirable for DSM is the charge rates that the batteries are capable of accepting. Taking the Leaf as an example, it is rated as having an 80kW motor and a 24kWh battery. This implies that the maximum rate of discharge the batteries encounter is 3.3C. As Lithium Ion batteries improve, safe charge / discharge rates are approaching 10C, which could mean a fast charge to 80% capacity in 5minutes or so. What this really highlights is the range of power consumption levels that EV chargers could employ when performing smart charging, and that the chargers are unlikely to be limited by the batteries, but by the power supply available.

In New Zealand the standard wall outlets have 10A circuit breakers, and high current outlets for HWC's, Dryers and Ovens are commonplace and range from 15-30A. Standard EV chargers are likely to sit just below 10A for an output of 2.4kW (note in Europe 13A output and 3.3kW is standard), and fast chargers are likely at the 7kW level using a 30A outlet.

As EV batteries are designed for locomotion, they can withstand lots of pulsing and variation in the charge rate. Even in the 7kW charging case, the rate is only 0.3C and nowhere near stressful for the battery. Again this means a well-designed EV smart charger could be constantly changing the charge rate over the full range of charge levels in response to a DSM signal without any expected additional battery wear.

The fact that charging will be electronically controlled means very little extra circuitry needs to be added to a charger in order for it to become a very powerful DSM device, with a high maximum power, fast response, fully variable power draw and a wide deferral window.

2.6.3 Benefits of DSM

As previously mentioned, many of the issues caused by renewable generation stem from the time mismatch between supply and demand, and time shifting loads can potentially solve this. The authors of [39] compared simulations both with and without DSM in the presence of higher levels of wind penetration, and found the wind was far more valuable when combined with DSM. The authors of [40] note that DSM is vital as Ireland aims to reach 40% renewable energy by 2020 and that the increased value of wind is greater than the cost of energy rescheduling.

A number of studies have also concluded that controlling the charging of EVs could significantly reduce the impact of the EVs on distribution transformers [20] [21] [22] [23] [24] [7] [41]. Furthermore, as controlled charging of EVs could be used to enable more renewable generation, EVs can be turned into a significant asset with respect to power grid operation as discussed in [42] and [43]. Many authors have also looked into the possibility of discharging EV batteries in order to return power to the grid at peak loading times in a technique termed vehicle to grid (V2G) [44] [25] [45] [46]. The full economic incentive for V2G is not yet proven due to the unpredictable nature of battery degradation.

It has also been proposed that DSM can provide better stability than equivalent generation, as large amounts of less reliable responsive loads provide better reliability that small amounts of more reliable generators [47].

In New Zealand there is also a direct economic incentive to reduce peaks within the transmission network, as Transpower charges parties connected to the national transmission network based on a half-hour average of either the 12 or 100 highest peaks (depending on location) over a year at the various grid exit points [48]. The authors of [49] looked into using DSM to reduce peaks in Portugal, and found reductions of 17.4% could be achieved by 2020.

A number of authors have looked into the economics of more general DSM systems. A simulation of the Italian power grid concluded that controllable load can compete with generators both in terms of performance and economics [50]. The authors of [51] performed a simulation on the economics of load shifting and battery storage in a micro-grid in the presence

of local renewable generation. They found that thermal storage devices are indeed economic to use, but that batteries may not be due to their high investment cost and short lifetime. The previously mentioned authors of [7] looked at the impact of EVs on the distribution system, and found that the NPV of network reinforcement was 20% better with controlled charging.

2.6.4 A Good DSM Implementation

The problem with DSM at the present time is the implementation. In order for DSM to work, there must by a method of deciding what the load behaviour should be, and then communicating the desired behaviour to the appropriate controllable loads. There are many options when it comes to DSM implementations and these will be covered in the next section, but first exactly what makes a good DSM system must be discussed. Response time, area of effect, cost, reliability and future proofing are all important metrics when evaluating a DSM system and will be expanded on here.

One of the major performance metrics is the speed of response. As has already been mentioned, the integration of renewable sources can be made easier with DSM. To accomplish this, the DSM system must be able to respond on a suitable time scale. A report investigating the impacts of a high amount of wind in the New Zealand power system predicted quite high variance on the 5 minute time scale [9], and as previously mentioned, variance on the low voltage network from residential installations was generally on a 5-12 minute time scale, with some fluctuations lasting as little as 5 seconds [17]. DSM systems could also be used directly for frequency regulation and emergency reserves. If instantaneous frequency responsive reserves are offered, then in New Zealand these operate on time scales of 0.4s, 6s and 15m [52]. The authors of [47] note that spinning reserve prices are typically 2-3 times that of non-spinning reserve prices. Overall, the faster the DSM system can respond the more services it can provide and the more valuable it is to the operation of the power system.

Closely related to speed of response is whether or not the system is automated. Some proponents of smart meters envisage efficiency savings from simply providing consumers with real time electricity prices and letting them alter their consumption accordingly. A Colorado University study on user participation with exactly this kind of system found that 90% of users barely used the information provided throughout the trial [53]. The GridWise project which tested a number of smart grid related technologies also concluded that automation was good for providing a consistent response [54]. In New Zealand, the majority of DSM that isn't implemented with ripple control or frequency measuring is done manually. A NZ Treasury

report from 2005 noted that no automated options were available at the time [31], and BuildingIQ, who sell energy management solutions have also stressed that the industry needs to move from manual DSM to optimised DSM [55].

Which communication technologies are used and where is a defining aspect of a DSM system. Furthermore, different systems can use different technologies for the various legs present in a power system. A simplified view of some of the originating and terminating points for communications are shown in Figure 2.1, along with common technologies for implementation. Less common links have been marked with dashed lines. Note that only one "System Operator" is shown here, whereas in actuality many different entities operate different parts of the power system. A Virtual Private Network (VPN) is used here as a catch-all for any custom communications infrastructure.

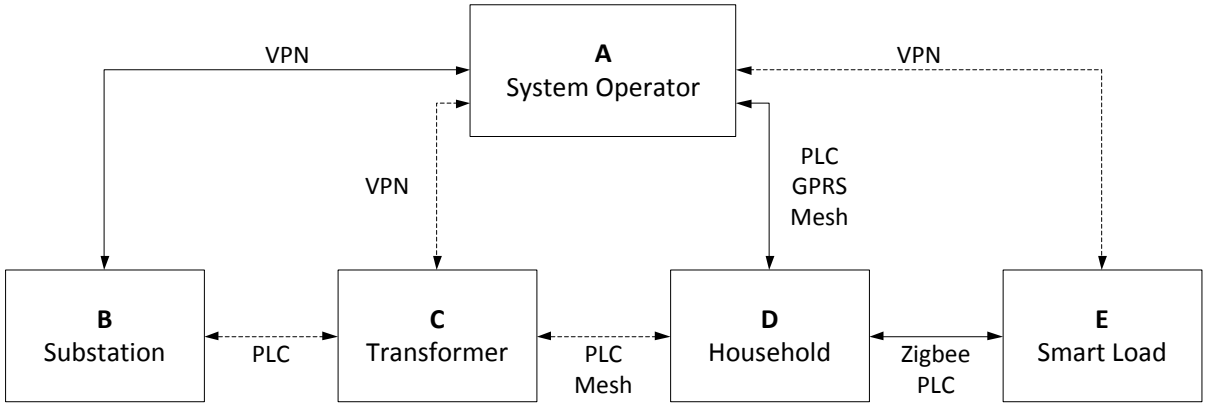

Figure 2.1 – Potential communication links in a Smart Grid

It is useful to have an idea of which technologies are suitable for which links, and also which systems leave parts unconnected. For example a smart meter enables communication from the operator to the premise (link AD), but unless it is equipped with additional hardware, it will not be able to communicate with smart loads. Also some smart thermostat systems use only the internet to connect the smart load directly to the operator (link AE) and in doing so bypass local hardware entirely. Which nodes DSM signals pass through and whether those nodes perform aggregation or modification of load control signals can impact how resource intensive a DSM system is and what area of effect DSM signals will have.

The area of effect is another metric referring to which power system constraints the DSM system can respond to. At the wider power system level, there is the overall instantaneous balance between supply and demand, but there are other constraints including the HV / MV / LV line loadings, and the loading of transformers between each of these that are just as

important, and require a very fast response. The more precisely a DSM system can be controlled in response to multiple constraints, the more useful the system is.

In addition to these application-specific performance metrics, there are also the more common metrics of cost, complexity and reliability, as well as harder to quantify metrics such as future-proofing. The authors of [18] note that whether technologies will be forward compatible with the future smart grid is a major consideration for utilities.

Another concern with any DSM system is that it could end up creating new problems if not designed correctly. One such problem is load synchronisation, where numerous loads all responding to a uniform steering signal all consume power at the same time, producing new peaks in the load profile. The author of [56] notes that under some circumstances, this could require even more line reinforcement than not using DSM at all. The authors of [27] mention a number of challenges that future smart grids may face including synchronisation. Also stated is the way in which a smart grid could give a profit motive for those wishing to exploit security vulnerabilities, and that new electronics could bring new failure modes into the power system. All of these potential issues mean that a DSM system for widespread use must be chosen and evaluated very carefully.

2.6.5 Present DSM Systems

Currently available or widely used DSM methods include ripple control, smart meters, dynamic demand control and internet connected devices. A defining aspect of a smart grid is the communication medium, and this is a good way of classifying the various systems. Ripple control is one of the oldest and most widely deployed DSM methods, and uses a simple one way ripple signal injected at the substation level onto the power line. Smart meters generally use some form of power line communications (PLC) or radio link in built up areas and GPRS modems in more rural areas [57]. Dynamic demand control uses the system frequency (and potentially voltage) to determine load behaviour, and internet connected devices use the consumers own internet connection for load signalling.

2.6.5.1 Ripple Control

Ripple control is a system whereby a higher frequency signal is injected into the power system at the substation or distribution level in order to shut off or limit the consumption of residential and commercial hot water cylinders during peak load periods. Ripple control is very common in New Zealand and is widely used for this purpose. For example in the Orion service territory which has around 190,000 customers, about 90% of residential electric hot water

heaters are controlled through ripple control [48], and ripple control or an equivalent is required on all new connections with hot water heaters over 1.2kW.

Ripple control is a very good system for the purpose it was designed for, namely managing overall peak loading over a wide geographic area. In the future, the prevalence of distributed generation and electric vehicles may result in a desire for a more precise demand control system that can take into account more localised conditions and react more quickly to transients. As ripple control takes between 5 and 12s to send out a message, it cannot really respond to events on a sub minute basis, especially if multiple messages were to be sent from a single transmitter in order to control multiple load areas. As mentioned at the start of this chapter, common time periods for distributed generation fluctuations range from tens of seconds to tens of minutes.

2.6.5.2 Smart Meters

The smart meter is one of the most obvious candidates for DSM functionality, though the majority are currently designed as advanced metering infrastructure (AMI), which allows the customers' consumption to be read remotely without the need for a physical meter inspection. For DSM concerns, there are two communication links of interest. The first is the connection from the utility to the meter (link AD in Figure 2.1), and the second is the connection from the meter to the controllable loads (link DE). The data connection to the utility is usually implemented with a PLC, but wireless systems can also be used. The authors of [58] give a thorough overview of the PLCs suitability for use in the smart grid. It is noted that some PLC types are well suited to the application, but many vendors are using the wrong type. In addition, there are numerous competing and incompatible standards. The authors of [59] investigate the reliability and cost of wireless smart meter communications, and conclude that the cost of smart meter communications failure is high, and therefore redundancies should be included. The authors of [60] investigated using the customers own broadband connection for connecting to the utility. A concern with this setup is that emergency signals may be significantly delayed, but it was predicted that with prioritised alarm traffic, node delays of ~11ms for emergency signals could be achieved.

Many smart meters do not have any DSM functionality, but those that do generally include a household wide cut-off switch, separate switchable circuits or a home area network (HAN) to enable digital communication with controllable loads. In 2009, the New Zealand Parliamentary Commissioner for the Environment raised the issue of smart meter deployments lacking HAN functionality [61]. A 2011 follow up stated that the 530,000 meters deployed thus

far in New Zealand had no HAN functionality, but that a planned future roll out of 750,000 meters were to be HAN functional [62]. As of early 2013 this second roll out is still in the planning and testing phase. Additionally, the EU is now requiring member states to roll out smart meters to 80% of consumers by 2020 [63], but the specified minimum functionality makes no mention of any DSM, making widespread inclusion unlikely.

The most common HAN for smart appliances is Zigbee, which for smart grid related applications uses the Smart Energy Profile (SEP) standard. While this does offer some DSM functionality, there is a problem with different versions of the SEP profile - SEP versions 1 and 2 are incompatible with each other. Currently deployed smart meters mostly use variants of version 1, while other utilities are waiting for version 2 to be finalised before beginning smart meter roll outs. Researchers using Zigbee have also found problems with security [64], network overhead as the number of nodes increase [65], and interference in the presence of Wi-Fi [66].

2.6.5.3 Time of Use and Real-time pricing

One of the simplest forms of DSM is Time of Use (TOU) pricing. In this scheme, there are predetermined prices allocated for certain periods of the day and these are not regularly changed.

TOU pricing is sometimes accomplished with two meters for different periods of the day, or a smart meter that is aware of the pricing scheme. There are multiple ways in which the demand response can be implemented. The simplest is where the end user knows or is reminded of the pricing level, and adjusts their behaviour accordingly. Simple demand response can also be achieved with a timed switch or time delay on an appliance. Alternatively, a smart meter may communicate with smart loads via a HAN as previously mentioned in order to perform the demand response. A number of cities in China have had good success with TOU pricing. In Beijing load shifting with a simple peak and off-peak pricing scheme has reached 200MW [67].

Real-time pricing (RTP) is quite similar to TOU, but the prices are updated anywhere from weekly to hourly, and they include a set window for which they are valid. A smart meter may receive this data via a dedicated communication channel or via the internet, and can then either relay this directly to the end user in a non-automated system, or directly signal loads as in an automated system, which as discussed earlier are much more desirable. Some NZ retailers do currently offer exposure to half hour wholesale prices with reduced margins.

The problem with an automated response to a price signal is that the signal needs interpreting. This requires either guaranteed pricing for the time period of interest (effectively TOU), or some statistical knowledge of the pricing signal. The authors of [68] proposed a predictive energy scheduling system for RTP that either uses guaranteed pricing data or makes predictions based on the last week of known data. The impacts of different time periods were simulated and compared to the base case, and the same simulation running with a simple TOU scheme. They found that the TOU system reduced costs to the end user by 5.7%, and that the RTP scheme with a 24 hour window reduced costs by 25%. Interestingly, when the pricing window is reduced to 2 hours and is supplemented with prediction, the costs are still 23.7% lower than the base case.

A RTP system could be a good solution to the growing proportions of renewable energy in the power system, but again it generally works over a wide geographic area and is not necessarily useful for keeping the LV network in balance.

2.6.5.4 Dynamic Demand Control

Dynamic Demand Control (DDC) is a method of DSM whereby the amount of power a controllable load takes is proportional to the system frequency. Since system frequency is an indication of the balance between supply and demand, adjusting the load in response to frequency should result in a stable operating point provided that the combination between latency and system inertia allows it.

DDC is very common in standalone micro-grids, where swings in the frequency are more pronounced and where there is an exact relationship between frequency and the supply/demand balance. A 1985 analysis of using frequency responsive integral cycle control showed it was a practical possibility for low-cost power systems, and highly advantageous in small stand-alone applications [69].

A number of groups have proposed the use of DDC controlled loads to integrate more wind into power systems. A simulation of DDC fridges in the UK showed that they could provide faster and more reliable response than current spinning reserve, and a good response during a sudden drop in wind power [70]. The authors of [32] and [71] also mention that frequency responsive loads can have a very fast response, and that EVs could provide <1s response times and can respond faster than traditional ancillary services. A simulation of Hokkaido Island in Japan found that if all hot water cylinders were controlled, the maximum allowed wind in the power system could go from 250MW to 675MW. The power demand of the island varies

between 2500MW and 5700MW. It was also found that the cost of the system would be 1/26 the cost of an equivalent battery storage system [72]. The authors of [33] looked into the amount of potential controllable load for integrating more wind in Belgium, finding that up to 358MW could be provided by residential appliances. The authors of [73] simulate using both frequency and voltage responsive loads to smooth power system fluctuations in the presence of wind. Perturbations in frequency, voltage, power and reactive power are all significantly reduced.

The authors of [74] also investigated using current smart meter technology for DDC, using the frequency monitoring already built into the meter. It was found that the response of the meters was too slow, taking between 3.3 and 4.6 seconds to respond to frequency changes. This is because delays of greater than 1 second are not at all useful and delays of less than 200ms are needed for primary response

DDC is a good potential solution to the issues of renewable generation and electric vehicle integration. A full analysis of DDC, including advantages and disadvantages is given in Chapter 3.

2.6.5.5 Automated Under Frequency Load Shedding

In New Zealand and many other power systems around the world, automated under-frequency load shedding (AUFLS) is used as a last resort emergency system to restore the power system balance under extreme loss of generation conditions. AUFLS works by fully disconnecting large areas of consumer demand when the frequency drops below a certain threshold.

In the North Island the AUFLS system is separated into two blocks each representing 16% of consumer demand. The first will trip at 47.8Hz and the second at 47.5Hz. On the 13th of December 2011, after an equipment failure at the Huntly power station, 870MW of generation was lost and the frequency dropped to 47.63Hz within a few seconds as shown in Figure 2.2. At this point ~30% of residential load was shed across the whole North Island of New Zealand to restore the frequency while other generation was brought online [52].

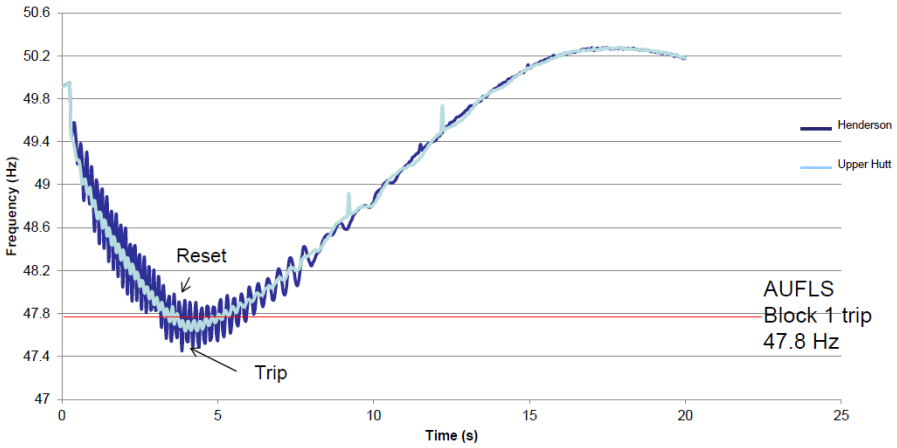

Figure 2.2 – Measured system frequency during December 2011 event

Obviously since AUFLS is such a blunt tool, it is almost not DSM and is really only useful for emergency situations. It is worth mentioning however, because any other DSM scheme that also took into account system frequency could produce the same results as AUFLS without impacting the end user at all, given more than 30% of residential load is deferrable.

2.6.5.6 Proprietary DSM Systems

In the absence of a universally accepted method of DSM, many companies have created their own proprietary systems. While these systems may solve an immediate need, they are generally tailored to a specific task (such as space heating and cooling) and cannot be extended to other devices, or work with other manufacturers products.

A number of utilities have partnered with smart thermostat providers to create their own custom demand response programs. For example three utilities in Ontario have partnered with Honeywell to offer free smart thermostats in exchange for periodic control of customers heating [75]. In New York, Con Edison gives customers a smart thermostat from Carrier and a $15 per month rebate in exchange for the ability to control their air conditioning periodically. The rebate is paid regardless of whether any load control is performed. They currently have 25,000 customers and up to 65MW of controllable load available [76].

Some of the company's manufacturing appliances have started creating their own standards for smart appliances, with Wi-Fi as the communication medium. Both Samsung [77] and Panasonic [78] announced separate smart appliance frameworks in 2012. Samsung updated theirs in 2014 with a new protocol and more features [79]. These appear unlikely to find widespread deployment as the smart appliances carry a major premium over standard models, and the smart appliance standards being created are incompatible with each other.

2.6.6 Summary

Overall the state of current smart grid and DSM technology is not nearly as smart as it could be, and no existing system provides the feature set that current communications and microcontroller hardware is capable of providing.

As mentioned in the previous chapter, a new DSM system is proposed in this thesis and will be described more fully in the coming chapters. The system was designed to have a fast enough response to integrate both small and large scale renewable generation, provide emergency reserves and frequency regulation, and have the ability to respond to power system constraints in both a localised and global manner.

2.7 Future DSM Options

Many researchers have also come to the conclusion that the current state of DSM technology needs to be improved, and have proposed both improvements and completely new systems. These will be summarised with a focus on whether these systems could lead to suitable DSM for the future power system.

2.7.1 New Algorithms for DSM Improvements

A number of authors have applied computational methods in an attempt to optimise energy consumption for a given scenario, assuming a suitable DSM system is already in place. While the results are generally promising, the actual method of implementing the DSM is not usually covered. The algorithms are of interest because they show what possibilities can exist in the presence of a high performance DSM system.

The authors of [80] present two algorithms for DSM, and compare the stability of each with varying levels of signal noise added to the system. The two algorithms have slightly different results, with one reducing cost to the minimum possible, and the other focusing on minimising component stress. The authors of [81] claim to have an optimal scheduling method for an environment with known electricity prices and controllable load requirements. Risk aversion for the consumer is also included. In [82], game theory is used assuming the presence of a two-way communication medium, and the peak to average ratio is successfully reduced.

As research in the area of the smart grid and DSM is highly active, a number of new works have appeared in the literature during the course of this research, with some including computational optimisations. The authors of [83] present an optimisation of EV charging taking into account both price constraints and user preference. The authors of [84] present another

optimal EV charging algorithm, where the EV charging requirements are aggregated and the charging times scheduled to reduce peaks while also ending within at least 4% of the desired charge level. An optimisation where distributed MPC is used to determine the operation of controllable loads and distributed generation is presented in [85]. Realistic data is used to show the systems potential. Realistic wind data is used in an optimisation presented in [86]. Here the optimised DSM significantly improves the system performance.

In [87], a novel DSM method including game theory is simulated in an isolated micro-grid with local wind and conventional generation. It is shown that energy costs can be reduced by 38% with good wind prediction, and an additional 21% with perfect knowledge of the wind.

The authors of [71] present a very interesting DSM solution. A combination of DDC, and digital communications are used in order to perform scheduling and frequency regulation. It is noted that other constraints could be taken into account but this has not been demonstrated.

2.7.2 Smart Grid Communication

As the communication medium is a defining aspect of a DSM system, a number of researchers have looked into the different methods that could be used for load control. As expected, the suitability of various solutions depends on the desired performance of the overall system.

In [88], the authors looked at implementing a very simple form of DSM, with required response times of 30 seconds. The conclusion was therefore reached that many technologies could provide that level of performance. The authors of [89] looked at systems for coupling just EVs and renewable generation, and also concluded that even very simple communications schemes could defer the network expansion EVs are likely to bring. It can be observed that where authors assume many communication technologies are sufficient, a relatively low level of performance is expected.

A number of researchers have looked into using internet based communications for DSM with varying conclusions. In [90] the authors argue that smart grid communication needs to be separate from the wider internet due to performance, reliability and security requirements. More specifically, the authors of [91] looked into using the DNP3 protocol over TCP/IP for all smart grid communication. The conclusion reached is that TCP/IP is not suitable and perhaps a profile based scheme based on IEC 61850 would be more suitable. As previously mentioned,

in [60] it was concluded that with prioritised alarm traffic, node delays of ~11ms for emergency signals could be achieved which would generally be an acceptable level of performance.

The authors of [92] did an interesting study on exactly what effect communication delays have in a micro-grid environment where coordination and control of local generation and loads is desired for smooth power flows. It was found that with 200ms delays, there is good performance, but with 2s delays there is some ringing, and with 4s delays, there is severe oscillation.

2.7.3 DSM Case Studies

There have been case studies performed where large scale DSM systems have been tested. Perhaps one of the most important is the previously mentioned GridWise project [54] which came to some useful technical and social conclusions. It was found that DSM could perform a number of practical tasks, such as managing a feeder constraint, integrating distributed generation and reducing the peak load. In terms of consumer response, it was found that residents would eagerly participate in price responsive contracts and shadow markets, and that automated response were the most consistent [54].

In [93] the authors optimised the actual HVAC system in a building in conjunction with a simulated battery. It was found that if a battery was used for load levelling, the addition of DSM reduced the storage requirements by 20%.

In [94] an energy optimisation was conducted on a hospitals' full energy supply, including local generation, boilers, gas supply and electricity supply. Consumption was estimated and optimised two days in advance. The total energy costs for the building were reduced by between 6% and 12% between 2004 and 2010.

2.8 DSM at the Transformer Level

As mentioned at the start of this chapter, one of the most significant concerns in the context of EV and renewable deployment is the loading of the local distribution transformer and the health of the LV network. From the overview of available DSM technologies, it has been shown that there are not yet any DSM systems available that include the transformer stress in determining load control actions. At the outset of this research, there was also limited discussion in the literature on this issue, but the topic has since become more common.

2.8.1 EV Charging Simulations

Some authors have simulated the effects of controlled EV charging on the distribution network. A study on the distribution system in Stockholm concluded that many areas could manage 100% EV deployment with the current grid capacity at the MV level, but this would not necessarily carry over to the residential (LV) level. Additionally, losses with controlled charging would be reduced by 0.5% [41]. A study on Guwahati City in India also found that controlled charging would reduce power loss and could flatten the voltage profile at the distribution level [95].

More authors have since begun discussing EV control at the transformer level as a solution. The simplest solution proposed is simply to limit the charge rate of the EV as discussed by [96] and [97]. In the later it was found that by limiting charging to 373W, 99.4% of EVs would be fully charged overnight and degradation of the LV transformer would be limited.

In [98], it is proposed that by simply delaying the EV charge cycle by 1.5 hours, transformer overloading can be prevented. This is contradicted by [22], where it is shown that delaying could actually increase the aging of the transformer, and a solution that takes into account transformer temperature is proposed as the ideal option. The authors of [43] also mention that when using EV charging to improve wind power utilisation, local conditions can be taken into account to improve system performance.

While these are all useful studies and are full of good ideas, the studies mentioned thus far are limited in scope to include only EVs and not the full range of household loads that could be utilised. Also due to the nature of these being simulations, no hardware or physical designs have yet been presented.

2.8.2 DSM Simulations

Recently, there have been some authors that have proposed systems that do go beyond just EV coordination. In [99] an optimisation scheme is proposed for commercial buildings whereby many loads including the HVAC system are controlled and utility network constraints are taken into account. The authors of [100] also discuss general load control at the transformer level. A piece of hardware was proposed that is placed at the transformer to coordinate controllable loads in a manner similar to what had already been proposed by the author and which is presented in this thesis. It is noted that the full hardware system does not yet exist, but the individual parts of the system are readily available.

An interesting piece of work is presented in [101] where an algorithm is proposed for making better use of distributed generation with local controllable loads. It is concluded that both in terms of communication and power flow, it is more efficient to consume power within the local grid than to share it with the wider grid.

Another interesting piece of work is the TRIANA system developed at the University of Twente. The design of an implementation as published in 2012 is well thought out. An energy management unit is placed in each household in order to measure, predict and signal all of the controllable loads within the house. The predicted load profiles are then sent to the local controller where they are aggregated and sent to the global controller. This creates a hierarchical structure that is extensible with more levels and also frugal with its communications needs. The global controller takes the aggregated load profiles and calculates the desired consumption patterns. Steering signals are then generated and passed back down the chain of controllers. The steering signals can be adjusted at each level in accordance with local conditions. The 2012 paper includes an overview of how the TRIANA system works as well as simulations covering the control and optimisation of heat pumps [102]. There are a couple of preceding works by the same authors focusing on smaller parts of the system. Work on basic cost optimisation in the presence of distributed generation, energy storage and controllable load was published in 2009 with peak shaving and power limiting investigated [103]. Further work on load prediction and optimising local generation was published in 2010 [104].

While no actual systems have been demonstrated, the proposed setup if implemented well could indeed provide a very capable DSM system. A number of these works have been summarised and compared with the design of LDC in Section 4.10.

2.8.3 DSM Hardware

A couple of authors have proposed hardware designs that are approaching the desired feature set by taking into account transformer conditions, but they are still relatively simplistic.

The authors of [105] demonstrate a load control system with smart plugs that sit in between the appliance and the power supply. The smart plugs monitor the system frequency and implement a DDC like action for real-time control. They can also have a HAN connection for performing other smart appliance related tasks. It is stated that a simple power line communication scheme is used for communication between the smart plug and the appliance, with few details given. A later work by the same author goes into more detail by proposing a

new low cost communication scheme for the so called “last meter” smart grid communications [106].

The authors of [107] also describe a smart plug based system that communicates with a local energy manager. The smart plugs either directly control the power supplied to the load, or communicate with it through a unspecified smart communication interface similar to the previous work. It is stated that the energy manager can receive “external load control signals” but other than the ability to be localised these are not expanded on further.

In a paper published in 2013 [108], the authors use voltage based droop control for DSM within a micro-grid. While the implementation method of using a multi-tap transformer for voltage control is questionable, the control methodology is similar to that proposed in this work. It does appear that the authors are attempting to solve exactly the same problem identified here.

2.9 Conclusion

The need for and benefit of a comprehensive demand side management system has been presented. The current state of DSM technology has been covered, as well as some of the features that are lacking in current systems.

Related research where improvements to DSM are proposed which aim to provide more features have also been summarised. It was found that while many authors have identified the similar shortcomings with current technology, very few have come close to creating a viable DSM solution that significantly improves on what is currently available.

Some of the work that has been published towards the end of this research is similar to the DSM system that will be discussed throughout this research, but no work has been found that is overly similar to the point that this work is no longer unique. Similar works have been summarised and compared in Section 4.10.

Chapter 3
Dynamic Demand Control

3.1 Introduction

In this chapter the concept of dynamic demand control (DDC) will be introduced as the first potential DSM solution to the power system issues caused by increased deployment of renewable generation and electric vehicles. The basic control methodology will be demonstrated, and then experimental results from a laboratory scale DDC system with simulated wind generation and an EV charger. A simulation where DDC is used in a slightly more complex power system will also be shown, and conclusions made about the expected performance of DDC in an actual power system.

DDC is a method of DSM originally patented by Fred Schweppe in 1982 [109]. In operation it works by simply adjusting the amount of load taken based on the measured system frequency. In simple terms, a power system can be considered as a rotating mass with various torques applied to it. Electricity generators apply positive torque so as to increase the frequency, and electrical loads apply negative torque which decreases the frequency. When these torques are equal, the frequency stays constant, and when there is an imbalance, the frequency will change. Normally a power system will have frequency responsive generators that constantly adjust power output in order to keep the frequency at the nominal value, which in New Zealand is 50Hz. As mentioned in Chapter 2, additional renewable generation can require additional frequency keeping as the generation output is often unpredictable.

Since system frequency is therefore a direct indication of the balance between supply and demand, adjusting the load in response to frequency should result in a stable operating point, provided that the combination between latency and system inertia allows it. One such example of a DDC scheme could be where a load draws zero power at 49Hz or below and full power at 51Hz or above as shown in Figure 3.1. As will be shown later, the stability of a DDC system is dependent on the system inertia. The lower the system inertia, the wider this frequency band has to be in order to maintain stability.

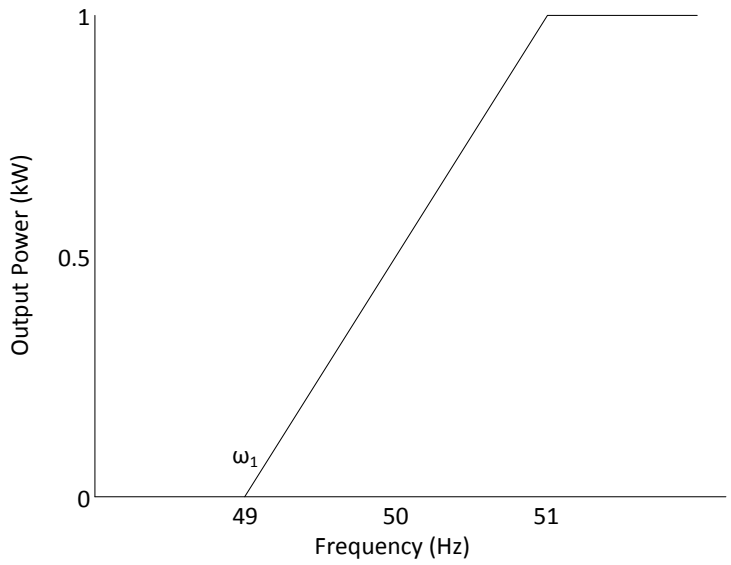

Figure 3.1 – Simple DDC controller response

If a significant amount of deferrable load was to be controlled in this manner, significant damping of the power system frequency would be provided. In turn, system wide fluctuations in generation from sources such as wind turbines would be absorbed without the need for extensive frequency regulation from other sources. This is a very desirable result given the attention that wind and other fluctuating generation sources currently receive.

3.2 DDC Control Structure

Consider a stable power system, operating with some set base generation and base load, as well as varying generation, varying load and frequency responsive load. If the fixed items are ignored, this can be simplified in the control diagram shown in Figure 3.2, with symbol definitions as given in Table 3.1.

Changes in either generation or load are represented by $\Delta\tau$, and the system inertia is given by J. Changes in torque will produce a change in the operating frequency of the system, represented by $\Delta\omega$, which here is the difference between the current frequency and the minimum operating frequency (ω_1):

$$\Delta\omega = \omega - \omega_1 \tag{3.1}$$

A frequency responsive load would draw no power at the minimum operating frequency, which in Figure 3.1 is 49Hz and is labeled as ω_1.

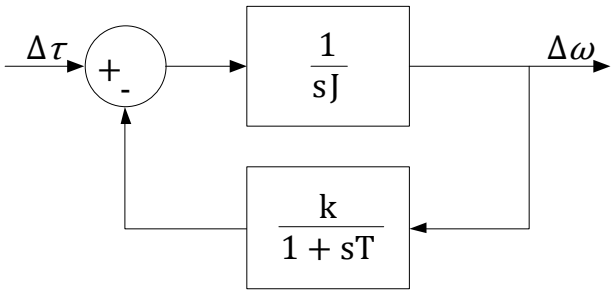

Figure 3.2 – Simplified DDC block diagram

Table 3.1 – Symbol definitions for DDC block diagram

Symbol	Definition
$\boldsymbol{\tau}$	Total system torque
$\Delta\boldsymbol{\tau}$	Change in input torque
J	System inertia
$\Delta\boldsymbol{\omega}$	Change in system frequency
k	Constant for conversion between frequency and torque
T	Filter time constant

In the diagram there is also a feedback path containing $\frac{k}{1+sT}$ which represents the controllable load. If the response of the controllable load was simply linear as given in Figure 3.1, the feedback path would just contain k, a constant of proportionality between frequency and power output, but a low pass filter is added for stability. Again for the example of a controllable load as shown in Figure 3.1, where a maximum power draw of 1kW occurs at 51Hz (labeled as ω_2), k can be calculated as follows:

$$P = k\Delta\omega = k(\omega - \omega_1) \tag{3.2}$$

$$k = \frac{P}{\omega - \omega_1} = \frac{P_{MAX}}{(\omega_2 - \omega_1)} = \frac{1000}{51 - 49} = 500 \tag{3.3}$$

The additional low pass filter has a number of benefits. Firstly, if inductance machines are being used to generate electricity, there are some non-linear effects that must be considered. Sharp changes in the torque applied to the machine cause a change in the phase of the voltage waveform due to the armature reactance, and the measured frequency can therefore shift temporarily. If small changes in measured frequency cause a large shift in power consumed (and therefore torque) by controllable loads then the system could become unstable. In addition, if very simple frequency measurement circuitry is used, then noise on the voltage waveform could produce erroneous measurements with unknown results. By adding a low pass filter, the impact of these effects can be dominated by the filter time constant and the system can be kept stable.

The addition of the controllable load adds negative feedback to the system, and the dynamics and stability of the system can now be calculated by determining the transfer function. The input of the system is $\Delta\tau$, and the output is $\Delta\omega$, such that the transfer function represents the response of the system frequency to random changes in load or generation.

The transfer function of the system with this low pass filter is calculated as follows:

$$\tau = \Delta\tau - \frac{k}{1 + sT}\Delta\omega \tag{3.4}$$

$$\Delta\omega = \frac{1}{sJ}\tau \tag{3.5}$$

So that

$$\frac{\Delta\omega}{\Delta\tau} = \frac{1 + sT}{s^2JT + sJ + k} \tag{3.6}$$

Which has a damping factor of:

$$\zeta = \frac{1}{2}\sqrt{\frac{J}{kT}} \tag{3.7}$$

The system performance is therefore dependent on the system inertia, filtering constant and available controllable load as given by (3.7). High gains and short time constants giving rapid response and high accuracy come at the expense of a low damping factor that is not acceptable. Thus the essence of control here is to have enough inertia in the local power system. Experimental validation has shown that using a digital low pass filter with a maximum delay of 0.16s gives an acceptable response for inertias of greater than 0.02 kG.m^2/2-pole kW. Since inertia increases as the 4th power of the machine diameter times the length, inertias are more readily achieved with larger machines. This is one of the reasons why DDC is easier to implement in large systems with high inertia.

3.3 DDC Experimental Performance

A classical DDC controller has been tested under laboratory conditions and by computer simulation. In the laboratory a controlled AC drive in a torque limited mode generated a string of random torques changing each second to simulate a variable wind regime. The AC drive (variable speed induction motor) was connected to a 3 phase alternator generating at nominally 50Hz to form a motor generator (MG) set. Two of the phases were on resistive loads, and the

third phase was passed to a DDC controlled load set up to emulate an electric vehicle battery charger. The VSD, loads and MG set are shown in Figure 3.3. Here the DDC controller adjusts the charging rate to keep the frequency at the nominal 50Hz.

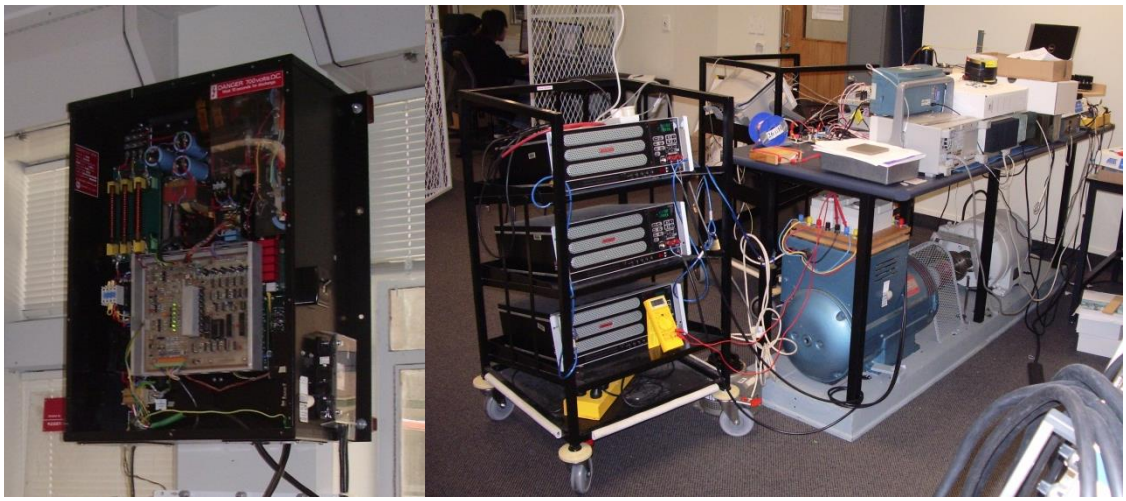

Figure 3.3 – Experimental setup of DDC system showing VSD (left), loads (center) and MG set (right)

The measured and computer simulated outputs are shown in Figure 3.4. An advantage of this experimental set-up is that the same random sequence can be used for all of the tests thereby allowing comparison of performance for different system parameters. In this experiment the operating frequency is allowed to vary and the system inertia is an important part of the damping factor for the controller.

In Figure 3.4, plot (a) shows the random torque signal, (b) shows the generator frequency, and plots (c) and (d) show the simulated and actual power consumption of the DDC load. As can be observed, the simulated output created with MATLAB Simulink® is a close fit to the measured data with the same average current and slightly less variation. This shows that the inertia figures for the experiment and the simulation are not quite identical and that the DDC controller behaved as expected to stabilise the system frequency. As previously mentioned, the stability of a DDC system is dependent on the system inertia. If DDC can be used to stabalise the operating frequency in a very small system, it will likely work even better in larger systems, especially as inertia scales exponentially with generator size.

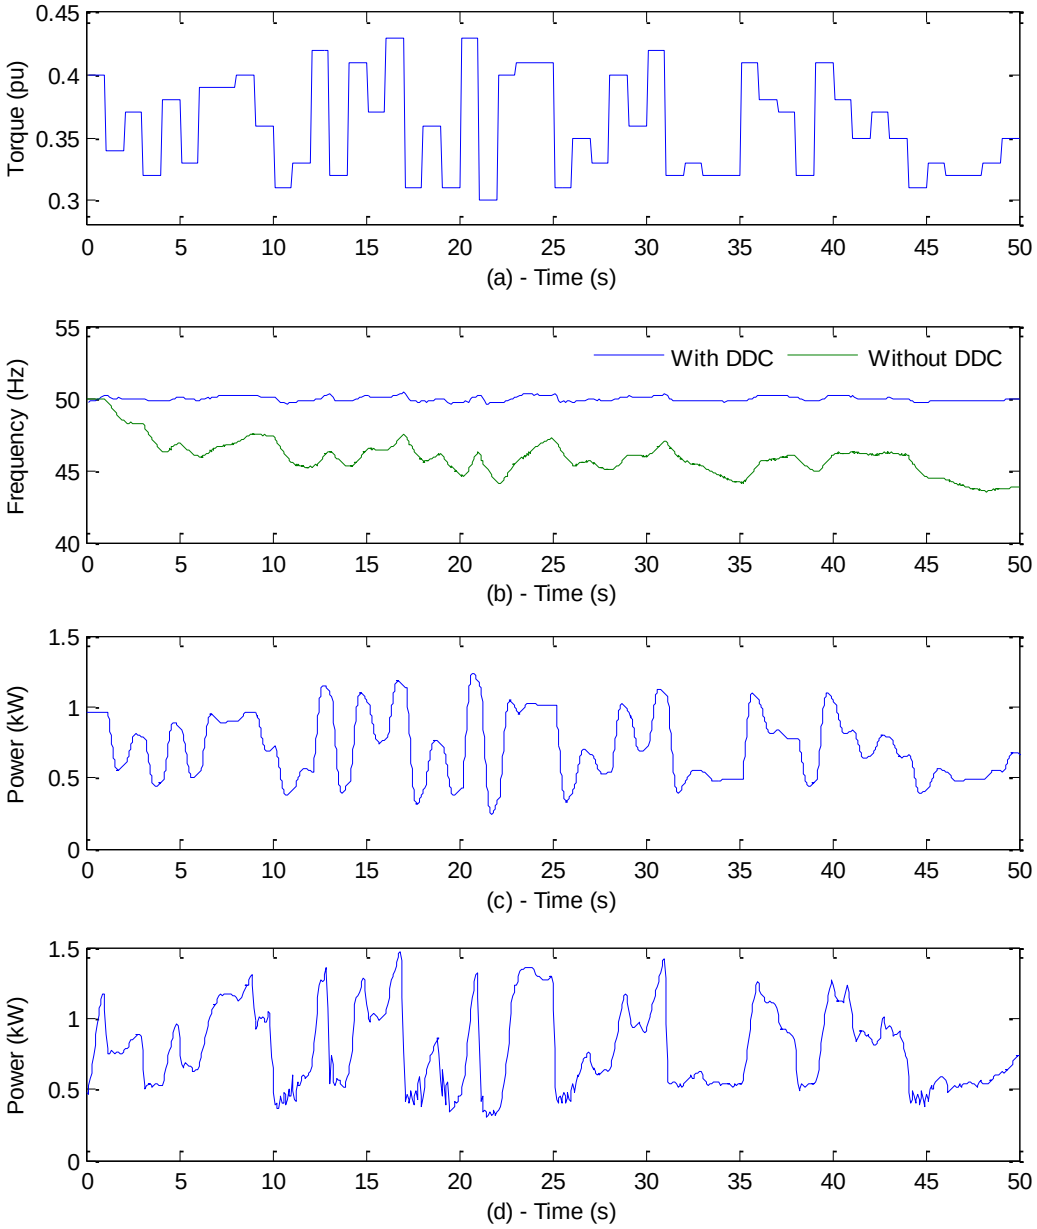

Figure 3.4 – Plot of DDC results showing (a) input torque, (b) system frequency, (c) simulated output and (d) actual output

3.4 Simulated Grid Performance

Further simulations were conducted in order to assess the performance of DDC in a grid environment. As mentioned in Chapter 2 there are a number of situations where a DSM system could assist in the operation of a power system. Of particular interest is the performance of DDC when integrating both renewable generation and electric vehicles. In the simulations here these were introduced both with and without DDC such that the results could be compared. The layout of the simplified power system simulated here is shown in Figure 3.5. Only 20 houses were simulated, each with a hot-water cylinder, electric vehicle charger, refrigerator and some base load.

