Location Awareness in Multi-Agent Control of Distributed Energy Resources

HARRIET HUTCHINSON

How to cite:

HUTCHINSON, HARRIET (2017) Location Awareness in Multi-Agent Control of Distributed Energy Resources. Doctoral thesis, Durham University.

Location Awareness in Multi-Agent Control of Distributed Energy Resources

Harriet Hutchinson

Thesis submitted towards the
degree of Doctor of Philosophy

School of Engineering and Computing Sciences
Durham University
United Kingdom

April 2017

Location Awareness in Multi-Agent Control of Distributed Energy Resources

Harriet Hutchinson

Abstract

The integration of Distributed Energy Resource (DER) technologies such as heat pumps, electric vehicles and small-scale generation into the electricity grid at the household level is limited by technical constraints. This work argues that location is an important aspect for the control and integration of DER and that network topology can inferred without the use of a centralised network model. It addresses DER integration challenges by presenting a novel approach that uses a decentralised multi-agent system where equipment controllers learn and use their location within the low-voltage section of the power system.

Models of electrical networks exhibiting technical constraints were developed. Through theoretical analysis and real network data collection, various sources of location data were identified and new geographical and electrical techniques were developed for deriving network topology using Global Positioning System (GPS) and 24-hour voltage logs. The multi-agent system paradigm and societal structures were examined as an approach to a multi-stakeholder domain and congregations were used as an aid to decentralisation in a non-hierarchical, non-market-based approach. Through formal description of the agent attitude INTEND_2, the novel technique of Intention Transfer was applied to an agent congregation to provide an opt-in, collaborative system.

Test facilities for multi-agent systems were developed and culminated in a new embedded controller test platform that integrated a real-time dynamic electrical network simulator to provide a full-feedback system integrated with control hardware. Finally, a multi-agent control system was developed and implemented that used location data in providing demand-side response to a voltage excursion, with the goals of improving power quality, reducing generator disconnections, and deferring network reinforcement.

The resulting communicating and self-organising energy agent community, as demonstrated on a unique hardware-in-the-loop platform, provides an application model and test facility to inspire agent-based, location-aware smart grid applications across the power systems domain.

Declaration

The work in this thesis is based on research carried out in the New and Renewable Energy Group, School of Engineering and Computing Sciences, Durham University. No part of this report has been submitted elsewhere for any other degree or qualification and it is all my own work unless referenced to the contrary in the text.

Parts of this work have been published in the following:

G Hutchinson, P C Taylor, G Coates, R Hair, "Multi-Agent System for Decentralised Control of Low Voltage Distribution Networks", *Universities Power Engineering Conference*, 2010.

G Hutchinson, R Hair, P C Taylor, "Leveraging Location Awareness for Distributed Energy Resources: Self-Organising Energy Communities", *International Conference and Exhibition on Electricity Distribution (CIRED)*, 2013.

Acknowledgements

This work was funded by EPSRC and E.ON. It would have not been possible without the support of supervisiors, technicians, colleagues and friends.

To my parents, to Phil, and to Margaret, without whom I would never have made it here. To all my loved ones, with thanks for their enduring patience, support and confidence.

Harriet Hutchinson
Durham, April 2017

Contents

List of Figures

List of Tables

Glossary

μCHP	micro-Combined Heat and Power.
AC	Alternating Current.
ACL	Agent Communication Language.
ADC	Analogue-to-Digital Converter.
ADMD	Average Daily Maximum Demand.
AMS	Agent Management System.
API	Application Program Interface.
BDI	Beliefs, Desires, Intentions.
CHP	Combined Heat and Power.
CLNR	Customer-Led Network Revolution.
CoP	Coefficient of Performance.
CPS	Cooperative Problem Solving.
CRES	Centre for Renewable Energy Sources and Saving.
DC	Direct Current.
DER	Distributed Energy Resource.
DF	Directory Facilitator.
DG	Distributed Generation.
DHCP	Dynamic Host Configuration Protocol.
DHW	Domestic Hot Water.
DNO	Distribution Network Operator.
DPS	Distributed Problem Solving.
DSM	Demand-Side Management.
DTI	Department for Trade and Industry.
EEGI	European Electricity Grid Initiative.
EMF	Electromotive Force.
EMTDC	ElectroMagnetic Transients including DC.
EMTP	ElectroMagnetic Transients Program.
EV	Electric Vehicle.
ExSSEZ	Experimental Small Scale Energy Zone.
FFT	Fast Fourier Transform.
FIPA	Foundation for Intelligent Physical Agents.
FIPA-ACL	FIPA Agent Communication Language.
FIPA-SL	FIPA Semantic Language.