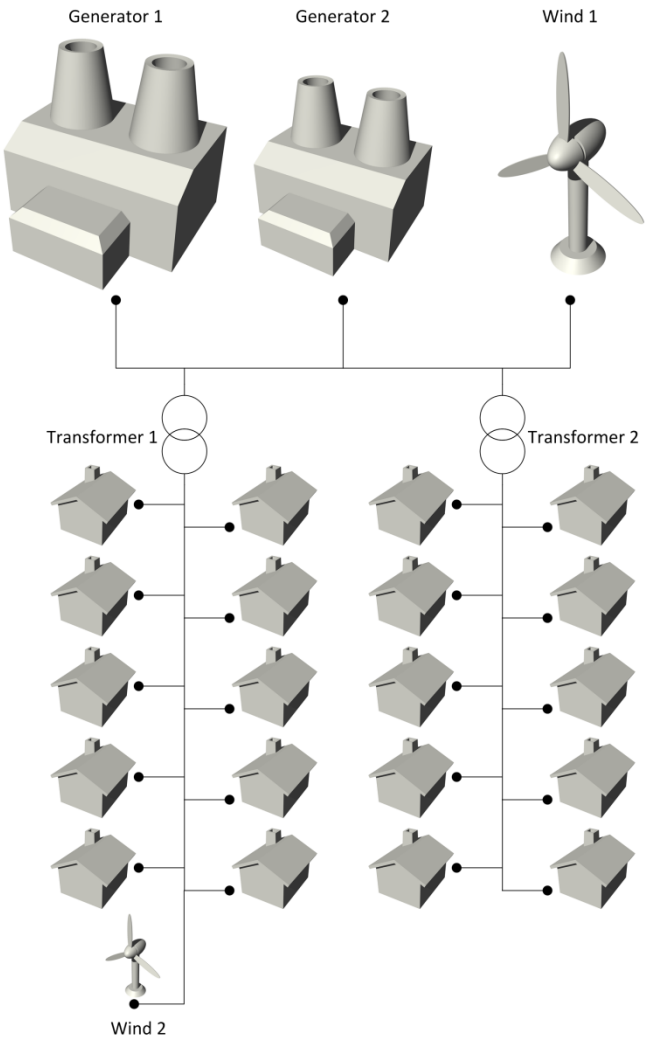

Figure 3.5 – Layout of simulated power system

There are two transformers, with 10 houses on each. There are two simulated fossil fuel generators – the first is a large generator with a slow response and a ramp-rate limited to 100W/s. It uses a simple integration based control scheme that constantly works to bring the frequency to a stable 50Hz. The second generator is a smaller, fast responding generator that can ramp its generation output almost instantly. There are also two wind generators included. In the first set of simulations, only the large wind turbine (Wind 1) will be used, and in the second simulation only the smaller turbine (Wind 2) is enabled. For simplicity, only details of the power system that were of interest were simulated: power flow and system frequency were modelled, while voltage was assumed to be constant and VAR and individual phase loading was ignored. These are of interest because DDC adjusts power flow in response to system frequency. Voltage stability is simulated in relation to electronic transformers in Section 7.5.

3.4.1 Integrating Renewable Generation

In the first set of simulations, fluctuating generation was added from Wind Turbine 1 up to a maximum of 30kW, potentially representing a large grid scale wind plant. The first simulation has only uncontrolled load, and the fast responding generation was increased by 30kW in order to achieve acceptable frequency stability. A one hour simulation of this setup is shown in Figure 3.6, with the wind generation, uncontrolled load, and fossil fuel generation in the first plot, and the system frequency in the second plot. There are a number of surges in wind generation spaced roughly 15 minutes apart, with the second gust having the most rapid change in output. The system frequency varies quite significantly, especially during the second gust ranging from 49.6 to 50.3Hz.

In the second simulation the fast responding generation (Generator 2) was disabled completely, and only the slow responding fossil fuel generation was left in, with the exact same wind profile as used previously. This time around the deferrable load was configured to respond to the system frequency, with an aggregate response similar to that of Figure 3.1 but with a 0.6Hz range. The result is shown in Figure 3.7 and it can be observed that the load now closely follows the changes in wind generation, and the frequency is far more stable, varying from 49.9 to 50.1Hz this time.

These results suggest that the frequency responsive load could provide similar frequency regulation to the fast responding generation currently used for this purpose. Interestingly enough, these results also align with those in the literature suggesting that an increased proportion of renewable generation in the mix requires an increase in the amount of frequency regulation. Traditionally this would imply more gas-fired plants but as these results suggest, frequency responsive load could also be suitable.

A third simulation was performed with quite different conditions. This time local generation was added (Wind Turbine 2), connected to a first transformer. The electric vehicles on the first transformer were removed, such that there is much less controllable load on the first transformer, and a high amount on the second transformer. The simulating results are shown in Figure 3.8, with the power demand on both transformers as well as the local generation output plotted. The generation ramps quickly up to 20kW, and it can be seen that the load on the first transformer goes negative for the majority of the generation pulse.

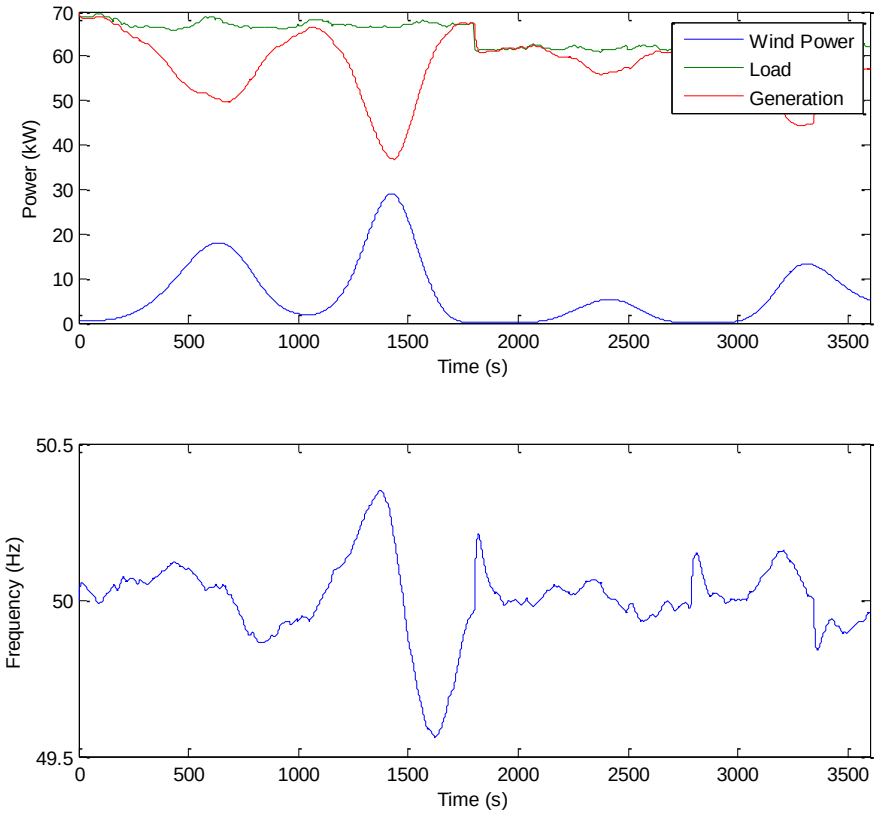

Figure 3.6 – Simulated response to fluctuating generation without DDC loads

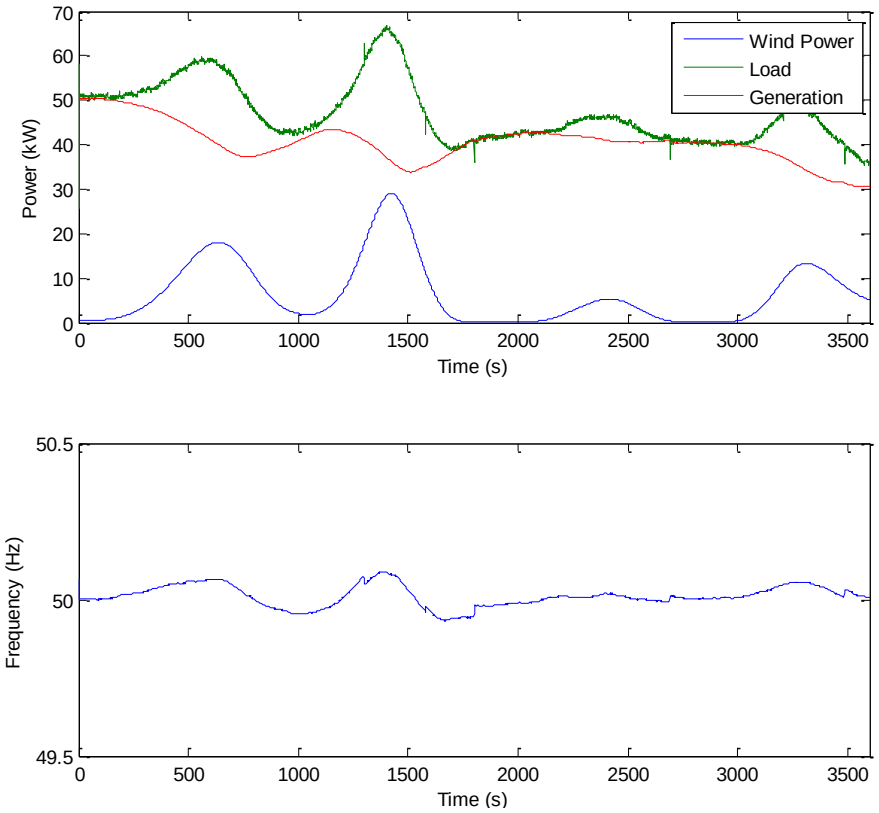

Figure 3.7 – Simulated response to fluctuating generation with DDC loads

This simulation highlights the way in which the system frequency is not representative of any distribution constraints in the power system, only the overall balance between supply and

demand. The first transformer is so lightly loaded that power is being exported back into the power system, while the second transformer has higher loading. Load on both of these transformers will respond to the same steering signal, even though there are quite different constraints on each. Of course the range of constraints not represented in the system frequency are not just limited to LV transformer ratings, but could also be MV line ratings, HV/MV transformers and even HV lines.

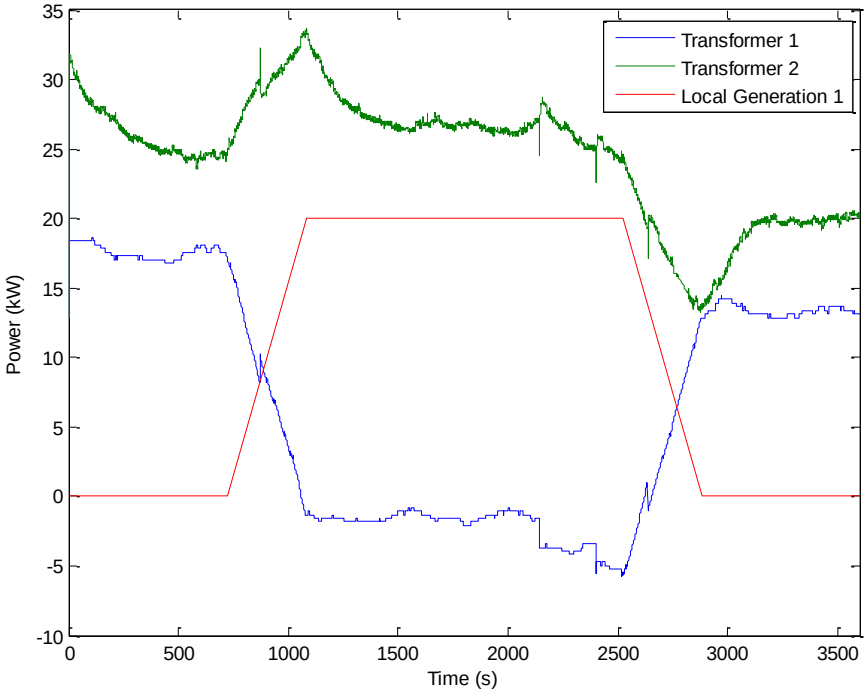

Figure 3.8 – Simulated response to a local surge in generation with DDC loads

3.4.2 Integrating Electric Vehicles

Another simulation was conducted in order to investigate the impact of DDC on the so called "EV switch-on" problem. This refers to the situation where commuters arrive home in the evening and all plug in their EVs at the same time, often when the grid is already heavily loaded. As introduced in Chapter 2, this could cause issues at the grid level, where there may not be enough generation to supply this extra demand at this time, as well as at the distribution level, with street level distribution transformers identified as being particularly vulnerable.

In this simulation, the entire EV load is added concurrently at t = 1200, but the simulation is configured such that there is not enough generation to cover the added load until extra generation is brought online at t = 2400. Generator 1 is limited to 45kW, and Generator 2 starts at 10kW, and provides an additional 20kW after t = 2400. Note that there is no wind generation added here.

The results are shown in Figure 3.9, with the second plot showing system frequency and traces on the first plot showing demand and the two generation sources. The EV switch on can be seen at t = 1200, with a brief spike in demand as the EVs are turned on and then back off as the frequency drops to a minimum of 49.9Hz. Within around 2 minutes the first generator has ramped output to its maximum and the frequency is left at a constant 49.95Hz. With the reduced system frequency the EVs stay at a reduced power consumption until the frequency is restored to 50Hz. With the sudden addition of the second generator, there is again another spike in frequency, and the demand also spikes up to keep the system stable. There is then a first order decay as the first generator reduces output in order to restore the frequency to a nominal 50Hz, which is achieved at t = 3000.

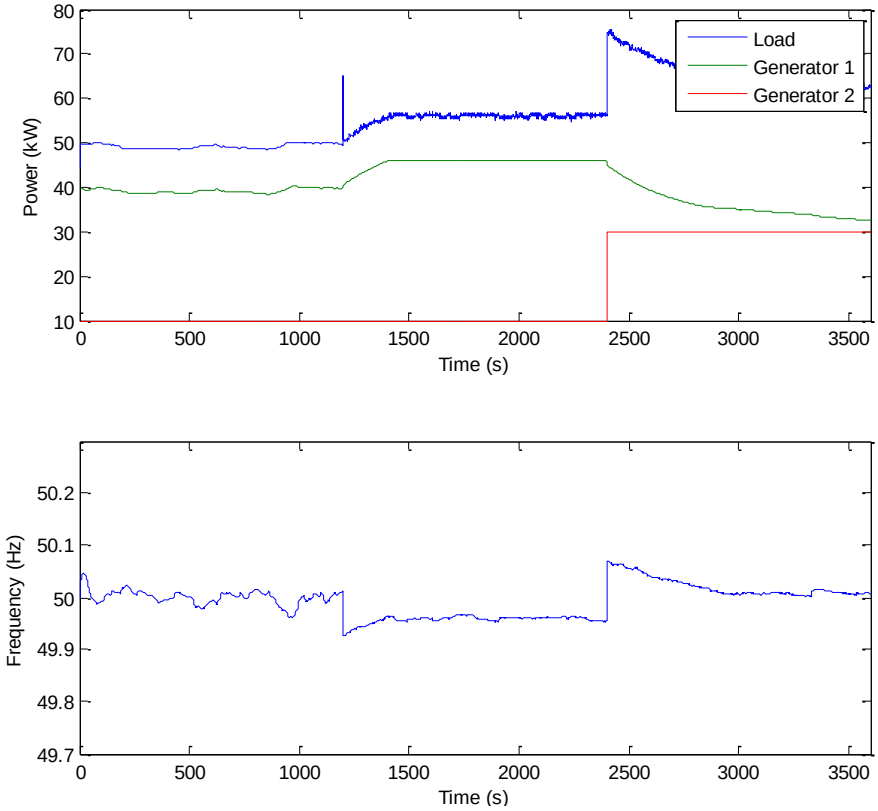

Figure 3.9 – Simulated EV switch-on scenario

Another interesting feature of these plots is that the frequency stability is significantly improved with the addition of EVs. Up until t = 1200, fast responding frequency regulation is provided only by hot-water cylinders and refrigerators, both of which are purely binary loads (either on or off). During this time the load is quite flat, while the frequency varies over roughly 0.1Hz. After t = 1200, the frequency has only tiny variations on top of the overall trend, and the load now shows what appears to be noise. This is because the EVs are simulated as having continuously variable power output over the range of 0-4kW. In practice, an EV charger may

not have quite such precise control over power output, but the aggregate response of a number of chargers will still be very linear. The charger used in the previous experimental work had controllable power consumption with roughly 8-bits (256 levels) of precision.

3.5 Performance of DDC

It has been shown in these simulations that DDC can be very effective when providing extra power system balancing and frequency regulation – which is particularly useful when integrating further grid scale wind and solar plants. What it does not provide is a response to any localised distribution constraints in the power system, such as individual line and transformer ratings.

There are also a number of practical considerations that make the actual deployment of DDC quite difficult. Because the control signal is the system frequency, and the areas of common frequency are generally very large, DDC needs to be deployed across the entire power system to have a meaningful effect. The fact that DDC cannot be deployed and tested in a small geographical area without significant system reconfiguration is a major deterrent. Deployment also necessitates that the system operator allows the frequency to vary over a slightly wider range than is allowable at the present time. This was seen in the fourth simulation where the frequency had a constant offset for the entire duration of the generation shortfall. These requirements make it unlikely DDC will be deployed within a large modern power system.

A combination of DDC and modern communication technology has potential, but this still has many of the same frequency variance and area of effect issues as normal DDC. Additionally, in this case the fast acting aspect of the system is responding to the system frequency, which generally does not vary that quickly, and other system metrics could benefit more from the fast response.

As DDC was found to be unsuitable for widespread DSM, a more appropriate system called LDC was designed. The rationale behind the design of LDC, as discussed next in Chapter 4, was to create a DSM scheme that had the same stability and operating characteristics as DDC, while also being a practical option for large power systems.

The original focus of this research was to create a wireless electric vehicle charger with a new higher operating frequency, as well as DDC functionality for improved grid integration. As the focus shifted to creating a new DSM scheme, less attention was placed on improvements

to wireless charging and general electronic loads were targeted. Some of the relevant work on improved wireless charging is still included in Chapter 7.

Chapter 4
Localised Demand Control

4.1 Introduction

In this chapter a new DSM technique called localised demand control (LDC) is introduced. The expected control structure, stability and transient performance are outlined and then the system is simulated under a range of conditions. The similarities between DDC and LDC are discussed, along with the expected performance of each in a practical grid environment. Finally, LDC is compared to other research found in the literature. Chapter 5 will continue discussing LDC with the introduction of the experimental system.

In Chapter 2 the difficulties of integrating electric vehicles, and further renewable generation, into the grid were outlined. The way in which demand side management could help to mitigate these issues was also discussed, and an overview of current DSM technology and research was also given. It was concluded that the current state of DSM could certainly be improved upon. In Chapter 3, DDC was introduced as a potential DSM solution, with the conclusion that the frequency variation required and the lack of a local focus make it unsuitable.

LDC is a new system designed to have much the same (or better) stability and transient performance than DDC while solving many of these previously mentioned drawbacks highlighted in Chapter 4. The goal of LDC is to create a DSM scheme that is suitable for deployment in a modern power system, and that adequately addresses the previously mentioned challenges that power systems are facing.

With LDC, a local controller continuously measures distribution constraints at a single distribution point in the power system. The controller then calculates whether that distribution point is within safe operating limits or not, and then sends a signal to all loads drawing power through that point of distribution accordingly. The most likely implementation points would be distribution transformers in large buildings or at the street level. The main distribution constraint measured is normally total power flow but other factors like system frequency, voltage, and transformer conditions can be considered as well.

This is the biggest difference between DDC and LDC: whereas DDC acts to maintain balance in the entire power grid, LDC is designed as a ‘local’ system and in this instance is

primarily concerned with the low voltage (LV) transformer and the local network health. With the version of LDC described here the power consumption is measured and regulated at the local distribution transformer. With this setup, a system is created that can be deployed in stages while still providing significant utility (lacking only grid supporting functionality as discussed later).

Ideally, the LDC system is designed in such a way that it can be retrofitted to existing transformers, extending their lifetime. The majority of LDC hardware required is placed at the transformer, including a power meter, an LDC controller and a control signal distribution system. The only additional hardware is a small low cost signal receiver/controller placed at each load or group of loads where control is desired. There are many benefits created by this system, and these will be discussed in detail in the following section.

The arrangement of a practical LDC implementation is shown in Figure 4.1. It can be observed that by adding in LDC, there is a forward information path to each of the loads, and the whole LV network is now a closed loop system allowing control of the power demand as seen by the wider grid.

For control, the power measurement made at the point of coupling to the grid is compared with a set operating point to generate an imbalance signal. Ideally this imbalance signal (kW) is zero but in practice the imbalance is integrated to produce an LDC control signal which has the dimensions of kWh. The LDC signal is then distributed to all loads drawing power through the point of coupling. As is the case in the DDC system, the LDC signal is then filtered and used directly for load control.

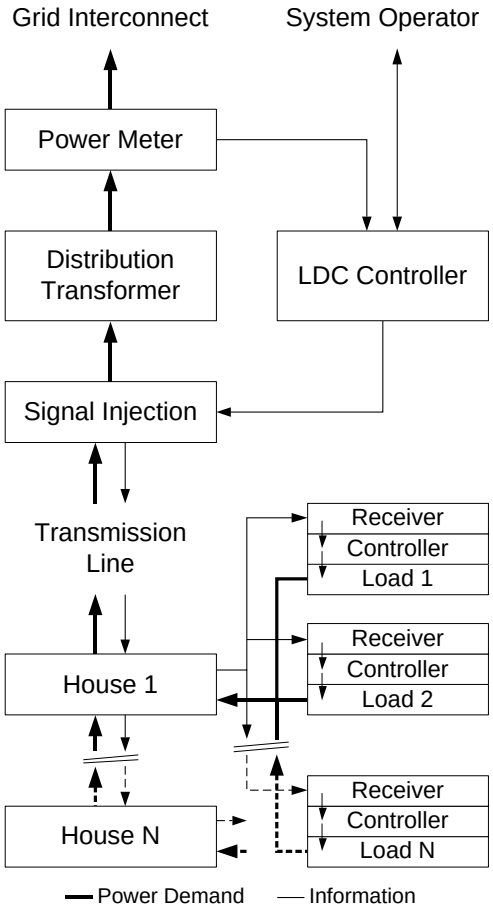

Figure 4.1 – LDC system structure

4.2 Controllable Loads

There is small amount of additional logic within each load controller that takes into account the internal state of the load, as well as a device priority that is either set at manufacture or by the user. For determining the load's switching behaviour, the most important constraint is keeping the load within the operating zone required by the end user, with switching within that zone then dependant on the LDC signal. For example a space heating or cooling system will first act to keep the temperature within the band set by the user, and then the load controller will be allowed to vary the temperature within that band. A small section of the lowest possible LDC signal could be reserved for signalling an emergency grid event, which could then override the device and user preferences.

It is a significant concern with any DSM system that the presence of a uniform steering signal will cause load synchronisation, which in turn could cause voltage dips when a large number of loads switch at the same time. By making the switching behaviour dependent on multiple metrics, it is unlikely that load synchronisation will occur, although this could be investigated in more detail. Furthermore, if a proportion of the controllable load is continuously

variable as is possible with only minimal modifications in EV chargers, heat pumps or hot water cylinders, then load synchronisation is even less of an issue as load changes will then be more incremental.

There are many practical considerations that need to be taken into account when performing load control. Topics including user comfort, scheduling and incentives are discussed in Chapter 7.

4.3 Benefits of LDC

There are numerous advantages to monitoring transformer conditions and regulating power through controllable loads. As previously mentioned, significant EV deployment could cause overloading of distribution transformers. If the transformer and EV chargers were compatible with LDC, the system could automatically delay the charge cycle to a later time when the load is reduced, and bring the load in progressively to avoid creating new peaks. This same load shifting can also help in instances of residential infilling, where new houses are built and the original distribution transformer is made to supply more homes than originally intended. Using these methods to protect the transformer from overload and burn-out is a very valuable feature, as these are expensive items with large capital and installation costs.

The challenges related to integrating renewable generation into the power grid have been covered previously in Chapter 2, and will only be summarised here. There are two main levels on which renewable sources are added to the power grid. Large scale installations are connected at the national grid level, where the generation output is generally well regulated and controlled so that the main concern is matching grid-wide demand with fluctuations in supply. Smaller generation sources such as those in residential areas are of more interest as they are connected at the distribution level and are likely to create a wide variety of issues. As discussed in Chapter 2, in the event of a grid fault or a power cut, a small solar inverter or wind turbine that lacks adequate protection could liven power lines, potentially endangering linesmen. Distributed generation in residential areas can also cause voltage stability issues, as there can be periods of significant mismatch between supply and demand. The unpredictable nature of these generation sources can also be problematic, as it can make the line livening and voltage issues appear randomly. In addition the varying power output can make load forecasting more difficult for network operators, as there is no guarantee that it coincides with consumer demand. LDC has the capability to solve or mitigate all of these issues. Firstly by simply monitoring and regulating the power flow through the transformer, local consumer demand can be more

closely matched with local supply, and the exporting of power into the wider grid can be prevented. The closer matching of supply and demand naturally reduces voltage stability issues, as the local network is adjusted to increase its load when there is local power production. Since the control signal is visible at all points in the local network, not only loads but also local generating units can be fitted with a signal receiver, allowing precise control and further mitigation of adverse effects. In this way the generation can be signalled to shut off for various reasons, including when there is a grid fault, when line voltages are too high or when there is too much power being exported into the wider grid.

If extra electronics were fitted to a distribution transformer, it would be sensible to add in a data uplink that connects the transformer to the appropriate system operator. This uplink would not be part of the real-time control loop governing transformer operation, but would simply be for monitoring transformer conditions and potentially updating operating parameters based on events happening in the wider grid. The amount of data transferred is low and latency is not of immediate concern so that the uplink could be realised with a simple GPRS modem. This is a relatively low cost way of providing simple two-way communication and will have signal coverage almost anywhere in the country. The data throughput is not necessarily high but is more than enough for basic LDC status and control info. Further to transformer monitoring, there are a number of other functions that would be enabled by a data uplink. Firstly, the power limit could be updated or set in advance for different times of the day. This could be used to shift load out of high demand periods to achieve load shaping, or it could be used in response to fluctuations in large scale renewable generation output in the wider grid by bringing on transient loads to absorb the power.

Much has been discussed in Chapter 3 on the ability of DDC to improve frequency stability and integrate large scale renewable generation through frequency responsive loads. If the LDC system also monitors the system frequency as one of its distribution constraints, then this same functionality is provided. Monitoring and responding to changes in system frequency or voltage in real-time is a valuable ancillary service, but this would not necessarily be the usual mode of operation. An experimental LDC system is discussed next in Chapter 5, which includes a section where the additional frequency responsive functionality has been added.

As also mentioned in relation to DDC, widespread deployment of LDC systems is required to provide meaningful levels of ancillary services. As shown, if the hot water cylinders in only 13% of NZ households were used for ancillary services, the entire contracted frequency

keeping reserve could be provided. Fortunately, LDC still provides many other benefits even without widespread deployment

4.4 LDC Control Structure

The control method (as described earlier) is as follows: The difference between the ideal power consumption (set-point) and current consumption is measured. This difference is then integrated to produce a control signal (with dimensions of kWh), and the control signal is then filtered and used directly for load control.

This setup of a simple LDC system with power feedback and controllable load is shown in schematic block diagram form in Figure 4.2, with symbol definitions as given in Table 4.1.

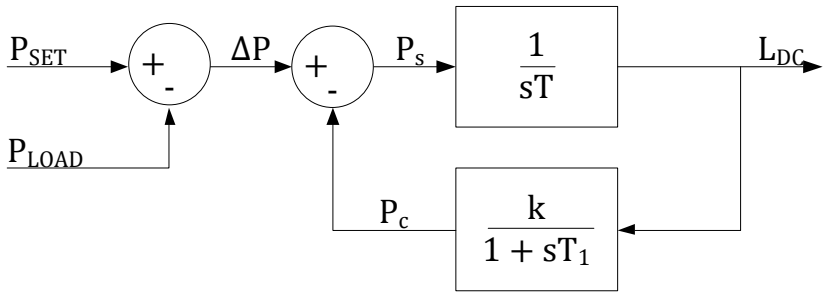

Figure 4.2 – Expanded LDC block diagram

Table 4.1 – Symbol Definitions for LDC Block Diagram

Symbol	Definition
ΔP	Power disturbance
P_{SET}	Power set-point
P_{LOAD}	Uncontrolled load
P_s	Power imbalance
P_c	Power consumed by controllable loads
T	Integration time constant that converts kW to kWh
L_{DC}	Control signal distributed to loads
k	Conversion constant between energy imbalance and power
T_1	Filter time constant

There are a number of power flows in a local grid, but for simplicity the set-point and local uncontrolled load can be combined into one power signal, represented as the power disturbance, ΔP as shown in (3.1). Under normal conditions, both these inputs are constant, and a change in either of them will produce a transient at the input that the system will respond to.

$$\Delta P = P_{SET} - P_{LOAD} \tag{4.1}$$

By combining the set-point and load into ΔP, the block diagram can be simplified as shown in Figure 4.3. The similarities between the DDC controller previously described and this LDC controller are now clear, but the controlling actions are actually quite different.

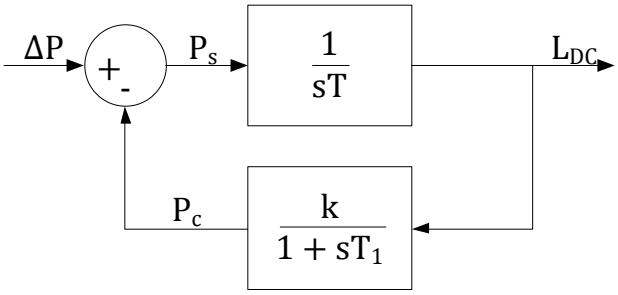

Figure 4.3 – Simplified LDC block diagram

Ideally the total power consumption in the local grid ($P_{LOAD} + P_c$) should be equal to the set-point P_{SET}, therefore from (3.1) the power disturbance ΔP should equal the power consumption of controllable load, P_c. This means the difference between these two, P_s is the item being regulated to zero and is the system output, resulting in equation (4.2) below. As the input of the system is ΔP, the transfer function of interest is $\frac{P_S}{\Delta P}$.

The feedback term $\frac{k}{1+sT_1}$ represents the response of the controllable load, where k is the amount of controllable load that is directly proportional to L_{DC} and the $\frac{1}{1+sT_1}$ term is a low-pass filter with a time constant of T_1. In a similar manner to DDC, a low-pass filter is required to dominate the effect of delays in generating, transmitting, receiving and responding to the control signal.

To calculate P_s the input power is subtracted from the controllable load as shown in (4.3),

$$P_s = \Delta P - P_c \tag{4.2}$$

$$P_s = \Delta P - \frac{k}{1+sT_1} L_{DC} \tag{4.3}$$

$$L_{DC} = \frac{1}{sT} P_s \tag{4.4}$$

By substituting L_{DC} from (4.4) into (4.3) and rearranging, the desired transfer function (4) can be obtained.

$$\frac{P_s}{\Delta P} = 1 - \frac{k}{s^2 T T_1 + sT + k} \tag{4.5}$$

Where the damping ratio is:

$$\zeta = \frac{1}{2}\sqrt{\frac{T}{kT_1}} \tag{4.6}$$

The final transfer function and damping ratio can be compared to the transfer function derived for DDC. The performance of the system is again dependent on the integral time constant, filtering constant and available controllable load. An important difference between this response and that of DDC is that the system inertia is not involved. In fact the inertia of the grid makes the whole network stable without having to add any extra inertia. As far as the controller is concerned the damping factor and hence the stability of the system is dependent on the controller gain and the integrator time constant and these are easily adjusted.

In the transfer function given by (4.5), there are two terms, first the leading "1", and then a standard second order response term. The "1" implies that changes in ΔP will be directly seen at P_s, and the second order response term then governs the control action that restores P_s to zero. This is the expected result given the structure of the LDC system. Any changes in power consumption or generation (ΔP) in the system are seen directly at the transformer where the imbalance (P_s) is measured.

In practice, this implies that transient spikes in the local power input are expected with changes in local load or generation, and that these will decay to zero with a response governed by the gain and time constants designed into the system. The predicted response of a simulated LDC controller is shown later in Figure 4.4, and an actual response measured in an experimental system can be seen in Chapter 5.

4.5 LDC Performance Requirements

To regulate local power demand, the LDC system simply needs to respond to local transients as quickly as is possible without any overshoot or oscillation. If ancillary services are to be provided to the wider grid, then the desired speed of response depends on the service being provided. If frequency responsive emergency reserves are offered, then in New Zealand these operate on time scales of 0.4s, 6s and 60s [52]. A response time of 0.4s is not achievable with the current LDC system. But since LDC is non-intrusive, load could be shed linearly as the system approaches a less extreme frequency such as 49.2Hz. In this way, a practicable 1s LDC response time would still have an overall faster response than emergency load shed at a lower frequency such as 48Hz. Generator governor action which is used for instantaneous frequency reserves also operates on a similar sub 1s time scale.

The only physical constraint governing the LDC system's speed of response is delays in generating, distributing, filtering and responding to the LDC signal. The filter time constant T_1 is designed to dominate these delays such that the operation of the system is predictable and stable, just as the DDC filter time constant was designed to dominate armature reactance effects to make that system stable.

The integral time constant T can be calculated using (4.6) to give an ideal response for a given amount of controllable load k as follows:

$$T = 4\zeta^2 k T_1 \tag{4.7}$$

For example, if there is 1kW of controllable load ($k = 1000W/kWh$), the required filtering time constant is $T_1 = 0.1s$ and a damping ratio of $\zeta = 0.8$ is desired, then an integral time constant of $T = 256$ will result.

From the above equation, another observation can be made. The filter time constant (T_1) and amount of controllable load (k) always appear together. This implies that if a larger filter time constant is to be used, the same system performance can be achieved if more controllable load is added. As expected, more controllable load makes the system performance better.

Figure 4.4 shows the per-unit step response of the calculated LDC system, with an equivalent DDC response added for reference. As can be observed both of these give an overshoot of 1% and a settling time to within 5%, of ~0.55s – which is well within the desired 1s response time. The differences in structure between DDC and LDC can also be seen. Whereas with DDC, additional torque results in a new stable operating frequency, with LDC additional power demand results in a brief step in observed demand then a steady decline as the load is reduced and the original conditions are restored.

As discussed later in this chapter and in Chapter 6, there are a number of options available to distribute the LDC signal. These include a radio signal, Wi-Fi, Zigbee, Ethernet, and power line communication.

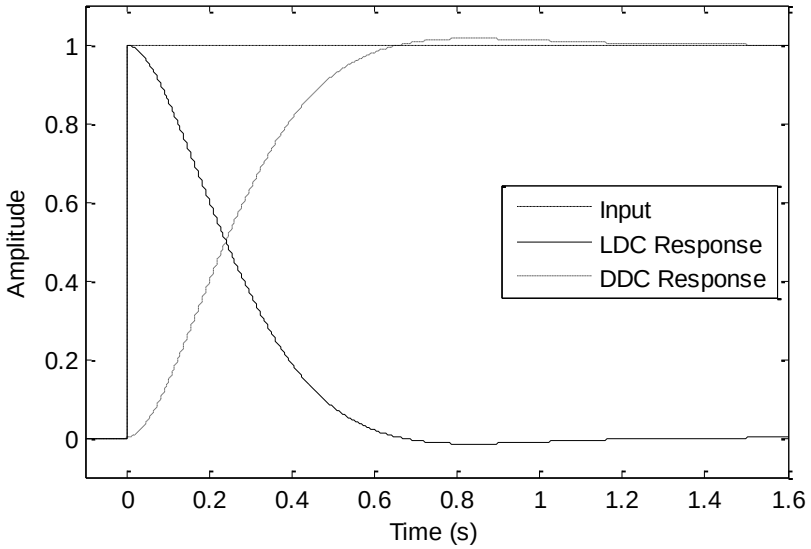

Figure 4.4 – Response of LDC and DDC systems to a per unit step response

4.6 LDC and DDC Comparison

It can be observed from the DDC and LDC system diagrams that they have very similar control structures. Since DDC is regulating $\Delta\omega$ and LDC is regulating P_s, the outputs are different but the underlying transfer functions (pole positions) are the same such that the dynamics are essentially identical. The advantage of LDC is that in addition to the filtering constants k and T_1 being configurable, the integral constant T can also be changed without affecting any mechanical machines. This means the speed of response and the stability are only dependent on the communication medium and electronics and not on any physical properties of the system as they are in DDC. The disadvantage however of LDC is that the control signal L_{DC} has to be manually distributed as opposed to DDC where the system frequency can be easily observed at all nodes. For good performance the L_{DC} signal must be distributed accurately, without corruption, and with low latency to every consumer on the local network. This is achieved here using a robust custom one way power line communication (PLC) system. A PLC is suitable as the signal only needs to be transmitted from the street level transformer to each house over the LV network – a maximum distance of a few hundred meters. Laboratory tests have shown no noticeable drop off in signal strength with 100m of cable using the communication apparatus described in Chapter 5 Chapter 6. It is also possible that in a future where electronic transformers are common the frequency and voltage can be used as the communication medium for the control signal, as these can be varied on each transformer independently of the wider power system. Preliminary simulations of this setup have been conducted in Chapter 7.

The biggest functional difference between LDC and DDC is that LDC attempts to regulate power on an individual distribution transformer or transmission line. The performance of LDC when performing this task is simulated in the next section. Localised power regulation has a number of advantages. Firstly, in the event of widespread EV deployment, street level transformers equipped with LDC would never exceed their maximum ratings. Thus LDC solves the so called "switch-on" problem of EVs which is a significant concern with current infrastructure. LDC is also resilient to controller failure. In the worst case where the controller fails with all controllable loads signalled on, this will simply be business as usual for the transformer, and no system damage should occur. Local power regulation also means the controller can make sure all local DG will be consumed locally and not fed back into the grid unless requested. DG units could also be configured to switch off when there is no LDC signal requesting power, further ensuring that energy cannot be exported and islanding cannot occur. As the authors of [17] mention, a scenario of widespread photovoltaic generation in a residential area could cause afternoon voltage peaks of up to 1.07pu at the distribution level. If the system also controlled based on the local voltage, these peaks could be avoided. These features make integrating LDC responsive DG more cost effective as expensive anti-islanding hardware and grid strengthening should not be required. In addition, many power systems do not provide feed in tariffs for DG, and many that do are constantly decreasing the tariff. In these systems energy exported into the wider grid is paid out at a fraction of the standard electricity rate. By exporting the bare minimum into the wider grid, the maximum amount of local consumption can be displaced, making the DG as cost effective as possible.

As LDC is a localised system, it could provide these benefits even when deployed in an isolated geographic location. Unlike DDC, LDC could be deployed and tested on a small scale, and then slowly rolled out with immediate results at each location. The system provides a simple, robust, low cost and easy to install method of DSM.

4.7 Simulated LDC System

A simulation of the power demand for a small community comprising twenty houses each containing some LDC compliant loads has been performed where the mains power is regulated to a nominal maximum amount of 20 kW (1 kW/house). In the first simulation local wind power is added as a random sequence, changing every ten seconds, with an average value of 70 kW. The load taken by each household averages 3.5kW, but can peak at up to 7kW. This system includes 77% wind penetration which is very high. A diagram of the system layout is

given in Figure 4.5, which shows the connection of the transformer, LDC controller, houses and local generation.

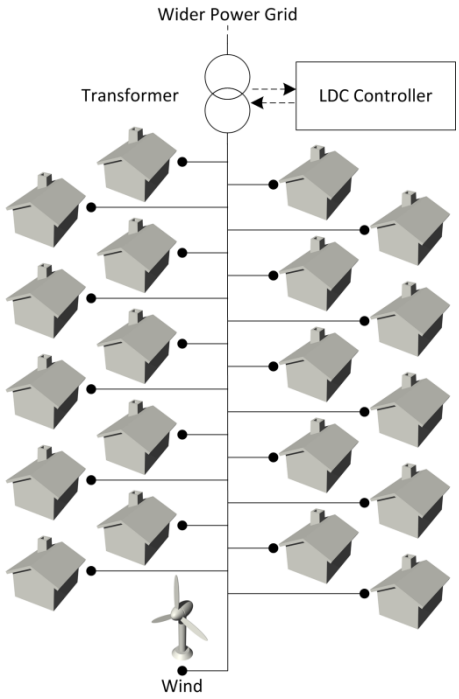

Figure 4.5 – Layout of simulated LDC system

Central to the system is the LDC controller which measures the power flow from the grid and compares this with a known set point. Internally, a simple integral controller outputs a power priority signal that varies from 0 to 10 in real time, attempting to keep consumption identical to the set point. The most important device is priority 1, whilst the least is priority 10. Consequently, devices with priorities below the signal will stay on whilst those above will be switched off. In this simulation each house consists of a number of LDC compliant (controlled) loads. These are listed in Table 4.2.

Table 4.2 – Simulated loads in each household

Load	Average Power (W)	Peak Power (W)	Priority
EV Charger	2000	4000	4-10
HWC	500	2000	4-9
Refrigerator	60	250	1-6
Base Loads (x4)	250	250	1,2,3,4

All loads except the base loads will vary linearly over their given priority range, consuming minimum power at a lower priority signal. The four 250W base loads are simply switched off if the signal goes below their given priority. A small random offset is given to each of the

house's priority signal so that not all house's base loads of equal priority switch at exactly the same time.

While having an EV charger in each and every house may appear unrealistic, the EV charger could equally be replaced with a heat pump or air-conditioning system with similar results.

As discussed earlier, deferrable loads are those that contain some kind of energy storage and which have some flexibility as when that energy needs to be stored. This means each load can be considered to have a state of charge (SoC) which refers to the current energy stored in the load as a proportion of the maximum energy that can be stored. While this is a common measure for items such as battery chargers, it can equally be applied to loads with thermal storage. The main difference is that with an EV charger, power is not consumed while the EV is charging. Contrast this with a HWC or space heater, where heat is constantly being removed from the system. For simplicity, the HWC and refrigerator have a constant amount of energy being extracted out, represented by the average power in Table 4.2. To increase the load diversity, the initial SoC of each load has been distributed linearly over the possible SoC range.

In these simulations, only details which are of interest have been included. So far only the power flows within the micro-grid have been simulated, and details such as voltage profiles, phase imbalances, and VAR loading have been ignored. The voltage profile and phase imbalance are investigated further in Chapter 7.

In addition, this simulation also assumes a perfect communication medium is in place for distributing the load control signal. The LDC controller generates the LDC signal, and this is available to all loads instantly, and with resolution only limited by the simulation software. The impact of using an imperfect communication medium is investigated in the Section 4.7.2.

4.7.1 LDC performance in the presence of wind

An example of the simulation output is presented in Figure 4.6. As shown, the wind varies significantly but the load on the system is kept in step with this varying wind. The power draw from the grid is regulated to 20kW. The probability density functions for the power taken from the grid and the power generated from the wind are shown in Figure 4.7. Note that the grid power is almost constant at 20 kW with deviations caused by loads switching on and off. The wind power is a roughly Gaussian distribution with a wide standard deviation – the ideal result would perhaps be a Weibull distribution p(x) where x is the wind speed, modified to x^3 to represent the power output.

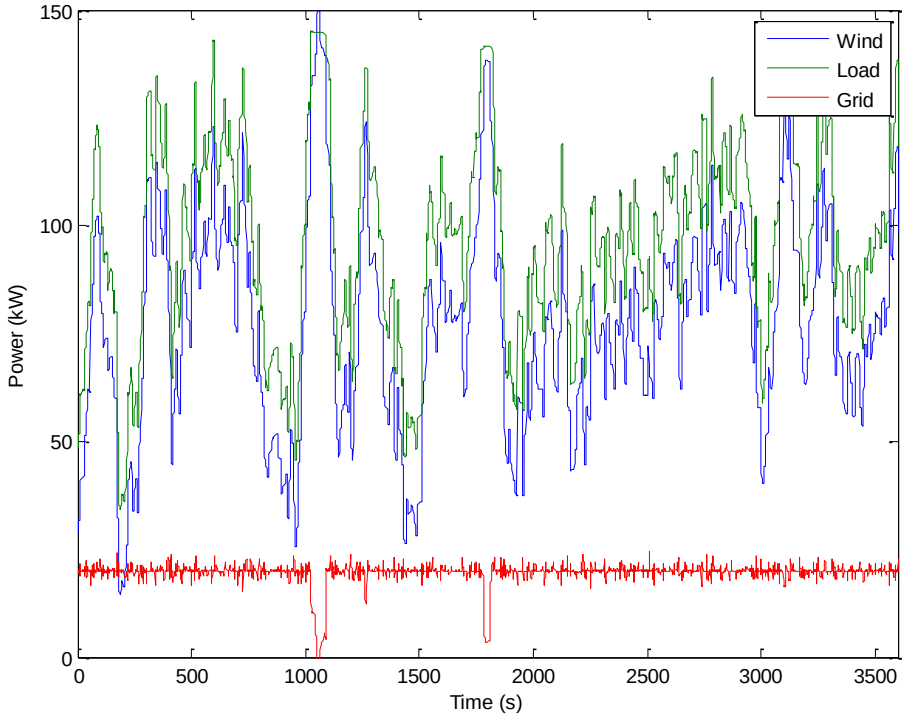

Figure 4.6 – Example plot of wind supply, grid supply and total consumption over a one hour period with a LDC controller.