FIR	Finite Impulse Response.
FIT	Feed-In Tariff.
GA	Generator Agent.
GIS	Geographical Information System.
GPIO	General Purpose Input/Output.
GPS	Global Positioning System.
HP	Heat Pump.
HV	high voltage.
I/O	Input/Output.
ICT	Information and Communication Technology.
ISP	Internet Service Provider.
JADE	Java Agent DEvelopment framework.
LA	Load Agent.
LSB	Least Significant Bit.
LV	low voltage.
LVDN	Low-Voltage Distribution Network.
MAS	Multi-Agent System.
MSW	Municipal Solid Waste.
MV	medium voltage.
OFGEM	Office of Gas and Electricity Markets.
OLTC	On-Load Tap Changer.
OS	Operating System.
PDF	Probability Density Function.
PHEV	Plug-in Hybrid Electric Vehicle.
PLC	Power Line Communication.
PV	photovoltaic.
RTDS	Real-Time Digital Simulator.
SA	Simulation Agent.
SSEG	Small Scale Energy Generator.
SSEZ	Small Scale Energy Zone.
SVC	Static VAR Compensator.
SVR	Static Voltage Regulator.
TLA	Thermal Limits Agent.
ToU	Time of Use.
UKGN	UK Generic Network.
V2G	Vehicle-to-Grid.
VPP	Virtual Power Plant.
VUF	Voltage Unbalance Factor.

Nomenclature

μ	Mean
σ	Standard deviation
τ	Phase shift
$\mathbf{I}$	Current, vector
$\mathbf{V}$	Voltage, vector
$\mathbf{Z}$	Impedance, vector
D	Decision
H	Hypothesis
I	Current, peak
i	Current, instantaneous
L	Inductance
P	Real power
Q	Reactive power
R	Resistance
S	Apparent power
V	Voltage, peak
v	Voltage, instantaneous
X	Reactance

Chapter 1

Introduction

The integration of technologies such as heat pumps, electric vehicles and small-scale generation into the electricity grid at the household level is limited by technical constraints. This work addresses this significant challenge in power systems by presenting a novel approach that uses a decentralised multi-agent system where the controllers learn and use the location of equipment within the low-voltage section of the power system. The benefits include improved power quality, better environmental and economic performance of small-scale generation, and a delay to network reinforcement of legacy electrical networks. The solution is implemented on a unique new test facility that combines real-time network simulation with embedded controllers to provide a electrical network/hardware/software-integrated multi-agent platform.

The embedded controller platform developed for this project has enabled multi-agent system testing of a type that was not previously feasible: it combines real-time simulation of arbitrary electrical networks with any geographical/electrical arrangement of equipment, and low-cost, general-purpose embedded controller hardware, running an industry-standard multi-agent system platform. This facility is a significant asset for research into fuller exploitation of agent capabilities in electrical power systems.

1.1 Structure of this Chapter

The research domain and background to this work are introduced in this chapter. The challenges presented by existing infrastructure are explained and the new technologies that create the problem are described. By examining both the technological and regulatory environment, the significance of the problem is illustrated as a driver to push forward from the existing work in the area.

The first section contains a description of the power system domain and how this work is a continuation of efforts into integrating new technologies. In the second section, the policy landscape that makes the research both timely and important is described. The low-voltage distribution network, as part of the power system, is described in more detail in the third section. Individual technologies (distributed

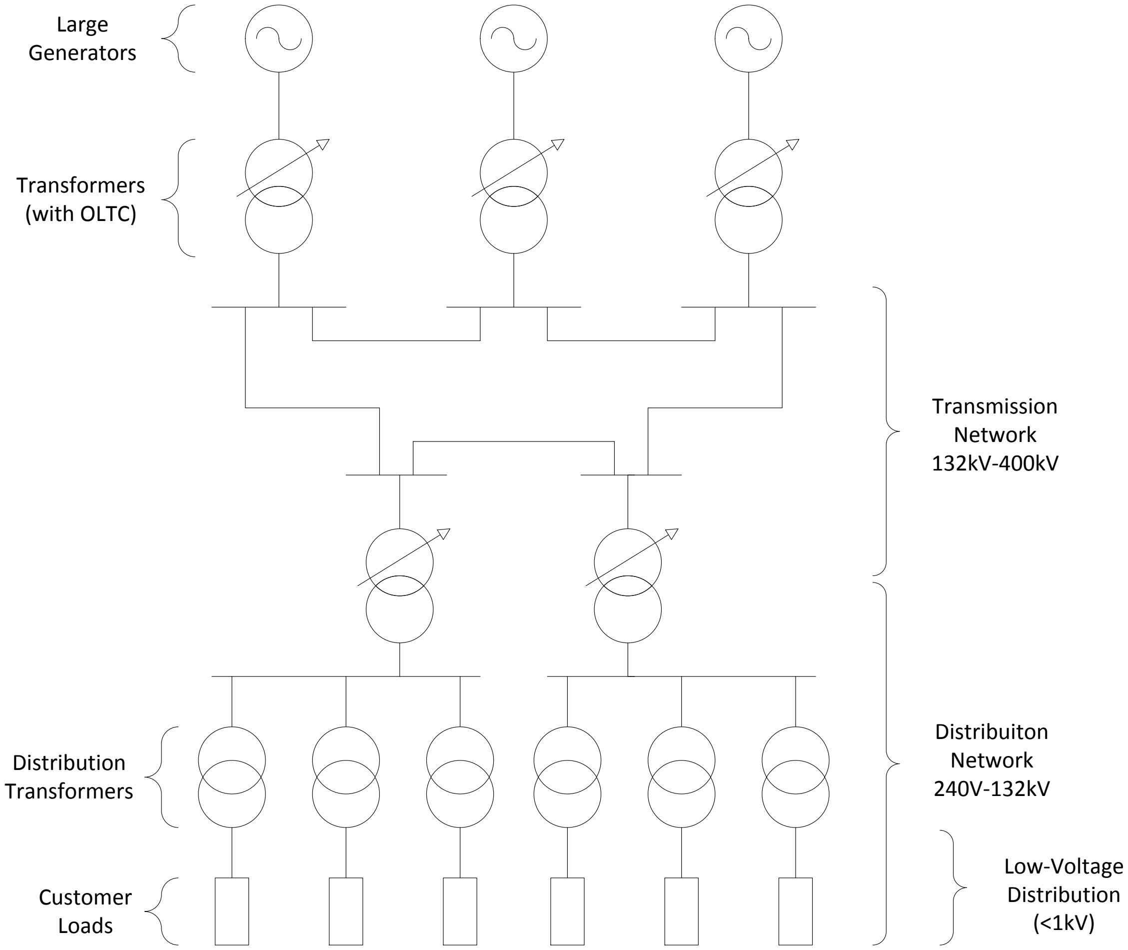

Figure 1.1: Structure of the UK electricity system: the diagram shows large, centralised generators at the top. Power then flows down through the transmission network. The bottom half shows the distribution network. Voltage levels are stepped down by transformers as the lines carry progressively less power to customer loads. Diagram adapted from [1].

generation, electric vehicles, heat pumps, demand-side management) are introduced with their role in creating the research challenge. An overview of the solution is provided in the fourth and fifth sections. The structure of this thesis is outlined in the final section of this chapter.

1.2 Research Background

The organisation of the system that delivers electricity from large-scale power plants to individual consumers is illustrated in Figure 1.1. The long-distance, high voltage (HV) (400kV to 132kV) *transmission* system delivers power away from centres of power generation. *Distribution* refers to the part of the network that delivers power from the transmission system to individual customers at voltages of 132kV and below. The low voltage (LV) section of the system is at the bottom of the diagram, between 230V single-phase (or 400V 3-phase) and 1kV in the UK.

Figure 1.1 shows all of the generation at the top of the diagram — conventionally, these are large thermal power stations: coal, gas or nuclear plants. However, smaller Distributed Generation (DG) (from a few hundred watts up to several MW) may be connected at the distribution level. These include backup diesel generators, Municipal Solid Waste (MSW) incinerators, large biomass plants, individual wind turbines or farms, some gas turbines and, more recently, large solar photovoltaic (PV) farms. Small Scale Energy Generators (SSEGs) of under 100kW installed capacity (rooftop PV, micro-wind, small biomass and combined heat and power, see below) may be connected to LV distribution networks at 230V/400V.

The current paradigm for Low-Voltage Distribution Networks (LVDNs) is one of *passive* control, where networks are designed to operate continuously without monitoring or adjustment. In contrast, measurement, online analysis and control devices in *active* networks allow characteristics such as transformer tap settings, equipment ratings or even network topology to be altered dynamically. This can be used as an alternative to network reinforcement or additional generation where capacity is reached; at the planning stage to reduce required infrastructure; or to improve the economic or technical operation of existing equipment.

Whilst active control techniques have gained acceptance in high-voltage transmission and distribution systems, their implementation in LVDNs is less widespread and usually confined to specially-designed remote or island power systems or demonstration "microgrids" (see Section 2.2.3). Active control has been increasingly viewed as a useful tool to solve network issues in distribution networks for some time: the UK DTI first recommended a move to active approaches in 2004[2]; since then some implementation has followed but active control remains a significant research area: the results from the Low Carbon Networks Fund projects included active techniques appropriate down to low-voltage networks, and the 2016 HubNet analysis of these projects recommended Active Network Management (ANM) for all DG connections[3].

Widespread integration of DG, Electric Vehicle (EV) and other technologies into LVDNs poses several challenges. Combining these technologies with active control and demand-side management techniques (see Section 1.4.5) renders them into *Distributed Energy Resources (DERs)*. DERs are resources at the domestic level that can be used in managing the electric network as well as for their consumer purpose. The liberalised electricity system and the regulatory framework in the UK add non-technical facets to the research landscape. The overall target for research is to propose and demonstrate the effectiveness of an active control approach at the LV level that can mitigate the technical issues to integration within the UK context.

This work forms a continuation of research into control of DG in LVDNs undertaken by Durham University, focusing on increasing the environmental, commercial and technical benefits of small-scale generation in evolving distribution networks. It is supported and funded by EPSRC and E.ON.