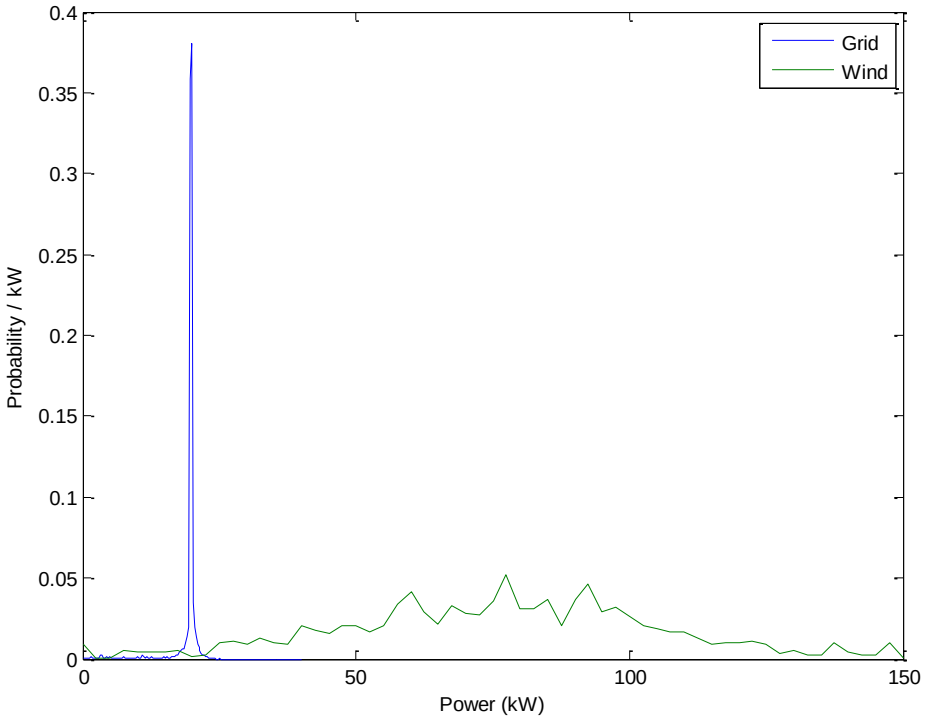

Figure 4.7 – Probability distribution of power consumed from the grid and provided by the wind turbine

Figure 4.8 shows the power usage over 1 hour for a single house. Here the power taken is quite volatile but when combined with all the other houses the % variation is improved

considerably. It can be seen that the fridge and hot water cylinder modulate their switching times to coarsely adjust demand, while the EV charger fills in the gaps. In this way a large load with a continuously variable control is seen to be an important adjunct to the controller strategy.

The response of the system to a step in wind speed is shown in Figure 4.9. The system adds 20kW of demand in approximately one second in a predictable first order response with no overshoot. It can be observed that this response is made of both large steps in load and more continuous load increase as a function of time. This is because the HWCs and EVs provide the majority of controllable load, and the HWCs have hysteretic temperature control whereas the EVs have continuously variable power output.

This again illustrates the value of having loads that are not just binary on/off loads but adjust power continuously based on the LDC signal. The high regulating performance of this simulation is mostly due to having an EV charger which is a relatively large continuously variable load. If LDC functionality was added to appliances through the use of a dongle in between the appliance and the wall socket, this would only provide binary control through hard switching of the appliance. Since some appliances could lose data or state information or even wear out if being switched on and off all the time, a compromise could be having a simple 5V signal that is fed into the appliance from the dongle and which gives an idea of how much power the appliance should consume. This could be based on a duty cycle signal or analogue voltage level. Again, the problem with this is that it creates another standard to agree upon, although most micro-controllers found in appliances should be capable of accepting a 5V duty cycle signal without significant extra circuitry. More details on the practical implementation of smart loads will be discussed in Chapter 7.

A simulation was also performed which was identical to that presented previously in Section 3.4.1, where in a DDC system a localised wind pulse on one of two transformers caused the exporting of power to the wider grid and the same network diagram shown in Figure 3.5 is applicable. Here LDC has been added to both transformers and the results are much improved. The plot is shown in Figure 4.10, and it can be observed that much more of the available capacity is utilised locally when compared to the results in Figure 3.8. Both LDC controllers are attempting to keep power draw from the grid at 15kW, and when the wind pulse arrives on Transformer 1, the load drops to around 8kW on average with all controllable loads drawing the maximum possible amount of power. This is much better than negative power flow as occurred in the DDC simulation.

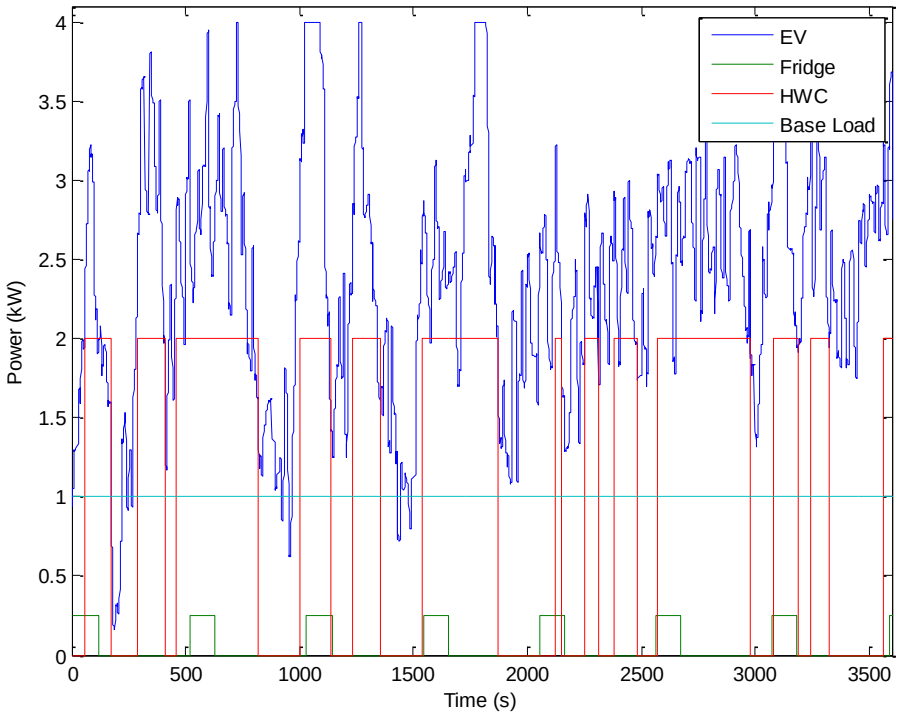

Figure 4.8 – Single household demand over one hour period

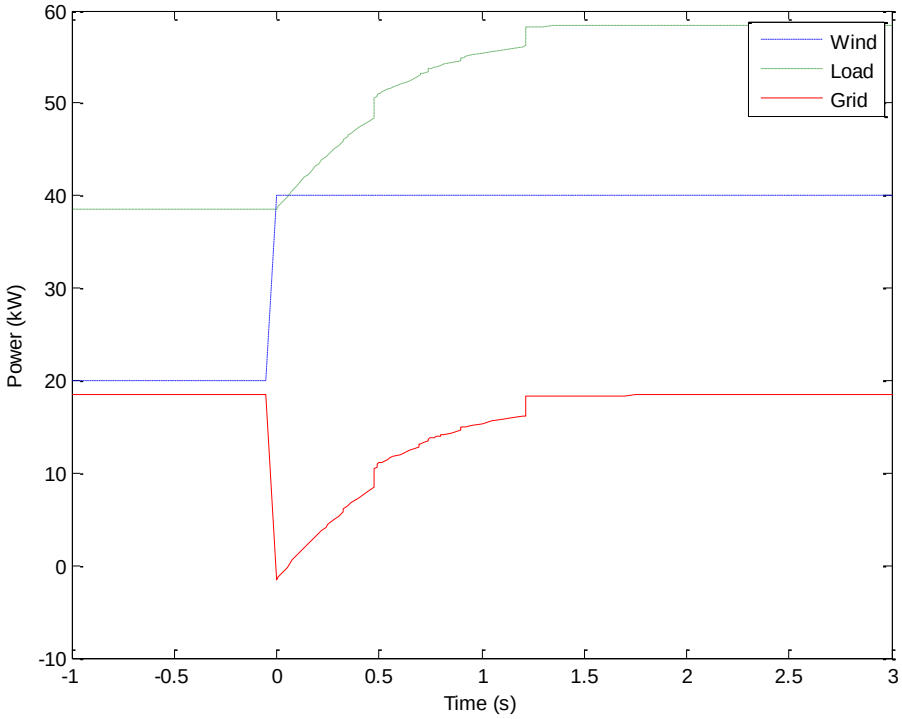

Figure 4.9 – Response of system to 20kW step in wind power.

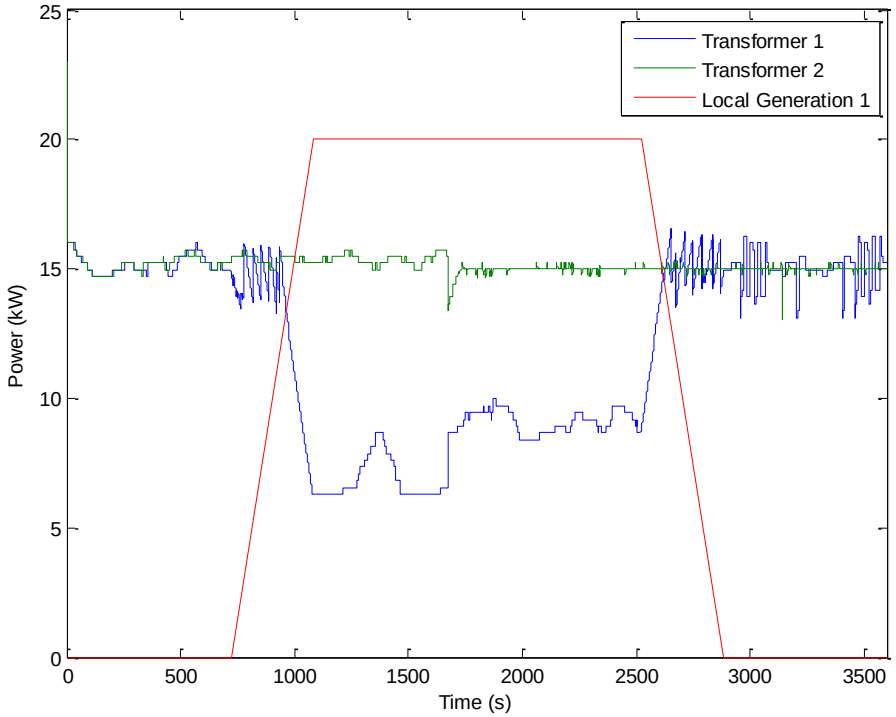

Figure 4.10 – Response of LDC enabled transformer to a local wind pulse

4.7.2 Effect of latency and sampling rate on LDC performance

As there are many different communication options to choose from, the exact level of performance required needs to be quantified. Depending on what transmission method is used, transport delays and sampling of the LDC signal may be unavoidable. These delays could be introduced by analogue filtering or the various delays and buffers in a digital communication scheme.

In order to quantify the effects of latency and sampling of the LDC signal, a modified version of the simulation mentioned above was run. The simulation was run for 15 minutes with a square-wave modulated wind turbine output. The turbine output changes between 20kW and 40kW four times during the simulation, and a single step response is shown in Figure 4.11. Two separate tests were done. In the first, an analogue delay was introduced that varied between 0s and 1s.

In the second test a sample and hold was added to the priority signal, with the sampling period also varying between 0s and 1s. In each, the system performance was measured with each change in delay or sampling rate.

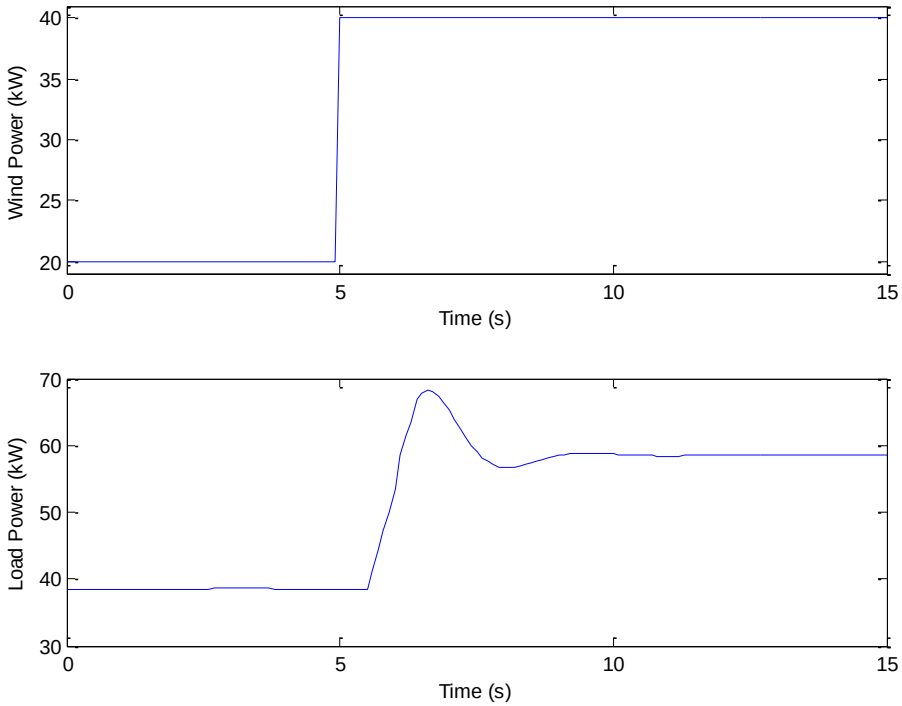

Figure 4.11 – Example output of a single step showing overshoot with 0.5s latency

As the LDC controller is trying to regulate the grid interconnect power to a specific level, the variation is a good measure of how well the LDC system is performing. In order to get the best performance from the LDC controller, the integrator gain was modified with each change in delay or sample time in order to avoid overshoots or oscillations caused by the transitions. The results of both tests are shown in Figure 4.12.

It can be seen that the ability of the LDC system to regulate power consumption is almost linearly related to any transport delay or sampling in the system. For this reason, the LDC signal needs to be transmitted as quickly as possible in order to get maximum performance. A delay of between 0.1 and 0.2 seconds is a realistic goal and this still yields good performance. It can also be seen that the sampling rate has less effect on system performance than transport delay. A sampling period of <100ms (>10Hz) is sufficient for a good performing system. These figures put limitations on the size of a network that might be served by a LDC controller in real-time. DDC is appropriate for levels up to grid scale but LDC may be best restricted to small islands within that DDC grid.

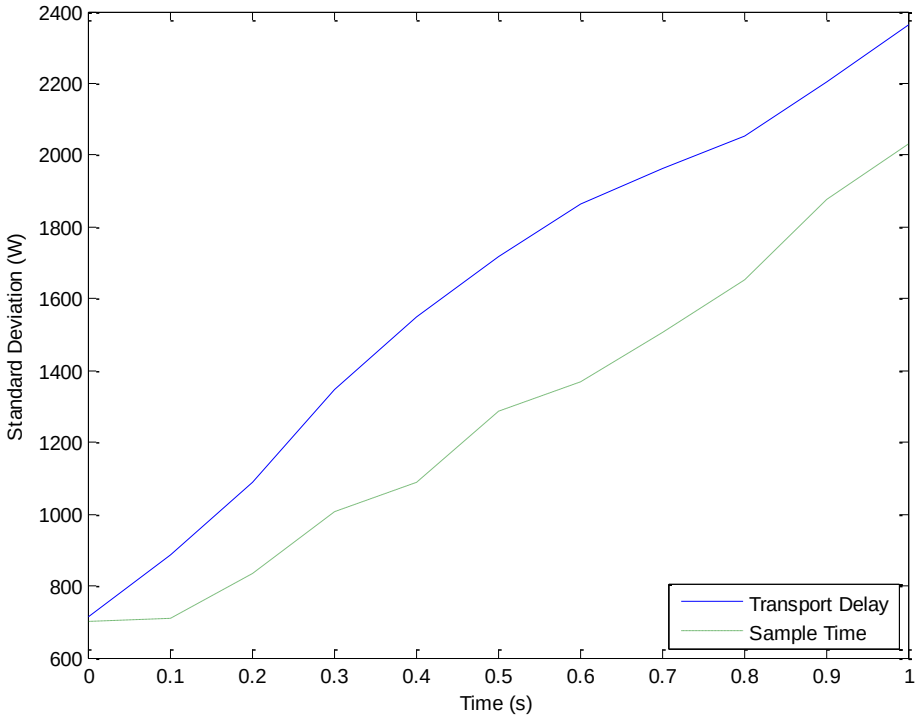

Figure 4.12 – Effect of transport delay and sampling on DDC system performance

Overall these simulations showed that using household loads to react to and smooth out a volatile power source is entirely possible. The system was able to respond to very large transients in around a second. The importance of having continuously variable LDC loads present in the system was also shown. The effect of using imperfect communication methods was also investigated, showing that good system performance could be achieved with realistic transmission methods. The required communication performance also lines up with that found in the literature [92] as discussed in Section 2.7.2.

A simple one-way PLC communication system has been designed and constructed for this work, and a general overview is given in Chapter 5, with full details and design covered in Chapter 6. The system was designed with the performance requirements established here for a fast and predictable response.

4.8 Load Scheduling Simulation

Some of the most common concerns when it comes to DSM are those of user impact and cost effectiveness. In order to show that there can be significant cost savings without any negative end user impact, some simple simulations have been performed.

These simulations show that a micro-grid equipped with the LDC system described earlier, operating in the presence of a real-time pricing scheme, creates a monetary benefit for shifting

loads out of peak times. The pricing profile was created by taking the shape of a typical load profile, and inverting it such that the times of highest load are the most expensive, and the times of lightest load are the cheapest.

4.8.1 Simulation Setup

In this simulation the micro-grid is made up of ten households each setup identically to those simulated earlier with the same loads as listed in Table 4.2. The difference with this simulation is that the micro-grid has a central community power manager added to the LDC controller that knows current grid consumption and the SoC of the large controllable loads within the micro-grid. The new network diagram with the added power manager, pricing signal and SoC information is shown in Figure 4.13. The LDC controller still distributes a single real-time priority signal to the entire community that represents the available power, but the power limit is now determined by the power manager. In these simulations the real-time price signal is received directly by the community power manager. This signal is updated once an hour and is not known precisely in advance. The only price information that is known in advance is the expected mean and standard deviation of the price over the 12 hour night cycle.

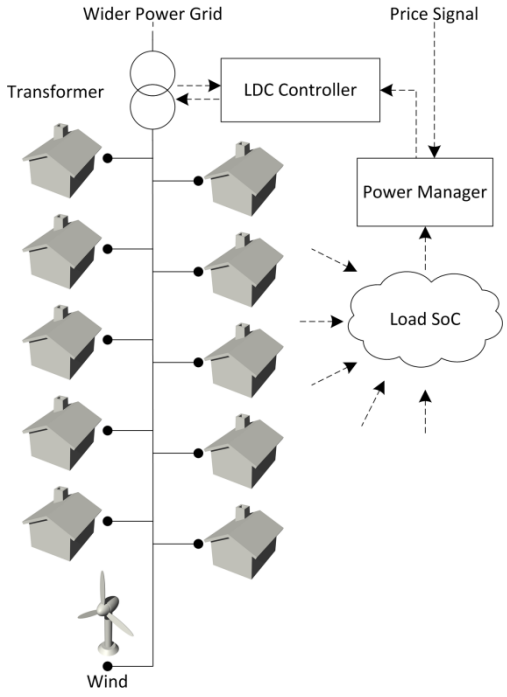

Figure 4.13 – Layout of simulated LDC system with added power manager

A novel addition here is that with just the mean and standard deviation, the power manager can interpret the price signal and determine how much electricity to consume at various price levels, in order to minimise the total cost of electricity. In this way the set-point can now be generated in an intelligent manner. The power manager continually updates the price threshold

and set-point throughout the night in response to any changing conditions. These could be unexpected extra generation from local sources or unexpected extra load. In these simulations the price thresholds and set-point are recalculated every half hour, with the SoC of devices in the community being polled at this same rate.

In terms of communication this means that an additional signal has been added to the system. As before, there is a one-way, high bandwidth and low latency signal for distributing the control signal around the community as discussed previously. Added is a path for the SoC and status of various loads within the system, which is only updated once every half hour in these simulations. In Figure 4.14 the communication links for the grid and a single house are shown. The solid line is power flow, and the dashed lines are various communication signals. The LDC signal goes from the controller to each controllable load, but now the SoC signal goes from each controllable load to the power manager.

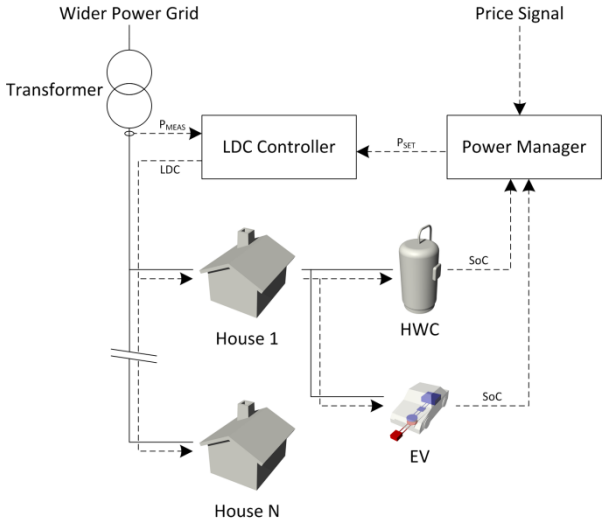

Figure 4.14 – Information flow in micro-grid for a single house

4.8.2 Simulated performance of LDC with price scheduling

The simulations undertaken here show the price and total power consumption (including local generation) throughout one 6pm to 6am cycle. For each set of conditions a simulation was run both with and without the community power manager.

The first simulation is a base case, with no power manager and no renewable generation. All loads simply consume power whenever it is required. With the pricing scheme as given the average cost of power delivered was 18.55c/kWh. The power consumption and pricing scheme is shown in Figure 4.15. It can be seen that power consumption is biased towards the peak pricing times and this will result in a higher average cost.

The second simulation is for the same conditions with the power manager enabled. The consumption with respect to time is shown in Figure 4.16. It can be observed that now much of the load has been shifted to the lower cost periods of the night, up to the maximum allowed 60kW. The 60kW constraint is a specification of the transformer and the power manager restricts the load in order to never exceed this level. This simulation resulted in an average power price of 16.892c/kWh.

In order to investigate the impact of adding renewable energy to the mix, a wind pulse was added which peaked at 11pm. As shown in Figure 4.17, without the power manager there is no impact on consumption, but as there is still large power draw at this time, most of the energy is utilised. The average cost of electricity is 17.351c/kWh for this simulation.

Figure 4.18 shows the same situation with the power manager enabled. There is now a large peak in demand that coincides with the wind gust. The average cost of electricity for this simulation is now 15.683c/kWh. Another interesting feature is that the total power consumption in this plot actually exceeds 60kW, even though this is set as the transformer limit due to the contribution from local generation.

A comparison of the four simulations is shown in Table 4.3. The simulations with the power management unit (PMU) enabled deliver less power overall. This is likely because the average and final state of charge values are kept lower. A lower final state of charge may result in energy use being shifted to higher priced periods, but this will only represent a small percentage reduction in efficacy. A lower average state of charge could actually improve efficiency slightly as standing losses will decrease.

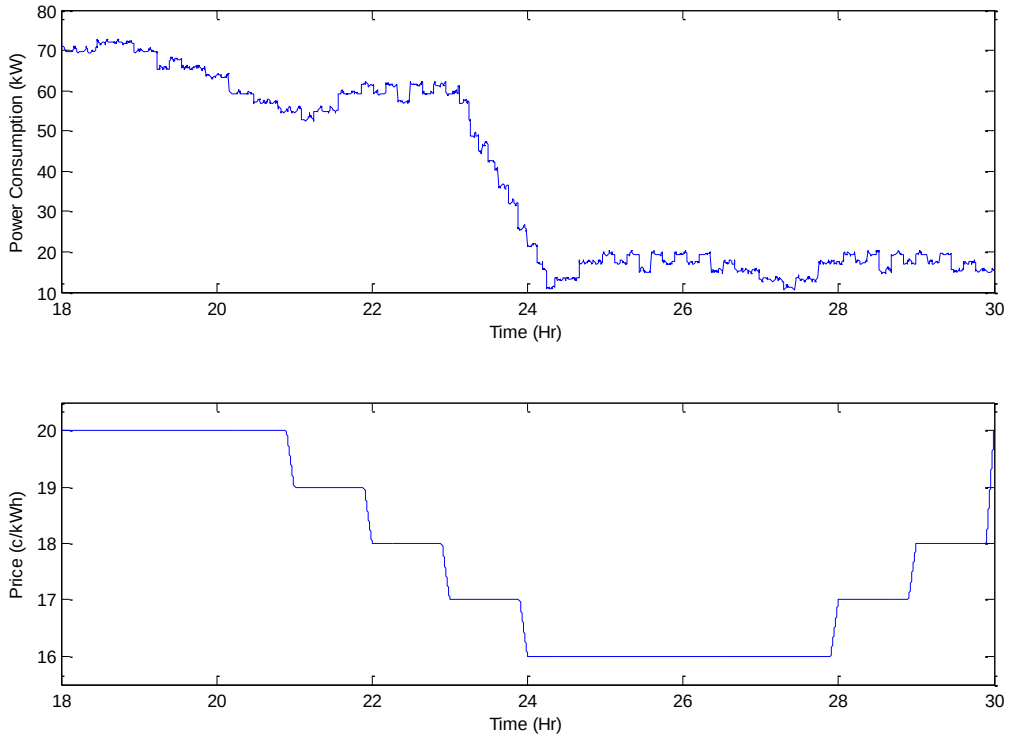

Figure 4.15 – Power consumption and price with no local generation and scheduling disabled

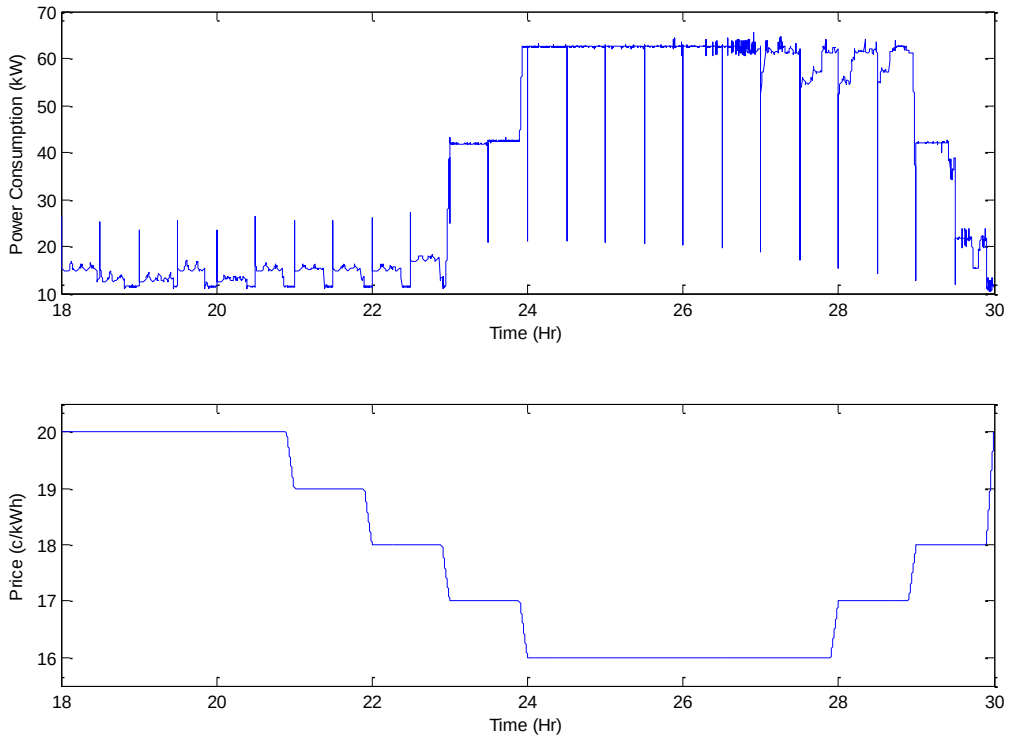

Figure 4.16 – Power consumption and price with no local generation and scheduling enabled

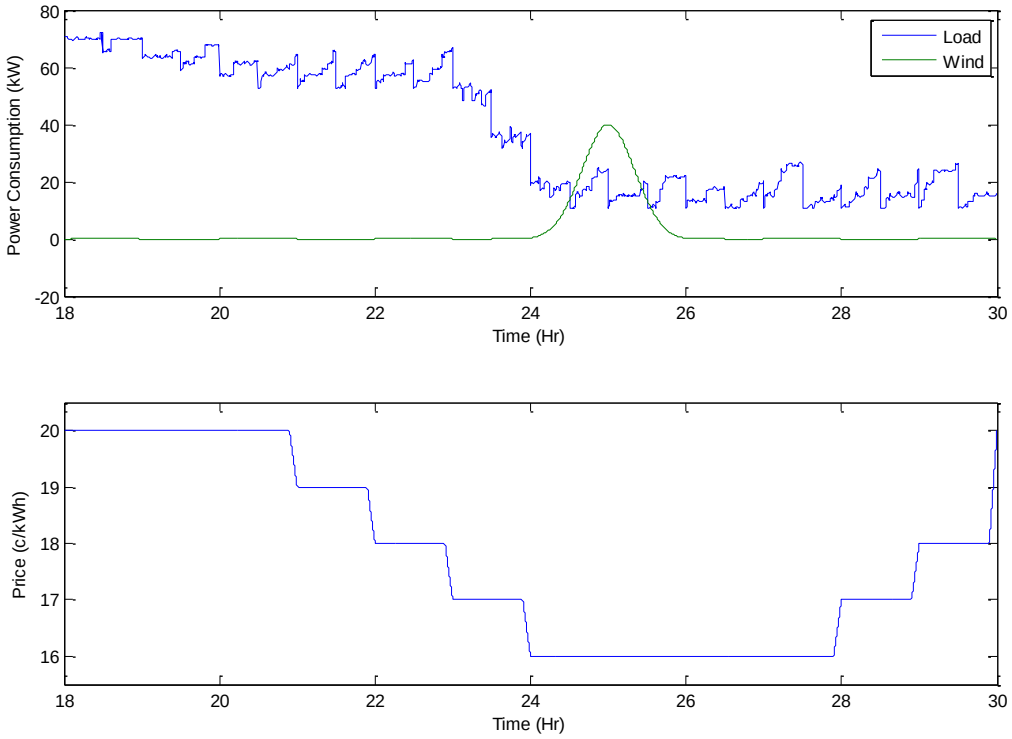

Figure 4.17 – Local generation without scheduling

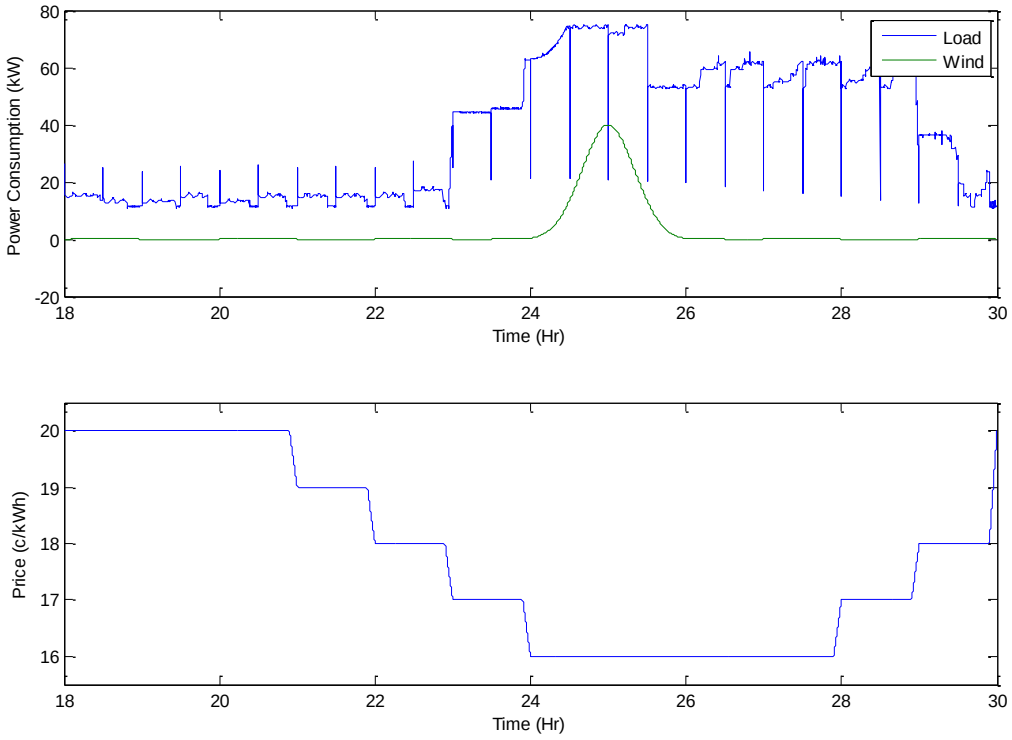

Figure 4.18 – Local generation with scheduling

Table 4.3 – Scheduling simulation results and cost comparison

Simulation	Total Cost ($)	Load (kWh)	Cost p/u (c/kWh)
No PMU, no wind	84.75	456.85	18.550
PMU, no wind	76.16	450.88	16.892
No PMU, wind	79.00	453.42	17.351
PMU, wind	70.21	451.06	15.683

While many simplifications have been made in order to carry out this simulation, the results show that an energy manager with a very small amount of feedback from the largest loads and general knowledge of the price of electricity could schedule loads with high efficiency and significant cost saving.

In this simulation, the state of charge of EVs and hot water cylinders only are considered in the power manager's calculations, and these details are only updated every half hour. This represents a very small amount of data being interchanged, but it has very effective results.

A pricing signal with a known mean and standard deviation may sound unrealistic, but these values could be generated by statistical inference or even provided by the utility company themselves. One of the biggest problems with real-time pricing (RTP) is that for the price signal to be effective, loads must be automated so they can respond immediately. But for automated loads to respond to a price signal, they must have a way of interpreting the price. By being supplied with or calculating an expected mean and standard deviation, making decisions based on the current price becomes quite simple.

An advantage of this scheme is that a supply or generation issue could be represented by a temporary hike in the real-time price, and the power manager would respond accordingly. As long as this increase was short, it would not affect the price statistics and the scheduling would still be effective.

4.9 LDC Performance Discussion

These simulations show the utility of LDC in a number of scenarios. In times of very high wind generation, the controllable loads can be engaged to prevent exporting of energy to the grid, and even without local generation the system could be used to schedule consumption for maximum benefit to the system operator and end users.

With any engineering solution there will always be trade-offs made, and in the case of LDC there are certain things that it is not designed to do which could be beneficial under some circumstances. The LDC system is designed to be placed at a single point of common coupling

(most likely a transformer) and then perform load control based on metrics at that point, by controlling loads drawing power through the point. Under some circumstances there may be metrics further down in the system that are ignored. Examples include individual line voltages and current ratings, including those that run into each premise. Even though voltage on each phase can be measured, when there is only a single steering signal the response will by definition be less effective.

Consider the situation where a street level transformer is to supply a medium amount of power, but one single phase has a large amount of solar generation, causing just that phase to rise significantly in voltage. When only a single steering signal is available, the LDC system can either tell all controllable loads to consume more power, or do nothing. If more power consumption is signalled, the phase with generation may be brought down to a desired level, but the other phases may then be loaded too much, or the transformer could become overloaded. If all phases have approximately equal generation and load this would not be a problem but this may not always be the case. If this is of significant concern, a control signal could be implemented for each individual phase but with added complexity and cost. It is possible that an electronic transformer could intrinsically provide the means of per-phase signalling as the transformer frequency and voltage are artificial and could be used for load signalling. This is discussed in the future work section of Chapter 7.

Another potential situation where LDC would not be beneficial is overloading of the line feeding an individual premise. A household with two high-power EV chargers could easily reach the maximum line current, especially when other household loads are also considered. While this may be considered highly unusual at the present time, an attempt must be made to design power systems with future needs in mind. A potential solution is creating a mini LDC system within the premise operating under the constraints of both the local line and local grid.

The speed of response of the LDC system is also the subject of a performance trade-off. In the design here, the system has been designed to respond in the shortest possible time frame, such that integrating any highly fluctuating generation or performing frequency regulation is possible. A rapid speed of response requires a low-latency communication system, which is likely to cost more. If it is assumed that the system will not respond on very short time frames (e.g. sub-second), then a slower and more generic communication system could be used, potentially including the internet.

4.10 Comparisons with other work in the literature

In Section 2.7 and 2.8, similar works to LDC, found in the literature were outlined, most of which have been published during the course of this research. The notable items will be expanded on here, and compared to the proposed LDC system.

Several works were cited where optimisation on EVs or smart loads was performed with a goal of better integrating renewable generation or improving the efficiency of the electricity network. Most of these simply support the rationale for a system like LDC as they are looking at the grid level energy balance, but one in particular was focused on a small micro-grid. In [87], a micro-grid with 50 users and both renewable and non-renewable generation was simulated. Generation scheduling and consumption was optimised using game theory to minimise the total energy cost for each individual. It was found that with coordination, costs could be reduced by 38% with good wind prediction, and a further 21% with perfect prediction. While this is an interesting study on integrating local renewable energy, it does not appear that any detailed local grid conditions were taken into account such as transformer loading and line voltages.

In Section 2.8.1, several works are summarised that control EV charging in order to limit the load on the distribution transformer. Most of these were very simplistic solutions, such as charge rate limiting [97] or delayed charging [98] but one study [43] came to the same conclusion that controlling EV charging was beneficial for both EV and renewable integration, and that taking into account local conditions could significantly improve the DSM performance. While this is of interest, no actual system or prototype was proposed.

A couple of works in Section 2.8.2 take this a step further and propose systems that perform DSM with a variety of controllable loads, while taking into account local conditions. Firstly, in [71] a system was presented where digital communications are used to efficiently schedule loads, and the system frequency is used for fast response frequency regulation. It is noted that local conditions could also be considered when scheduling, but this has not yet been demonstrated. In [99] a load optimisation scheme was proposed for commercial buildings whereby many loads including the HVAC system are controlled and utility network constraints are taken into account. The simulated system is a simplified, linear programming problem for economic load dispatch. No actual system structure or communication method has been proposed.

The authors of [100] proposed a system that is quite close to the work described here. They conclude that a good DSM system needs to be distributed in order to be efficient with communication and data storage resources. A device is proposed that is placed at the distribution transformer and measures current, voltage and temperature in order to determine whether or not the transformer is overloaded. The device then signals controllable loads via a wired or wireless communication medium in order to bring the transformer into a safe operating state. A queue of device states is used to determine in which order loads should be signalled.

While this is indeed very similar to the operation of LDC, there are some significant differences. Firstly, no sampling rate is given, so the speed of response of the system and general system dynamics are unknown. Secondly all loads are assumed to be simple binary loads, instead of continuously variable loads that can provide a much smoother response, better stability and less load cycling. Thirdly, it does not appear that any work has been done on the actual hardware, communication method, or physical load control. It is an interesting work but it is light on details, making direct comparisons with LDC in terms of performance, cost, complexity, reliability etc. difficult.

In [101], a proposed system structure is presented where the distribution network is divided up into multiple levels of smaller micro-grids, with a controller at each level. The work is focused on the communication needs of the system, with the conclusion that a distributed hierarchical arrangement is more resource efficient than a centrally controlled model. The solution does have some similarities to LDC, with energy managers co-ordinating generation and load autonomously at a local level, while also responding to higher level system commands. Since only the communication needs have been simulated, performance comparisons with LDC are again difficult to make.

The TRIANA system is a very interesting and well thought out solution to the problems of EV and renewable energy integration. An energy manager is placed in each home to perform load forecasting as well as controlling loads. It is noted that due to the use of forecasting, smart load communication is not necessarily required, with devices that simply switch discrete loads on and off providing adequate performance. Predicted load profiles are aggregated at each level and then sent to the controller at the next highest level. This creates a hierarchical structure that is extensible with more levels but also frugal with its communications needs. The global controller takes the aggregated load profiles and calculates the desired consumption patterns.

Steering signals are then generated and passed back down the chain of controllers. The steering signals can be adjusted at each level in accordance with local conditions.

This TRIANA system again has similarities to LDC but is also hard to compare directly in terms of performance. Comparisons can be made in terms of complexity, as each house now has a device that monitors load behaviour and performs prediction, which is likely a computationally heavy task, and each controllable load will have a device that can perform load switching and communication. In a simple LDC system, there are only devices at the transformer level and load level, which communicate directly with each other, so that the household device is unnecessary. If the TRIANA household device also performs metering, then premise-specific load control can be performed, which at the present time is something LDC cannot do. Overall the TRIANA system is a good potential solution, and results showing actual system performance would be of interest. As load control signals would be carried via a digital communication medium, the speed of response and stability would depend on how good the communication system is.

All of the works mentioned previously do not include any hardware proposals or experimental results. There are a few works where hardware systems have been proposed in Section 2.8.3, but most are quite simple and none that have results from a working or laboratory scale system. In [105] a dongle is presented that can monitor frequency, communicate with a HAN and can perform load control in a way similar to the system proposed in [71]. In [107] another smart plug is presented that can respond to “external load control signals”, but little detail is given. Both of these works could be useful as part of a larger system but as of yet, the full functionality has not been demonstrated.

In [108], a system very similar to LDC is proposed, but it uses voltage droop instead of an artificial load control signal. A multi-tap transformer is used which changes between different voltage levels in order to signal controllable loads, while the proposed control scheme is quite similar to that of LDC. It is assumed in the work that these transformers are already common and therefore retrofitting of transformers is simple, but at least in New Zealand, this is not the case. In practice, multi-tap transformers only give a few differing levels of voltage in order to perform load signalling, so therefore the fidelity of load control is likely quite poor. In contrast, the communication system currently used for LDC & described more fully in Chapter 6, uses a frequency tone that can vary by 70-80Hz and changes of less than 1Hz can be discerned. A

voltage droop based system could be effective in an electronic transformer, and this is expanded on in Chapter 7.

In summary, while many researchers are active in the field of smart grid development, as of yet very few works have been discovered that provide similar features to that of LDC, and those that do are still theoretical systems, with no demonstrable results as of yet.

4.11 Conclusions

In this chapter the potential theoretical benefits of LDC have been discussed and many of these have been demonstrated in computer simulations. These include more effective integration of renewable generation and electric vehicles, as well as better management of the low voltage network and distribution transformer. The similarities in control structure between LDC and DDC have also been discussed, as well as the benefits LDC offers over DDC. The implications of the simulated LDC performance have been discussed, as well as situations where LDC would not be able to mitigate grid issues.

Chapter 5
LDC Laboratory System

5.1 Introduction

In this chapter the experimental LDC system constructed as part of this PhD project is presented, including both the LDC controller and the signal transmission system. The system layout described in Chapter 4 is prototyped at laboratory scale, consisting of a power meter, LDC controller, signal transmitter and signal receivers. The power meter measures total power flow into (or out of) this micro-grid, and the power measurement is passed to the LDC controller which then calculates the LDC signal that is passed onto the control signal transmission system. This system then injects a small voltage ripple that can be picked up by load controllers throughout the micro-grid.

A basic overview of the signal transmission system is given, however further detail is covered in Chapter 6. Following a detailed description of the hardware developed, the performance of the system under a number of conditions is evaluated against the simulations of Chapter 4. The LDC system is then augmented with some of the functionalities of DDC, and both simulated and compared experimentally.

5.1.1 Power Meter and LDC Controller

The LDC controller and power meter are integrated into one unit and are shown in Figure 5.1 along with the transmission system. The controller includes a three phase four wire power meter, an Atmel microprocessor, and isolated com ports for interfacing with the transmission system and a user terminal.

A simplified diagram of the power meter and controller is shown in Figure 5.2. The power consumption is measured by three analogue multipliers that each receive a signal from three voltage transformers and three LEM (current) sensors. With this setup, the power meter will work correctly with both three and four wire three-phase systems. The outputs of the three multipliers are then added with an op-amp addition circuit. Since a negative power flow will produce a negative voltage, the second op-amp of the pair is used to generate a second inverted power signal. Both of these power signals are then positive and can be accepted by a microcontroller with a single sided supply. These two power signals, as well as the three original voltage signals are then passed to the LDC control circuitry.

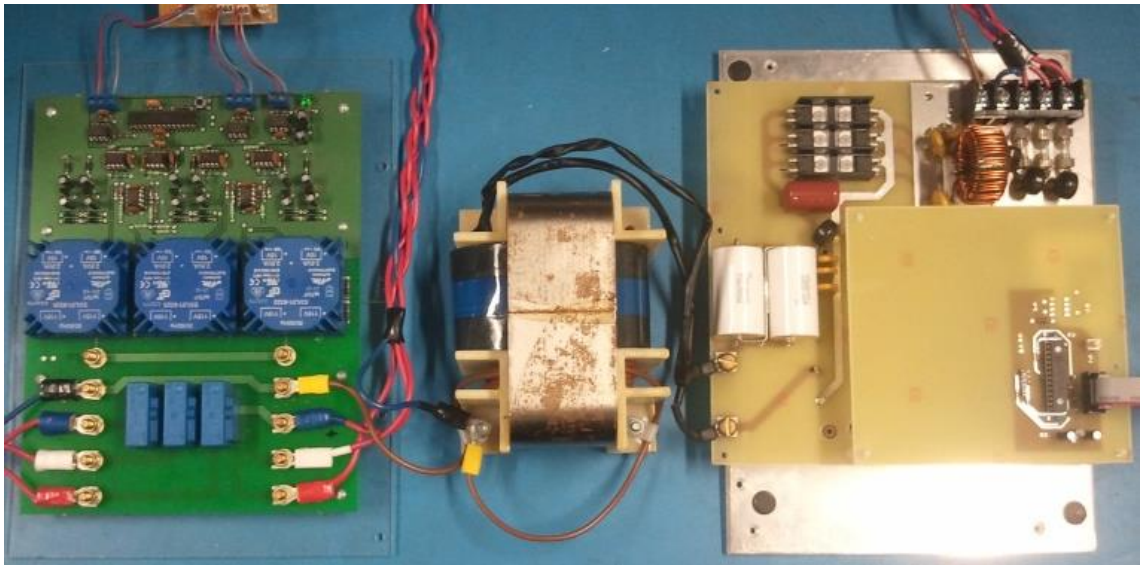

Figure 5.1 – Photo of LDC test hardware showing control unit (left) along with signal injection transformer and electronics (center & right)

In order to make the signals safe for connecting to a microcontroller, each power signal is passed through a precision rectifier with a 5V supply, such that the maximum output voltage is between 0 and 5V. To measure line voltage and system frequency, the voltage waveforms are stepped down to around the 5V level, then passed through a passive peak detector, and then stepped in half again. One phase is passed through a Schmidt trigger to produce a 50Hz square wave for measuring the system frequency. The two power signals and three voltage signals are then each fed into an ADC channel on the central microcontroller, and the 50Hz square wave is passed into a timer capture pin.

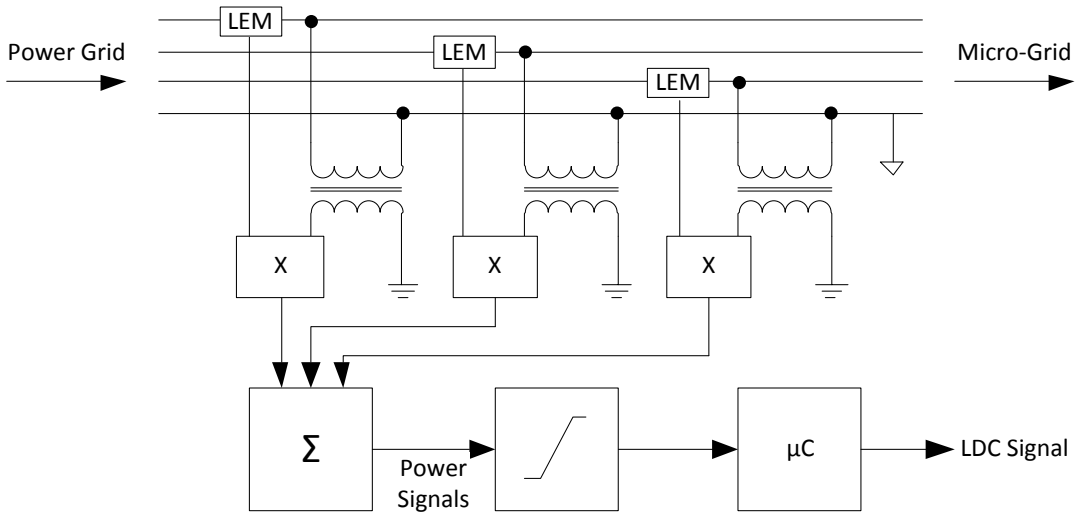

Figure 5.2 – Power meter block diagram

The input power is sampled by the microcontroller 4000 times per second and an averaged sample used for control is created at 50Hz. Since an entire cycle is averaged, oscillations in the power measurement caused by phase imbalances are reduced.

To implement the integral control action described earlier, the difference between the internal power set point and the input power is calculated. This difference is then multiplied by the controller gain and added to a variable representing the LDC signal level. The level is then sent to the signal distribution system via an isolated pulse width modulated (PWM) output.

Through the use of a terminal interface, a whole host of variables can be set and observed. Variables such as the controller gain and current power set point can be adjusted and internal variables such as the input power and the LDC level can be observed. The controller can also be put into one of three modes: normal mode, silent mode and test mode. The normal and silent modes run the LDC system in the standard manner, but the silent mode does not output anything to the terminal unless requested. This is useful when the system operation is being logged via MATLAB, and a MATLAB program polls for the LDC system status at set intervals. The test mode enables the system to output a constant predetermined LDC signal for testing the communication system.

5.1.2 Control Signal Transmission System

While the operating philosophy of the transmission system is discussed further in Chapter 6, with a focus on the signal filtering design, the basic design will be outlined here in order to clearly convey the experimental setup. The block diagram of the transmission hardware is shown in Figure 5.3, and the actual system is shown previously in Figure 5.1.

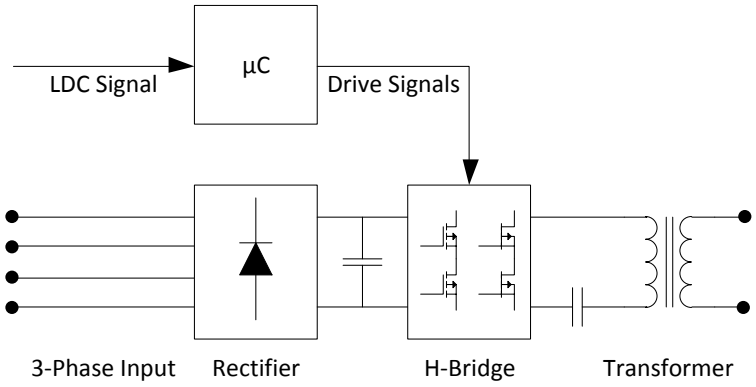

Figure 5.3 – Block diagram of LDC signal injector

As shown the input three-phase power is rectified into a DC bus. A simple H-bridge then drives the signal injection transformer, with a series capacitor for removing any DC bias. The H-bridge is driven by a simple programmable system on chip (PSoC) design, which takes in a square wave from the power meter and uses it to drive the first leg of the H-bridge. The second leg of the H-bridge is driven from this same generated square wave, delayed in time.

5.1.3 Load Controller

A small piece of circuitry is used to receive the control signal and consequently determine the switching behaviour of the controllable load. A block diagram of this load controller is shown in Figure 5.4.

The mains signal is taken in and split off, with one set of connections going to generate the 5V supply, and the other going to a resistor divider and high-pass filtering network. The output of the high-pass filter is then biased about 2.5V before being fed into the PSoC. Inside the PSoC, a four pole band-pass filter is used to fully extract the LDC signal, before it is fed into the microcontroller core and filtered in software. The design and testing of the filtering circuitry is the main subject of Chapter 6. Once filtered, the signal can then be used to determine the load switching behaviour, which is sent to the PWM and UART outputs.

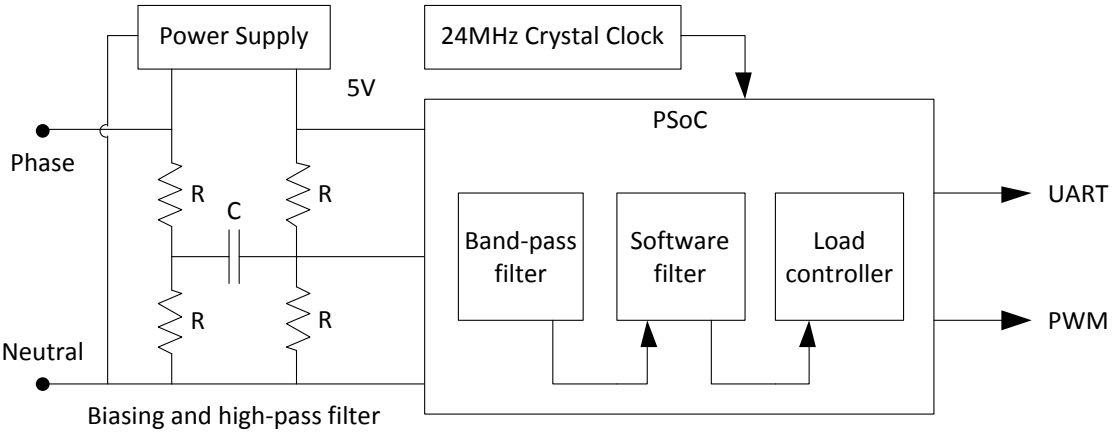

Figure 5.4 – Block diagram of load controller

The final version of the load controller is shown in Figure 5.5. It is a very simple piece of circuitry, with most of the components used to provide a 5V power supply from the $230V_{RMS}$ mains. A small network of resistors and capacitors can be seen on the right side of the figure. These make up the high-pass filter and biasing circuitry. Left of these is the PSoC IC and a crystal clock source. There are headers for PWM, UART and for programming the device.

The PWM and UART outputs are the two methods by which the board can perform load control. If a digital programmable load is used, the desired impedance can be sent to the load via the UART. This gives a continuous and smooth response, but adds a little delay because of the time taken by the load to process each new value. Care must also be taken to not flood the load with new values, as it is possible to overload it and cause it to stop responding. In addition, the power limits on the load must be set correctly, as it is possible for an error in the program

or UART link to cause a 0Ω to be received, possibly in connection with the flooded UART receive – this is of course highly undesirable.

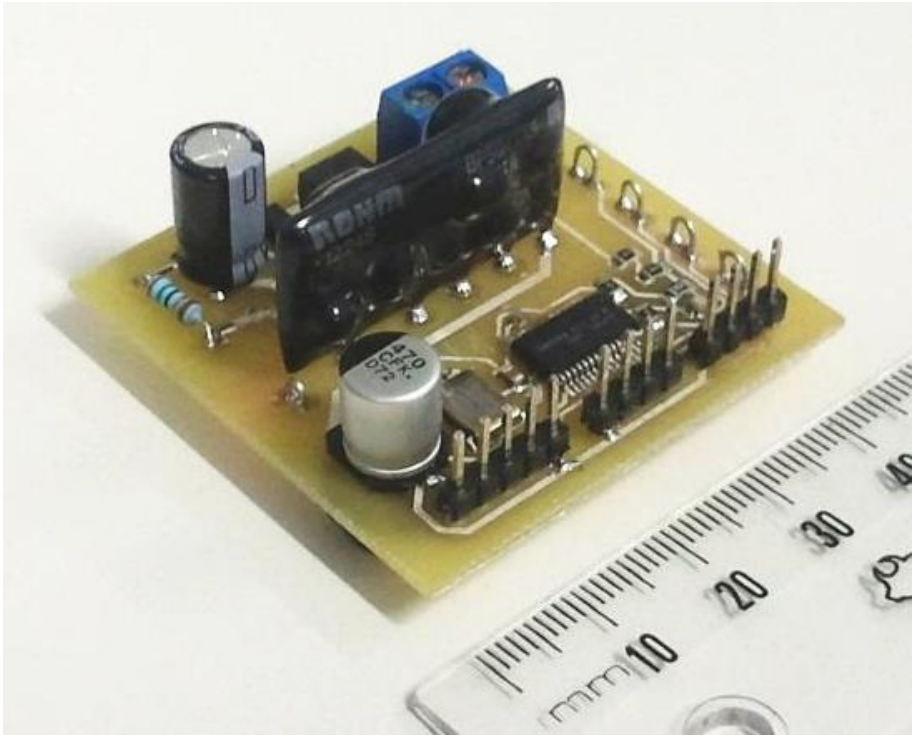

Figure 5.5 – Photo of LDC signal receiver and load controller

The second method is the connecting of the PWM output to a solid state relay. The response of the relay is much faster than that of the controllable load, but due to its binary nature, it is not nearly as smooth. In practice, the UART controlled load provides the best response. These methods model the two different ways in which smart loads are likely to be implemented. In the first case the load has some internal intelligence within its circuitry and can be told how much power to consume through either a standardized smart load interface or internal LDC receiver. In the second case a 'dumb' load has its power supply switched on and off as required by the LDC receiver – which could be achieved by a "dongle" placed inline at the wall outlet.

5.2 LDC Experimental Performance

5.2.1 Experimental Setup

A 3kW laboratory scale system has been constructed to verify the system performance. The system models a small micro-grid with 3 loads, some local generation and a grid connection. The layout of the experimental setup is shown in Figure 5.6. The supplied power is measured at the input to the system (P_{MEAS}) and a control signal is generated and distributed (LDC) to (ideally) keep this at a 400W set point. The system has three loads, the first two of which are

controllable, varying from 0-600W while the third is set at a constant 500W for the entire experiment.

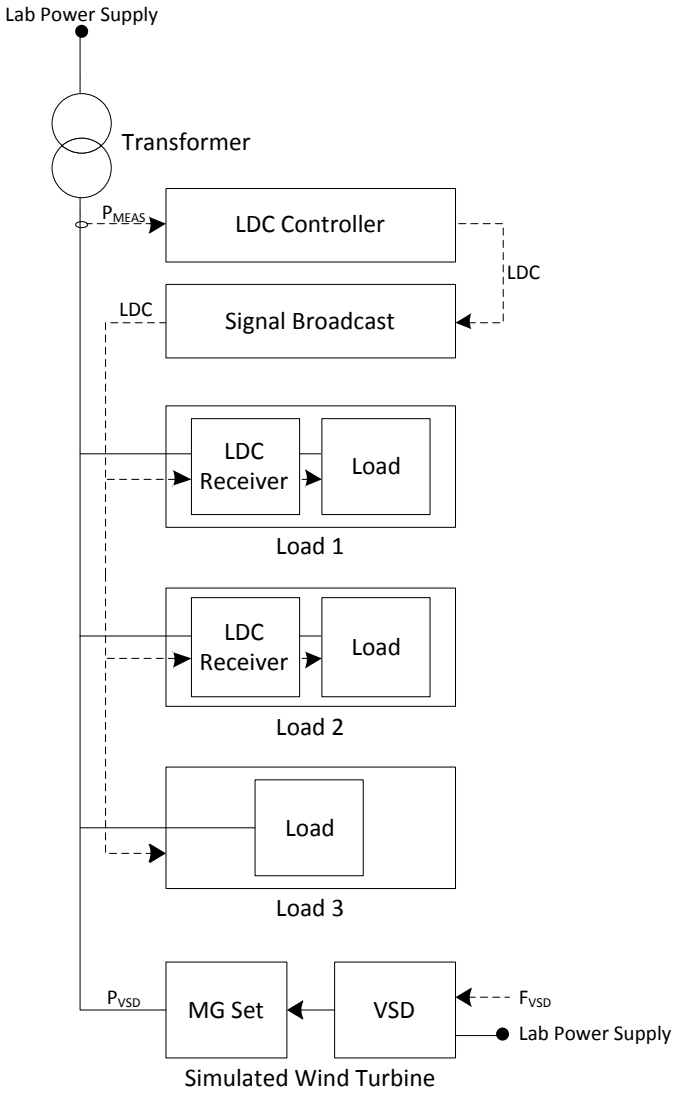

Figure 5.6 – Network diagram of experimental setup of LDC system

The experimental rig includes local generation simulated by a motor generator (MG) set: A variable speed drive (VSD) controls the motor with a predetermined 'random' frequency pattern (F_{VSD}). This widely fluctuating pattern could represent the output from a small wind turbine with a highly variable wind speed. As induction machines were used, small changes in input frequency caused changes in slip and consequently motor torque. These changes in torque caused step changes in the power supplied by the MG set (P_{VSD}), and these can be seen in the LDC system response. This frequency reference was only used within the MG set to produce a random torque, as opposed to the system frequency which was the mains supply at a nominal 50Hz.

A MATLAB program has been written and additional circuitry created to automate the experiments. A small Atmel ATMega8 board has been added which uses a PWM output to control the VSD. The output is filtered and fed into the VSD as an analogue voltage signal. The VSD has also been configured to output its current power consumption as an analogue signal, which is sampled by the ATMega8. The ATMega8 receives the desired frequency and can report back the current power via an RS232 interface. By using a number of USB to RS232 convertors, the power meter, VSD controller, and digital loads, can all be connected to the one computer, and operated from the same MATLAB program.

To run the experiment, the program opens serial connections to all of the devices, sets up the desired initial parameters and then starts sending the different frequency reference values to the VSD. Power consumption is then polled every second. Ideally a faster poll rate would be used but the controllable loads can have a delay of up to a second when responding to requests for power consumption. In practice it was most effective to send a request for power, then perform all the other functions required each cycle before checking the returned power value. The received values from each device are logged to an experiment file, ready to be plotted.

5.2.2 Power Regulation Performance

The results of an experiment running over three minutes are shown in Figure 5.7. A varying frequency reference (F_{VSD} in Figure 5.6) is supplied to the VSD and is shown in plot (a). The output power (P_{VSD}) of the VSD in response to this is shown in plot (b). The calculated DSM signal (LDC) that is broadcast to each load is shown in plot (c). Plot (d) shows the consumption for the two controllable loads, while the input power to the system (P_{MEAS}) is shown in plot (e).

The set point was a nominal value of 400W and the power input shows good regulation around this value, with a standard deviation of 33W. Transient spikes are visible on the power input when there are sudden changes in local generation, but these are expected given the layout and nature of the system as discussed in Chapter 4.

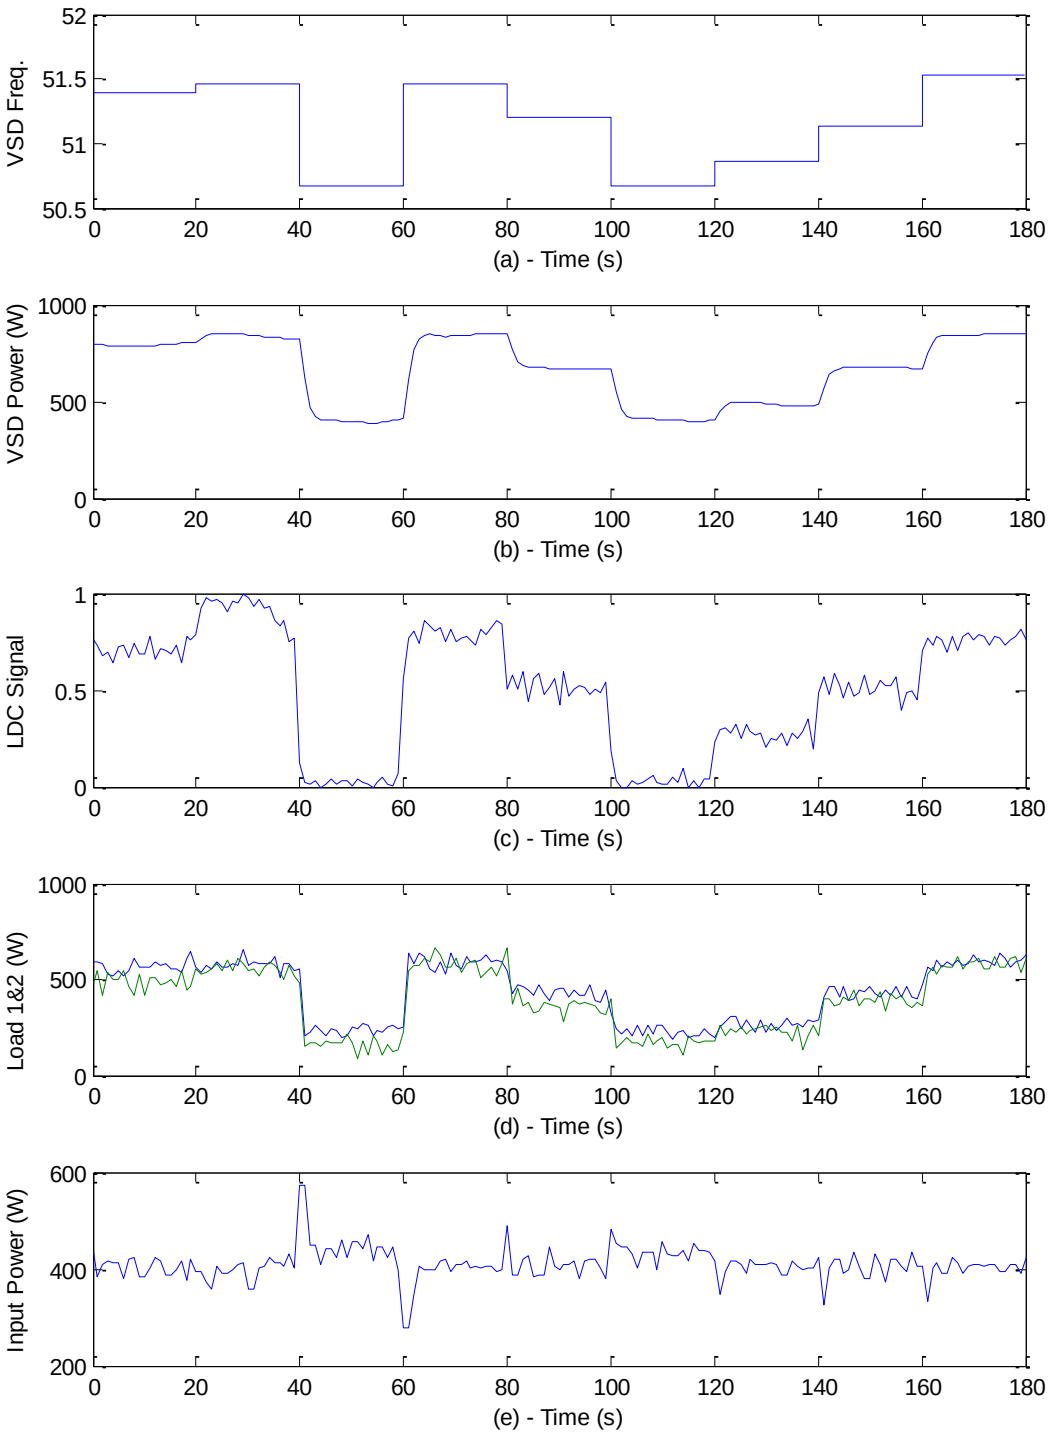

Figure 5.7 – Plot of LDC results showing (a) input frequency, (b) input power, (c) LDC signal, and (d-e) power flow

While it is expected that the two controllable loads would have exactly the same consumption, in practice there is an observed difference of 10-20W. This is most likely caused by slight differences in clock speed of the LDC receiver microprocessor. As minute changes in a frequency tone are being sampled, variations in clock speed will cause variations in the LDC signal observed and in the receiver response. Although crystal references have been added in later versions of the hardware, these variations are not necessarily a disadvantage as it leads to

more diversification of the response. In addition, LDC signal noise and slight variations in RMS line voltage will also cause minor differences.

5.2.3 Step Response Performance

In order to more accurately quantify this transient response, the analogue power signal (P_{MEAS} in Figure 5.6) measured during a step transient is shown in Figure 5.8. Here roughly 500W of fixed load is switched out manually, and the LDC system brings 500W of controllable load back online in less than one second in a well-controlled manner.

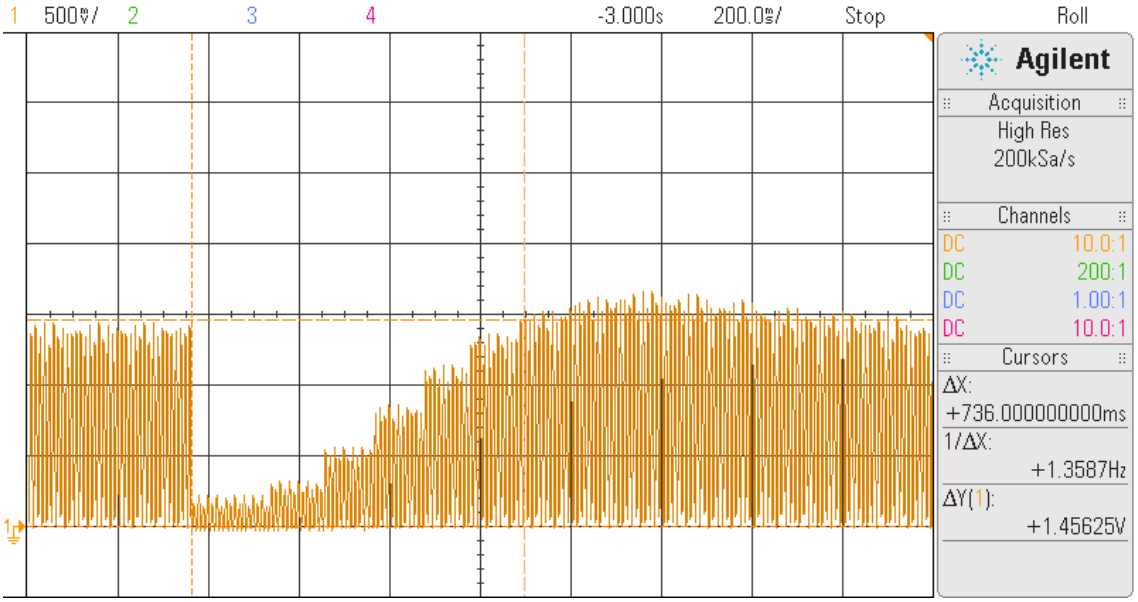

Figure 5.8 – Step response of LDC controller showing power measurement regulation.

Overall the power regulation of the system is very good even with only two controllable loads. It is expected that a larger system with a greater number of controllable loads will behave more smoothly due to the averaging effect of the combined responses. The effect of the controller gain on the speed of response and the system stability will now be looked at in more detail.

5.2.4 Effect of Controller Gain on Power Regulation

Within the LDC controller, there is a digital value representing the controller gain, which affects how fast the measured power imbalance is integrated to produce the LDC control signal. If too high a gain is used, the system can become unstable and will oscillate. High gains will also cause large control actions and overshoots.

In Section 4.4, the control structure of LDC was introduced, and the integrator was shown to have a time constant T, which controlled the speed of integration. Within the actual LDC controller, the power imbalance is multiplied by the gain, so there will be an inverse

relationship between the digital gain (G) and the time constant in the transfer function (T). By looking at the various sampling rates within the microcontroller, the relationship between the digital gain (G) and integrator time constant (T) has been determined to be:

$$T = \frac{1.34 * 10^6}{G} \tag{5.1}$$

With this relationship, it is possible to predict the behaviour of the LDC system from the equations derived in Section 4.4 for a given digital gain. The damping ratio of the LDC system was given as:

$$\zeta = \frac{1}{2}\sqrt{\frac{T}{kT_1}} \tag{5.2}$$

Here, T is the integration time constant, k is the controllable load and T_1 is the filter time constant. The communication system as configured for these experiments had a filter sample length of 256, which with a nominal operating frequency of 800Hz, has a nominal time constant of $T_1 = 0.32$. With a controllable load of k = 700W/Wh the time constant and damping ratio can be calculated for a given digital gain. Here three gains are given that are integer powers of two. These are used to keep multiplication within the microcontroller simple. The predicted damping ratio and overshoot for each digital gain is shown in Table 5.1.

Table 5.1 – Predicted performance for varied digital gain values

Digital Gain (G)	Time Constant (T)	Damping Ratio (ζ)	Overshoot (%)
2048	654.3	0.8545	0.56
4096	327.1	0.6043	9.23
8192	163.6	0.4273	22.7

From the calculated damping ratios, it can be predicted that with G = 2048, the response will be good with almost no overshoot. With G = 4096, a fast response should be observed with overshoots of around 10%. With G = 8192, there should be a very fast response with large overshoots and significant control action.

An experiment was performed where the same load and generation conditions as in the previous sections, but with these three different controller gain values. In the same manner as in the previous section, a random frequency reference is fed to the VSD, and the system is run with three loads, two of which are controllable. The output of the system with gains of 2048, 4096 and 8192 are shown in Figure 5.9, Figure 5.10 and Figure 5.11 respectively. In each

figure, plot (a) shows the VSD frequency reference frequency, plot (b) shows the LDC signal and plot (c) shows the input power to the system (labelled as P_{MEAS} in Figure 5.6).

In Figure 5.9 it can be observed that the LDC signal is quite smooth, but the input power is quite noisy, with a standard deviation measured separately at 30.65W. In Figure 5.10, the LDC signal shows a little more control action and the input power is more stable with a standard deviation of 20.04W. Finally in Figure 5.11, the LDC signal is very noisy, but the input power is now very flat, with a standard deviation of 14.51W.

These observations fit with the expected results for the damping ratios calculated earlier. It can be observed that a gain of 2048 was stable but a little slow, 4096 was fast and a little erratic, and 8192 was very fast but bordering on instability.

In these experiments that sampling rate was too slow to gauge whether the percentage overshoot matched the predicted value. The step response performance shown previously in Figure 5.8 was sampled at a higher rate and performed with a gain of 8192, which does indeed show approximately 20% overshoot as predicted.

As expected the regulating performance with a gain of 8192 is very good, but the LDC signal varies quite sharply, showing that significant control action is now being applied to produce the highly regulated input power. In practice, this means that with this high gain, loads would be switching on and off more frequently to get just the right level of power consumption. A gain of 4096 was used for general experimentation (except in Figure 5.8), as it was a good trade-off between power regulating performance and LDC signal stability.

In equation (4.6), the system stability is shown to be dependent on the integral time constant, filter time constant and amount of controllable load. The amount of controllable load is likely to change throughout the day, as different loads are either removed or reach operating limits. An ideal LDC system would know exactly how much controllable load is available and adjust the time constant / controller gain accordingly. Alternatively, the gain could be configured for a good response with the maximum expected controllable load, such that with fewer loads the response would simply be slower, and the system would always be stable.

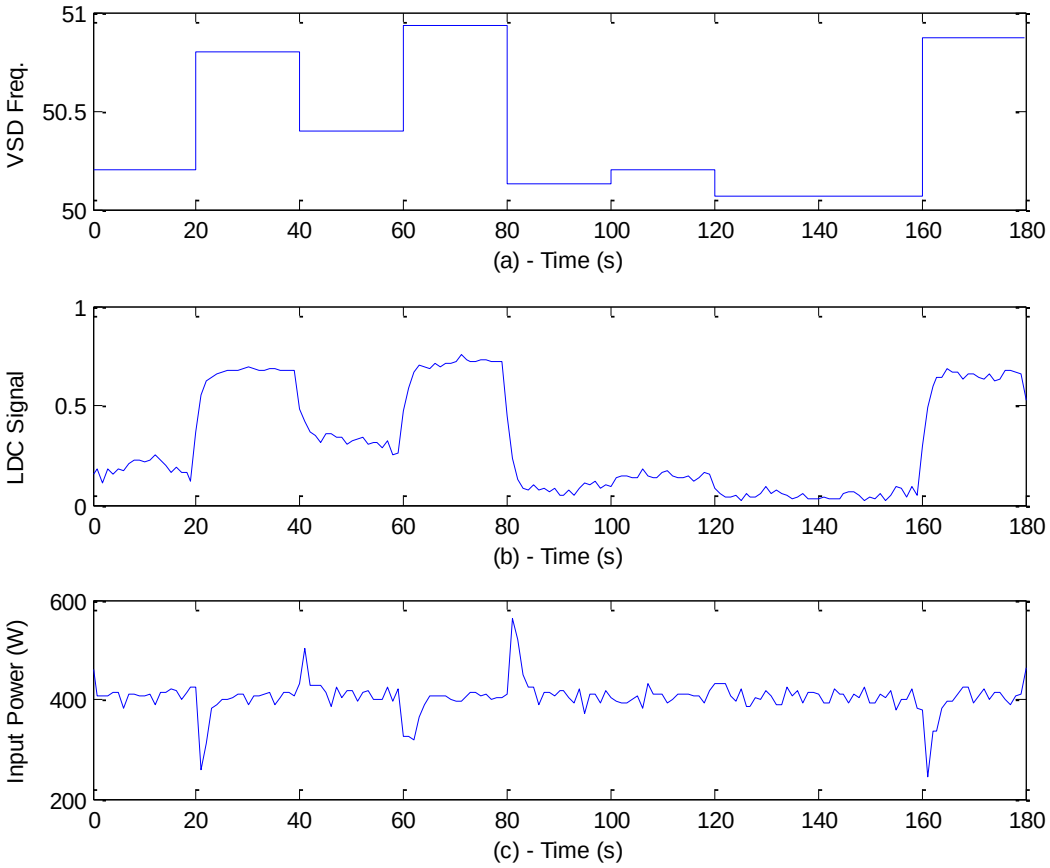

Figure 5.9 – Power regulation with controller gain of 2048

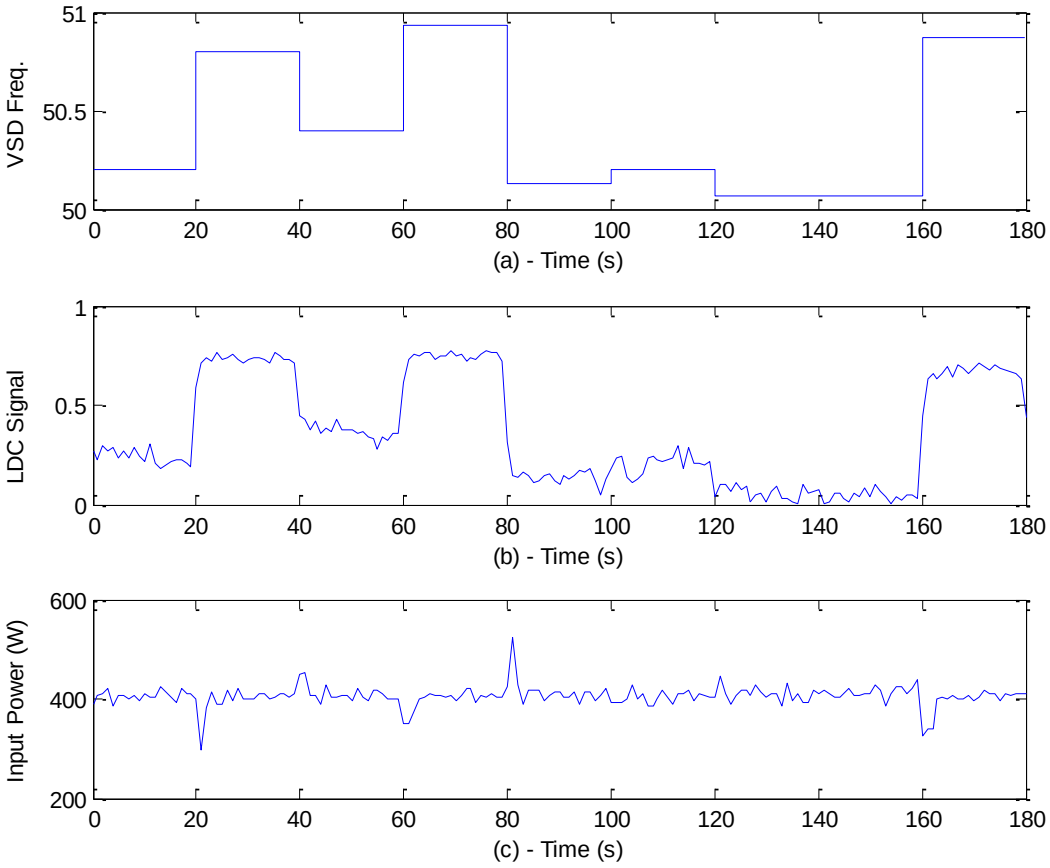

Figure 5.10 – Power regulation with controller gain of 4096

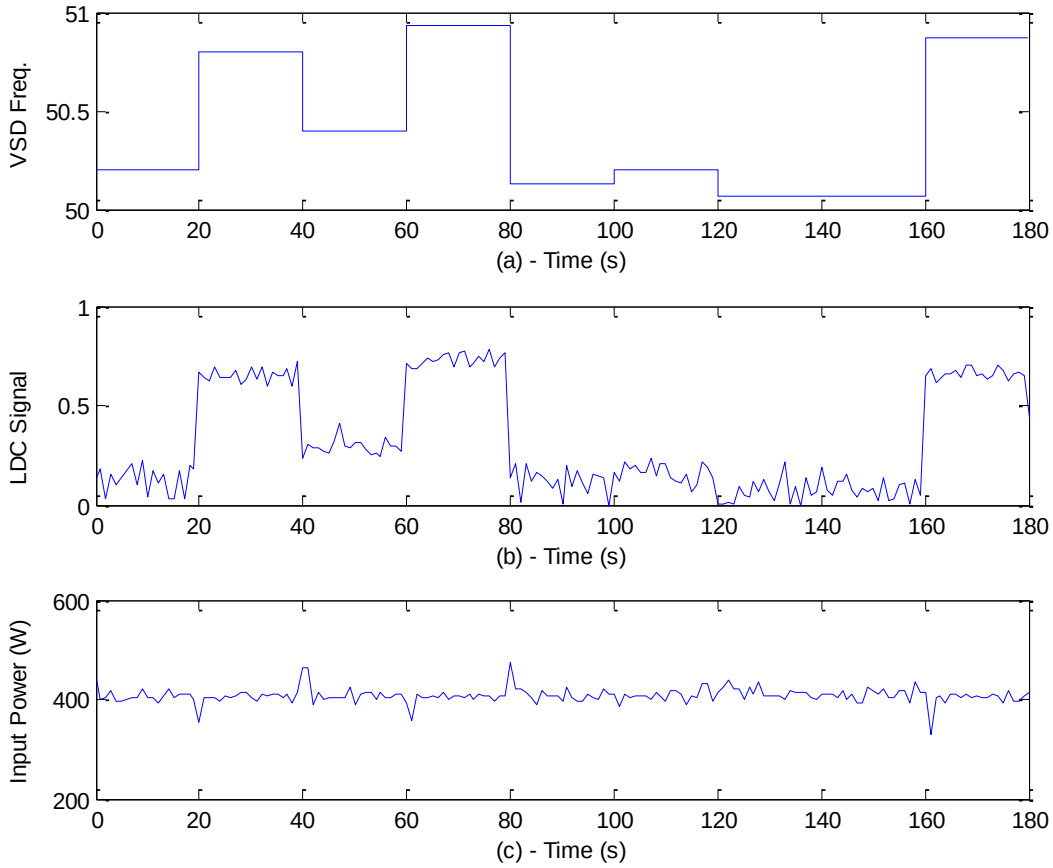

Figure 5.11 – Power regulation with controller gain of 8192

5.3 A combined DDC and LDC system

Dynamic Demand Control is a method of DSM whereby loads adjust the amount of power taken based on the measured system frequency. A full overview of DDC has been given in Chapter 3. Localised Demand Control (LDC) is a new system designed to have much the same operating principles as DDC while being able to deal with smaller scale areas of operation, and being suitable for modern power systems. A full overview of the design and performance of LDC is given in Chapter 4.

By adding additional circuitry to the LDC controller, system frequency and voltage can be monitored as seen at the distribution transformer where the controller is located. In this way the system can automatically respond to significant frequency events, dropping controllable load as the frequency drops, and adding additional load during an over frequency event.

This scheme has a number of advantages over other emergency response methods such as automated under-frequency load shedding (AUFLS) and ripple control. With AUFLS, when a frequency event is detected, whole suburbs simply have their power supply removed, irrespective of how important each load is. With ripple control, a message can take from 5.2 to

6.6 seconds for the transmission only, which precludes it from any fast acting load response situations where the response time should be less than 1s.

With the current LDC communication system, the response time is much less than one second, and only large controllable loads such as water heating, HVAC and vehicle charging would be affected. Two methods of frequency response are envisaged; the first is where controllable load is switched when the system frequency is detected to have moved outside a set band, and the second is where the LDC power set point is adjusted directly in response to system frequency for a linear, highly damped response. Since a linear response will likely provide the best performance, this method is investigated here.

5.3.1 Simulated DDC and LDC system

A simulated LDC micro grid has been created in MATLAB Simulink, and a diagram of the network layout is shown in Figure 5.12. It consists of ten houses running off one distribution transformer with an LDC controller. Each house has a hot water cylinder, EV, fridge, and fixed base load – the same as in previous simulations. Aside from the base load, all of these are controllable. The average daily power consumption is 2kW, half of which is controllable. This is higher than might typically be found but includes 600W for an EV.

A simulated power grid feeds the micro-grid. It consists of a simulated fixed (uncontrolled) load and three generation sources. Generator 1 is a large generator with a slow responding frequency-keeping characteristic. Generator 2 is a medium sized generation source that is used to simulate faults. Generator 3 is a small frequency-keeping generator, which has a very fast response but small overall capacity. Note that the power flows shown in the expanded figures do not sum to zero as the fixed load has been excluded.

An example of the simulated output for half a day is shown in Figure 5.13. Here the LDC controller is performing some load shifting, moving the bulk of controllable load into the early morning with operation as discussed in Section 4.8. The load shifting is conducted in such a way that user convenience is not impacted. A large frequency event is simulated at 11:45pm and this can be seen in expanded form in subsequent figures.

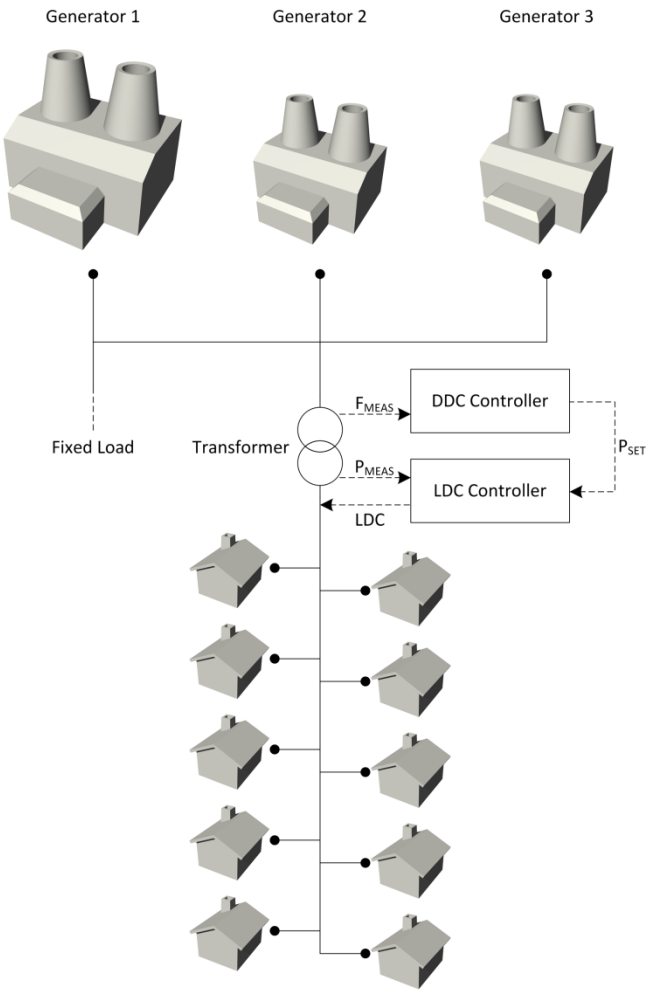

Figure 5.12 – Layout of simulated power network with both DDC and LDC

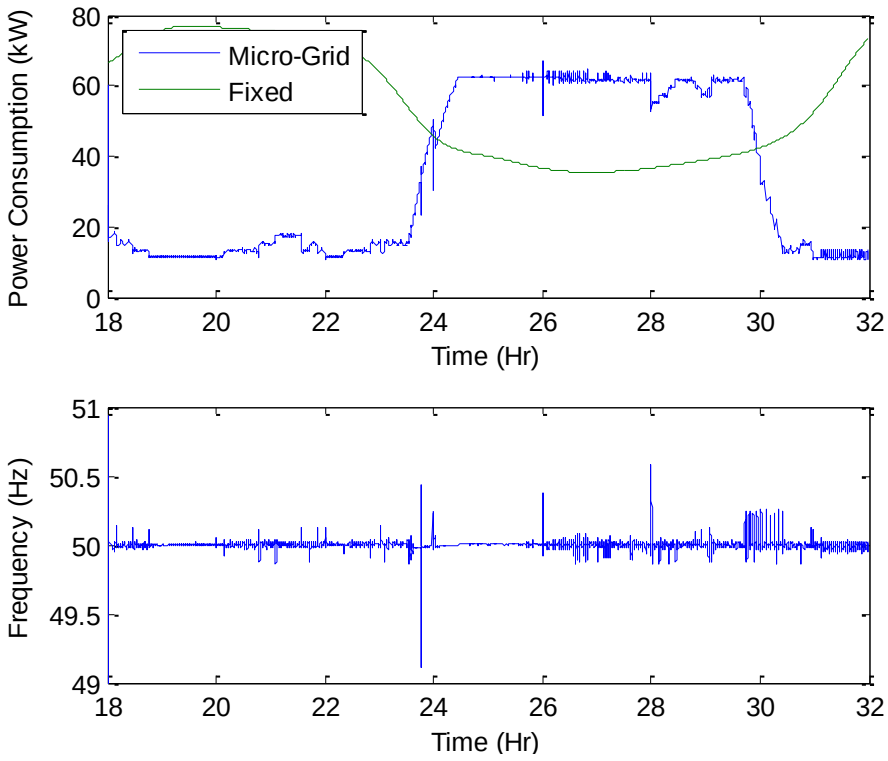

Figure 5.13 – Simulated grid and micro-grid output for a 14 hour overnight period

A DDC frequency monitor was added to the LDC controller which adjusts the LDC power set-point (P_{SET}) linearly in response to changes in the system frequency (F_{MEAS}). The simulated grid has transmission delays of 0.5s between the generation sources and the LDC micro-grid in order to simulate the effects of transmission line reactance and delays in filtering the frequency signal.

5.3.1.1 Results with sudden loss of generation

A power failure was simulated by dropping the output from Generator 2 from 20kW to zero. One half of the generation was brought back after 15s and the other half after 30s. When the system is drawing 60-120kW, this is a significant amount of generation to lose. The responses with both no DSM, and linear DDC were simulated.

As shown in Figure 5.14, with no DSM there is a significant drop in frequency to around 48.3Hz, and then a spike in frequency to above 52Hz when the generation is brought back online. This is not a desirable result.

The linear DDC response is shown in Figure 5.15. Here the controller adjusts its power set point by ±20kW over the frequency range 49.3-50.7Hz with a small dead zone in the middle. The out of band response is still present but is not needed. The response shown has far better damping and the frequency is kept within a tighter band, without any oscillation.