1.3 Regulatory Framework and Non-Technical Context

1.3.1 UK Regulations

The UK electricity system is "liberalised": that is, since the Electricity Act (1989)[4], it has been operating as a mixture of competing and collaborating stakeholders. Generation plant, transmission networks and distribution networks are owned and operated by separate legal entities, with electricity then retailed to customers by companies who compete on price by trading within wholesale energy futures and settlement markets.

The 1989 Act was supplemented by the Energy Act (2008)[5], which added the Feed-In Tariff (FIT), and the Energy Act (2013) [6], which added a Capacity Market (CM) and Contracts for Difference (CfD), and these schemes run alongside the wholesale market[7]. Further details on these market structures and their implications are included in Section 1.3.3.

Large energy companies may incorporate supply, distribution and retail roles as subsidiaries; these *vertically integrated* organisations are regulated to ensure they operate as several competing, independent businesses.

In the UK, the regulator is the Office of the Gas and Electricity Markets (OFGEM), which grants regional monopoly licenses to network operators. It promotes competition and ensures that network operators adhere to the Distribution Code[8].

In order to facilitate growth in distributed generation, Engineering Recommendation G83/1[9] was introduced as part of the Distribution Code. These guidelines for both users and operators of a distribution network create simple, favourable conditions for owners of small generators. Currently, DG which outputs at less than 16A per phase may be connected to existing LV networks following the "fit-and-inform" principle in which there is no centralised planning of how or where DG is connected by consumers. The procedure outlined in [9] specifies how a generator may be connected and describes limits on operating conditions (e.g. mandatory disconnection after 1600ms of a low-voltage event) and a Distribution Network Operator (DNO) cannot refuse the connection without a good technical reason. Currently, the regulations do not permit small generators to be assessed as making any contribution to system security.

1.3.2 The Broken Value Chain

The requirements of the multiple actors in the electricity network are quite different and sometimes in conflict: generation plant owners may be seeking revenue from an investment and hence maximise output at all times; Combined Heat and Power (CHP) owners may only wish to export when there is a large local heat demand; a DNO may wish to limit loads or energy sources to protect equipment and ensure stable system operation; and a trader will wish to exploit patterns in electricity consumption to improve their market performance. Generation plant owners may be able to provide services to stakeholders not represented in the current market structure. A workshop in November 2010 involving the Department for Energy and Climate Change, industry representatives and Durham Energy Institute described the lack of a clear and direct link as a barrier to the transition to a low carbon grid, summarising, “It is not clear who the ‘customer’ is for smart grid technologies”[10]. Even in 2016, a whitepaper by Policy Exchange levelled significant criticisms at this disconnect. The policy thinktank directly called for simplification of the ancillary services markets, and commented,

> “in order to further decarbonise the power system, it will need to become smarter and more flexible.... Overall, it is clear that cleaner forms of flexibility such as demand response and storage face a number of policy and regulatory barriers.”[11]

There is a continued “broken value chain” between grid impacts, services, and stakeholders, which adds complexity into any attempt to fully exploit the integration of DERs in the UK.

1.3.3 Market Drivers

The 2013 Energy Act added the Capacity Market (CM) and the Contract for Difference (CfD) to the wholesale market[6]. In a CM, ‘firm’ energy suppliers receive a payment for making energy available[1]. This allows a system that includes significant intermittent renewable energy: if intermittent sources do not deliver and there is a shortage of energy, the firm suppliers offered contracts through the CM are required to make up the shortfall. The CM payment is only for *availability* of generation— it does not pay for the energy supply itself.

The CM is paired with Contracts for Difference. This incentivises intermittent generation by guaranteeing a price for energy supplied. If the wholesale market energy price drops below the guaranteed ‘strike’ price, the scheme tops up the payment; if if the market price is above the strike price, the generator must pay back the difference. The contracts are awarded competitively as the strike price is determined by sealed-bid auction to determine true market prices. At the same time as reducing the risk in investment in intermittent generation, the budget allocated for CfDs is allocated in order to keep costs within the budget of the Levy Control Framework, a policy for managing government spending on renewables and decarbonisation targets[12]. The strike price is adjusted by technology to reflect the differences in cost of competing low-carbon generation.

Onifade[7] offers insight into the effectiveness of the CM + CfD policy instrument for incentivising the uptake of renewables. However, these two support mechanisms apply mainly to larger facilities and backup generators. At the small scale (and in particular at the domestic scale), the market driver is still the Feed-in Tariff for generation below 5MW in capacity.

---

[1] In this case, the suppliers include conventional generators and also “proven” and “unproven” demand-side response; storage technologies; and interconnectors that can inject power into the UK grid.

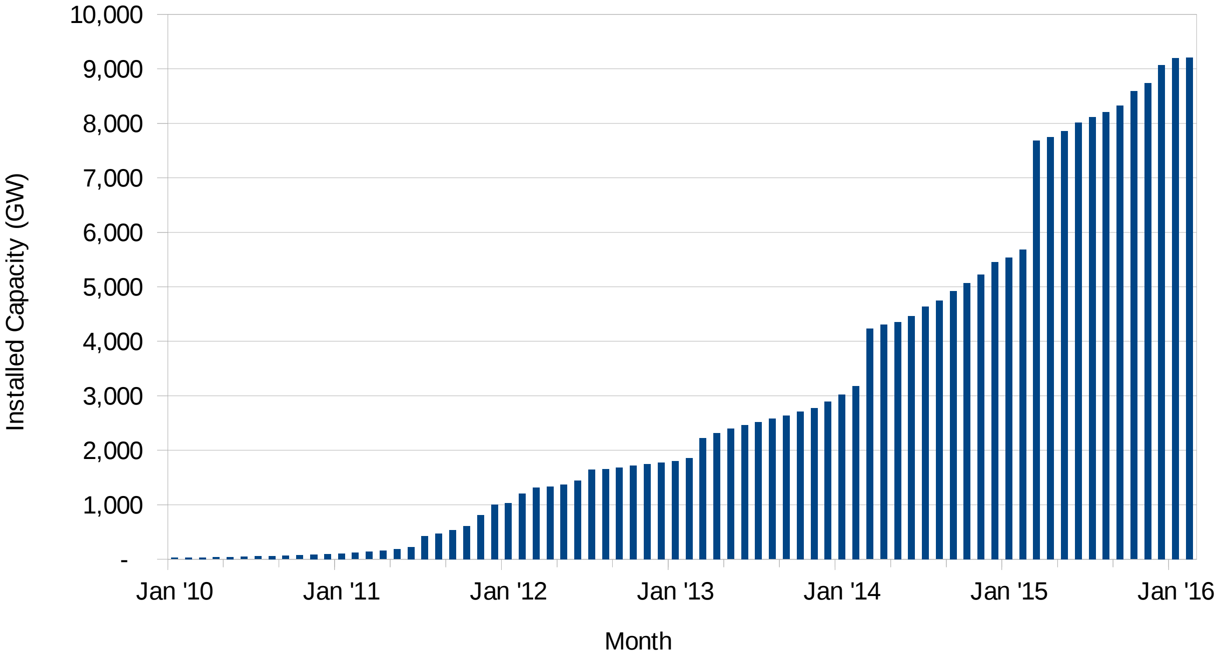

Figure 1.2: Growth in UK photovoltaic generation
The graph shows the cumulative capacity of PV installations since 2010. Data from[14].

After the passing of the 2008 Energy Act, Feed-In Tariffs (FITs) were implemented in March 2010[13]. These guarantee a fixed electricity price to the owners of microgeneration technologies. The price is composed of two parts: the *generation tariff* subsidises generation whether the energy is used locally or not; there is an additional *export tariff* payment if the electricity is exported to the grid. The generation tariff is dependent upon the type of technology.

Initially, PV was more expensive to install and attracted a higher price than micro-Combined Heat and Power (μCHP); the generation rate for PV (43p/kWh) drove significant investment in new installations, at the same time as its capital costs reduced significantly. PV gave a rate of return in excess of many conventional financial investments. As well as consumers claiming the subsidy, many companies started to rent roofs (amongst other agreements) for installations. The feed-in tariff has been a significant driver to increasing microgeneration in the UK. As shown in Figure 1.2, the UK installed capacity is over 9GW in 2016, up from 0.015GW in 2010.

However, the longer-term rate of PV uptake may reduce considerably in the due to changes in FIT policy, in keeping with Onifade's observations on uncertainty caused by governmental changes in FIT schemes[7]. The export tariff was recently cut to 4.91p/kWh for all technologies; the PV generation tariff dropped to 4.32p/kWh and theμCHP generation tariff was set to the comparatively higher rate of 13.45p/kWh from April 2016[15]. Some uncertaintainty is reduced by the projection of tariff rates to March 2019.

As of August 2016, National Grid has concluded a tendering process for new frequency response services[16]. These can be provided by large storage facilities, but also through demand response and aggregation. Whilst aggregators were able to bid in this process, none were successful in this round; in the future, DERs may play a role in this market.

In addition to the increasing cost of petrol creating market conditions for cost-effective EV use, UK government subsidies directly incentivise their uptake. More than 51,000 EVs were bought with government support between up to December 2015[17]. The subsidy scheme has been extended to March 2017, with support for EVs worth £4500 per vehicle from March 2016[18].

1.4 Low Voltage Distribution Networks

1.4.1 Characteristics of LVDNs

Figure 1.1 shows power flow from centralised generation through the transmission system to loads within the distribution system. The low voltage (LV) distribution system consists of the lower-voltage part of the network, taking power to individual customers connected at 230V/400V (single/three-phase).

Household electrical power demand varies between users, depending upon precisely when particular appliances are used. However, the average demand can be assessed by taking an average across multiple households. Figure 1.3 shows this average demand profile for the UK over a 24-hour period.