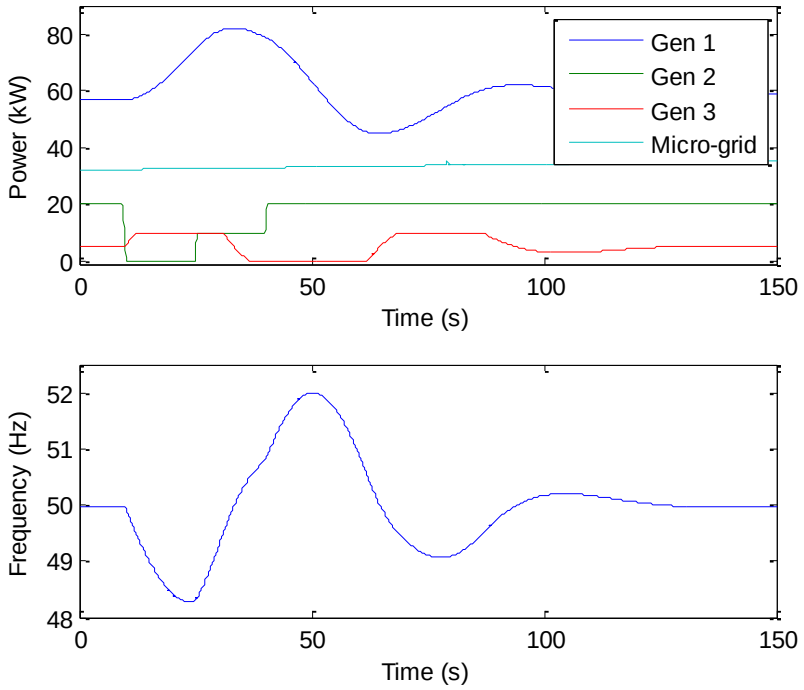

Figure 5.14 – Response to sudden decrease in generation with no DSM

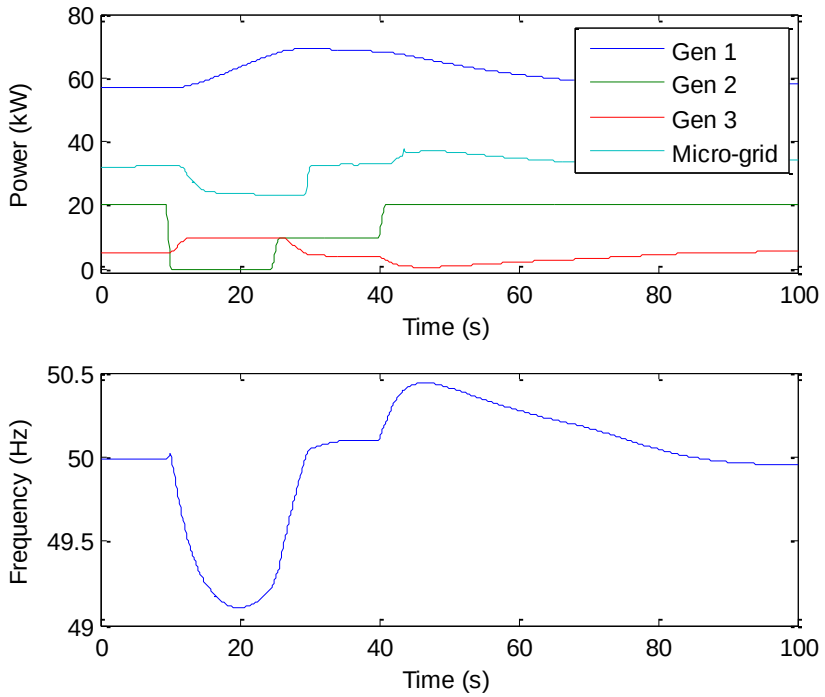

Figure 5.15 – Response to sudden decrease in generation with linear DDC

5.3.1.2 Results with sudden increase in generation

A sudden increase in generation was simulated in a similar manner by adding 20kW of generation to Generator 2. One half of this is dropped within 15s and the other half in 30s. This could simulate a sudden gust of wind driving a wind turbine. A similar response would also be seen in the event of sudden load disconnection. Figure 5.16 shows the response of linear DDC. Again load is added and removed linearly for a far smoother response.

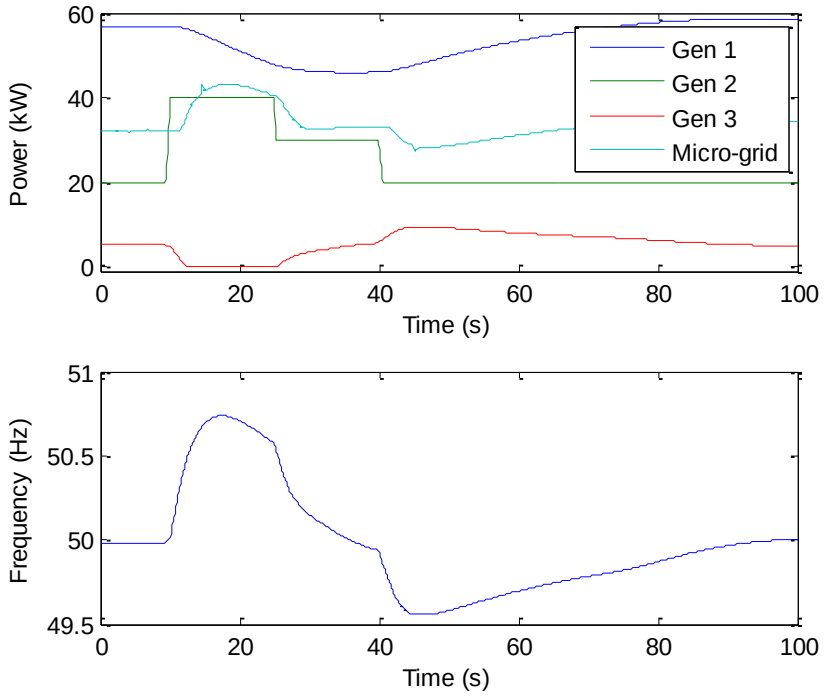

Figure 5.16 – Response to sudden increase in generation with linear DDC

These results show that a combination of LDC and DDC could perform well in regulating system frequency during system disturbances. This frequency regulating service could even be aggregated and sold as an ancillary service to the system operator.

5.3.2 Experimental DDC + LDC system

An LDC system previously created for verifying the performance of LDC has been modified with frequency measurement functionality. As shown earlier in Section 5.2, the LDC system has good performance when regulating the power consumed by a micro grid to a given power set point. By linearly adjusting the set point based on the measured system frequency, DDC functionality can be integrated.

The added circuitry to provide frequency measurement is minimal, consisting of a comparator operating off the line voltage, producing a square wave that is fed into an interrupt generation pin of the microcontroller. Better performance could have been achieved if the timer input capture pin was used instead, but this was not available without fully redesigning the system.

A diagram of the experimental setup is shown in Figure 5.17. In order to provide the ability to vary the system frequency, a motor generator set was used as the main power source. A large VSD drove the MG set, which consisted of a 15kW induction machine driving a 3-phase synchronous generator.

Three loads were put on the system, with one on each phase. The first is controllable, receiving a signal from the LDC controller and adjusting resistance accordingly. The second two loads were fixed and combined consumed a constant 600W, while the controllable load could be varied from zero up to around 600W.

The LDC controller had a nominal set point of 850W, with ±200W added or removed over the frequency range 49.5-50.5Hz. The frequency response was not limited to this range though, so further load dispatch would be attempted under more extreme frequency events, with the same slope of 400W/Hz.

The VSD was placed in a torque limited mode, with three different torque values used. The torque limit is shown as τ_{VSD} in Figure 5.17, and was adjusted during the experiment using a terminal interface. The torque values corresponded to three operating frequencies of 49.5Hz, 50Hz, and 50.5Hz with the chosen LDC and load setup.

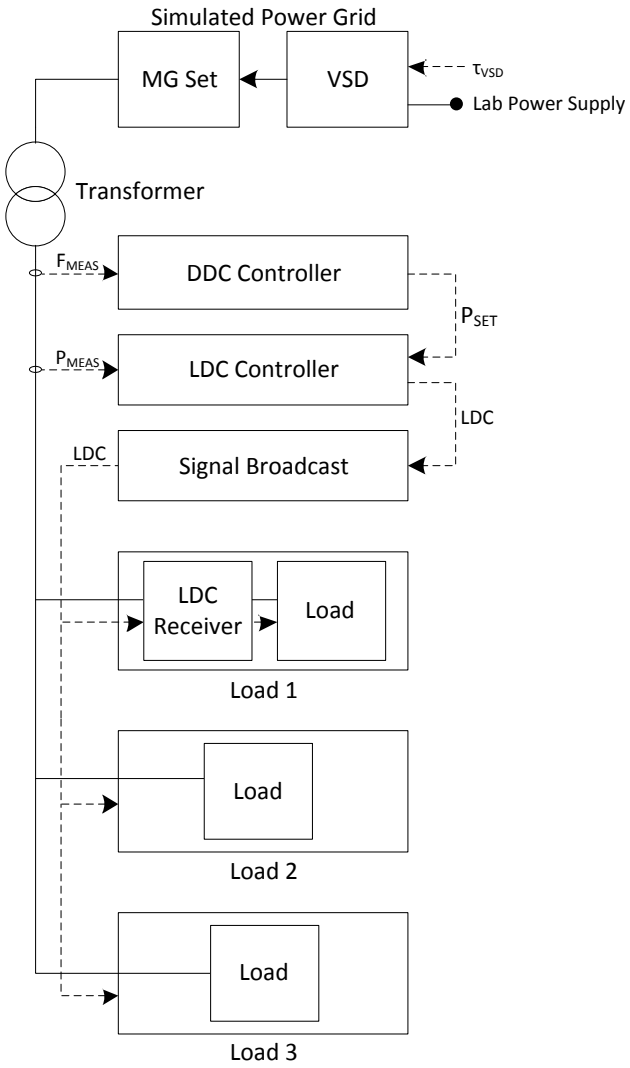

Figure 5.17 – Network diagram of experimental setup of combined DDC and LDC system

The operation of the system under the three different torque limits is shown in Figure 5.18. Plot (a) shows the system frequency, while plot (b) shows the LDC signal, and plot (c) shows the system power consumption. The LDC signal is a scalar that indicates to controllable loads the level of power available. In this case, the controllable load will vary its power consumption linearly in response to the LDC signal. Note that errors in the measured system frequency are generated due to the non-ideal measurement setup that involves interrupt driven timer capture.

It can be seen that the system responds well to changes in system torque, arriving quickly at a stable operating point. It is also worth noting that the constant operating frequency during periods of constant torque is only possible due to demand side regulation, further demonstrating good system performance. Without this regulation the generator quickly reaches its upper or lower frequency limit as defined within the VSD.

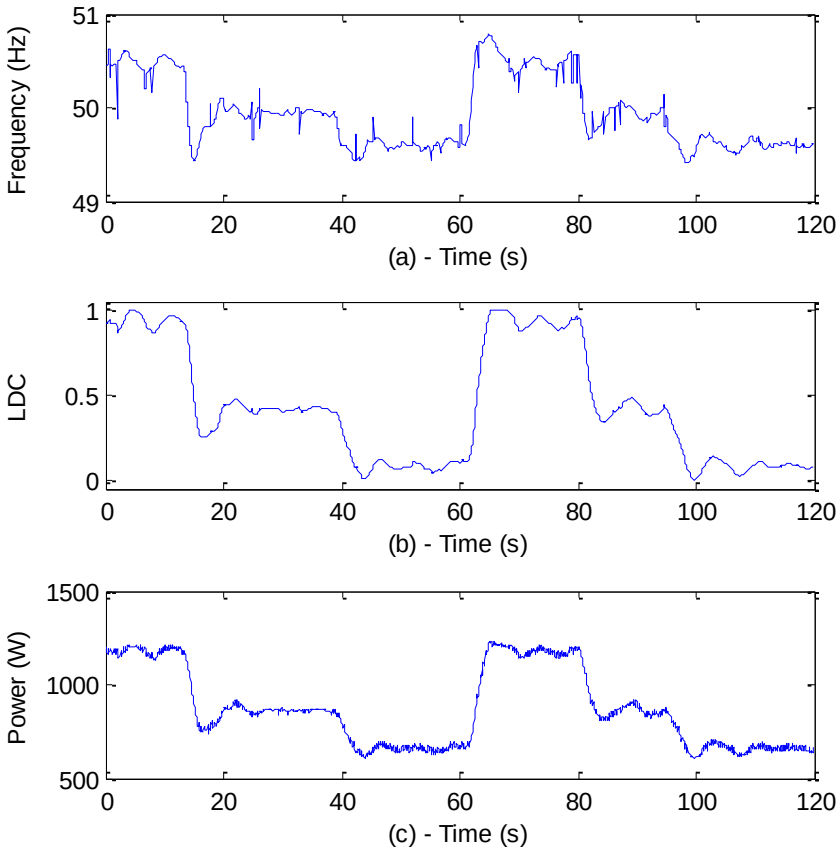

Figure 5.18 – DDC operation of LDC system under varied torque conditions showing (a) system frequency, (b) LDC signal, and (c) system power

An enlarged view of the system during a step torque reduction is shown in Figure 5.19. Due to the large capacity (and therefore inertia) of the generator relative to the power being consumed, the system frequency does not shift rapidly. It can be observed though, that the system power consumption does nicely track the system frequency, with a small amount of overshoot and ringing.

Similarly, Figure 5.20 shows the system response to a step increase in generator torque. Again, the LDC system tracks the system frequency without any noticeable delay and a small amount of overshoot and ringing.

These results show that an LDC system operating in a DDC like frequency responsive mode can provide valuable frequency regulation. A similar result would be expected if the power meter was additionally monitoring for drops in the system voltage.

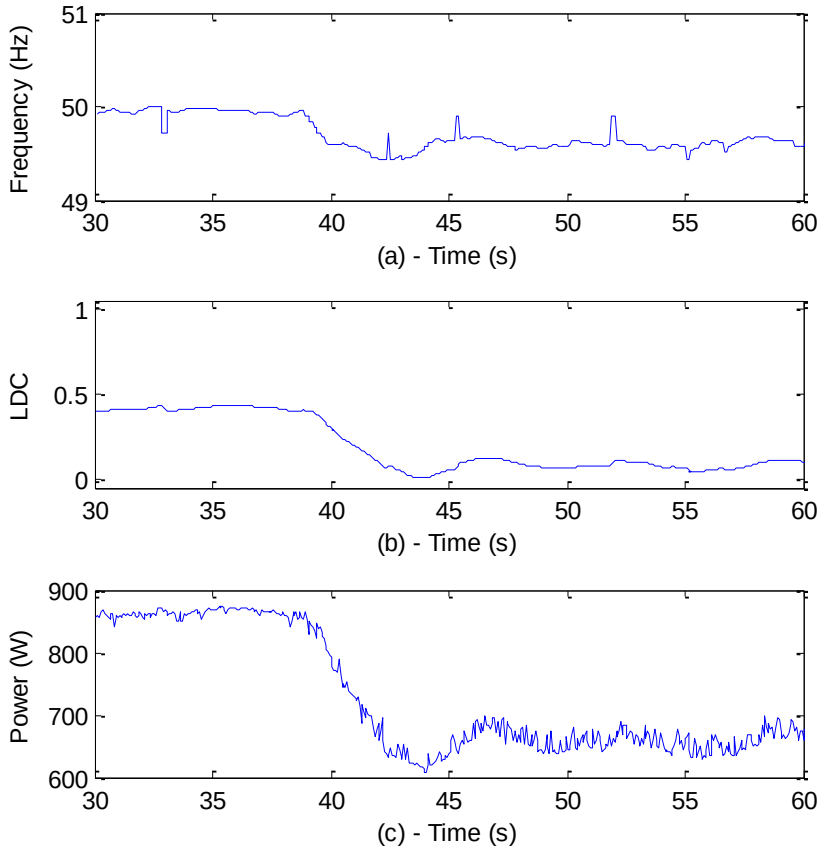

Figure 5.19 – System response during sudden loss of generation showing (a) system frequency, (b) LDC signal, and (c) system power

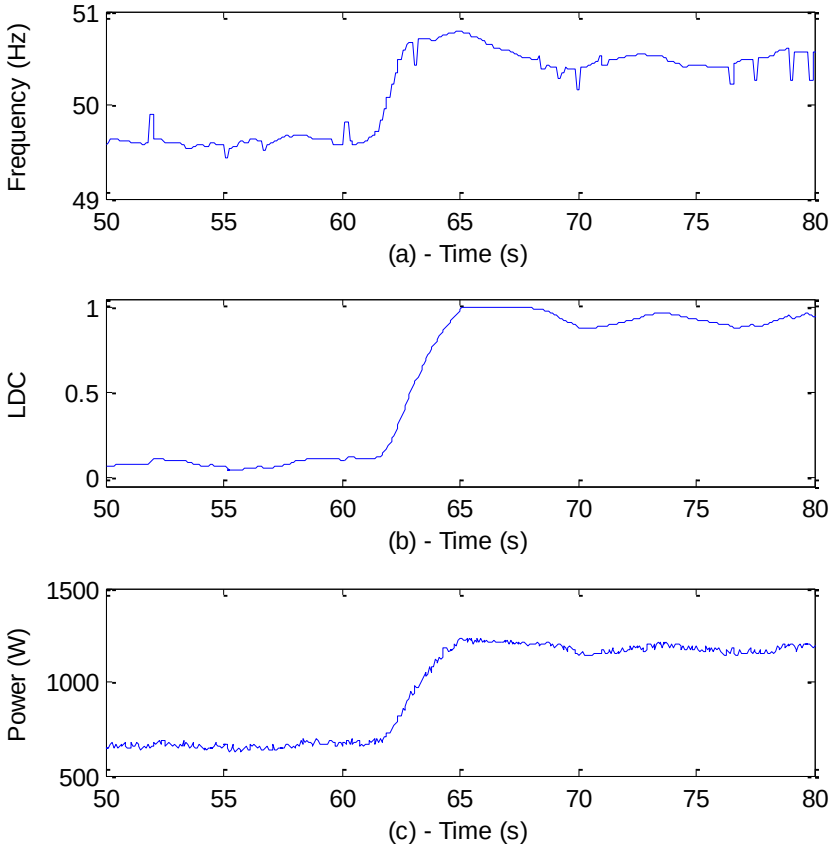

Figure 5.20 – System response with sudden increase in generation showing (a) system frequency, (b) LDC signal, and (c) system power

5.4 Conclusions

In this chapter the experimental performance of the LDC system has been verified. Firstly the response to step changes in power consumption was demonstrated, followed by the effect the controller gain has on power regulation. The performance of the system when combined with system frequency monitoring was then demonstrated both in simulation and with experimental results. It was found that the combined LDC and DDC system could provide very good frequency regulation.

These experimental results agree well with the simulated results shown in Chapter 4, showing that the system dynamics, speed of response and stability are as predicted. This in turn lends some further credibility to the more advanced simulations where twenty houses and around one hundred loads were simulated in Chapter 4. In practice, a full scale LDC system should be able to integrate large variations in renewable energy generation and significant numbers of electric vehicles, while providing distribution transformer protection and ancillary grid services.

Chapter 6

Control Signal Transmission System

6.1 Introduction

As mentioned previously in Chapter 5, the LDC system requires a load control signal to be transmitted to all controllable loads within the micro-grid. This design requires the signal to be unidirectional, of medium resolution (≤ 8-bits) and have very low latency.

To this end, an analogue communication system has been designed in order to simply distribute the LDC control signal around the micro-grid. A system whereby a nominal 800Hz tone is injected at the star point of a local distribution transformer and picked up and filtered at each load has been created. The system runs on the $230V_{RMS}$, 50Hz mains as found in New Zealand. An example of the spectrum seen by any load in the system when the signal injection system is running is shown in Figure 6.1.

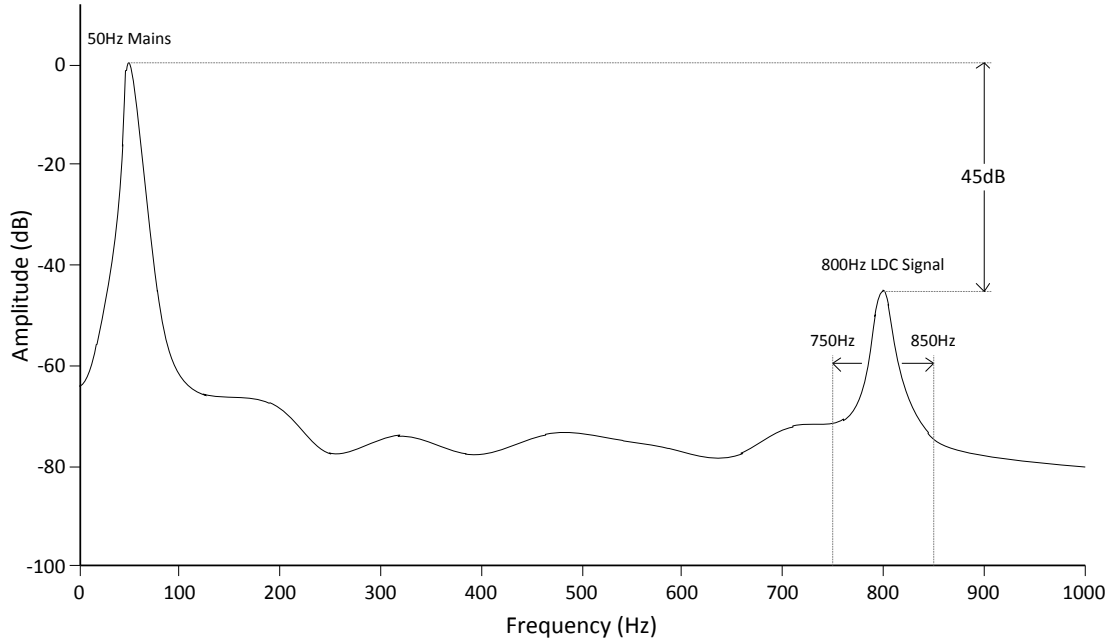

Figure 6.1 – Annotated frequency spectrum showing 50Hz mains and 800Hz LDC signal

There is initially a -45dB signal to noise ratio (SNR) between the 50Hz and 800Hz components, and the two signals are separated by 1.2 decades. In order to represent the maximum and minimum LDC signal values, the tone is varied in frequency – in this case by ±50Hz. If the tone is at 750Hz or below all controllable loads are switched off, if it is at 850Hz

or above, all user loads may be switched on, and between these two extremes loads can be switched in a priority sequence.

It should be stressed that this is only one possible solution to the problem of effectively distributing a load control signal. There are of course many other technologies and variants of the design used here that would provide the desired level of performance. A future where electronic transformers are common within the power system would mean that no extra communication medium is required, and this is discussed in the future work section of Chapter 7. This chapter details the design of the apparatus that distinguishes between the two signals (at 50Hz and 800Hz) and performs the load control function. The goal here is simply to create something that will give realistic results in a laboratory environment, so the LDC system performance can be investigated.

The frequency range of 750Hz - 850Hz was chosen as it is in between the 15^{th} and 17^{th} harmonic of the mains, is far enough away from 50Hz to be filtered and yet is low enough in frequency to still propagate well through standard wiring. Frequencies in the low kHz are quite common for power line communication. The hardware used here could easily be configured to operate at a higher frequency with only software changes, but for a number of practical reasons, an 800Hz center frequency was convenient and simpler to use.

Firstly the microcontroller processing is kept simpler, as a higher frequency requires more calculations per second for any sort of processing or digital filtering. Secondly the filter that is used in the final system is digital, with a maximum sampling rate of around 47kHz. This gives an over-sample ratio of ~60 for 800Hz, and an output waveform that is quite sinusoidal - a higher frequency will create a more distorted waveform. Better clock accuracy would also be required for a higher operating frequency. The current system can work fine using a microcontroller with an internal clock source, but if a 100Hz band is used in the 1-2kHz range, the accuracy requirement is much higher and a crystal clock source may be required.

6.2 Star point injection system

As shown in Figure 6.2, to inject a tone in the 750Hz-850Hz region at the star point of the transformer, an inverter and a transformer are used. The inverter consists of a 3-phase rectifier, DC Bus, H-Bridge and a 300:1 transformer for isolation and some output filtering. The signal is small to the point where it has no effect on electrical loads.

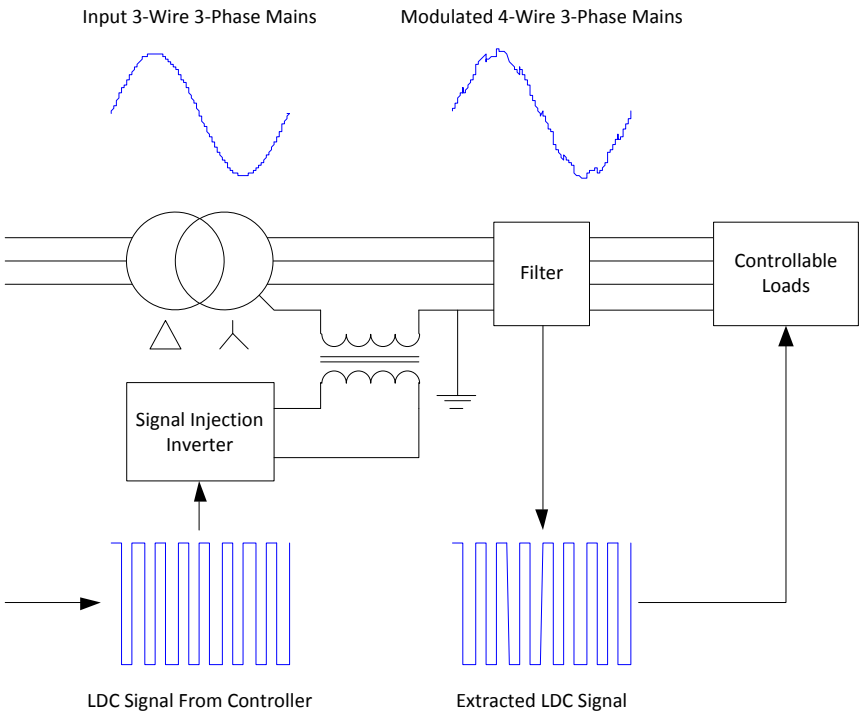

Figure 6.2 – Arrangement of signal transmission system

One side of the injection transformer secondary is connected to the neutral / ground of the micro grid and the other to the star point of the local transformer. In this way the tone can be picked up at any outlet within the system. Since a single turn secondary can be used, the transformer can be realised with a simple clip on transformer such that no wires need cutting. An example of this is shown in Figure 6.3. The core shown could be realised with a 'C' core, or with two halves of a circular core that can be joined together. As the system utilises a transformer with a high turns ratio, the transmission system is tolerant of very high fault currents, as the resulting current seen by the electronics is reduced down by the same ratio.

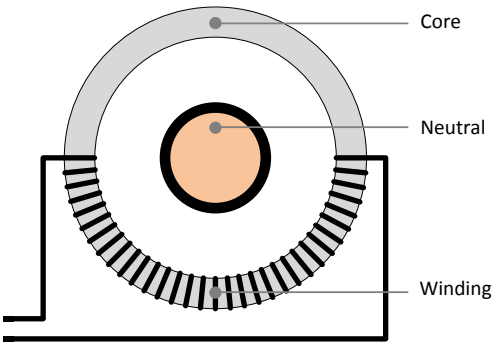

Figure 6.3 – Expanded signal injection transformer

It is important to note that with the Delta-Star transformer used the tone is a common mode and cannot propagate on the delta side of the transformer. Thus all local micro-grids connected to the same grid are independently controlled and there is no leakage from one micro-grid to another.

6.3 Filter Design Criteria

The communication scheme requires that each controllable load has circuitry for filtering out or detecting the 750Hz - 850Hz signal added to the 50Hz mains. Since the signal is used for real-time control, this circuitry needs to be fast acting while still being low cost. As some logic will need to be applied in converting the LDC signal into a control action, the signal needs to be digitised quickly by a microcontroller which will perform the final load control.

The signal is added to the distribution lines by an inverter with bus voltage of ~570V_{DC} and an H-Bridge driving a 300:1 step down transformer, generating a ~1.3V_{RMS} signal. Ideally the injected signal would be a sine wave, but as it is created by an H-Bridge voltage sourced inverter it will have the shape shown in Figure 6.4. Here D represents the duty cycle of the waveform, with D = 1 creating a perfect square wave. The period of the signal is 2L and the amplitude is A. While the waveform may have a measured RMS value of 1.3V, the magnitude of the fundamental harmonic is of more interest, and can be calculated from the Fourier transform of the signal.

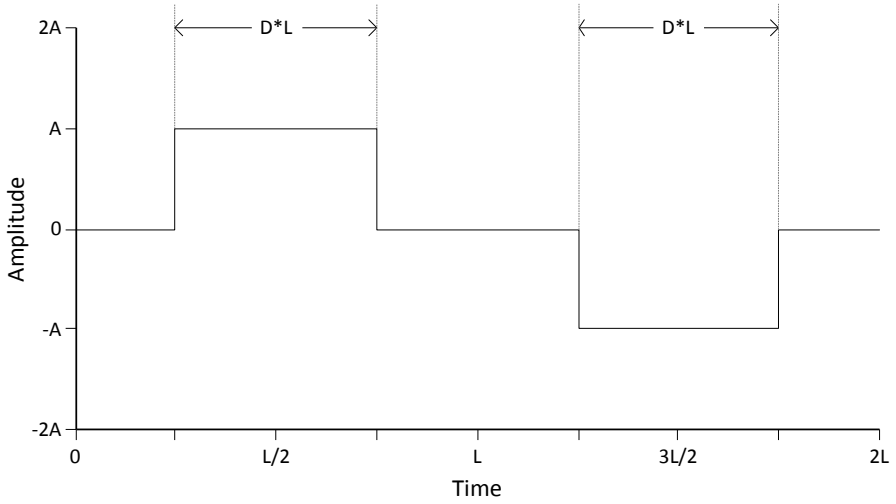

Figure 6.4 – Annotated control signal waveform

Because $f(x)$ is odd, and the integer cycle integral of an odd function times an even function is zero, then $a_n = 0$ as in (6.1). b_n can be calculated as in (4.5).

$$a_n = \frac{1}{L}\int_0^{2L} f(x)\cos\left(\frac{n\pi x}{L}\right)dx = 0 \tag{6.1}$$

$$b_n = \frac{1}{L}\int_0^{2L} f(x)\sin\left(\frac{n\pi x}{L}\right)dx \tag{6.2}$$

Again as $f(x)$ is odd, the integration of both the first and second half of the signal will be equal, so the integration between 0 and L can be taken and doubled. The signal will be equal to A from $\frac{L}{2}(1-D)$ to $\frac{L}{2}(1+D)$and zero elsewhere resulting in (6.3).

$$b_n = \frac{2A}{L}\int_{\frac{L}{2}(1-D)}^{\frac{L}{2}(1+D)} \sin\left(\frac{n\pi x}{L}\right) dx \tag{6.3}$$

$$b_n = \frac{4A}{n\pi}\sin\left(\frac{n\pi}{2}\right)\sin\left(\frac{n\pi D}{2}\right) \tag{6.4}$$

The peak value can be found by substituting $n = 1$, $D = 0.5$, and $A = \frac{570}{300} = 1.90\text{V}$ into (6.4).

$$b_1 = \frac{4(1.9)}{\pi}\sin\left(\frac{\pi}{2}\right)\sin\left(\frac{\pi}{4}\right) = 1.90 * 0.90 = 1.71\text{V}$$

In consequence the RMS value of the fundamental is $1.21\text{V}_{\text{RMS}}$. Given that the system voltage is 230V_{RMS}, this gives an initial SNR of approximately -46dB. Therefore to reliably pick up this LDC signal, the filter needs provide much more than 46dB attenuation of the 50Hz component.

6.4 Chosen Filter Implementation

There are numerous filtering implementations that could be used to provide this level of performance: passive networks, active filters and digital filters were all considered. The chosen implementation uses a combination of a passive filtering network and a digital filter inside a Programmable System on Chip (PSoC) from Cypress Systems. The PSoC contains configurable analogue switched capacitor blocks that are capable of making up to eight two pole filter stages. By combining a number of blocks, 2,4,6 and 8-pole band-pass filters can be created. While the PSoC is slightly more expensive than other microcontrollers that could implement the load control logic, the fact that it can provide high performance filtering stages as well as load control logic means it is a good fit.

The PSoC runs off a 5V supply, and expects signals below this level. An RC high-pass filter is used to step down and bias the input signal by ~ 2.5V_{DC} in order to be accepted by the PSoC. This has the added benefit of attenuating the mains component significantly more than the LDC component.

6.4.1 High-pass Filter Design

The input is first stepped down using a resistor divider network and then passed through a RC filtering network as shown in Figure 6.5 before being fed into the PSoC.

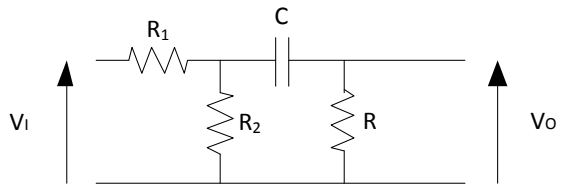

Figure 6.5 – Resistor divider and RC High-pass filter

In order to analyse this circuit, the input resistor divider can be converted to a Thevenin equivalent source as shown in Figure 6.6.

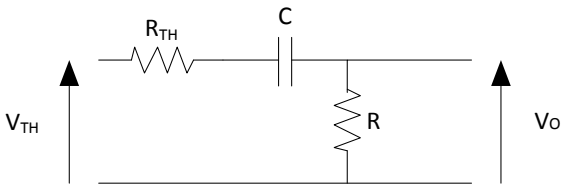

Figure 6.6 – Thevenin equivalent High-pass filter

The parameters R_{TH} and V_{TH} can be calculated as follows:

$$R_{TH} = R_1||R_2 = \frac{R_1 R_2}{R_1 + R_2} \tag{6.5}$$

$$V_{TH} = \frac{V_I R_2}{R_1 + R_2} = V_I \frac{R_{TH}}{R_1} \tag{6.6}$$

The transfer function is therefore:

$$\frac{V_O}{V_{TH}} = \frac{R}{R_{TH} + \frac{1}{sC} + R} = \frac{sR}{s(R_{TH} + R) + \frac{1}{C}} = \frac{R}{R_{TH} + R} \frac{s}{s + \frac{1}{(R_{TH} + R)C}} \tag{6.7}$$

The original voltage, V_I can be substituted back in to get:

$$\frac{V_O}{V_I} = \frac{1}{1 + \frac{R_1}{R} + \frac{R_1}{R_2}} * \frac{s}{s + \frac{1}{(R_{TH} + R)C}} \tag{6.8}$$

This is a standard high-pass filter transfer function, preceded by two linear gain terms. Of note is that the corner frequency is dependent on both R_{TH} and R, which implies that the source impedance has a significant effect on the filtering performance. Since the source impedance is

$R_1||R_2$, and a step-down of roughly 1/10 is desired, $R_1||R_2$ will be only slightly greater than R_2.

There are a number of practical considerations when choosing the values for this network. Firstly, the combination R_1 and R_2 have the potential to consume significant power and therefore must be high in impedance. R_1 will likely be made of two resistors due to both voltage and power rating limits. For cost saving reasons, R_1 and R_2 should step down the voltage enough so that standard voltage capacitors can be used: a target of 1/10 reduces the signal to around 30V. By taking each of these constraints into account, all four values can be optimised for the best filter performance. In addition to these constraints, the filter should have the maximum base gain, maximum attenuation of the 50Hz component and realistic component values.

Firstly by taking into account the capacitor voltage rating, the ratio of R_1 and R_2 can be calculated. The peak mains voltage will be $230\sqrt{2}$ so the following gives the ratio:

$$M = \frac{R_2}{R_1 + R_2} = \frac{30}{230\sqrt{2}} = 0.0922$$

By setting a maximum desired power consumption, the exact values of R_1 and R_2 can be found. Here there is a trade-off between base gain and power consumption, so an acceptable level of 250mW is chosen. This is small enough to not be of significant concern in terms of standing loss, but large enough to not attenuate the signals too severely. In a $230V_{RMS}$ system, the calculation is as follows:

$$P = \frac{230^2}{R_1 + R_2}$$

By substituting in M from above and rearranging we get:

$$R_2 = \frac{230^2}{\left(\frac{1 - M}{M} + 1\right)P} = 19.5\text{k}$$

$$R_1 = \frac{R_2(1 - M)}{M} = 192\text{k}$$

If resistors from the E24 series are chosen by rounding the above values up, then values of $R_1 = 200\text{k}$ and $R_2 = 20\text{k}$ can be used. This still gives a slightly lower power consumption and capacitor voltage than the desired limit, but is acceptable.

From the overall network transfer function derived earlier in (6.7), the corner frequency will be as follows:

$$\omega_0 = \frac{1}{(R_{TH} + R)C}$$

Since the lowest frequency of interest is 750Hz, by placing the corner frequency here there will be maximum attenuation of the 50Hz component and a minimum attenuation of frequencies 750Hz and above. As R_{TH} is already set, sensible values of R and C need to be chosen to give this corner frequency. A larger R and smaller C will give less base attenuation for the whole network, so R is chosen as large as possible while still being predictable in a RC network. By using $R = 120k$ and $C = 1.5nF$, the corner frequency is set at 768Hz, so there will be a maximum attenuation of ~3dB in the 750-850Hz range.

A bode diagram which gives the response of the whole RC network is shown in Figure 6.7, with 50Hz and 750Hz marked.

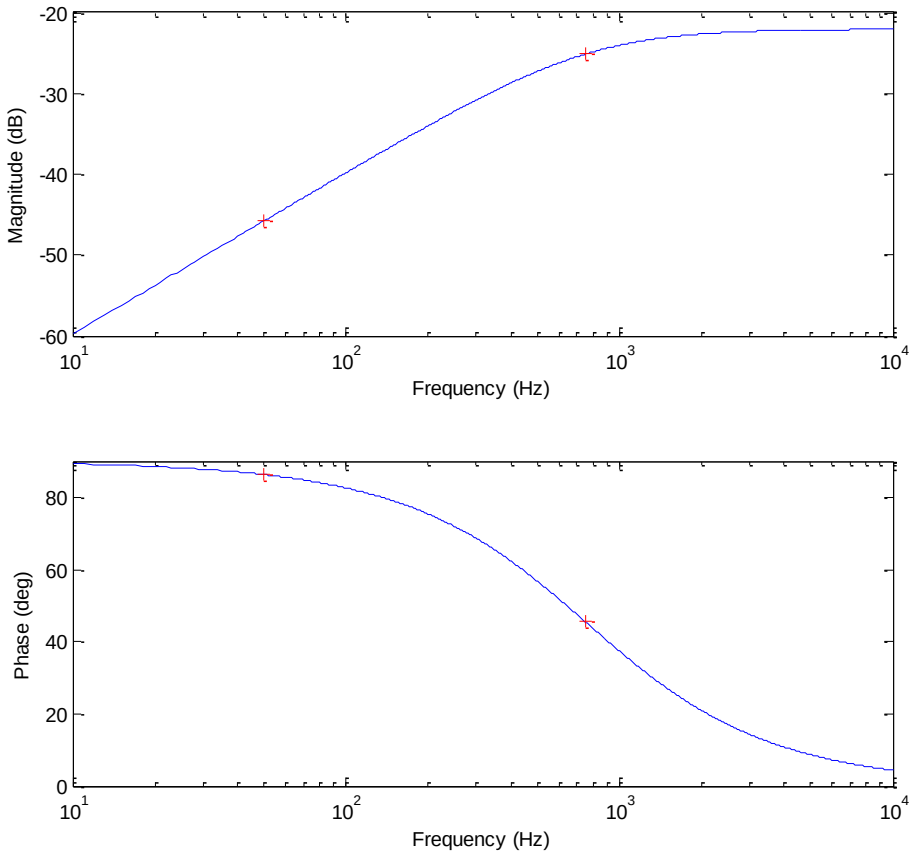

Figure 6.7 – Bode diagram of high-pass filter response with 50Hz and 750Hz marked

The final attenuation is -45.8dB at 50Hz and -25.2dB at 750Hz. This means the 50Hz component is attenuated by 20.6dB more than the 750Hz component. In order for the signal

not to saturate the PSoC input, the 50Hz $230V_{RMS}$ component needs to be less than 5V peak to peak – an attenuation of at least 42.3dB which is indeed provided by the calculated implementation.

6.4.2 Band-pass Filter Design

Following the above pre-processing, a band-pass filter can be created within the PSoC to perform the rest of the filtering and produce an output that can be digitised. Since the signal needs to be used for control in real-time, simple zero-crossing detection is used. While this may not be the most accurate method, it is fast and can be done with the limited resources available within the PSoC.

The band-pass filter is intended to provide the remainder of filtering such that the waveform can be digitised. In order to specify the required filter performance, the relevant spectral components need to be quantified. In addition to the $230V_{RMS}$ 50Hz signal, there are also odd harmonics present in the power waveform (including at 750Hz and 850Hz). The relevant components and their amplitudes before entering the band-pass filter are shown in Table 6.1.

Table 6.1 – Signal amplitudes of interest for BPF design

Component	Initial Amplitude (dBV)	After HPF (dBV)	Normalised (dB)
50Hz	47.2	1.44	0
750Hz	-6.60	-31.8	-33.2
800Hz	1.30	-23.6	-25.0
850Hz	-13.8	-38.4	-39.9

The ideal filter design would be a notch filter, with unity gain between a band such as 755Hz and 845Hz and zero gain outside this band such that harmonics are rejected and the control signal is passed. Obviously this is not possible, but the goal of the filter design and implementation is to get close to this result, and a band-pass filter is a good approximation. Given that the 800Hz component starts with a SNR of -45dB with respect to the 50Hz mains and the mains is then attenuated by 20.6dB by the high-pass filter, it will be at -25dB when entering the band-pass filter. If just a two-pole filter is used, then given the 1.2 decade separation between the two signals, there will be at least -24dB attenuation of the 50Hz component – more if the filter has a high Q. As this will put the 50Hz component and control signal at approximately equal amplitudes, at least a 4-pole filter should be used to effectively differentiate between the two signals. If the control signal SNR can be raised to +24dB with

respect to the 50Hz mains, then the impact of the mains can be ignored and only the harmonics need to be considered.

Since the band around 800Hz has been chosen, the main spectral components of interest are 750Hz and 850Hz. These were measured on the mains signal at the University of Auckland, and while this measurement may not be representative of all mains harmonics, it is likely to be on the high end as the University has a significant electronic load that causes these harmonics. The 750Hz and 850Hz harmonics were measured as listed in Table 6.1. More comprehensive simulations with different levels of harmonics is carried out later in the Section 6.6.

The filter must therefore be designed to attenuate the harmonics as much as possible while still letting the control signal through. Ideally the band over which the control signal can vary should also be as wide as possible. With an 8MHz clock, an 800Hz signal will be 10000 clock cycles. A band of 1024 ($\pm$512) around this point will give a band from 761Hz to 843Hz. Using even powers of two is useful for keeping the mathematics within the microcontroller simple, and for this reason the band from ~760Hz to ~840Hz will be attempted.

The PSoC software provides a graphical user interface which can be used to design the filter. The interface has seven parameters that can be adjusted: center frequency, bandwidth, sampling frequency, filter type, gain, and two digital capacitor values.

With the filter type, center frequency and bandwidth chosen, the sampling frequency, gain and digital capacitor values can then be adjusted for best performance. The sampling frequency was chosen to be the maximum allowed for the band-pass filter block. A sample rate of 47kHz gives an oversampling ratio of around 60, and in consequence a digital analysis of the filter does not need to be performed.

For the filter type, a Butterworth was chosen because a flat in-band phase response is desirable. This is especially true as timed zero crossings will be used for digitisation and changes in phase can distort the reading. A center frequency of f_0 = 810Hz was chosen so as to attenuate the 750Hz component slightly more than the 850Hz component as in practice it was found that the 750Hz caused more interference. A bandwidth of 60Hz was chosen as a good balance between attenuating harmonics and a wide pass-band, and the sample rate was close to the maximum that the filtering system could utilise. The filter could not be implemented without overall gain across the spectrum, but this proves to be useful as all signals have been significantly reduced in amplitude and will need boosting in order to be accurately

digitised. An increase of 25dB means the control signal will be large enough for a comparator with some hysteresis, but not so large as to saturate the filter output.

The value of C2 can now be chosen in order to get the actual filter performance close to the desired parameters. There are only 16 possible values that each can take. The two values were separated slightly so as to get a wider pass band and a steeper descent into the stop band. Each pole pair is effectively a two-pole band-pass filter, with a transfer function of:

$$\frac{V_O}{V_I} = \frac{G\frac{\omega}{Q}s}{s^2 + \frac{\omega}{Q}s + \omega^2} \tag{6.9}$$

Each of the values G, ω, Q are determined by the capacitor values, and the equations that govern the relationships between them can be found in the PSoC documentation. In this implementation, the values are as follows:

First pole pair: $\omega = 2\pi(789.8),\ Q = 20.19,\ G = 6.40$

Second pole pair: $\omega = 2\pi(832.4),\ Q = 19.18,\ G = 5.33$

These two filter stages combine to give the following 4th order transfer function

$$\frac{V_O}{V_I} = \frac{2.289 * 10^6\, s^2}{s^4 + 518.6s^3 + 5.205 * 10^7 s^2 + 1.344 * 10^{10} s + 6.736 * 10^{14}}$$

The expected magnitude and phase response of the filter is shown in Figure 6.8. As desired, there is more attenuation of the 750Hz component than the 850Hz. It can also be observed that the 50Hz component is reduced extensively. An enlarged view of the pass-band is shown in Figure 6.9. Here it can be seen that the band between 760Hz and 840Hz has a suitable amount of gain and this band is suitable for containing the control signal.