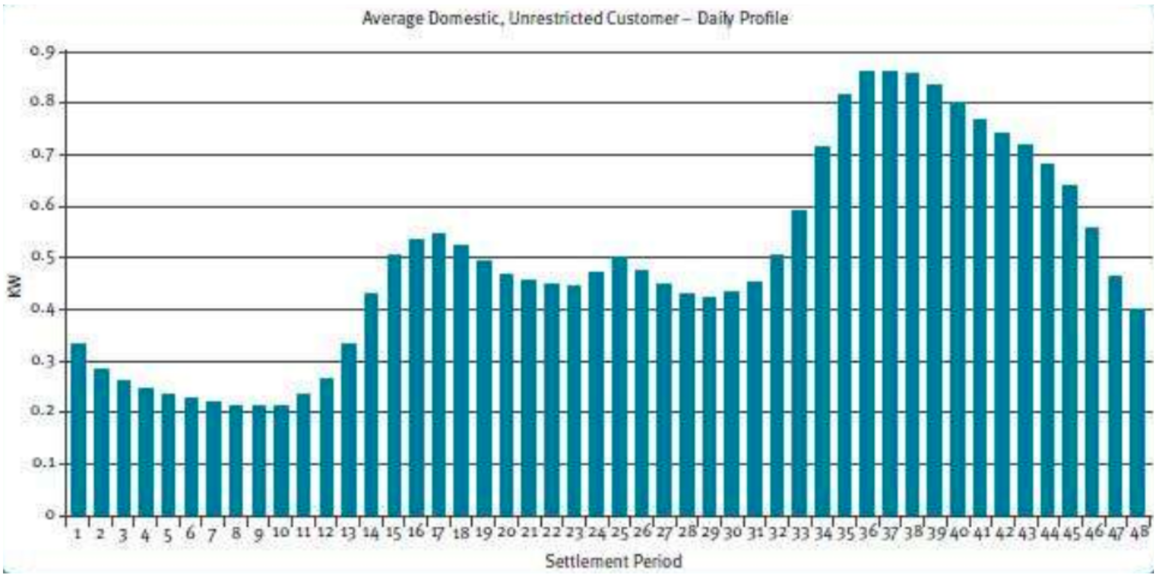

Figure 1.3: Daily demand profile of a house
The diagram shows how the power demand of a UK household varies with time. The diagram illustrates a typical unrestricted (i.e. not Economy 7) customer, and is one of 8 domestic profiles used by suppliers to estimate demand. The data are averaged to show the after-diversity demand (in kW) over a 24-hour period in half-hourly intervals. Image from [19].

The figure shows low demand during early morning hours, followed by an increase when people wake. There is a small peak around lunchtime, then another peak after 6pm, then the demand drops as people sleep. Demand is season-dependent. There is considerable variation in power use — demand on summer nights may be less than a quarter of the winter peak.

The demand profile is an average across many users and not all customers will be using appliances in the same way at the same time; the coincidence curve in Figure 1.4 shows the decreasing ratio of cumulative peak demand and the likely peak demand as the number of customers (and hence *load diversity*) increases. This probabilistic method allows for specification of infrastructure that is below the cumulative load peak whilst remaining cost-effective. The exact shape of the demand and diversity curves is determined by the profile of the users on a particular network segment.

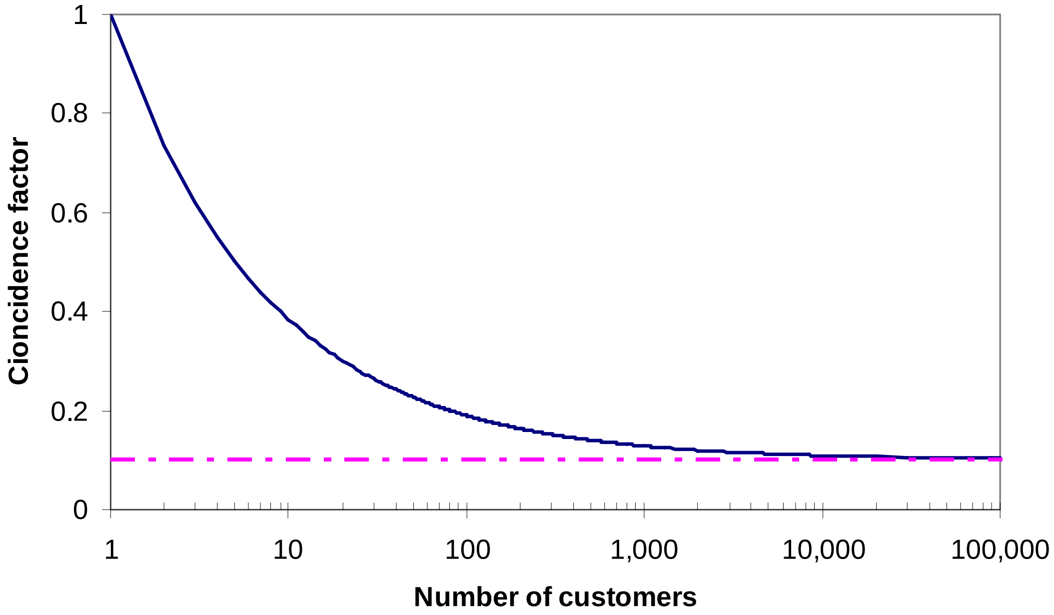

Figure 1.4: Coincidence curve
The ratio of cumulative peak demand to probable peak demand reduces as the number of users increases[20].

Conventional LV network planning takes into account growth in household use and number of customers on a feeder; however, many networks have been planned without taking into account the new, significant requirements for technologies such as microgeneration and electric cars, both of which will have significant impacts on existing systems. These technologies and impacts are explored in Sections 1.4.2-1.4.5 and in Chapter 2 in Section 2.2.1.

1.4.2 Distributed Electricity Generation

The role of small-scale DG has developed significantly over the past two decades. Diesel and gas generators and renewable technologies such as photovoltaic cells can be used to supplement a weak grid or provide a replacement for it in remote locations. DG has grown in importance as it is deployed for new reasons: as well as having uses in securing reliability of supply and high quality power, it also forms part of the UK strategy to provide low-carbon electricity using renewable energy sources and local CHP production. The generation of electricity near its point of use reduces losses in transmission and distribution networks and can provide cheaper energy than that bought from a supply company. Wind, CHP, biomass, anaerobic digestion/gas turbine systems and photovoltaic generation are all connected at the distribution level. In particular, large wind turbines and farms can be connected to the distribution network if there is sufficient capacity for the extra power. Wind turbines are connected across the complete range of voltage levels in the distribution network, although micro-wind has not been extensively deployed at domestic voltages. The payback period (both economically and environmentally) for micro-wind is typically unfavourable due to the combination of installation costs compared to capacity and poor wind characteristics in urban contexts. Power system studies are required to assess the impacts of connecting larger DG, but low-power DG is not examined in the same way and has the potential to be disruptive (see Section 1.3.1). PV and μCHP can be connected at low voltage at the household scale and have the potential to be deployed extensively.

1.4.2.1 Combined Heat and Power

Combined Heat and Power (CHP) generation uses a thermal energy source (such as gas or biomass) to supply both electrical and heat demands, improving efficiency of power generation. In 2003, the UK Department for Trade and Industry (DTI) set out plans for 10GW of installed capacity of distributed CHP by 2010 [21], although by 2009 this had only resulted in a modest increase from 4.8GW to 5.5GW[22]. In 2011, the government recommitted to increasing CHP capacity, saying: “The Government will continue to promote the development of good quality CHP in the UK”[23]. It subsequently legislated for this to be the policy and strategy under the Energy Act (2013)[6].

The proportion of non-domestic small-scale CHP (of under 100kW electrical power) remains low, at only 454 installations constituting 28MW. Between 2011 and 2014, 477 μCHP units — below 2kW capacity each, a total of less than 1MW — were installed in the domestic sector in the UK[24]. CHP uptake is driven by a combination of reduced energy costs and increased efficiency, with the impact of reducing carbon emissions. There are market benefits since μCHP is included in the feed-in tariff scheme described in Section 1.3.2. The operation of CHP plant is typically driven by a local heat demand, so benefits of increased efficiency are often dependent upon a coinciding electrical load.

1.4.2.2 Photovoltaics

Photovoltaic (PV) panels can be fitted to roofs to supplement a building’s energy needs. The use of PV has expanded rapidly in the UK due to the introduction of FITs for generation and export (see Section 1.3.2). Householders installing a new PV array are guaranteed a subsidy for any energy production. As well as for reducing electricity bills (and electricity-related CO_2 emissions), consumers and others are using PV arrays as stable investments. FITs have resulted in the formation of companies willing to buy or rent roof space for PV arrays in order to take advantage of the profitable subsidies. Figure 1.2 shows the recent rapid increase in UK PV capacity. Sustained growth of this magnitude will require interventions to ensure safe electrical network operation; the implications of this are discussed in more detail in Section 2.2.1.

Daylight hours — and the power output of PV — may not match consumer use patterns, particularly in cases where a PV array is fitted to a house where its occupants are absent during the day. Arrays may be oversized for a building in order to take advantage of the FIT incentive, so there may be a considerable proportion of energy exported. In addition to the challenges associated with more predictable distributed generation, the electricity network must tolerate the intermittent nature of solar energy which can result in power flow swings and harmonic distortion.

1.4.3 Electric Vehicles

The transport sector produces 22% of UK annual greenhouse gas emissions[25] and road transport is the largest component of this. Personal transport in petrol/diesel cars forms the majority of fuel consumption and is a major target for decarbonisation. A move away from combustion engines to EVs and Plug-in Hybrid Electric Vehicles (PHEVs) has been part of the UK low carbon strategy since 2009[26]. This includes government subsidies and investment in infrastructure. EVs need to be combined with a low-carbon electricity supply in order to have a beneficial impact on CO_2 emissions. The additional generating capacity required for this transition (including personal transport and light goods vehicles) could be as high as 40GW by 2050[25].