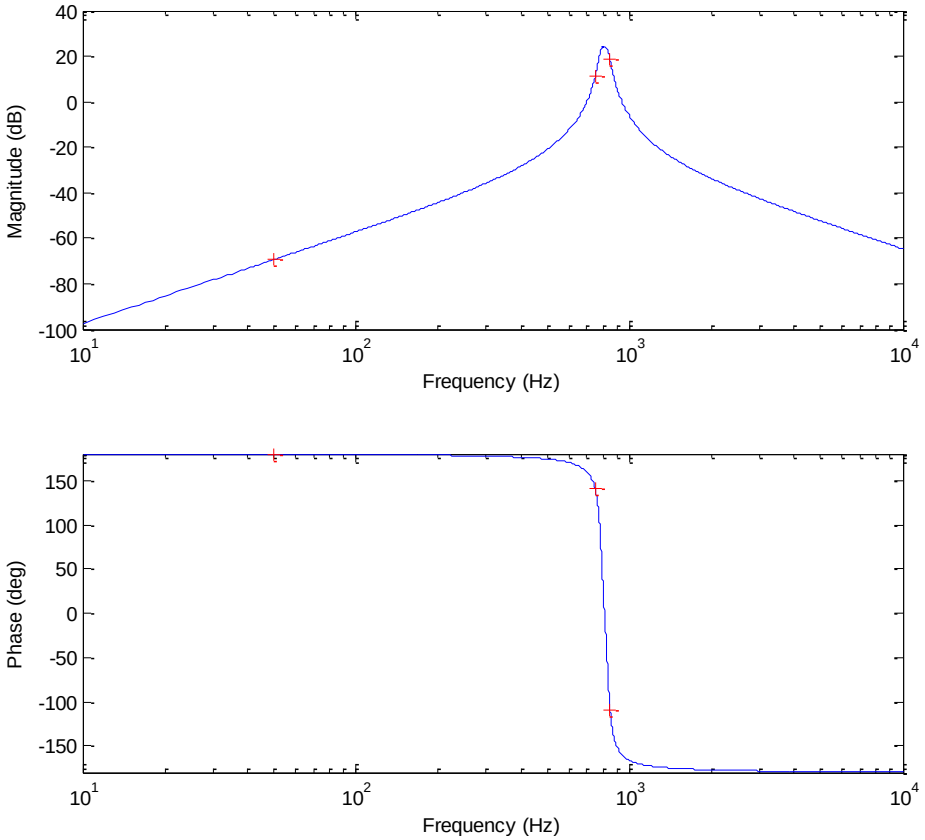

Figure 6.8 – Bode plot of high-pass filter with 50Hz, 750Hz and 850Hz marked

In Table 6.2 the expected signal component strengths at the filter output are shown, with at least 10dB of difference between the desired signal and expected noise. It can also be observed from Figure 6.9 that the signal strength at 760Hz and 840Hz will be 10dB and 3dB less than at 800Hz respectively. With the 750Hz harmonic at -22.5dB, there would still be +12.5dB SNR with the control signal at 760Hz.

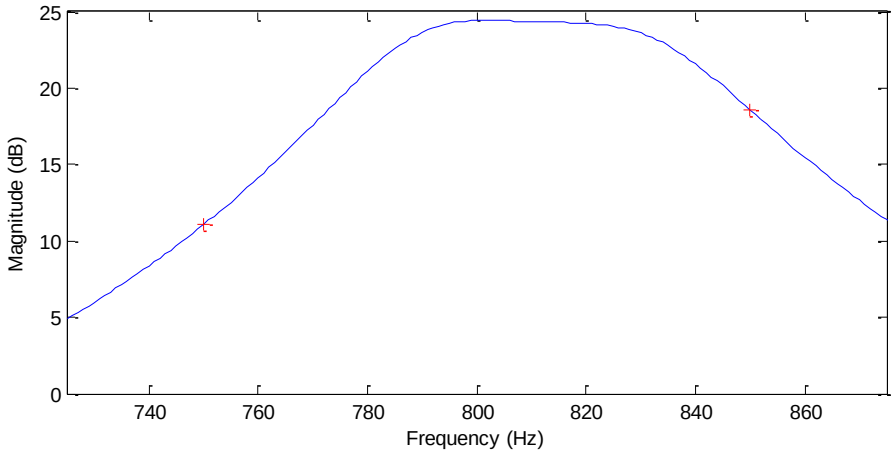

Figure 6.9 – Detail of high-pass filter pass band

These results shown that given the expected levels of noise in the system, the 760Hz to 840Hz band should be usable with the current filter implementation. If more noise is added to the system, the band may have to be reduced or the signal amplitude may need increasing.

Table 6.2 – Signal amplitudes of interest after BPF

Component	Initial Amplitude (dBV)	After HPF (dBV)	After BPF (dBV)	Normalised (dB)
50Hz	+47.2	+1.44	-68.0	-69.8
750Hz	-6.60	-31.8	-20.7	-22.5
800Hz	+2.30	-22.6	+1.81	0.00
850Hz	-13.8	-38.4	-19.9	-21.7

6.5 MATLAB Design Verification

The transfer functions mentioned previously for the high-pass and band-pass filter have been realised in MATLAB/Simulink in order to further verify the system performance. The input signal containing both the 50Hz mains and 800Hz LDC signal is shown below in Figure 6.10.

It can be observed that the 800Hz signal is barely noticeable on the outline of the mains waveform, with small peaks and troughs just visible on close inspection. A frequency spectrum of this input is also shown, and the main signal components are of course the 50Hz mains and the 800Hz control signal. It can be seen that the 800Hz signal is at -46dB with respect to the 50Hz mains.

Note that the FFT plots in this section were created in MATLAB using the standard FFT function. A sampling frequency of 4kHz was used and in each simulation exactly 4s of data was simulated. This means for the 50Hz component there is an integer number of cycles, and for the higher spectral components of interest, there will be many thousands of cycles sampled, which will give an accurate indication of component strength.

The high-pass filter output is shown in Figure 6.11, and the control signal is now at -22dB with respect to the 50Hz component, which is in line with what was calculated previously. The band-pass filter output is shown in Figure 6.12, and it can be seen that the 50Hz component has been attenuated significantly, and the SNR is now -70dB. Again this is exactly the same as the analytical value calculated in the previous section. These first simulations show that the expected filter performance is almost exactly equal to that of the predicted analytical performance.

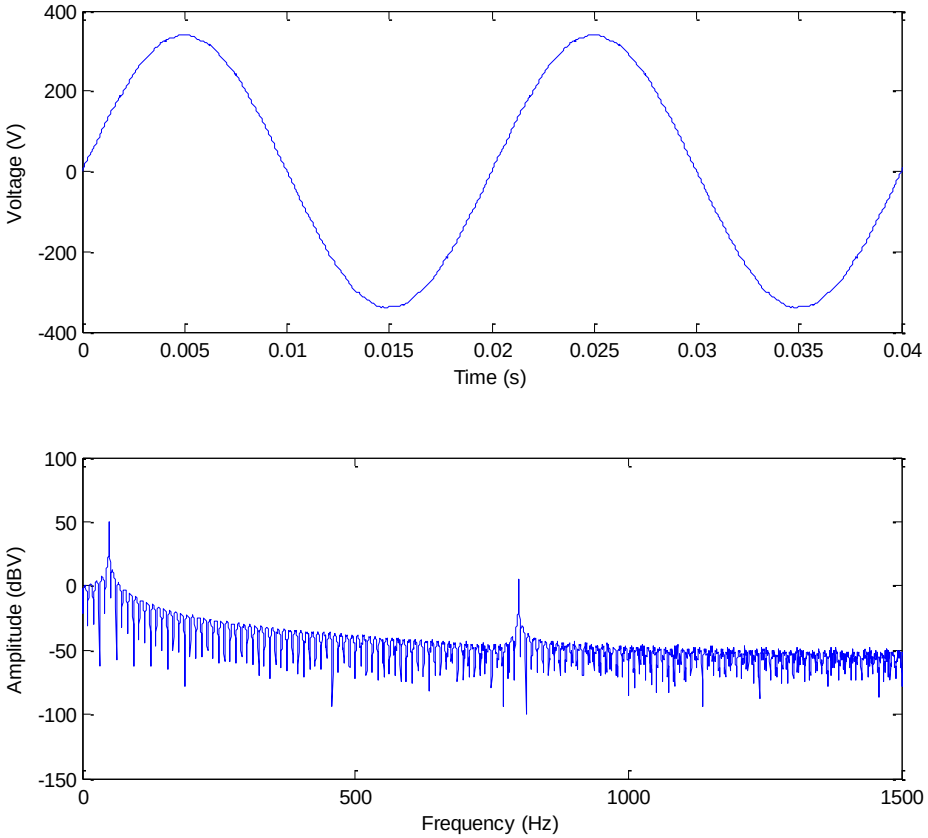

Figure 6.10 – Time and frequency domain plot of mains waveform with superimposed LDC signal

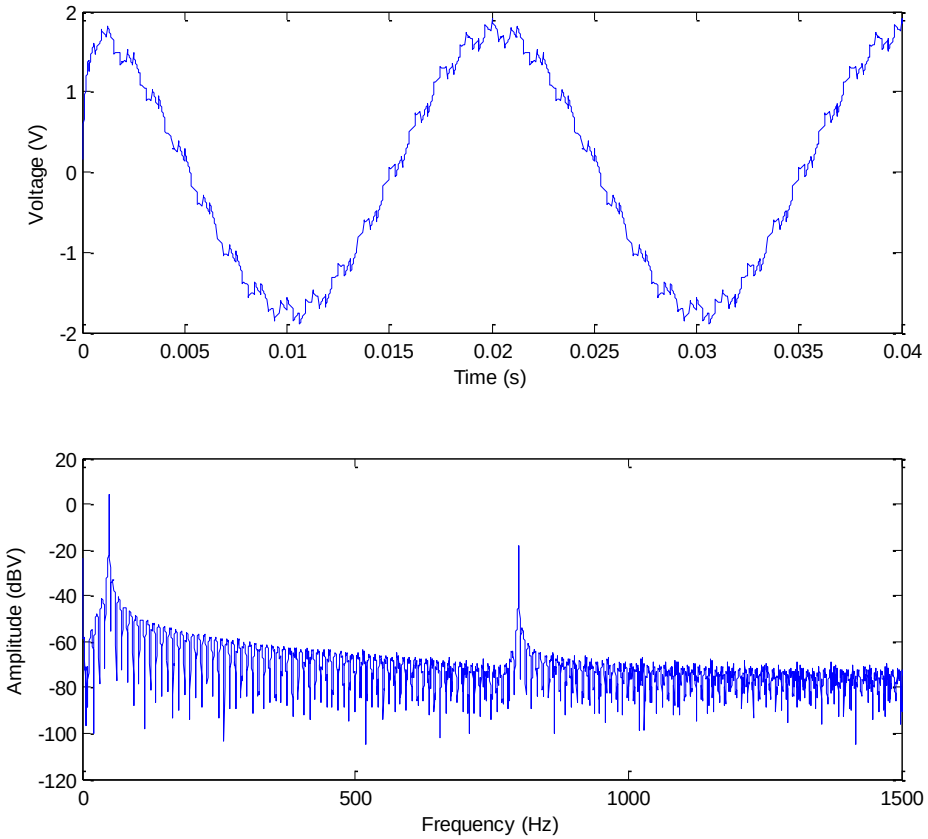

Figure 6.11 – Time and frequency domain plot of high-pass filter output

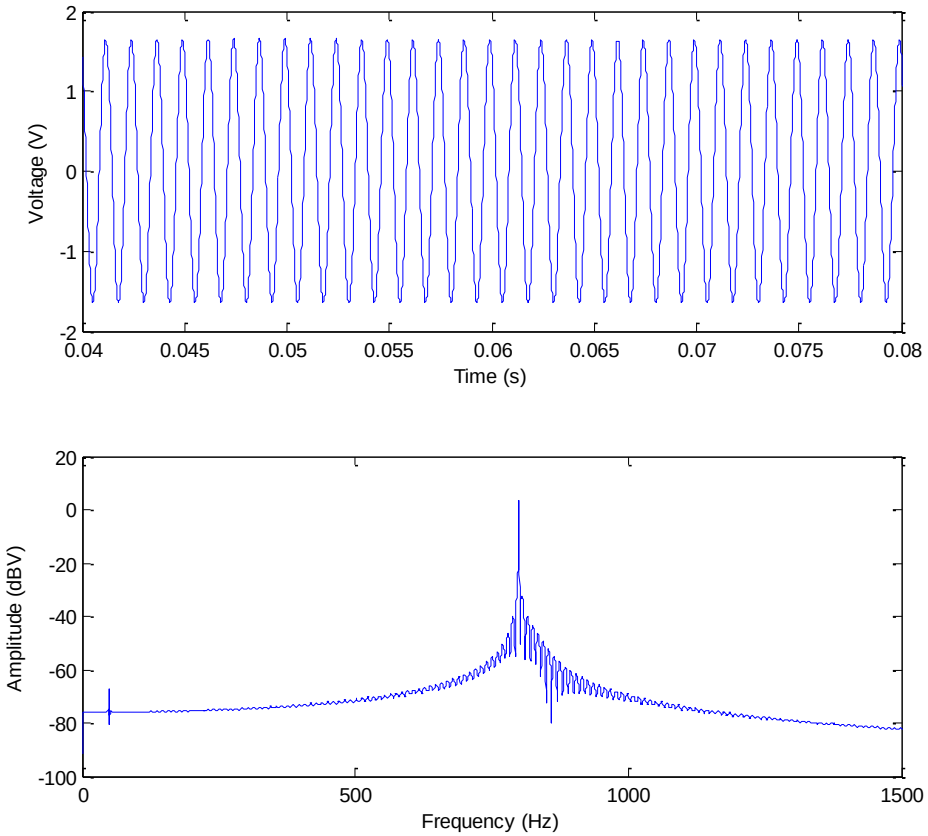

Figure 6.12 – Time and frequency domain plot of band-pass filter output

6.6 MATLAB Noise Rejection Simulation

While demonstrating correct performance in a perfect environment is one thing, determining that the system will work in a non-ideal environment is also important. To this end, the system performance has been measured with various levels of noise added to the system. To generate the expected frequency measurement, componentry similar to that which could be implemented in a microcontroller was simulated. The final band-pass filter output is put through a comparator to create a digital signal which can then be timed and filtered.

6.6.1 Filter performance with expected noise levels

The first noise simulation was created by adding a number of sine and cosine waveforms to the mains waveform for the expected odd and even harmonics present on standard mains. An FFT of the mains was investigated and the harmonics were approximated from 100Hz – 1100Hz (the 2^{nd} – 22^{nd} harmonics). The signal with added noise was then passed through the HPF and BPF filter implementations described earlier, before being put through a software comparator and timer. The input spectrum and band-pass filter output are shown in Figure 6.13. The most important harmonics are at 750Hz and 850Hz as these are closest to the control signal and measure in at -10dBV and -20dBV respectively.

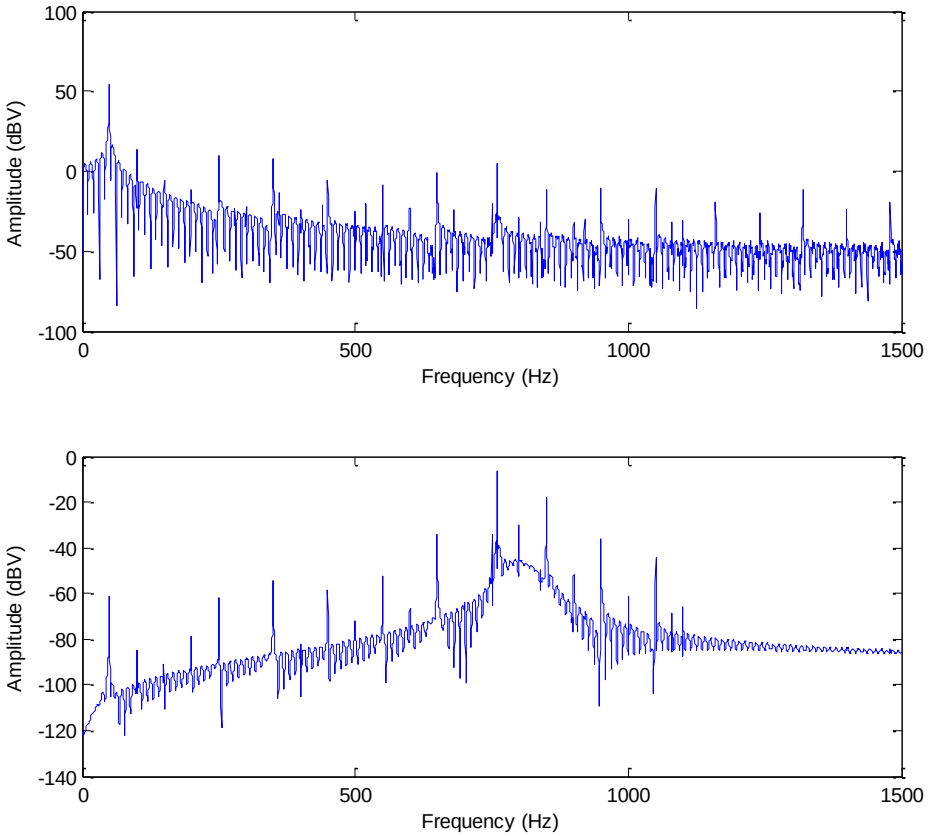

Figure 6.13 – FFT of input mains and band-pass filter output with added noise

From the plot it is clear that at both the input and output, the control signal should be relatively easy to isolate and measure. The simulation was run with the control signal at 760Hz as this is considered the worst case scenario for a number of reasons. Firstly 760Hz and 840Hz are the edges of the desired operating band, and are therefore both on the edges of the filter pass-band and have lower gain than filters in the middle of the band. Secondly the control signal is injected using a standard H-bridge design, but the conduction angle is fixed for simplicity. This means that while the effective duty cycle is set at 50% for the center operating frequency of 800Hz, the duty cycle drops to 47.5% at 760Hz and increases to 52.5% at 840Hz. As a result there is slightly lower signal amplitude at 760Hz than at 840Hz. The high-pass filter also has the effect of attenuating the 750Hz harmonic and 760Hz signal more than the 850Hz harmonic, such that at high noise levels, the 850Hz causes more interference with a 760Hz signal than that of the 750Hz harmonic.

The comparator output produced a series of measurements with a mean of 760.7Hz and a standard deviation of 23.14Hz. Since a comparator is a poor method of digitising frequency signals, software averaging is required within the microcontroller to produce a stable measurement. There is of course a trade-off between how many samples are averaged (increased response time), and how accurate the measured frequency is. The performance of

the software filter with various sample lengths is shown in Table 6.3, as well as signal accuracy which will be discussed next.

Table 6.3 – Software filter mean, standard deviation, delay and accuracy for various sampling lengths

Sample n	Mean (Hz)	Standard Dev. (Hz)	Time Delay (ms)	6-bit Accuracy (%)	7-bit Accuracy (%)
1	760.7	23.14	1.25	3	1
2	760.7	21.08	2.5	3	1
4	760.7	14.41	5	4	2
8	760.7	1.83	10	29	16
16	760.7	1.40	20	36	20
32	760.7	1.24	40	40	22
64	760.7	0.87	80	51	30
128	760.7	0.26	160	91	70
256	760.7	0.16	320	98	86

Here the time delay is approximately equal to the filter time constant (T_1) used to calculate the predicted LDC performance. A delay of the order of 0.1s has been shown to produce acceptable results with the design of the LDC system as discussed in Section 4.7.2. From the table it can be observed that 64 or 128 samples would produce time delays on the order of 0.1s.

As mentioned at the opening of this chapter, a signal accuracy of $\leq$ 8-bits is desired, so a good result would be accurate recreation of a 7-bit signal – a simple simulation is performed at the end of this section showing the impact of signal accuracy on LDC performance. If the frequency range used is 760 to 840Hz, then this is an 80Hz band. For 128 signal levels (7-bits), the measured frequency needs to be accurate to within 0.625Hz. If 128 samples are used, the band of $\pm$0.3125Hz is 1.2 standard deviations from the mean, which results in an estimated 70% of measurements being correct to 7-bits (and 91% accurate to 6-bits). If 256 samples are used, the limit will be 2 standard deviations which results in an estimated 86% of measurements that will be correct to 7-bits (and 98% will be correct to 6-bits).

Comparing these two options, 128 samples will provide a faster response (T_1 = 0.16) but may have slightly more unnecessary load cycling (because readings will vary a little more), and 256 samples will provide a slightly slower response (T_1 = 0.32) but likely with less unnecessary load cycling.

The measurements taken at 760Hz may not provide the best indication of overall transmission system performance. These previous accuracy calculations can be extended to find the optimal band for best control signal accuracy with various signal resolutions and

sampling lengths. Widening the band results in a lower accuracy requirement, but also means measurements will be made closer to the harmonics which will in turn be less accurate. Sampling lengths of 64, 128 and 256 were tested with signal resolutions of 6-bits, 7-bits and 8-bits to determine expected system performance, with the measurement accuracy shown in Table 6.4.

Table 6.4 – Measurement accuracy with various resolution and sampling lengths

Bits	Signal Levels	64 Samples (%)	128 Samples (%)	256 Samples (%)
8	256	42	66	90
7	128	67	89	99
6	64	89	99	100

The optimal frequency band for the majority of these results was 772.5Hz to 830Hz, showing that performance at 760Hz was not indicative of overall system performance.

What level of signal accuracy is required can be determined from the LDC simulations. The exact same system simulated in Section 4.7 has been modified with a discrete LDC signal, and the regulating performance for different signal resolutions and filter sample lengths is shown in Table 6.5 for a rough idea of the performance reduction. A slight reduction in wind was used to allow the LDC system to properly regulate input power such that the standard deviation is a meaningful measure of performance. A 10-bit signal and 64 samples gives approximately the same level of performance as a perfect communication medium. If 7-bits are used, the regulating performance is 6% worse, and if 6-bits are used the performance is 22% worse. It can also be observed that increased filter sample length has a lesser impact on regulating performance compared to reduced signal accuracy.

Table 6.5 – Standard deviation of input power with reduced signal accuracy and increased sample length

Bits	Signal Levels	64 Samples (W)	128 Samples (W)	256 Samples (W)
10	1024	492	522	604
8	256	494	517	634
7	128	521	549	653
6	64	598	588	675

From these two tables it can be predicted that the performance reduction when changing from 10-bit signal resolution and small filter length to 256 samples and a 7-bit signal is only 30%, and that this performance is easily achievable (99% accuracy) with this level of signal noise.

6.6.2 Filter performance with maximum possible noise

In order to thoroughly investigate the system performance in the presence of noise, an even noisier signal was created. In this instance the maximum allowed harmonic distortion for power systems in New Zealand was added to the signal. In NZECP 36:1993 [110] the maximum level of harmonics allowed at any point of common coupling is specified. The levels for each harmonic range from 2.3% to 0.2% of the nominal voltage for harmonics numbered 2-50. The harmonics of most interest here are the 15^{th}, 16^{th} and 17^{th} harmonic, which have a limit of 0.5%, 0.2% and 0.4% respectively. The result of this added noise is shown in Figure 6.14 with the same control signal used previously at 760Hz. It can be seen that there are now a number of signals with similar or greater amplitude than that of the 760Hz component. Indeed when an attempt at digitising the signal is made, the result varies widely but most commonly tends to 800Hz.

In order to determine what the usable system bandwidth now is, a sweep was performed where the control signal was varied and comparator output recorded in 2.5Hz increments over the frequency range 750Hz to 850Hz. At each frequency increment, 4s of signal was simulated, generating around 3,200 samples. The resulting mean error and standard deviation of the measurements is shown in Figure 6.15. The mean error is calculated by subtracting the input control signal frequency from the mean measured output frequency. There are two traces on each plot, the first is with the standard Butterworth implementation used previously but centered at 800Hz, and the second is with a 1.0dB Chebyshev implementation. The mean gives an indication of what the usable range of the filter is, and the standard deviation gives an indication of how many samples are required for a stable measurement. From the error plot it can be observed that the usable band in the presence of this worst case noise is from around 773Hz to 833Hz for the Butterworth, and from 770Hz to 835Hz for the Chebyshev. As expected the variation in measurement using the Chebyshev is improved compared to that of the Butterworth.

A similar calculation as done previously can be performed to again find the optimal band for best control signal accuracy with various signal amplitudes and sampling lengths. For example in Figure 6.15 the optimal band is 777.5-837.5Hz, which with 512 samples produces a 128 level (7-bit) signal accuracy of 66%. The overall result with this noise level is quite poor, so the system was simulated with three levels of increased control signal amplitude, which correspond to a 300:1, 200:1 and 100:1 transformer step down. The measured accuracy for 7-

bit signal resolution is shown in Table 6.6 with the three different amplitude levels and three different sampling lengths.

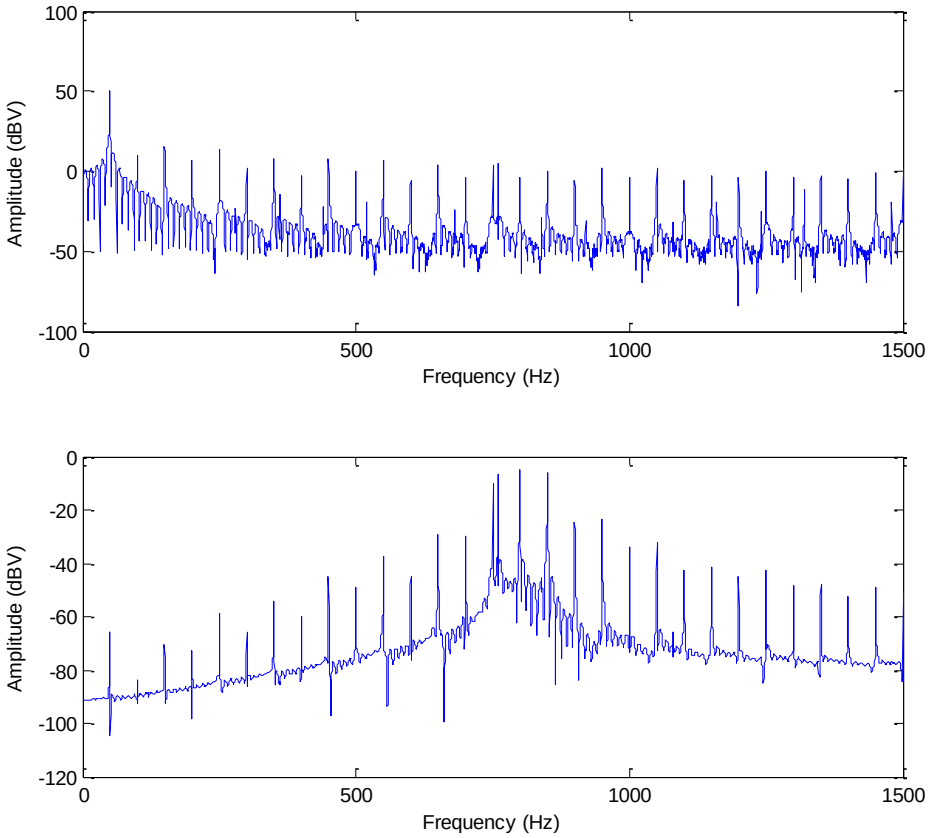

Figure 6.14 – FFT of mains waveform and band-pass filter output with maximum allowable noise

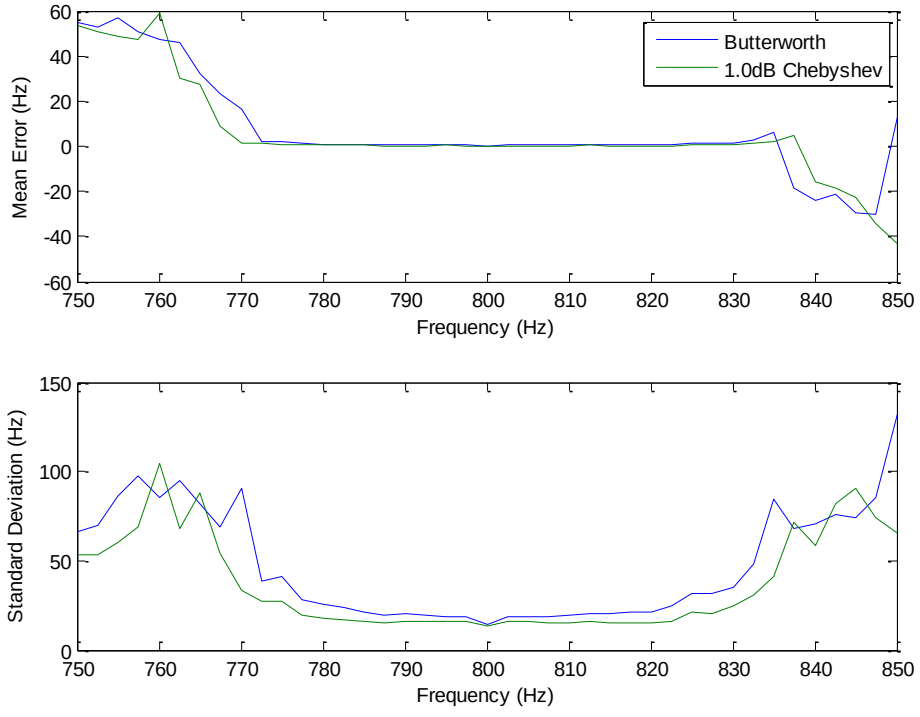

Figure 6.15 – Error and spread in frequency measurement with respect to frequency

Table 6.6 – 7-bit measurement accuracy with varied signal amplitudes and sample lengths

Amplitude (V_{RMS})	128 Samples (%)	256 Samples (%)	512 Samples (%)
1.71 (300:1)	8%	45%	66%
2.56 (200:1)	24%	64%	83%
5.13 (100:1)	61%	86%	98%

From the table it can be observed that with a 200:1 transformer and 512 samples, an accuracy of 83% would be obtained, and 6-bit accuracy was calculated separately at 97%. This result implies that in environments with very high levels of noise, an increase in signal amplitude or reduction in accuracy would be required for reliable signal reception.

6.7 Filter Performance in LDC System

A practical system has been tested within a laboratory scale micro grid. The signal injection setup and filter implementation is as described previously in Sections 6.2 and 6.4. A 300:1 turn injection transformer was used, with the injection inverter running off the same voltage as the rest of the system. The full implementation of the LDC system is covered in more detail in Chapter 5.

The input to the filtering stage with an additional FFT plot overlaid is shown in Figure 6.16, running with the signal injection system on at 800Hz. Here there is 46.8dBV of 50Hz, and 0.199dBV of 800Hz, which equates to $220V_{RMS}$ of 50Hz and $1.02V_{RMS}$ of 800Hz. The numbers from the FFT are slightly lower than those indicated by simply taking the RMS value of the waveform, as these were measured separately at $230V_{RMS}$ and $1.30V_{RMS}$ respectively. For the control signal, this is almost exactly as expected, as the calculated value of the 800Hz harmonic in the control signal is $1.11V_{RMS}$.

A plot of the filter output with the control signal turned off is shown in Figure 6.17. There is a noise signal with maximum amplitude of around 1V peak-to-peak. This determines the noise floor, and the comparator is set to have a hysteresis band just slightly larger than this noise signal, such that noise will be rejected and only legitimate signals will be picked up.

The output of the high-pass filter is shown in Figure 6.18, again with the added FFT. The 50Hz component is now at 1.25dBV, and the 800Hz signal is at -26.8dBV, which compares well to the 1.44dBV and -23.4dBV predicted in Table 6.1. The signal is also contained between 0V and 5V and can therefore be routed directly to the PSoC inputs. The time domain signal is also very similar to the simulated output shown previously in Figure 6.11.

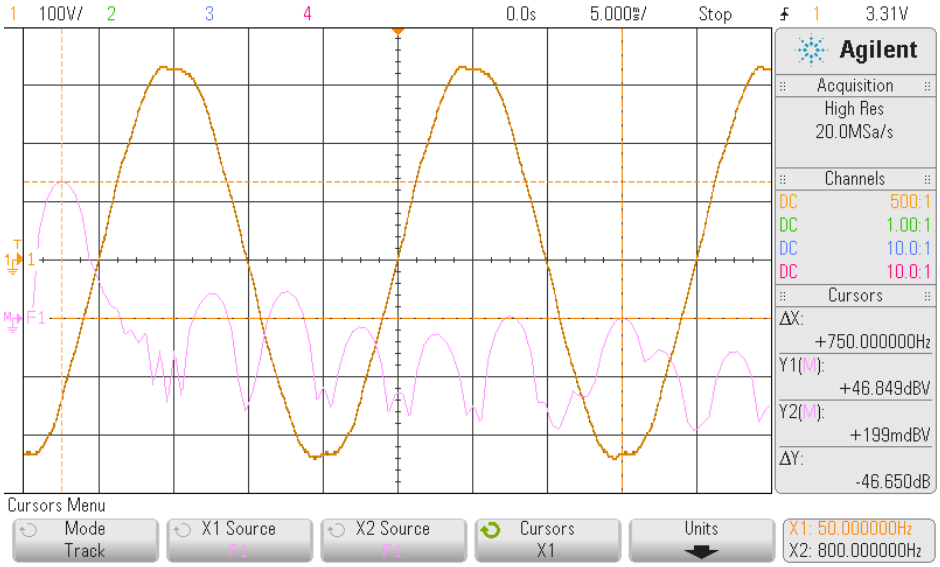

Figure 6.16 – Mains waveform with control signal at 800Hz

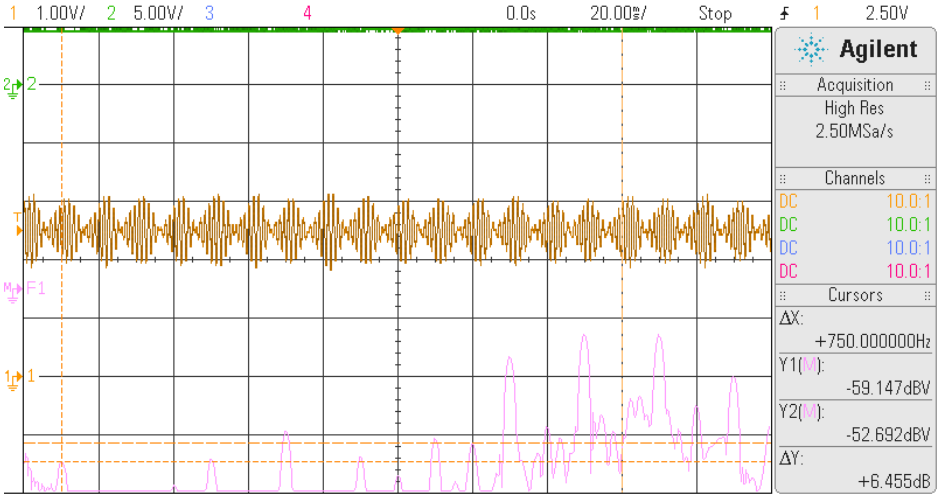

Figure 6.17 – Output of band-pass filter and comparator with no control signal

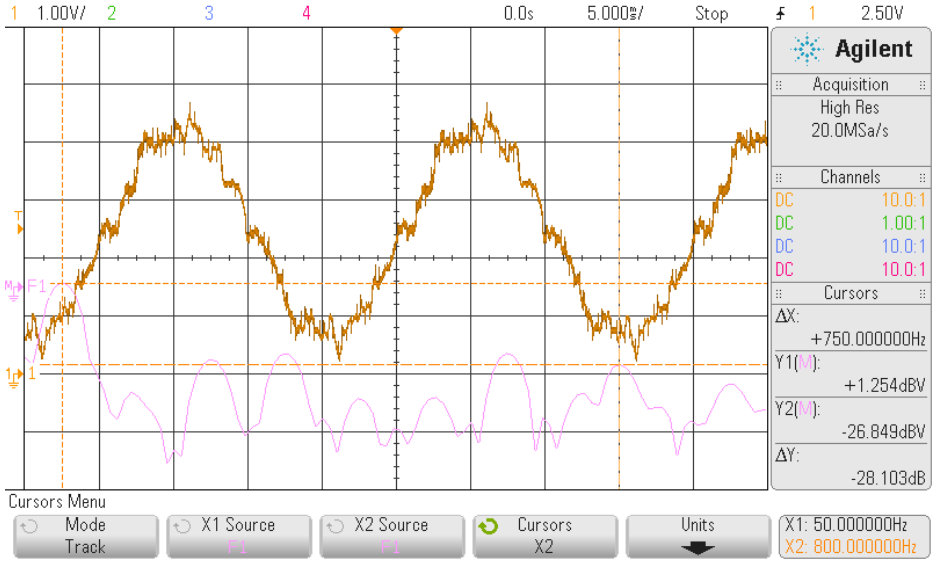

Figure 6.18 – Output of high-pass filter with control signal at 800Hz

The output of the band-pass filter is shown in Figure 6.19, along with an FFT of the signal and the output from a comparator block used for digitisation. A number of features can be observed in these waveforms. Firstly the amplitude of the output is quite varied, likely due to interactions with other spectral components. At times the signal drops below the threshold of the comparator, causing the signal to disappear – the implications of this on LDC system performance will be discussed next. It can also be observed from the FFT that the harmonic components are around 10dB lower in amplitude than the control signal.

A similar plot is shown in Figure 6.20 with the control signal at 761Hz. Again the signal varies in amplitude, and gaps in the comparator output are slightly more common. When compared to the signal at 844Hz shown in Figure 6.21, it can be seen that the 844Hz signal has slightly higher overall amplitude and less gaps in comparator output, as expected. The reason for this difference has been mentioned earlier: the filter center frequency and operating band was moved up slightly in frequency.

Overall the performance of the laboratory system matches that predicted by the analytical and simulation results, with similar imperfections in the measured signal. There are a number of software filtering features inside the load controller that reduce the impact of signal gaps on the operation of the LDC system. Firstly, the controller is configured to outright reject any measurement that is more than 10% outside the expected range. Additionally, when no signal is received, the controller continues to use the previous measurement, so no sudden jump in controllable load will occur if the LDC signal is removed or not received for a period of time. Ideally, the controller might slowly move the load to the middle signal level if nothing is received for a long enough period of time, but this was not needed in these tests.

Also as mentioned previously, a number of samples are taken to determine the actual control signal – usually 128 or 256. This means even if an erroneous measurement is taken that is far from the correct measurement (but within the usable range), it will have minimal effect on the final determined value.

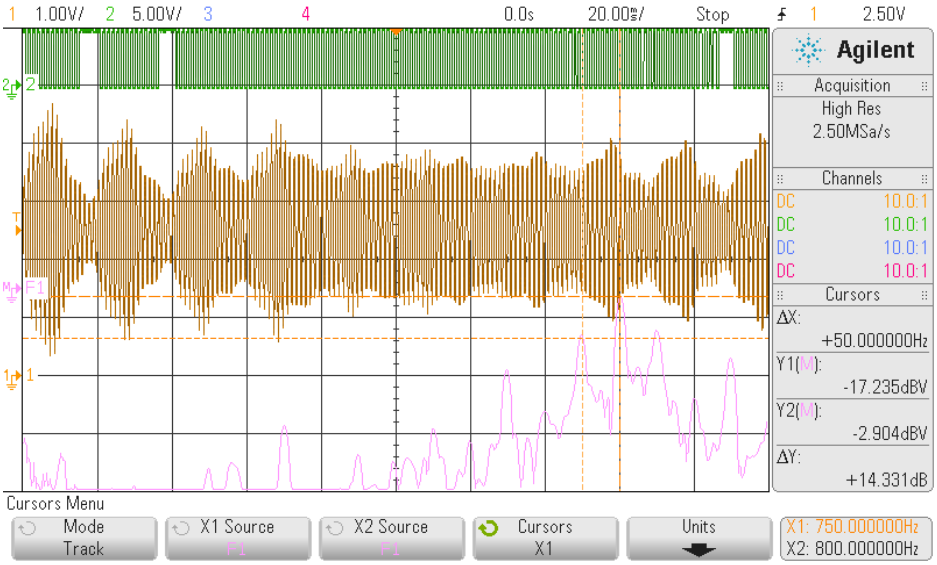

Figure 6.19 – Output of band-pass filter and comparator with control signal at 800Hz

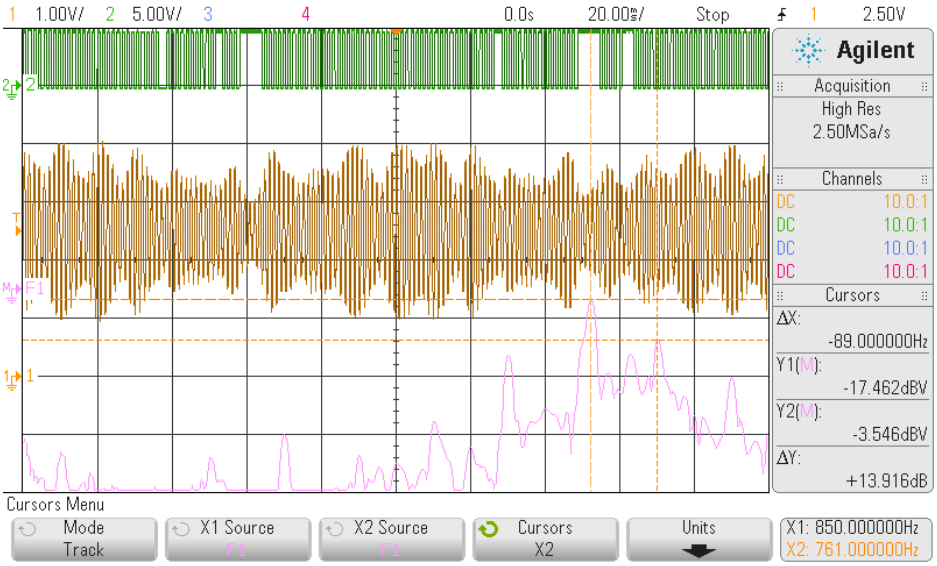

Figure 6.20 – Output of band-pass filter and comparator with control signal at 761Hz

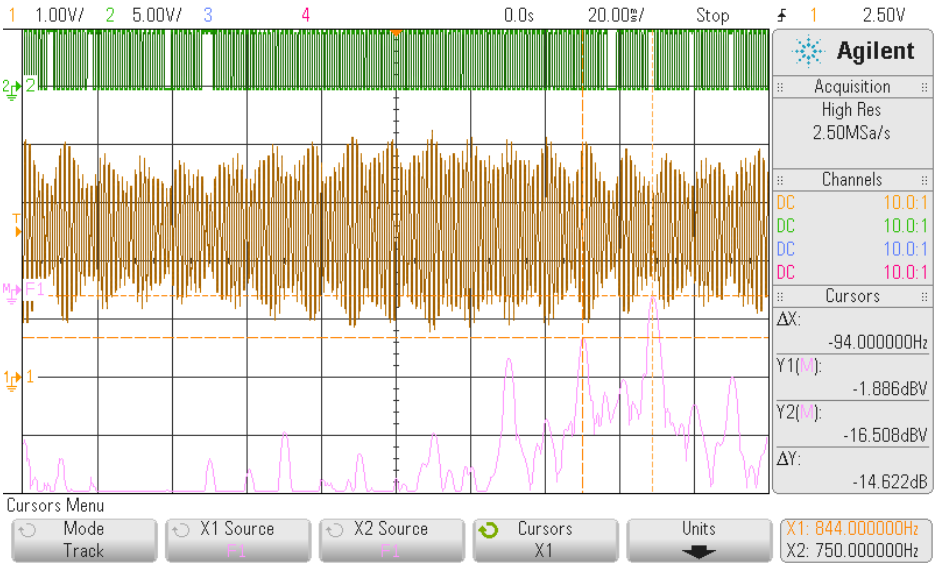

Figure 6.21 – Output of band-pass filter and comparator with control signal at 844Hz

Gaps in the received signal could also affect the timeliness of load response, but this will likely be minimal as there are a large number of measurements available per second. In Figure 6.20 with the control signal at 761Hz there are 13 missing measurements over a 200ms time period. This represents 8.5% of the expected 152 measurements made in this time. If the filter length is 256, then even if 10% of all measurements were missed, the filter time constant would only change from 0.32 to 0.352. It should also be considered that the result at 761Hz is right on the limit of correct filter operation, so the likely impact is much less than this. Indeed, the usable band was found to start closer to 772.5Hz in simulation.

6.8 System Performance Discussion

These results show that within the filtering band, the LDC signal can be picked up reliably and used for local demand control in the presence of normal noise levels. From the simulation with maximum allowed levels of noise, it was shown that under these circumstances some changes would be required for reliable operation. Such changes could include increasing the amplitude of the control signal, further optimizing the filter or moving to a different frequency where there is less harmonic content.

If normal levels of noise are planned for, there will be a small reduction in LDC performance compared to that of a perfect communication system. In one simulation this was calculated to be around a 30% reduction. The power regulation of the LDC system is already very good, so a 30% performance reduction can be managed. Under most noise conditions it was found that at least 256 measurement samples are required for acceptable signal accuracy. While this only has a small impact on overall regulating performance, it will have a larger impact on response time. 256 measurements gives a filter time constant of $T_1 = 0.32$, which in a critically damped system will give a settling time of more than 1s (at around 2s). This means performance when responding to some emergency signals or very short surges in wind turbine or solar output is reduced.

In addition to noise tolerance, the system also has good fault tolerance. Because of the large ratio transformer, high fault currents in the neutral wire are not an issue. Furthermore, the communication system can also be used to signal grid faults. Anti-islanding hardware in distributed generation is generally quite expensive, but if this detection was included in the LDC system, the signal (or absence of) passed to distributed generation could be used to stop the exporting of power and livening of lines during a fault in a potentially more cost effective fashion.

The goal of this communication system was to create a system where the real performance of LDC could be assessed in a laboratory environment. Many other designs and technologies could potentially provide this level of performance, and the design presented here is just one option.

In particular an increase in operating frequency could be investigated for better performance, as more samples per second and less harmonic noise is likely to be beneficial as long as the signal can still propagate throughout the LV network effectively.

Chapter 7
Suggestions for Future Work

7.1 Introduction

There are a number of topics that have been investigated as part of this research, but which still require further work or are beyond the scope of this thesis. These include work on practical smart loads, a high-frequency IPT supply, and electronic transformers. In addition, there are some interesting studies that could be done on the impacts of deploying LDC or building on it further. Initial discussion and results relating to these topics will be presented here, with comments on how these works fit with the other ideas in this thesis, and where further work might be undertaken in the future.

In Section 7.2, a number of practical considerations that need to be taken into account when creating smart loads and deploying DSM systems, are presented. Much of this is in relation to user acceptance and integration of DSM functionality into smart loads. These considerations provide the rationale for the functionality and design goals of the smart appliance interface presented in Section 7.3. Recommendations for both the capabilities and the physical design of a smart appliance interface are presented, with the comment made that many options could be used for the physical design, as long as the same essential functionality is provided.

In Section 7.4, the preliminary design of a high-frequency IPT power supply is presented. Ideally a full IPT system would have been completed, with a prototype smart appliance interface built into the system. Only the power electronics stage was completed, using new prototype IGBTs from Infineon that are designed for low switching losses, and which may allow the IPT system to run at a higher operating frequency. There are many potential advantages when increasing the operating frequency, and these are also discussed in this section. Although a full smart appliance interface has not been created, one was simulated by using an analogue voltage to represent the DSM signal, and the power regulation of the supply in response to the signal is presented.

In Section 7.5, some initial simulations on electronic transformers are shown. Electronic transformers are a very new technology, and mostly exist as theoretical models and research prototypes. Though still very new, the implications of electronic transformers for DSM systems are large. Electronic transformers are looking to become economical because power electronic

devices are improving at such a rapid pace. This is a similar reason as to why high power IPT systems for EV charging have moved to much higher frequencies than earlier thought possible, and in so doing reduced the volume of material required. In Chapter 6, one design for a DSM signal transmission system was given, where the goal was to distribute the signal quickly to the entire LV network. An electronic transformer uses power electronics stages to change between the different voltage levels in a power system, and it is assumed that the output frequency does not need to match the input frequency at all times. This feature means the line frequency could be adjusted in a DDC like manner to perform this task and no additional transmission system would be needed. In addition, it is also assumed that on each phase, the voltage can be controlled independently of the other phases and input voltage. This could be used for per-phase load signalling and phase balancing. Some simple simulations of these features showing line voltage regulation when there is significant imbalance between the phases are presented.

7.2 Practical Smart Loads

As was introduced in Chapter 2, the concept of DSM involves utilising controllable loads for the benefit of power system operation. In order for loads to be controllable, a number of factors must be considered. First the load must be deferrable – either it is storing energy or performing an energy intensive task, with flexibility as to when the energy is consumed. Secondly, the loads must be capable of receiving a DSM signal, and as discussed in Chapters 2 and 4, there are a number of methods that can be used to achieve this. Thirdly, the operating constraints of the load must be considered with respect to device limits, user comfort and scheduling. Consideration must also be given to remuneration or incentives passed on to the end user, who in many cases will be the ones paying for and affected by smart loads. These topics will now be developed in more detail.

7.2.1 Communication

As discussed in Chapter 2, there are many communication links in a smart grid, such as those connecting utilities, substations, distribution transformers, households and smart loads. Here, technologies for providing the communication link that terminates at the smart load will be discussed.

As covered in Chapter 5, there are two general ways in which the final load control can be performed, a device could be placed inline with the load's power supply, to simply turn the power supply on and off in response to DSM signals. The alternative is that the load be capable of receiving and responding to DSM signals directly.

Obviously the first method is far easier to retrofit to existing devices, but provides much lower performance when compared to direct DSM integration. Cutting off power to a device could cause it to lose state information or even cause damage in some cases. In addition, no information about the operation of the device can be fed back into the DSM system.

Integrating the DSM receiver directly into the device may be more costly, but would provide much higher performance as device operation can be considered when responding to signals, and more precise load control could be performed. Direct integration requires extra circuitry with added complexity and space, but if a purpose built chip is used, the overall cost may in fact be lower. There is also the issue of standardisation, given limiting the capability of the smart load to only one particular communication technology would preclude other technologies that could arise in the future. This is particularly relevant given smart grid technology is still new and undergoing constant development.

A third option could also be considered, where the DSM functionality is built into the device only for the so called "last meter" smart grid communication. With this method, an interface is built into the device where the goal is to provide communication between the smart load and an intermediate device for the lowest possible cost. The interface could utilise power line communication or an additional wired interface. The advantage of this option is that if a suitable and low cost last meter protocol can be agreed upon, end users can then use any advanced DSM technology within the home, and replace this technology without replacing or upgrading the firmware of their appliances.

In [106] one option for providing last meter communication is presented, where a very simple power line communication scheme is used to communicate between an inline device and the smart appliance. The scheme is designed with low cost as the primary goal. The data rate from the inline device to the smart appliance is relatively low at around 200Bps, but the data rate from the appliance to the inline device is even lower – at around 50Bps. The data rate to the appliance is potentially not quite fast enough for the performance levels discussed in this thesis. A preliminary design for a simple wired smart appliance interface is presented in Section 7.3.

7.2.2 User Comfort and Device Limits

The operating constraints of the smart load are also an important consideration when performing load control. Temperature based appliances have limits relating to user comfort and safety, as space heating is generally set up to keep the temperature within a set range desired

by the end user, a fridge must keep food below a certain limit for food safety, and hot water must be kept above a certain temperature to avoid bacteria growth, but not so high as to be a hazard.

Generally, a DSM system will always keep devices within set operating limits, but depending on the incentives provided to the end user, they may be willing to sacrifice some degree of comfort in exchange for better incentives. Even with these limits on device operation, as discussed previously there is still a large amount of flexibility that can be exploited with DSM. All the simulations performed throughout this thesis had hard device operating limits that could not be overridden by the DSM system.

Further to operating limits, there may also be scheduling limits placed on smart loads. For example, the limits on an EV charger may include a condition that there must always be a minimum level of charge within the vehicle, but also that it must be near fully charged every morning. Scheduling could also be put in place for better operation of heating loads. For example, hot water cylinders could be scheduled to be at the maximum allowed temperature before a known period of peak power system load and high consumption. If scheduling information can be communicated to a higher level power manager such as is simulated in Chapter 4, more effective DSM can be performed as was demonstrated there. For this reason, being able to obtain load status and operating information is an important feature that a smart grid communications scheme must provide.

7.2.3 Incentives

No DSM scheme will be implemented without a financial incentive to do so. Adding extra hardware to appliances adds cost and complexity in an industry with very narrow margins. As such, a user should either receive a rebate when purchasing a smart appliance, or receive a discount on electricity when the smart appliance is operating.

Under circumstances when the end user is receiving an incentive for operating the smart appliance, it must be assured that the appliance is actually operating as desired by the system or line operator. With LDC or DDC, one solution would be to have the power meter record how much power is consumed at various ranges of the LDC signal or system frequency, and the user could be rewarded accordingly.

User acceptance and incentives are beyond the scope of this research, but nevertheless are very important to consider, as there have been many examples of good DSM systems put in

place that ended up providing a poor user experience or miscommunicated benefits. An example of this is the backlash against smart meters [111], which likely have no impact on the end user yet due to miscommunication can face strong opposition.

7.2.4 Summary

There are a number of considerations that must be taken into account when looking at technologies for implementing DSM. Those relating to communication methods, device limits and incentives have been summarised here. For the best results when implementing a DSM system, information relating to scheduling and device operation needs to be available and this is the driver behind a number of the functionality requirements presented in the next section. If the LDC system were to be deployed commercially, further work in these areas would need to be undertaken.

7.3 A Proposed Smart Appliance Interface

In order for demand side management to become widespread, suitable appliances will need to have smart load functionality built in when they are manufactured. While some devices could be retrofitted or simply have their power supply regulated at the input, this option would not be suitable for the vast majority of loads.

If a manufacturer has the desire to integrate smart functionality at the present time, there are only limited options available. The most viable is the Zigbee smart energy profile (SEP). SEP has support for a number of DSM functions including the following:

- Load scheduling
- Temperature set-points & offsets
- Randomised start & end times
- Profiles for thermostats, water heaters, lights, pool pumps, EVs etc.

SEP was originally designed to work solely with Zigbee communications hardware. A renewed version of the profile, termed SEP 2.0 has now moved to IP based communication. As a consequence, SEP 2.0 is not backwards compatible with SEP 1.x. While this makes SEP more versatile, it also increases the complexity, cost and power consumption of SEP. As an example, the CC2538 SOC from Texas Instruments costs $6-10 depending on the level of bulk pricing. As appliance manufacturers often operate on very slim margins, a $6 increase in component cost is likely unreasonable for high volume products.

Some manufacturers are simply creating their own smart appliance frameworks. One such example is Samsung, who with their wide range of consumer appliances are well placed to create a whole family of products that work well together. One problem with this approach is that it creates "lock-in" and manufacturer incompatibility. Once a consumer has a range of Samsung devices that all work well together, they are forced to keep buying Samsung even if other manufacturers have built better performing or more cost effective devices.

Since much of the drive towards smart appliances is to create devices that save energy and are generally more environmentally friendly, the long term utility and inter-device compatibility is very important. For this reason, choosing an interface with good compatibility and a wide feature set is very important.

A low level smart appliance interface is proposed here that contains all the functionality required for smart loads at the lowest possible cost, including the functionality discussed in the previous section. The idea here is that any manufacturer can include this interface for very little cost, and a customer can then plug a Zigbee, Homeplug or other adapter into the interface as they so desire.

7.3.1 Smart Load Classes

Ideally, a low level interface should support every possible function required by current and future smart loads. Obviously this is a hard task and is likely impossible, but a good place to start is classifying different types of smart loads into groups, and then listing what functionality those groups need.

So far, three main types of smart loads have been identified. These are as follows:

1) Pure energy storage e.g. Battery Charging.
   An exact amount of energy is required before a known time, storage can occur at any point before this time. None of the stored energy will be consumed while the device is connected.
2) Hysteretic energy storage e.g. HWC, HVAC, Fridge/Freezer.
   There is a known amount of energy stored and limits on the upper and lower levels that can be reached under normal operation. Storage and consumption can and does occur at any time.
3) Device cycle with energy intensive periods e.g. washing machine, dishwasher.
   The device has started a cycle and at a known time will require significant power for a set period. Power intensive period can be delayed and potentially interrupted.

7.3.2 Functionality Requirements

Most deferrable loads can fit into the classes listed above. A comprehensive feature set for DSM loads has been outlined below along with an indication of which device classes the feature relates to in brackets.

- DSM Control
  - Proportional load control (1,2)
  - Binary load control (2,3)
- DSM Characteristics
  - DSM Mode of operation (1, 2 or 3)
  - Power consumption levels available e.g. continuously variable, 500W increments (1,2)
- Status & Scheduling
  - State of charge (1,2)
  - Target state of charge level (1,2)
  - Target state of charge time (1,2)
- Error Reporting
  - Error code
  - Error description
  - Error severity / level
- Basic Info (static)
  - Device max power consumption
  - Device max energy storage (1,2)
  - Device Model
  - Device Name
  - Device Type
- Interface Info
  - Max command frequency
  - Load response time
  - Interface version
  - Command list

Whichever technology is implemented for smart appliance communication, if it provides all of this information, a powerful scheduling and optimisation system is possible with minimal input from the end user. Additionally, as long the base interface can provide this information, other convertors can change between different smart appliance protocols while still providing good system performance.

7.3.3 Proposed Physical Design

There are a number of considerations that must be taken into account when determining the exact type of interface to use. These generally come down to cost, performance and reliability.

One possible solution to this problem is presented here that attempts to achieve a realistic trade-off between these factors.

To keep the cost down, a commonly available digital interface such as UART, SPI, or RS232 could be used. The proposal here will focus on a UART solution, as the majority of microcontrollers have these available, with some even providing multiple implementations on the same chip. A UART is simple to use and well understood by the majority of electrical engineers.

As has been discussed in Chapter 4, an ideal smart load would have a very fast response to load control signals. In order to cater for devices that may not be able to process UART commands in real-time, a PWM input in addition to the UART is envisaged. This has the added benefit of working for 'dumb' loads which lack in-built intelligence, but could process a PWM signal – an example of such could be a heating unit.

Safety and reliability are very important when it comes to domestic appliances, and this should be a driving consideration for any interface design. There are likely to be both wired and wireless devices connected to the interface. In order to keep costs down, it could be assumed that any wired device will have its own power supply, and any wireless device will not need supply isolation, only fusing. With these assumptions, a simple quad opto-coupler IC could be used as the only hardware required by the interface. Wireless devices would have access to a fused local 5V rail, and wired devices would only interact via the opto-coupler.

This results in a 6 pin interface with the electrical circuitry inside the smart appliance as shown in Figure 7.1. On the left side are the connections that go to a physical header for routing signals outside the device, and on the right are connections that go to the supply and microcontroller circuitry within the smart appliance. The component cost here is very small, with the optocoupler available for less than 50c with bulk pricing, and a small number of resistors and fuses also included.

The biggest manufacturing cost here may in fact be providing a port and wiring that routes the 6 electrical lines outside of the device. For this reason, although the component cost is quite low, the overall cost could in fact be greater than a PLC system as proposed in [106]. As previously mentioned, as long as the chosen interface used can provide the correct information for device control and scheduling, convertors between different smart grid communications technologies can be used.

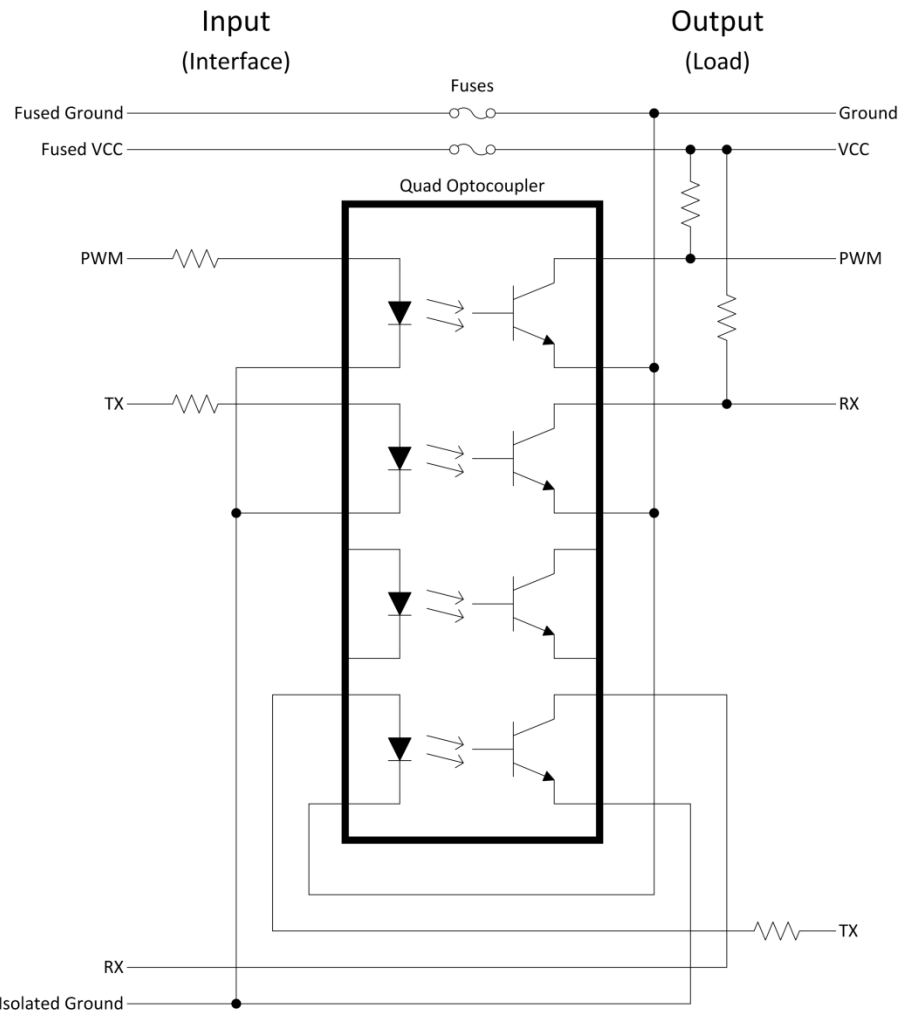

Figure 7.1 – Electrical connection of smart appliance interface

7.3.4 Summary

A proposed design for a smart appliance interface has been presented, including the functionality the interface should provide, as well as an option for the physical hardware interface.

7.4 High Frequency IPT Power Supply

Much of this research was centered around the impact of EV chargers on the power system, and possible ways of improving EV grid integration. In addition to the work on DSM and the design of LDC, the beginnings of a high frequency IPT power supply that integrates with the LDC system was created.

The goal of building this supply was to create an IPT power supply where the performance at a higher operating frequency could be measured, and an informed decision on frequency could be made, while also integrating the supply into the LDC system to determine likely real-time response features for the simulations used in earlier chapters. Only the power electronics

section of the power supply was completed, which includes the high frequency gate drive and H-bridge circuitry.

IPT, which stands for inductive power transfer, is a technology whereby significant power can be transferred over an air-gap with fully galvanic isolation and good efficiency. A full background of IPT is beyond the scope of this chapter, but in [112] the authors present a very good background and history of IPT systems and development. Current IPT systems created at the University of Auckland can transfer anywhere from 3-10kW over a 300mm air-gap at greater than 90% efficiency.

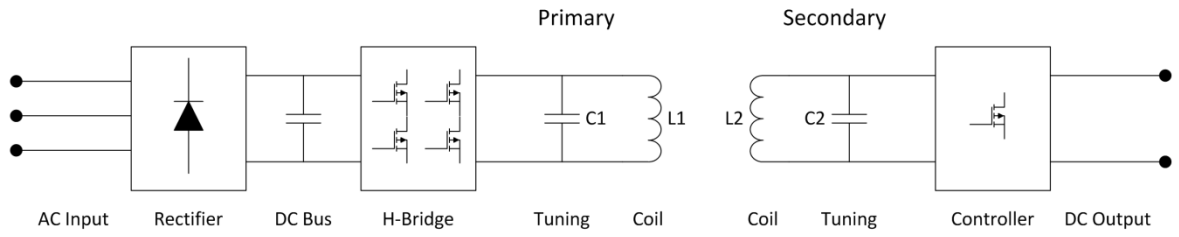

Figure 7.2 – Layout of typical IPT system

They layout of a typical IPT system is shown in Figure 7.2 – it works by first converting input AC electricity to DC, and then using an H-Bridge to step up the frequency into the kHz range, after which it is fed into a specially designed coil or pad. The coil is effectively one half of a transformer, and is mostly made up of ferrite and copper windings. There are many effective coil designs with quite varying magnetic structures and the design of these is an area of significant research. A second coil, usually of similar or identical design to the first, is placed such that there will be good magnetic flux linkage between the two. A current through the first (primary) coil will induce a current in the second (secondary) coil, and power can be transferred. In order to increase the current in the primary coil, a tuning stage is usually added to counteract some or all of the coil reactance at the desired operating frequency. Tuning is also added on the secondary side for similar effect.

Existing IPT systems generally operate at a frequency of between 10kHz and 45kHz. It is anticipated that there could be advantages to running the system at higher frequencies, and a system with an operating frequency of around 85kHz was desired. The main advantage relates to power transfer for a given field strength. With IPT, energy is transferred by a magnetic field, and the strength of the magnetic field is regulated. For a given field strength, a higher frequency will allow more power transfer. Since there is one set of field regulations for frequencies under

100kHz, and another for those above, the operating frequency of 85kHz is close to the maximum that can be used without moving into the next set of regulations.

It was also considered that there could be advantages in reactive component sizes as for a given impedance, a higher frequency will require a lower capacitance or inductance. This could results in smaller magnetic designs and smaller capacitors.

There are of course going to be trade-offs when attempting to raise the operating frequency. Firstly, all other things being equal, doubling the frequency would likely double the switching loss, as a set amount of energy is lost each switching cycle. In addition, conductors need to be improved, as the skin depth of wires will decrease. Some components will also have significantly reduced ratings at higher operating frequencies, and this needs taking into account.

One of the reasons that moving to a higher frequency is being considered is that the power electronics devices are improving at a rapid rate, and switching loss is becoming a smaller part of overall system losses. If the same amount of power can be transferred using a lower current and higher operating frequency, it is likely the overall efficiency would be improved. Modern silicon carbide (SiC) devices and the latest generation IGBTs now have very low switching and conduction losses, and these are put to good effect here.

7.4.1 Single switch gate drive performance

Designs were attempted with both SiC JFETs and the new generation of Infineon IGBTs. The JFETs proved very difficult to design effective gate drives for, and the final design used the IGBTs. Here acceptable gate drive was also difficult to achieve due to Miller capacitance issues, but a reference design with a fast gate driver chip and high performance decoupling capacitors was found. A single switch was tested at 400V and 2.5A. Turn on and off times of 15.4ns and 27ns were measured as shown in Figure 7.3 and Figure 7.4.

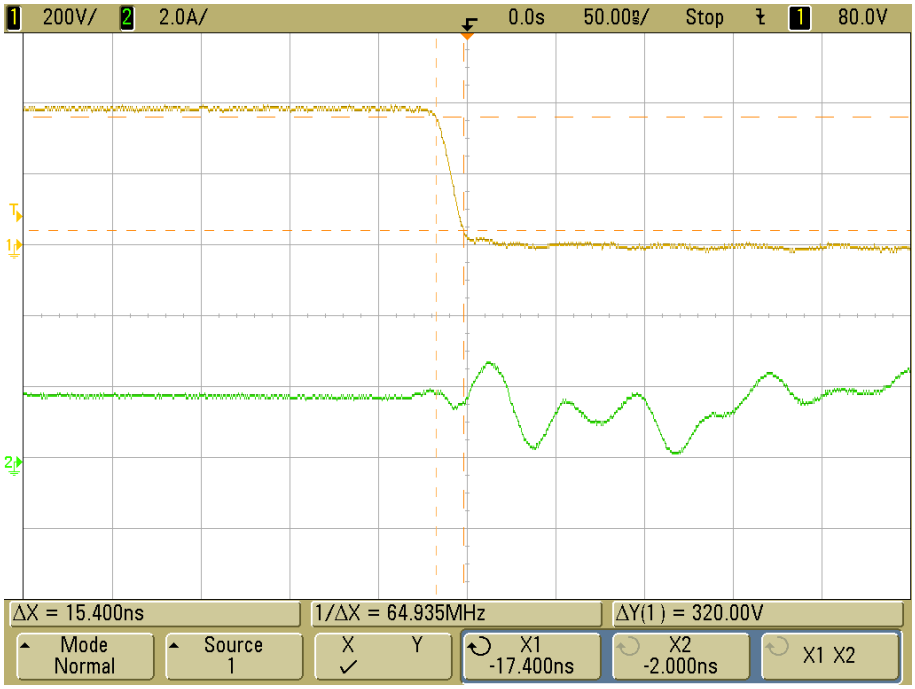

Figure 7.3 – Infineon IGBT turn on

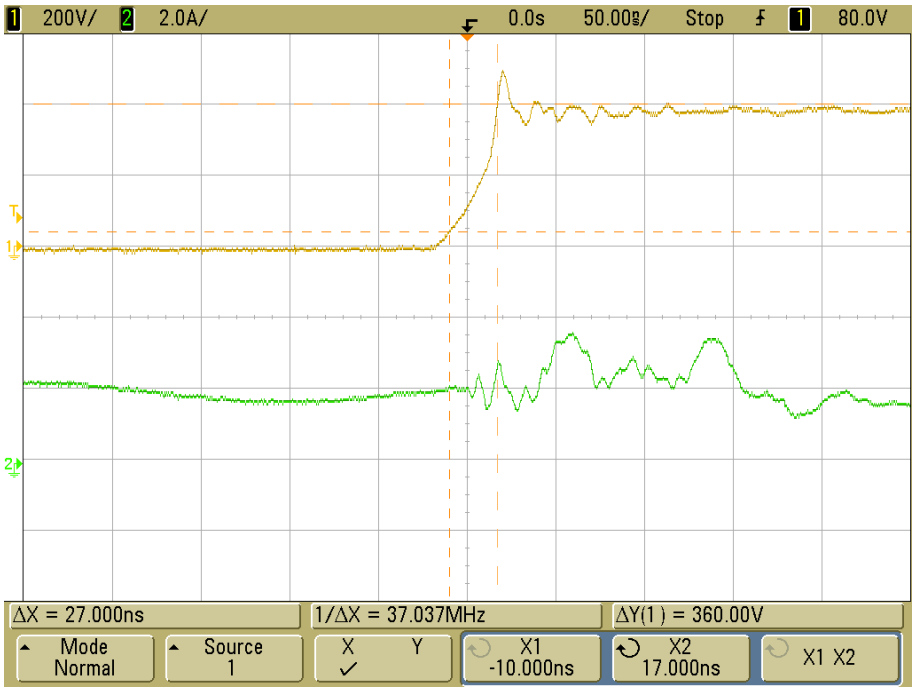

Figure 7.4 – Infineon IGBT turn off

7.4.2 H-bridge performance

With the working gate driver design, an H-bridge was made utilising new Infineon IPW50N65FFD switches. The H-Bridge is shown in Figure 7.5 with heat-sink and cables removed, and in Figure 7.6 with heat-sink, cables and control board connected. The goal of this design was keep the bridge as small as possible, minimise loop area of tracks with high current, and minimise distance between the gate drive and IGBT. It can be observed that the switching devices occupy almost the minimum possible board area, and the gate drive circuitry is very

close to the power devices. Space has also been left for additional anti-parallel diodes, as the inherent body diode in the IGBT package is not always fast enough for these applications.

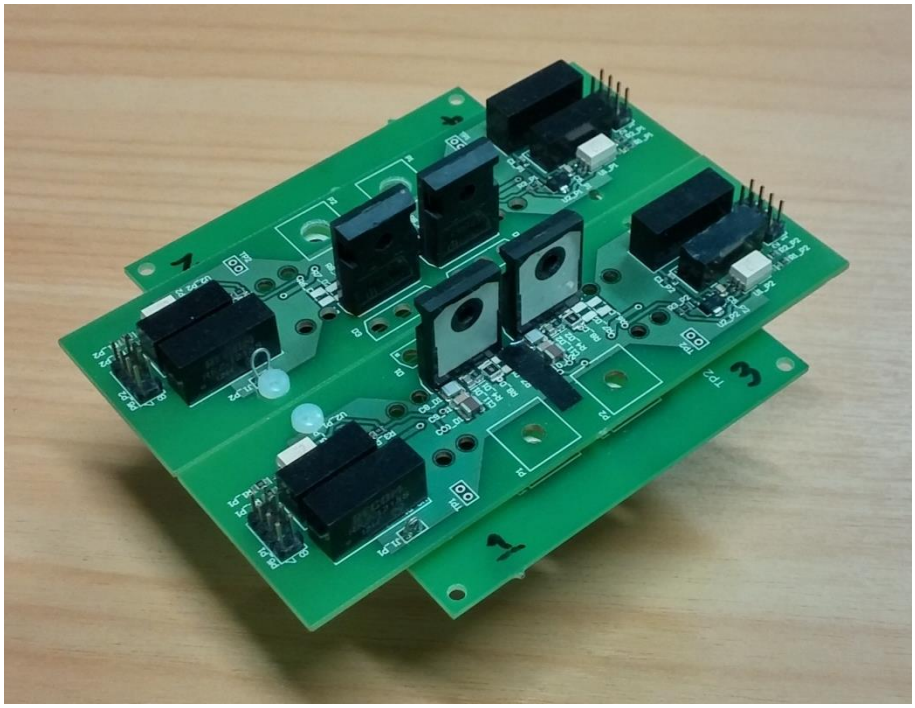

Figure 7.5 – H-Bridge with heat-sink and cables removed

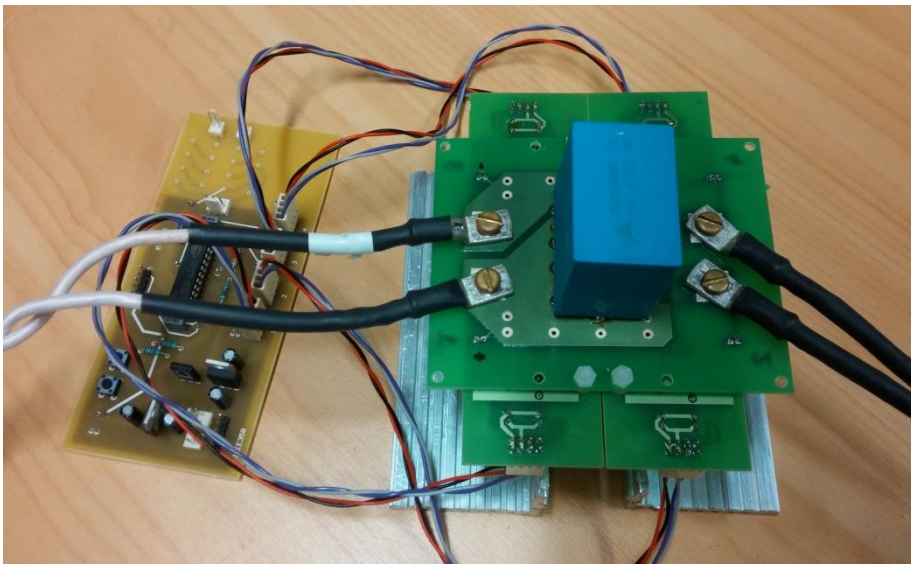

Figure 7.6 – H-Bridge in place with heat-sink and control board

The H-Bridge was tested to over 3kW, driving into a resistive load and a waveform of the bridge output during this test is shown in Figure 7.7.

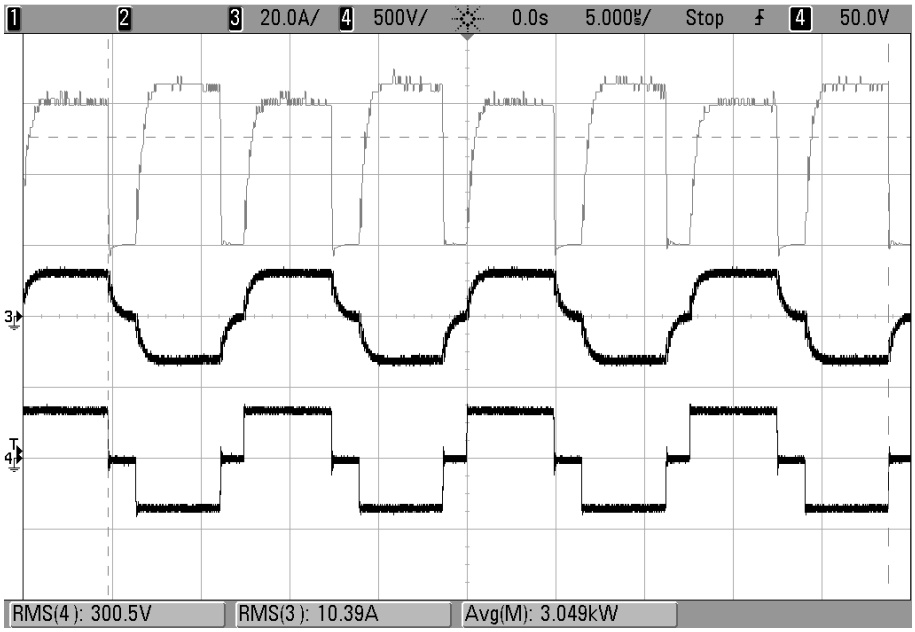

Figure 7.7 – Voltage and current output of H-Bridge, with calculated output power

Here the top trace is the math channel of the oscilloscope, which is multiplying the current and voltage together. The average of this waveform is calculated at 3.049kW which should be a relatively accurate measure of power throughput. The middle trace is the current, with an measured RMS value of 10.39A, and the bottom trace is the voltage, with an RMS value of 300.5V.

7.4.3 Comparison between current IXYS and new Infineon devices

With a working H-bridge, a comparison was made between the current design used and the bridge with the new devices. The existing power supply has a rectifier at the input and a tuning stage at the output, which the prototype bridge does not have. In order to test both under identical conditions, the tuning stage was removed and voltage was measured at the input to the bridge so the rectifier is ignored.

Even though the new bridge is designed for 85kHz, it was run at 46kHz so as to match the current bridge. An oscilloscope with power measurement software was used to calculate the efficiency of each H-bridge. The measured performance of the Infineon bridge and the IXYS bridge is shown in Table 7.1 and Table 7.2 respectively. It can be observed that the new Infineon device is slightly more efficient than the current IXYS device. Note though that the new bridge tested does not have any ancillary circuitry powered from the input, whereas the old supply is a complete unit including its own low voltage DC supply. The new bridge was consuming 6W to supply the gate circuitry, but even if 10W of extra input power is added to

the efficiency calculation, the performance is still better than the current design. This is a very good result as the existing IXYS devices were chosen for their high efficiency.

Table 7.1 – Infineon IPW50N65FFD bridge performance

Overall Efficiency Attribute	Value
Vin RMS	253.593 V
Iin RMS	4.242 A
Real Power (Input)	1.07 kW
Vout RMS	214.914 V
Iout RMS	4.906 A
Real Power (Output)	1.034 kW
Apparent Power (Output)	1.054 kVA
Reactive Power (Output)	203.143 VAR
Power Factor (Output)	0.981
Vout Freq	46.089 kHz
Overall System Efficiency	0.967

Table 7.2 – IXYS IXGH30N60C3C1 bridge performance

Overall Efficiency Attribute	Value
Vin RMS	249.97 V
Iin RMS	4.218 A
Real Power (Input)	1.053 kW
Vout RMS	214.394 V
Iout RMS	4.76 A
Real Power (Output)	990.672 W
Apparent Power (Output)	1.02 kVA
Reactive Power (Output)	244.806 VAR
Power Factor (Output)	0.971
Vout Freq	46.124 kHz
Overall System Efficiency	0.941

It should be noted that each of these measurements were taken under identical conditions with the same instruments, as such comparisons between the two systems are valid, but the absolute measurements are not necessarily 100% accurate as this would require very precise instrument calibration. The current probes used were HIOKI 3274 and 3275 units, which are rated to within $\pm 1\%$, and voltage probes were Tektronix P5200 units which are rated to within $\pm 3\%$. Both of these have ambient temperature and warm-up requirements which were met for these measurements. This implies an absolute accuracy of $\pm 4\%$ when voltage and current are multiplied together, and $\pm 8\%$ for the efficiency.

Also of interest was the efficiency of the new switches when operating at 85kHz. In Table 7.3 the measured performance of the H-bridge under the same conditions as the previous experiments, but with the frequency increased to 85kHz is shown. It can be observed that the efficiency has now dropped to 91.2%, which is a significant reduction in performance. Depending on the advantages that are found when operating at this higher frequency, it may prove uneconomic to attempt systems at 85kHz unless better devices are found.

Table 7.3 – Infineon IPW50N65FFD bridge performance at 85kHz

Overall Efficiency Attribute	Value
Vin RMS	253.788 V
Iin RMS	4.359 A
Real Power (Input)	1.105 kW
Vout RMS	215.345 V
Iout RMS	4.833 A
Real Power (Output)	1.008 kW
Apparent Power (Output)	1.041 kVA
Reactive Power (Output)	260.838 VAR
Power Factor (Output)	0.968
Vout Freq	84.301 kHz
Overall System Efficiency	0.912

7.4.4 Response of H-Bridge to a DSM signal

The goal of building an IPT power supply was to create something that could be integrated into a DSM system and perform very fast and precise load control. While a complete IPT system was not completed, a test has been performed showing the response of the H-Bridge to an external load control signal.

The bridge controller was configured so that the power output of the bridge would be changed in response to an external analogue voltage signal. This was achieved by using an analogue to digital convertor to receive the signal, and by adjusting the bridge phase angle accordingly. For this experiment, a 100Hz sine wave created by a function generator was used as the DSM signal, and a purely resistive load was placed on the bridge output. A current probe measured the output current so that the input and output waveforms could be compared. A waveform showing system operation is shown in Figure 7.8, where the top trace is the simulated DSM signal, and the bottom waveform is the measured H-Bridge output current. It can be observed that the envelope of the current waveform is varying by around 3A from around 1.5A to 5.5A. A phase delay of around 1ms can be seen between the two signals, and

small steps in the current waveform can be observed on close inspection. An enlarged view of the steps in the current waveform is shown in Figure 7.9, where the measured period for each step is 176μs, with a change in peak current of 137.5mA.

This experiment shows that electronically controlled loads could indeed provide fast and precise response to DSM signals if configured correctly. If just the sampling period is considered, the response time of 176μs is indeed very fast. If the observed phase delay is also taken into account, the delay of 1ms is still fast, and much quicker than the 0.1s time constants and propagation delay considered earlier in the design of LDC.

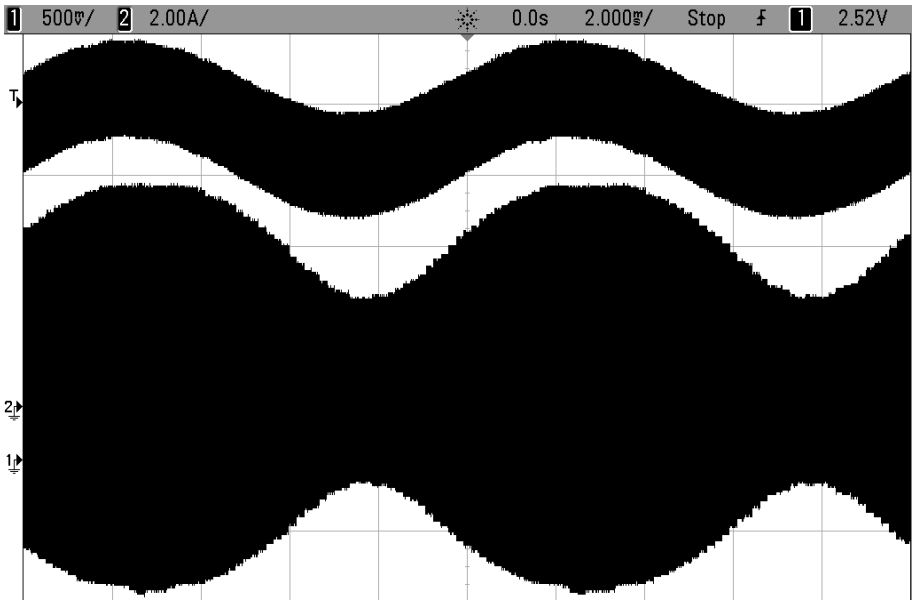

Figure 7.8 – Simulated DSM signal (top) and modulation of H-Bridge output current (bottom)

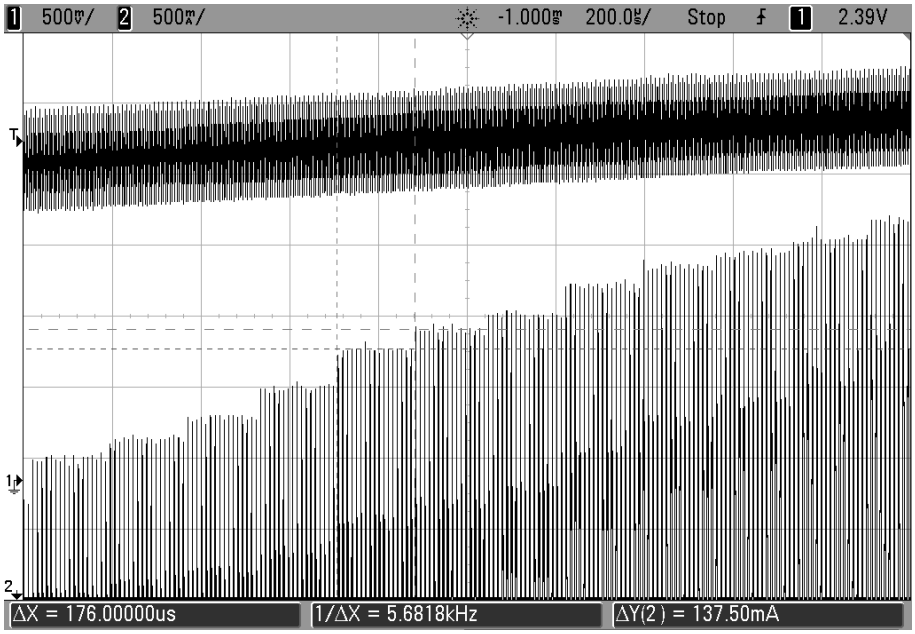

Figure 7.9 – Enlarged view of simulated DSM signal (top) and H-Bridge current with steps in amplitude (bottom)

Each step change in output is caused by a new sample of the DSM signal from the ADC, which is configured to sample at 5860Hz – a period of 171µs. In the enlarged view, peaks on the current waveform can be distinguished, with around 18 peaks per step in the output giving a rough estimate of 102kHz as the operating frequency. The actual frequency of operation in this experiment was closer to 98kHz.

While there is a significant amount of noise in the measured DSM signal waveform, it should be noted that there is a capacitor right next to the ADC pin so this noise is unlikely to propagate into the ADC measurement. In the enlarged waveform it can also be observed that the noise on the simulated DSM signal is close in frequency to the operating frequency of the IPT system – likely because it originates from the switching signals of the H-Bridge. The noise looks especially extreme in Figure 7.8 due to the large frequency difference between the 100Hz sine wave and 98kHz IPT system and aliasing of the oscilloscope.

7.4.5 Summary

An H-bridge designed for use in an IPT supply has been prototyped and compared with the current system. The overall efficiency of the new H-bridge with the Infineon power devices was measured to be around 2% better than the existing design. A test of the response of the bridge in response to a simulated load control signal was also performed, which showed electronically controlled loads could provide an exceptional response to DSM signals. To fully prove the power supply performance, a tuning stage and IPT pad needs to be added to the system so that actual wireless power transfer can be attempted. Further work could also be undertaken to integrate the power supply into the LDC experimental system.

7.5 Electronic Transformers

As has been mentioned throughout this thesis, it is quite likely that in the future, distribution transformers will be mainly electronic devices, which have a number of advantages over standard transformers. These solid state transformers (SST) – also called power electronic transformers or intelligent universal transformers – work by using power electronic circuitry to step up the operating frequency, so that a physically smaller transformer can be used to step down the voltage. While this configuration is more complicated than a line frequency transformer, the entire assembly is much smaller and lighter. As a consequence, SSTs which may be more expensive to purchase can end up being cost effective when installation costs are considered, as these are a significant proportion of the overall transformer replacement cost.

If a SST is implemented using a topology with a (large) capacitive bus, then the transformer can improve power quality by performing VAR compensation directly. The transformer could also dynamically adjust voltage in response to changing grid conditions, and use frequency droop for load control. The utilisation of voltage and frequency droop on a SST will be investigated more closely here.

As LDC requires the distribution of a real-time load control signal, additional hardware must be added to the transformer in order to perform this task. If an SST is used, where the frequency can be set artificially to be different from the wider power system, then the frequency could be used to distribute the load control signal and no extra hardware would be needed.

As mentioned during the discussion on the performance of LDC in Chapter 4, LDC does not necessarily prevent phase imbalances in the presence of unbalanced renewable generation or load, but if the individual phase voltages can be controlled independently, these can be used to perform load control on a per-phase basis while also working to keep phase voltages within specification.

7.5.1 Simulation Setup

A system where the overall load signalling was communicated using the local frequency, and where per-phase adjustments were performed with line voltage has been simulated under a number of conditions. In this simulation the transformer has two phases with a common frequency but independently controlled voltages. On each phase there are five houses and one standalone generating unit as shown in Figure 7.10. The same control scheme as in LDC is used, but now the LDC signal is encoded in the local frequency, with 49.5Hz signalling to controllable loads to avoid using any power, and 50.5Hz signalling maximum power. Each phase also has a voltage regulator which measures the voltage at the very end of the transmission line, and adjusts the phase voltage within a set range in an attempt to keep the end line voltage at 230V. In addition to measuring the frequency, the controllable loads also measure voltage, and use voltage droop to adjust the allowable load draw as signalled by the frequency accordingly. These three feedback mechanisms will be tested in various combinations here, where LDC refers to overall transformer load control, voltage feedback refers to the transformer level voltage regulation, and voltage droop refers to individual load based voltage regulation.

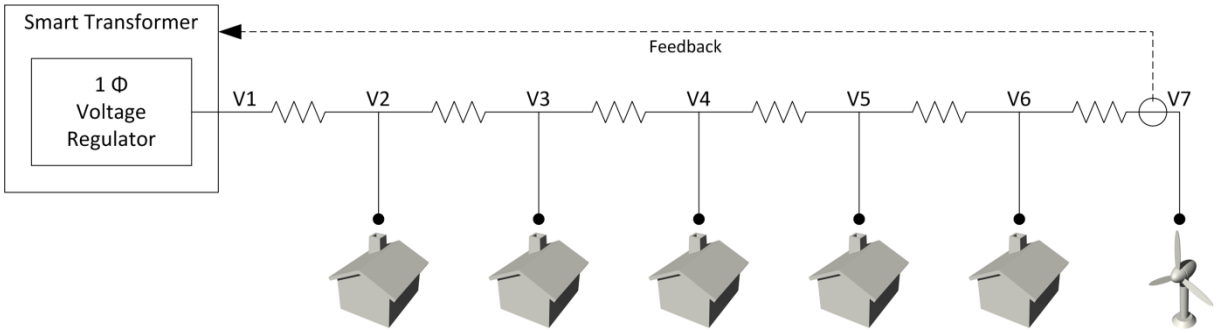

Figure 7.10 – Layout of an individual phase showing each voltage node and voltage feedback path

7.5.2 Simulation Results

The first simulation presents the base case showing the voltage stability issue that can occur in the low voltage network with local renewable generation. The generating unit at the end of the transmission line starts producing 8kW, but only on phase one. Here the line impedance is configured to be quite high (0.1Ω per segment) to amplify voltage stability issues.

A plot of the voltage profile along the transmission line is shown in Figure 7.11. The x-axis is the voltage node, labelled one through seven to represent each of the voltages labelled in Figure 7.10. There are three main traces on the plot, representing phase one without distributed generation enabled, and then both phases with generation enabled.

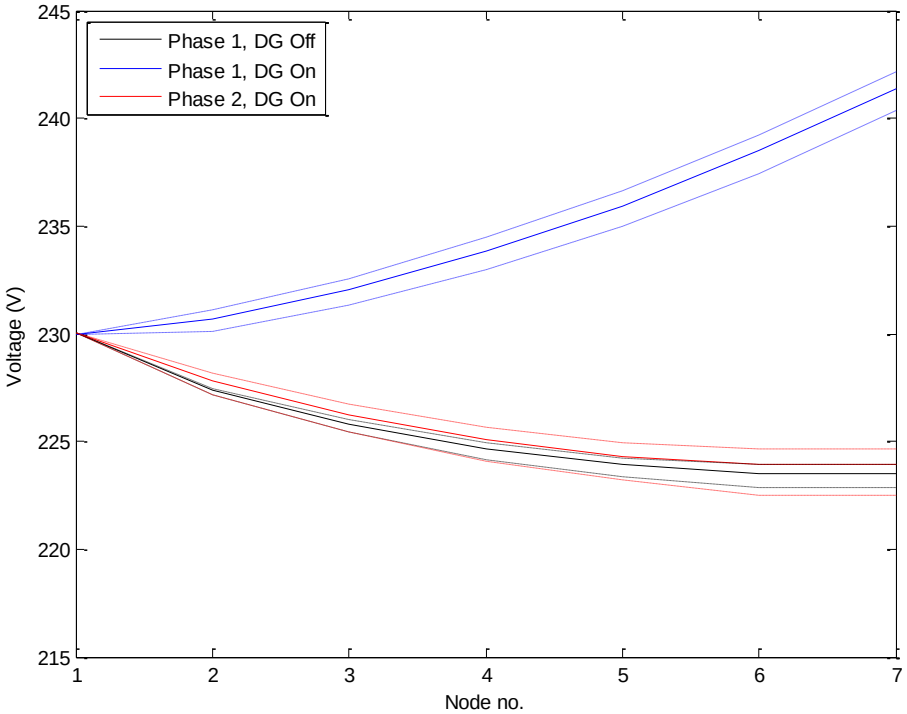

Figure 7.11 – Voltage profile with LDC, voltage feedback and voltage droop disabled

The dotted lines above and below each trace are the maximum and minimum voltage measured over a roughly 30min period. As expected, enabling the generation has no effect on

the voltage profile of phase two, so it is very similar to the profile of phase one with generation disabled. The voltage of phase one with generation enabled increases with the distance from the transformer, reaching a maximum of around 243V, when the nominal voltage is 230V. The maximum variation in voltage occurs at the end of the line, and is around 3V.

A simulation where LDC is enabled is shown next in Figure 7.12. The LDC controller measures the drop in consumption (increase in local generation) and signals more controllable load online. This has the effect of bringing the voltage on phase one to within safe limits, but the voltage on phase two now drops below 215V as it is similarly loaded but has no local generation. As is apparent, this is not a good solution and while this scenario is quite contrived, it is not beyond the realm of possibility.

If voltage droop is also enabled, then ideally the controllable loads will reduce consumption when the voltage is low, and the phase two voltage profile should be improved. This has been simulated and the results are shown in Figure 7.13, where the profile of phase two now matches phase one, and phase two is also kept closer to the nominal voltage. It can also be seen that the fluctuation in voltage on each line has reduced from around 3V to 1V. This result could be achieved in a standard LDC system without an electronic transformer, where voltage sensing is also added to the load control devices and used accordingly.

A simulation where just LDC and voltage feedback are enabled and voltage droop disabled has been performed, and the results are shown in Figure 7.14. Now the voltage at the transformer can range between 225V and 240V in order to regulate the end of line voltage to 230V. This method also helps to keep the voltages closer to the nominal 230V under all conditions.

Finally a simulation with all the functionality enabled is shown in Figure 7.15. The voltage is even more closely regulated and is again closer to the nominal 230V, showing the maximum performance of all these systems operating together.

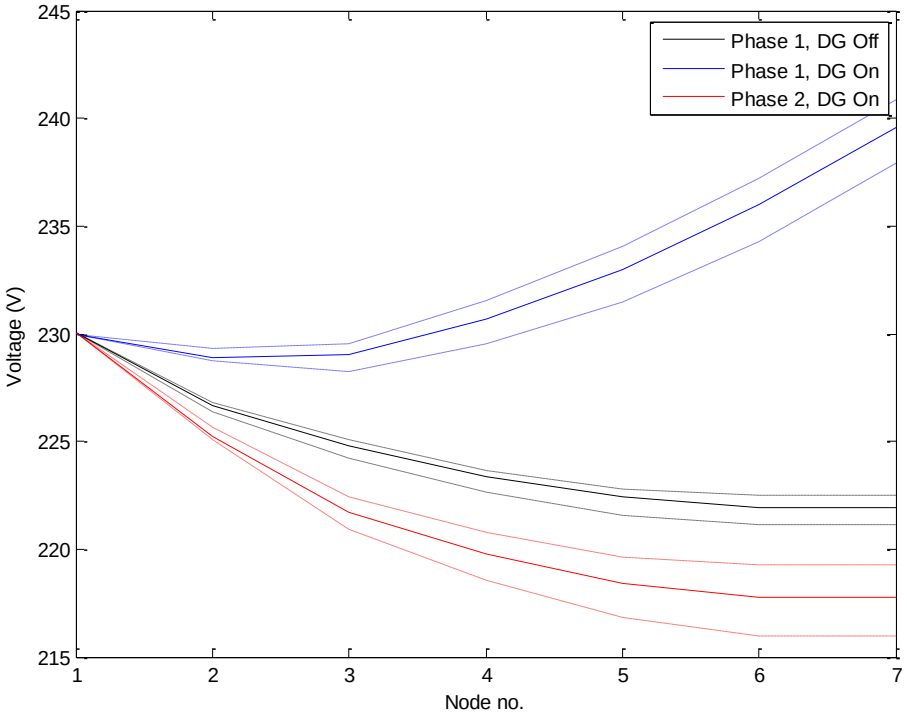

Figure 7.12 – Voltage profile with LDC enabled, and voltage feedback and voltage droop disabled

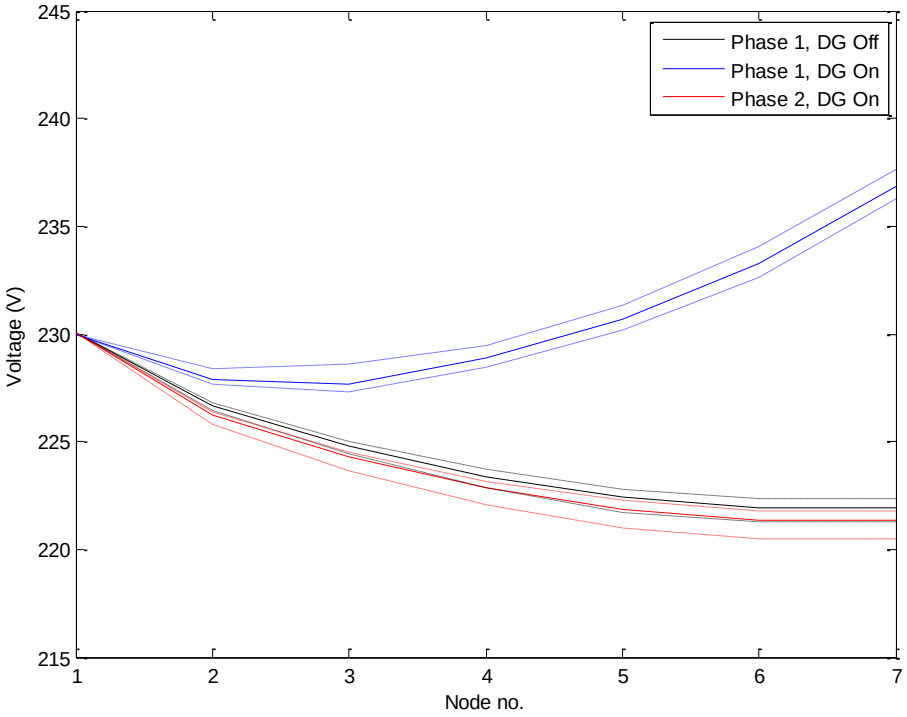

Figure 7.13 – Voltage profile with LDC and voltage droop enabled, and voltage feedback disabled

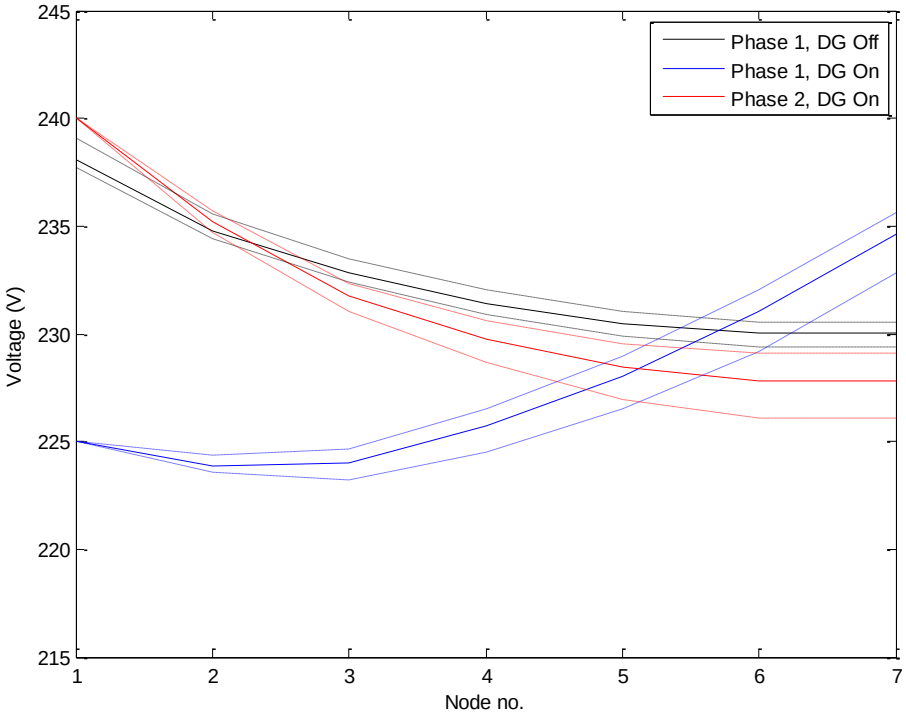

Figure 7.14 – Voltage profile with LDC and voltage feedback enabled, and voltage droop disabled

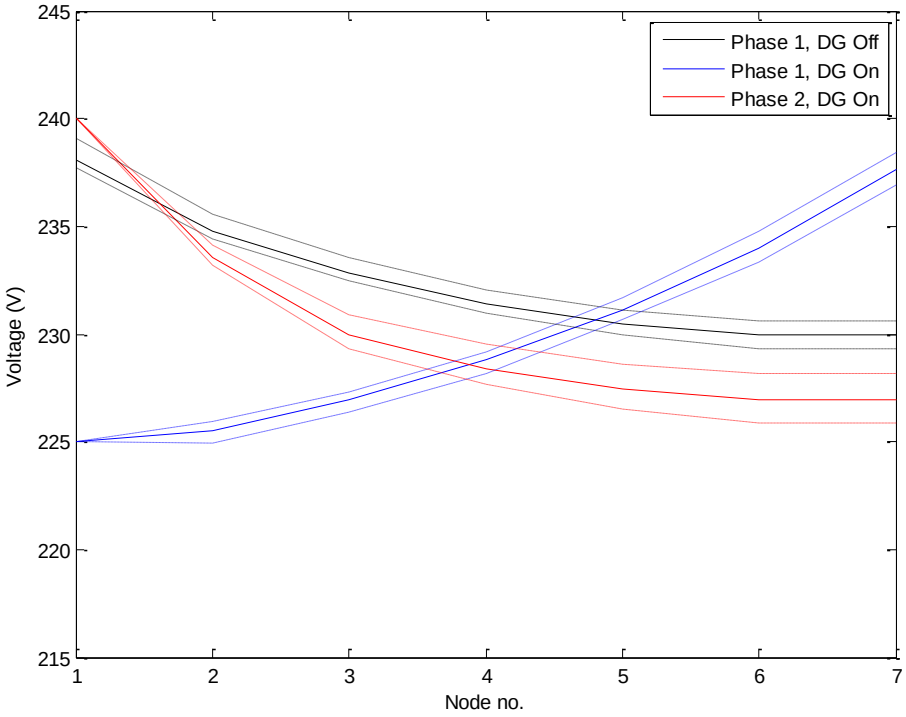

Figure 7.15 – Voltage profile with LDC , voltage feedback, and voltage droop enabled

7.5.3 Summary

In this section some preliminary simulations have been performed on a simple low voltage network with some embedded renewable generation on a single phase. In the base case, local generation causes significant voltage rise towards the end of the line, with some houses seeing voltages outside the desired operating range. Due to the single phase nature of the generation,

the addition of LDC does bring the voltage down on that phase, but at the same time causes a further drop on the phase without generation. Adding voltage monitoring to controllable loads improves the results, as does per-phase voltage control at the transformer level. The system with LDC, voltage monitoring and voltage control provides the best performance. Electronic transformers are a very large area of research, and far beyond the scope of this research, but testing of an LDC like load control system on an actual electronic transformer could produce very interesting results.

A potential issue with this control scheme is that houses (and therefore smart loads) connected at different nodes on the transmission line would measure different voltages on average. If the measured voltage is then used for load control, then houses closer to the transformer or local generation would receive more power on average than those further away. Ideally each smart load would have a slightly different voltage reference, which is representative of the range of expected voltages at that particular node. This could be determined statistically by each smart load or the smart meter at each household, or set manually as a constant within the smart load when the DSM system is installed. Whether or not this would have significant effect on electronic transformer DSM, and what could be done to mitigate these effects would be an interesting topic for future work.

7.6 Further Development of LDC

A number of areas could be investigated in relation to both the new possibilities that a system such as LDC opens up, and specific scenarios where it could be of significant benefit. Once such study that is relevant to New Zealand is the benefits LDC could provide when integrating significant levels of rooftop solar and air conditioning in areas such as Auckland. There is likely a good correlation between solar generation and air conditioning use, but while the total generation and consumption may match, the variance and effect of cloud movement could cause issues at the distribution level. The improvements that LDC could provide as well as the characteristics of solar generation in New Zealand would both be interesting studies.

Another interesting topic would be investigating where LDC can be beneficial in a move to a truly intelligent power system. With the current trends in the cost of computing power and communication, power systems could start making use of distributed intelligence when responding to system constraints and faults in order to reorganise the system and improve resilience. There is already much research on the smart grid and distributed intelligence, but

looking into what can be done with the extra information and control that LDC provides would be quite interesting.

7.7 Conclusions

In this chapter, a number of topics where further work could be undertaken have been discussed. Practical considerations when integrating DSM functionality into smart loads were outlined, and the motivation for using a simple, low-cost interface was given. Once such option for this interface was presented, along with a comprehensive outline of the functionality the interface needs to provide. The initial design of a high-frequency IPT supply using prototype Infineon IGBTs was presented, and showed efficiency improvements over the current IXYS IGBT based design. The bridge was also demonstrated in a DSM responsive mode of operation to determine the response of this load in a real network, and it showed very good performance in terms of response time and fidelity. Additionally, some initial simulations on electronic transformers were presented, and it was shown that making use of per-phase voltage regulation could be beneficial for improving voltage stability in the low-voltage network. Finally, some areas where research could be undertaken in relation to the impact and further development of LDC were suggested.

Chapter 8
Thesis Outputs

8.1 Summary of Conclusions

In this thesis the new DSM method LDC was presented. Simulations showed that with the addition of LDC, the power taken by a micro-grid could be regulated well, and that any renewable generation could be consumed locally, without exporting energy to the wider grid. In one simulation, the power input from the grid was regulated to 20kW, while local wind gusts of up to 150kW were consumed locally with no export of power to the grid. Scheduling simulations showed that the system could provide this functionality while still keeping controllable loads within the limits desired by the end user, while also minimising the overall cost of electricity by shifting as much consumption as possible to periods of lowest electricity cost.

The performance of the lab-scale LDC system was shown to agree well with the results of the simulation. Local wind generation was simulated with a small motor-generator set, and the LDC controller managed to consume the fluctuating generation locally, without exporting any power back to the wider grid. An experiment was also performed where the LDC system adjusted power draw to regulate the grid frequency, and the ability of the LDC system to provide good frequency regulating performance was demonstrated. Both of these features improve the grid integration of renewable energy. Local consumption of local generation reduces power system loading and improves stability both locally and at the grid level. By regulating the system frequency, fluctuations in generation from large scale renewable generation can be absorbed, further aiding integration. In addition, frequency regulation is helpful during grid emergencies where there is either a generation or line failure. Load will automatically be reduced (or added) as needed to reduce the impact on power system operation.

A power line communication system for performing load signalling was also presented, and simulations showed that the communication system should still work well in a practical grid environment, with expected levels of signal noise. Practical suggestions for the deployment of LDC and final communication link to controllable loads were also made, as well as the impact on LDC operation if solid state transformers were to be used.

The initial design of an IPT EV charger has also been presented and it was shown that if DSM functionality is fully integrated into the charger, then very good DSM performance can be provided. This result combined with the earlier simulations shows that EVs are well suited as controllable loads, and that with DSM, grid integration of EVs can be improved.

In summary, LDC was shown to be a good solution for mitigating many issues facing power systems, including those caused by increased deployment of grid-scale and localised renewable generation. In addition, LDC can regulate electric vehicle charging, turning what could be a burden for current power system operation into a valuable resource.

8.2 Innovations

The central innovation of this thesis is the act of performing demand side management at a distribution transformer level, which was indeed quite novel at the beginning of this research. Furthermore, it is argued that of the many technologies and topologies that could be used, this is a very good option with regards to the high level of performance and low complexity and implementation cost. A number of further innovations have been made, most of which build upon this initial idea.

With regards to what is present in the literature, it appears one of the most innovative aspects of this research is that a system has been designed and simulated, and then actually tested with real hardware in a laboratory setting. Additionally the laboratory results match very well with the simulated performance of the system.

One of the main goals of LDC is to improve the integration of renewable generation. The idea that fluctuating renewable energy can be smoothed out by intelligent electric vehicle charging is not completely novel. Though demonstrating that it is possible for the renewable generation to be smoothed almost completely, such that a constant load is seen by the wider power system does indeed appear to be novel. A related innovation is that in the presence of LDC, lower cost distributed generation can be used, as the expensive grid protection can be moved from each individual unit into one LDC controller.

The communication system described previously is also quite innovative. It is a simple design, yet it is robust and has good fault tolerance. The use of digital filters also means it can be reconfigured for different operating environments and frequencies with only a firmware change.

8.3 Research Outputs

A number of publications and patents were produced from the work undertaken in this thesis. The initial work on the simulated LDC system and the requirements for the communication system was presented at IECON 2011 [113]. Prof. John Boys was the primary author of a paper presented on the initial ideas behind LDC at INTELEC 2011 [114]. The adding of DDC functionality into the LDC system was published at ICRERA 2013 [115], and the overall design, comparison with DDC and experimental performance has been accepted for publication in IEEE Trans. on Smart Grid [116].

In addition, two patents have been applied for by Boys, Covic and Lee, and these are currently in the national phase [117] [118].

References

[1] 'Key World Energy Statistics', International Energy Agency, 2012.

[2] IEA, *Medium-Term Renewable Energy Market Report 2012*. Paris: OECD Publishing, 2012.

[3] 'Gordon Hughes: The Impact of Wind Power On Household Energy Bills | The Global Warming Policy Foundation (GWPF)'. [Online]. Available: http://www.thegwpf.org/gordon-hughes-the-impact-of-wind-power-on-household-energy-bills/. [Accessed: 03-Apr-2014].

[4] 'High Power Wall Connector | Charging Solutions | Tesla Motors'. [Online]. Available: http://www.teslamotors.com/roadster/charging/high-power-wall-connector. [Accessed: 30-Jan-2014].

[5] '2013 Nissan LEAF® Specs', *Nissan USA*. [Online]. Available: http://www.nissanusa.com/electric-cars/leaf/versions-specs. [Accessed: 30-Jan-2014].

[6] 'Nissan LEAF Global Sales Reach 100,000 Units', *Nissan Online Newsroom*. [Online]. Available: http://nissannews.com/en-US/nissan/usa/channels/us-united-states-nissan-models-leaf/releases/nissan-leaf-global-sales-reach-100-000-units. [Accessed: 30-Jan-2014].

[7] R. A. Verzijlbergh, M. O. W. Grond, Z. Lukszo, J. G. Slootweg, and M. D. Ilic, 'Network Impacts and Cost Savings of Controlled EV Charging', *IEEE Trans. Smart Grid*, vol. 3, no. 3, pp. 1203–1212, 2012.

[8] 'New Zealand Energy Data File 2012', Ministry of Economic Development 2012, 2012.

[9] G. Hassan, 'Wind power variability and forecast accuracy in New Zealand', Mar. 2007.

[10] 'System Operator: Annual Review and Assessment 2011/12', Transpower New Zealand Ltd, Nov. 2012.

[11] 'System Operator: Annual Review and Assessment 2013/14', Transpower New Zealand Ltd, Nov. 2014.

[12] 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. IEEE*, vol. 99, no. 1, pp. 214–232, 2011.

[13] M. Nicolosi, 'Wind power integration and power system flexibility-An empirical analysis of extreme events in Germany under the new negative price regime', *Energy Policy*, 2010.

[14] E. A. DeMeo, W. Grant, M. R. Milligan, and M. J. Schuerger, 'Wind plant integration [wind power plants]', *IEEE Power Energy Mag.*, vol. 3, no. 6, pp. 38– 46, Dec. 2005.

[15] D. Lew, D. Piwko, N. Miller, G. Jordan, K. Clark, and L. Freeman, 'How Do High Levels of Wind and Solar Impact the Grid? The Western Wind and Solar Integration Study', National Renewable Energy Laboratory, Technical Report NREL/TP-5500-50057, Dec. 2010.

[16] C. L. Masters, J. Mutale, G. Strbac, S. Curcic, and N. Jenkins, 'Statistical evaluation of voltages in distribution systems with embedded wind generation', *Gener. Transm. Distrib. IEE Proc.-*, vol. 147, no. 4, pp. 207–212, 2000.

[17] A. Woyte, V. Van Thong, R. Belmans, and J. Nijs, 'Voltage fluctuations on distribution level introduced by photovoltaic systems', *IEEE Trans. Energy Convers.*, vol. 21, no. 1, pp. 202–209, 2006.

[18] H. Farhangi, 'The path of the smart grid', *Power Energy Mag. IEEE*, vol. 8, no. 1, pp. 18–28, Feb. 2010.

[19] X. Mosquet, M. Devineni, T. Mezger, H. Zablit, A. Dinger, G. Sticher, M. Gerrits, and M. Russo, 'Powering Autos to 2020', Jul-2011. [Online]. Available: http://www.bcg.com/documents/file80920.pdf. [Accessed: 01-Aug-2011].
[20] N. Anglani, F. Fattori, and G. Muliere, 'The impact of electric vehicles on the grid for local energy models', *IECON 2012 - 38th Annu. Conf. IEEE Ind. Electron. Soc.*, pp. 3000–3006, 25.
[21] K. J. Yunus, M. Reza, H. Zelaya-De La Parra, and K. Srivastava, 'Impacts of Stochastic Residential Plug-In Electric Vehicle Charging on Distribution Grid', in *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, 2012, pp. 1–8.
[22] A. D. Hilshey, P. D. H. Hines, P. Rezaei, and J. R. Dowds, 'Estimating the Impact of Electric Vehicle Smart Charging on Distribution Transformer Aging', *IEEE Trans. Smart Grid*, vol. 4, no. 2, pp. 905–913, 2013.
[23] J. C. Gomez and M. M. Morcos, 'Impact of EV battery chargers on the power quality of distribution systems', *IEEE Trans. Power Deliv.*, vol. 18, no. 3, pp. 975–981, 2003.
[24] P. Richardson, D. Flynn, and A. Keane, 'Impact assessment of varying penetrations of electric vehicles on low voltage distribution systems', *Power Energy Soc. Gen. Meet. 2010 IEEE*, pp. 1–6, Jul. 2010.
[25] S. Huang and D. Infield, 'The potential of domestic electric vehicles to contribute to Power System Operation through vehicle to grid technology', 2009.
[26] M. Kintner-Meyer, K. Schneider, and R. Pratt, 'Impacts assessment of plug-in hybrid vehicles on electric utilities and regional US power grids part 1: Technical analysis', *Pac. Northwest Natl. Lab.*, 2007.
[27] S. E. Widergren and H. Kirkham, 'Smart grid - transforming power system operations', pp. 1–5, Jul. 2010.
[28] Litos Strategic Communication, *The Smart Grid: An Introduction*. DOE Office of Electricity Delivery and Energy Reliability, 2008.
[29] K. Tomsovic, D. E. Bakken, V. Venkatasubramanian, and A. Bose, 'Designing the Next Generation of Real-Time Control, Communication, and Computations for Large Power Systems', *Proc. IEEE*, vol. 93, no. 5, pp. 965–979, May.
[30] N. Isaacs, K. Saville-Smith, M. Camilleri, and L. Burrough, 'Energy in New Zealand houses: comfort, physics and consumption', *Build. Res. Inf.*, vol. 38, no. 5, pp. 470–480, 2010.
[31] 'Electricity Demand-side Management — The Treasury - New Zealand'. [Online]. Available: http://www.treasury.govt.nz/publications/research-policy/tp/edm. [Accessed: 08-May-2011].
[32] A. Brooks, E. Lu, D. Reicher, C. Spirakis, and B. Weihl, 'Demand Dispatch', *IEEE Power Energy Mag.*, vol. 8, no. 3, pp. 20–29, Jun. 2010.
[33] S. De Rijcke and J. Driesen, 'Balancing Wind Power with Demand-side Response', Feb. 2010.
[34] S. Kreutz, H.-J. Belitz, and C. Rehtanz, 'The impact of Demand Side Management on the residual load', in *Innovative Smart Grid Technologies Conference Europe (ISGT Europe), 2010 IEEE PES*, 2010, pp. 1–5.
[35] N. Isaacs, M. Camilleri, and L. French, 'HEEPS OF DOMESTIC HOT WATER'.
[36] Lisa Burrough, 'Heat Pumps in New Zealand Houses', presented at the 5th Australasian Housing Researchers Conference, Auckland, New Zealand, 2010.
[37] Ministry of Transport, New Zealand, 'New Zealand Household Travel Survey 2009–2012', Mar-2012. [Online]. Available: http://www.transport.govt.nz/assets/Import/Documents/NZ-Household-travel-survey-driver-travel-April-2013.pdf. [Accessed: 03-Jul-2014].

[38] Z. McDonald, ‘Texas Energy Provider Gives Free Power at Night to Electric Car Owners | PluginCars.com’, 31-Jul-2013. [Online]. Available: http://www.plugincars.com/free-electricity-night-texas-ev-owners-127861.html. [Accessed: 13-Dec-2013].
[39] V. Hamidi, F. Li, and F. Robinson, ‘Responsive demand in networks with high penetration of wind power’, in *Transmission and Distribution Conference and Exposition, 2008. T&# x00026; D. IEEE/PES*, 2008, pp. 1–7.
[40] P. Finn, C. Fitzpatrick, M. Leahy, and L. Relihan, ‘Promotion of wind generated electricity using price responsive Demand Side Management: Price prediction analysis for imperfect energy storage’, in *Sustainable Systems and Technology (ISSST), 2010 IEEE International Symposium on*, 2010, pp. 1–5.
[41] A. Karnama and V. Knazkins, ‘Scenario-based investigation of the effects of Plug-in Hybrid Electric Vehicles (PHEVs) in 11 kV substations in Stockholm’, in *Energy Market (EEM), 2010 7th International Conference on the European*, 2010, pp. 1–6.
[42] J. E. Hernandez, F. Kreikebaum, and D. Divan, ‘Flexible electric vehicle (EV) charging to meet renewable portfolio standard (RPS) mandates and minimize green house Gas emissions’, in *Energy Conversion Congress and Exposition (ECCE), 2010 IEEE*, 2010, pp. 4270–4277.
[43] A. Schuller, J. Ilg, and C. van Dinther, ‘Benchmarking electric vehicle charging control strategies’, in *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, 2012, pp. 1–8.
[44] Y. Mu, J. Wu, J. Ekanayake, N. Jenkins, and H. Jia, ‘Primary Frequency Response From Electric Vehicles in the Great Britain Power System’, *IEEE Trans. Smart Grid*, vol. 4, no. 2, pp. 1142–1150, 2013.
[45] P. Kadurek, C. Ioakimidis, and P. Ferrao, ‘Electric Vehicles and their impact to the electric grid in isolated systems’, in *Power Engineering, Energy and Electrical Drives, 2009. POWERENG ’09. International Conference on*, 2009, pp. 49–54.
[46] K. Shimizu, T. Masuta, Y. Ota, and A. Yokoyama, ‘Load Frequency Control in power system using Vehicle-to-Grid system considering the customer convenience of Electric Vehicles’, in *Power System Technology (POWERCON), 2010 International Conference on*, 2010, pp. 1–8.
[47] B. J. Kirby, ‘Load Response Fundamentally Matches Power System Reliability Requirements’, 2007.
[48] D. Crossley, ‘Orion Network DSM Program - New Zealand’, New Zealand, DC08.
[49] P. S. Moura and A. T. de Almeida, ‘The role of demand-side management in the grid integration of wind power’, *Appl. Energy*, vol. 87, no. 8, pp. 2581–2588, Aug. 2010.
[50] D. Cirio, G. Demartini, S. Massucco, A. Morim, P. Scalera, F. Silvestro, and G. Vimercati, ‘Load control for improving system security and economics’, Geneva, 2003.
[51] Z. Xu, X. Guan, Q.-S. Jia, J. Wu, D. Wang, and S. Chen, ‘Performance Analysis and Comparison on Energy Storage Devices for Smart Building Energy Management’, *IEEE Trans. Smart Grid*, vol. 3, no. 4, pp. 2136–2147, 2012.
[52] H. Heath and N. L. Young, ‘13 December AUFLS Event - System Operator, Transpower New Zealand Ltd’, 01-May-2012. [Online]. Available: http://www.systemoperator.co.nz/sites/default/files/bulk-upload/documents/AUFLS%20event%2013%20Dec%202011%20report.pdf. [Accessed: 05-Jul-2014].
[53] ‘Colorado University finds Prius plug-in hybrids a hit; online assistance, not so much’, *AutoblogGreen*. [Online]. Available: http://green.autoblog.com/2012/10/23/colorado-university-finds-prius-plug-in-hybrids-a-hit-online-as/. [Accessed: 23-Oct-2012].
[54] D. J. Hammerstrom, R. Ambrosio, T. A. Carlon, J. G. DeSteese, G. R. Horst, R. Kajfasz, L. L. Kiesling, P. Michie, R. G. Pratt, M. Yao, and others, ‘Pacific Northwest

GridWise™ Testbed Demonstration Projects; Part I. Olympic Peninsula Project', Pacific Northwest National Laboratory (PNNL), Richland, WA (US), 2008.
[55] M. Zimmerman, 'The industry demands better demand response', in *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, 2012, pp. 1–3.
[56] B. Gwisdorf, S. Stepanescu, and C. Rehtanz, 'Effects of Demand Side Management on the planning and operation of distribution grids', in *Innovative Smart Grid Technologies Conference Europe (ISGT Europe), 2010 IEEE PES*, 2010, pp. 1–5.
[57] V. C. Gungor, D. Sahin, T. Kocak, S. Ergut, C. Buccella, C. Cecati, and G. P. Hancke, 'Smart Grid Technologies: Communication Technologies and Standards', *IEEE Trans. Ind. Inform.*, vol. 7, no. 4, pp. 529 –539, Nov. 2011.
[58] S. Galli, A. Scaglione, and Z. Wang, 'For the Grid and Through the Grid: The Role of Power Line Communications in the Smart Grid', *Proc. IEEE*, vol. 99, no. 6, pp. 998–1027, Jun. 2011.
[59] D. Niyato, P. Wang, and E. Hossain, 'Reliability analysis and redundancy design of smart grid wireless communications system for demand side management', *IEEE Wirel. Commun.*, vol. 19, no. 3, pp. 38–46, June.
[60] F. Gomez-Cuba, R. Asorey-Cacheda, and F. J. Gonzalez-Castano, 'Smart Grid Last-Mile Communications Model and Its Application to the Study of Leased Broadband Wired-Access', *IEEE Trans. Smart Grid*, vol. 4, no. 1, pp. 5–12, 2013.
[61] J. Wright, 'Smart electricity meters: How households and the environment can benefit', NZ Parliamentary Commissioner for the Environment, Wellington, Jun. 2009.
[62] J. Wright, 'Advanced Metering Infrastructure: Nomination of the MEP and access to data - Submission to the NZ Parliamentary Commissioner for the Environment', Jun-2011. [Online]. Available: www.pce.parliament.nz/assets/Uploads/PCE-EA-Submission.pdf. [Accessed: 05-Jul-2014].
[63] 'Directive 2009/72/EC of the European Parliament and of the Council of 13 July 2009 concerning common rules for the internal market in electricity and repealing Directive 2003/54/EC', *Off. J. Eur. Union*, no. L 211, p. 55, Aug. 2009.
[64] H. Chaudhry and T. Bohn, 'Security concerns of a plug-in vehicle', in *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, 2012, pp. 1–6.
[65] D. Bunyai, L. Krammer, and W. Kastner, 'Limiting constraints for ZigBee networks', *IECON 2012 - 38th Annu. Conf. IEEE Ind. Electron. Soc.*, pp. 4840–4846, 25.
[66] G. Stanciulescu, H. Farhangi, A. Palizban, and N. Stanchev, 'Communication technologies for BCIT Smart Microgrid', in *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, 2012, pp. 1–7.
[67] Jin Zhong, Chongqing Kang, and Kai Liu, 'Demand side management in China', in *Power and Energy Society General Meeting, 2010 IEEE*, 2010, pp. 1–4.
[68] A.-H. Mohsenian-Rad and A. Leon-Garcia, 'Optimal Residential Load Control With Price Prediction in Real-Time Electricity Pricing Environments', *IEEE Trans. Smart Grid*, vol. 1, no. 2, pp. 120–133, Sep. 2010.
[69] J. M. Elder, J. T. Boys, and J. L. Woodward, 'Integral cycle control of stand-alone generators', *Gener. Transm. Distrib. IEE Proc. C*, vol. 132, no. 2, pp. 57–66, 1985.
[70] D. G. I. Joe A. Short, 'Stabilization of Grid Frequency Through Dynamic Demand Control', 2007.
[71] C.-T. Li, C. Ahn, H. Peng, and J. Sun, 'Integration of plug-in electric vehicle charging and wind energy scheduling on electricity grid', in *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, 2012, pp. 1–7.
[72] J. Kondoh, 'Autonomous frequency regulation by controllable loads to increase acceptable wind power generation', *Wind Energy*.

[73] M. Tokudome, K. Tanaka, T. Senjyu, A. Yona, T. Funabashi, and Chul-Hwan Kim, 'Frequency and voltage control of small power systems by decentralized controllable loads', in *Power Electronics and Drive Systems, 2009. PEDS 2009. International Conference on*, 2009, pp. 666–671.

[74] K. Samarakoon, J. Ekanayake, and N. Jenkins, 'Investigation of Domestic Load Control to Provide Primary Frequency Response Using Smart Meters', *IEEE Trans. Smart Grid*, vol. 3, no. 1, pp. 282–292, 2012.

[75] 'Honeywell Helps Canadian Utilities Increase Energy Efficiency and Improve Grid Performance With peaksaver PLUS Programs'. [Online]. Available: http://honeywell.com/News/Pages/Honeywell-Helps-Canadian-Utilities-.aspx. [Accessed: 29-Aug-2012].

[76] 'Association for Demand Response and Smart Grid - Con Edison CSI'. [Online]. Available: http://www.demandresponsesmartgrid.org/Default.aspx?pageId=1409851. [Accessed: 27-Sep-2012].

[77] 'SAMSUNG and Patricia Heaton Will Launch Groundbreaking New Line of Smart Home Appliances to Help Moms Work Smarter, Not Harder', *Samsung Electronics America*, 20-Jun-2012. [Online]. Available: http://www.samsung.com/us/news/20193. [Accessed: 02-Apr-2014].

[78] 'Panasonic expands smart home appliance line, adds Android Smart App, cloud services', *Engadget*. [Online]. Available: http://www.engadget.com/2012/08/21/panasonic-expands-smart-home-appliance-line-adds-android-smart/. [Accessed: 21-Aug-2012].

[79] 'SAMSUNG Unveils New Era of Smart Home at CES 2014', *Samsung Electronics America*, 05-Jan-2014. [Online]. Available: http://www.samsung.com/us/news/22331. [Accessed: 02-Apr-2014].

[80] A. Zeman, M. Prokopenko, Y. Guo, and R. Li, 'Adaptive Control of Distributed Energy Management: A Comparative Study', in *Second IEEE International Conference on Self-Adaptive and Self-Organizing Systems, 2008. SASO '08*, Oct., pp. 84–93.

[81] Jin Xiao, Jae Yoon Chung, Jian Li, R. Boutaba, and J. W.-K. Hong, 'Near optimal demand-side energy management under real-time demand-response pricing', in *2010 International Conference on Network and Service Management (CNSM)*, 2010, pp. 527–532.

[82] A. Mohsenian-Rad, V. W. S. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia, 'Autonomous Demand-Side Management Based on Game-Theoretic Energy Consumption Scheduling for the Future Smart Grid', *Smart Grid IEEE Trans. On*, vol. 1, no. 3, pp. 320–331, 2010.

[83] D. Ban, G. Michailidis, and M. Devetsikiotis, 'Demand response control for PHEV charging stations by dynamic price adjustments', in *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, 2012, pp. 1–8.

[84] H. Yano, K. Kudo, T. Ikegami, H. Iguchi, K. Kataoka, and K. Ogimoto, 'A novel charging-time control method for numerous EVs based on a period weighted prescheduling for power supply and demand balancing', *Innov. Smart Grid Technol. ISGT 2012 IEEE PES*, pp. 1–6, 16.

[85] G. K. H. Larsen, N. D. van Foreest, and J. M. A. Scherpen, 'Distributed Control of the Power Supply-Demand Balance', *IEEE Trans. Smart Grid*, vol. 4, no. 2, pp. 828–836, 2013.

[86] K. Ogimi, K. Uchida, A. Yona, T. Senjyu, and T. Funabashi, 'Optimal operation method of wind farm with demand response', in *2012 International Conference on Renewable Energy Research and Applications (ICRERA)*, 2012, pp. 1–6.

[87] C. Wu, H. Mohsenian-Rad, J. Huang, and A. Y. Wang, 'Demand side management for Wind Power Integration in microgrid using dynamic potential game theory', in *2011 IEEE GLOBECOM Workshops (GC Wkshps)*, 2011, pp. 1199 –1204.
[88] S. Mak and D. Radford, 'Communication system requirements for implementation of a large scale demand side management and distribution automation', *Power Deliv. IEEE Trans. On*, vol. 11, no. 2, pp. 683–689, Apr. 1996.
[89] M. Kuss, P. Denholm, and T. Markel, 'Communication and control of electric drive vehicles supporting renewables', in *Vehicle Power and Propulsion Conference, 2009. VPPC '09. IEEE*, 2009, pp. 27–34.
[90] M. Daoud and X. Fernando, 'On the Communication Requirements for the Smart Grid', *Energy Power Eng.*, vol. 3, no. 1, pp. 53–60, Feb. 2011.
[91] X. Lu, W. Wang, and J. Ma, 'An Empirical Study of Communication Infrastructures Towards the Smart Grid: Design, Implementation, and Evaluation', *Smart Grid IEEE Trans. On*, vol. PP, no. 99, pp. 1–14, 2013.
[92] Q. Shafiee, J. C. Vasquez, and J. M. Guerrero, 'Distributed secondary control for islanded MicroGrids - A networked control systems approach', *IECON 2012 - 38th Annu. Conf. IEEE Ind. Electron. Soc.*, pp. 5637–5642, 25.
[93] D. Kawashima, M. Ihara, Tianmeng Shen, and H. Nishi, 'Real-time simulation of cooperative demand control method with batteries', *IECON 2012 - 38th Annu. Conf. IEEE Ind. Electron. Soc.*, pp. 3588–3593, 25.
[94] P. Stluka, D. Godbole, and T. Samad, 'Energy management for buildings and microgrids', in *2011 50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*, 2011, pp. 5150 –5157.
[95] M. Singh, I. Kar, and P. Kumar, 'Influence of EV on grid power quality and optimizing the charging schedule to mitigate voltage imbalance and reduce power loss', in *Power Electronics and Motion Control Conference (EPE/PEMC), 2010 14th International*, 2010, pp. T2–196–T2–203.
[96] S. J. Gunter, K. K. Afridi, and D. J. Perreault, 'Optimal Design of Grid-Connected PEV Charging Systems With Integrated Distributed Resources', *IEEE Trans. Smart Grid*, vol. 4, no. 2, pp. 956–967, 2013.
[97] H. Turker, A. Hably, and S. Bacha, 'Dynamic programming for optimal integration of Plug-in Hybrid Electric Vehicles (PHEVs) in residential electric grid areas', *IECON 2012 - 38th Annu. Conf. IEEE Ind. Electron. Soc.*, pp. 2942–2948, 25.
[98] C. Desbiens, 'Electric vehicle model for estimating distribution transformer load for normal and cold-load pickup conditions', in *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, 2012, pp. 1–6.
[99] J. Berardino, M. Muthalib, and C. O. Nwankpa, 'Network constrained economic demand dispatch of controllable building electric loads', in *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, 2012, pp. 1–6.
[100] A. Vukojevic, N. Parwal, and C. Yasko, 'Distribution transformer loading with deferred loads', in *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, 2012, pp. 1–7.
[101] C.-H. Lo and N. Ansari, 'Decentralized Controls and Communications for Autonomous Distribution Networks in Smart Grid', *IEEE Trans. Smart Grid*, vol. 4, no. 1, pp. 66–77, 2013.
[102] H. A. Toersche, V. Bakker, A. Molderink, S. Nykamp, J. L. Hurink, and G. J. M. Smit, 'Controlling the heating mode of heat pumps with the TRIANA three step methodology', in *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, 2012, pp. 1–7.
[103] A. Molderink, V. Bakker, M. G. C. Bosman, J. L. Hurink, and G. J. M. Smit, 'Domestic energy management methodology for optimizing efficiency in Smart Grids', 2009, pp. 1–7.

[104] A. Molderink, V. Bakker, M. G. C. Bosman, J. L. Hurink, and G. J. M. Smit, 'Management and Control of Domestic Smart Grid Technology', *IEEE Trans. Smart Grid*, vol. 1, no. 2, pp. 109–119, Sep. 2010.
[105] A. Ricci, B. Vinerba, E. Smargiassi, I. De Munari, V. Aisa, and P. Ciampolini, 'Power-Grid Load Balancing by Using Smart Home Appliances', in *Consumer Electronics, 2008. ICCE 2008. Digest of Technical Papers. International Conference on*, Las Vegas, NV, 2008, pp. 1–2.
[106] A. Ricci, E. Smargiassi, D. Mancini, I. De Munari, V. Aisa, and P. Ciampolini, 'Wr@p: A "last-meter" technology for energy-aware networked appliances', in *2011 IEEE International Symposium on Power Line Communications and Its Applications (ISPLC)*, 2011, pp. 193–198.
[107] M. Pipattanasomporn, M. Kuzlu, and S. Rahman, 'Demand response implementation in a home area network: A conceptual hardware architecture', in *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, 2012, pp. 1–8.
[108] T. L. Vandoorn, W. Willems, J. D. M. De Kooning, J. Van de Vyver, and L. Vandevelde, 'Contribution of a smart transformer in the local primary control of a microgrid', in *Innovative Smart Grid Technologies Europe (ISGT EUROPE), 2013 4th IEEE/PES*, 2013, pp. 1–5.
[109] F. C. Schweppe, 'Frequency adaptive, power-energy re-scheduler', 4317049, 23-Feb-1982.
[110] Ministry of Commerce, New Zealand, 'NZECP 36:1993 NEW ZEALAND ELECTRICAL CODE OF PRACTICE for HARMONIC LEVELS'. 18-Mar-1993.
[111] D. C. Lineweber, 'Understanding Residential Customer Support for – and Opposition to – Smart Grid Investments', *Electr. J.*, vol. 24, no. 8, pp. 92–100, Oct. 2011.
[112] G. A. Covic and J. T. Boys, 'Inductive Power Transfer', *Proc. IEEE*, vol. 101, no. 6, pp. 1276–1289, Jun. 2013.
[113] J. R. Lee, J. T. Boys, and G. A. Covic, 'Communications requirements and performance of a distributed demand control system', in *IECON 2011 - 37th Annual Conference on IEEE Industrial Electronics Society*, 2011, pp. 2711–2716.
[114] J. T. Boys and J. R. Lee, 'Power quality with green energy, DDC, and inductively powered EV's', in *Telecommunications Energy Conference (INTELEC), 2011 IEEE 33rd International*, 2011, pp. 1–8.
[115] J. R. Lee, J. T. Boys, and G. A. Covic, 'Transient Response of a Localised Demand Control System', in *2nd International Conference on Renewable Energy Research and Applications*, Madrid, Spain, 2013, pp. 1–5.
[116] J. R. Lee, J. T. Boys, and G. A. Covic, 'Improved Grid Dynamics using a Localised Demand Control System', *IEEE Trans. Smart Grid*, vol. In Press, Accepted Manuscript, 2014.
[117] J. T. Boys and G. A. Covic, 'Local demand side power management for electric utility networks', WO2012138235 A2, 11-Oct-2012.
[118] G. A. Covic, J. T. Boys, and J. R. Lee, 'Local demand side power management for electric utility networks', WO2014038966 A1, 13-Mar-2014.