EVs could have a role as active network components, rather than just passively consuming electricity until fully charged. Their batteries can be used as energy storage devices for supporting network operation. This approach where a car battery supplies energy, called Vehicle-to-Grid (V2G), allows aggregation of large numbers of vehicles to provide ancillary services, for example, for frequency or voltage support by supplying stored energy during periods of insufficient generation. This can help the integration of intermittent renewable sources[27]. In this case, market-driven or specifically tailored pricing structures could result in a vehicle charge controller buying electricity at cheap rates and selling it again at peak times, informally known as “carbitrage”. At present, there is a lack of clear structure for settlement with vehicle owners. Consumers may be reluctant to participate where this may reduce their utility of having a car available on demand or at the cost of battery life reduction.

However, until these issues are addressed, they remain useful in other ways: they can be used to smooth electricity demand profiles by adjusting charging power or delaying or advancing charging cycles over time. Aside from the 40GW+ increase in the generation required to power large numbers of electric cars, the classic network problem is one of capacity and use patterns: a large number of users return to their homes after work and plug in their cars to charge, creating a large number of high loads that draw power for a long time. As the literature in Section 2.2.1 will demonstrate, with extensive EV use, the combination of reduced diversity from typical load patterns and the large current draw for these devices will be sufficient to overload the existing infrastructure[1]; these network issues are discussed in more detail in Section 2.2.1.

1.4.4 Heat Pumps

Heat pumps are important to the decarbonisation of energy for space and Domestic Hot Water (DHW) heating. The Coefficient of Performance (CoP) of a heat pump determines the efficiency with which energy from an external source is converted to useful heat. The UK Energy Mix is described in [28], but including transmission and distribution losses the average UK grid carbon intensity was around 367gCO_2/kWh in 2015. [29]. Consequently, a Heat Pump (HP) with a CoP of around 3 produces heat at 122gCO_2/kWh, which is comparable to efficient condensing gas boilers; as the UK energy mix becomes greener, this comparison will only become more favourable. This makes it a low-carbon technology eligible for the Renewable Heat Incentive, a government subsidy scheme.

Similarly to EVs, a HP is a large load with a long cycle. The network impacts are consequently similar in terms of thermal overloading and voltage drop. Consequently, electrical network reinforcement will be required[30] if HPs are to contribute significantly to the UK's decarbonisation efforts. However, the electrification of heating provides opportunities for Demand-Side Management (DSM). Existing static Time of Use (ToU) tariffs encourage owners of immersion heaters and electric space heating to charge overnight. However, models have shown that thermal storage (for space heating or via a tank for DHW) can allow staggered startup to mitigate the effects of large loads operating simultaneously and delay network reinforcement (see [31]). The Office of Gas and Electricity Markets (OFGEM)-funded Customer-Led Network Revolution (CLNR) project[32] has demonstrated HPs with thermal storage in large field trials with successful load-shifting for voltage support.

1.4.5 Demand-Side Management

Demand-Side Management (DSM) encompasses a range of techniques in network operation that alter customer use of a system, rather than assets owned by a network operator. This comprises both technical and behavioural interventions: for example, a DNO may send a signal to a customer instructing an appliance to switch off (*demand response*), or electricity pricing structures may discourage system use at particular times of day. As well as instructing or incentivising the use of customer loads in a way that is beneficial to the network, microgeneration may also be regarded as "demand-side" for control purposes; small generators are modelled as a negative load on the customer's side of the network.

1.4.5.1 Benefits of DSM

OFGEM calculated that capability to load-shift 10% of UK customer electricity would avoid capital costs of £536million for new generation and £28million for network reinforcement annually[33], as well as reducing wholesale electricity costs. Strbac [20] outlines the advantages of adoption of DSM as:

- Reduction in generating capacity margin.

  Generation capacity in the region of 20% above peak consumption is regarded as sufficient headroom for system security. However, the cost of plant for capacity margin purposes may be £250-£400/kW or more; DSM may provide a more cost-effective alternative to standby generation.

- Improving effectiveness and economic efficiency of distribution networks.

  Figure 1.5 shows the intended network impacts of DSM for load shaping, where DSM techniques postpone or advance electricity consumption. Improved load profiles benefit a DNO through delayed network reinforcement, lower equipment ratings and reduced distribution losses. In addition to increasing average utilisation of assets, DSM can be used in outage management and for power quality improvements. Whilst Strbac mentions improved electrical efficiency and carbon reduction, DSM does not itself significantly reduce electricity use; Shaw *et al.* demonstrated that load-shifting is likely to reduce consumption by 0.02% since a flatter load profile has lower overall I^2R losses[35]. This is not significant compared to other system losses, leading to the conclusion that DSM is of higher value for peak reduction, rather than directly improving network efficiency. Shaw draws no conclusions about the value of matching local demand and generation for reducing distribution system losses, but Strbac notes that the potential 25-40% reduction in distribution network losses from widespread use of μCHP is dependent upon matching generation with peak load[36].

- Balancing supply and demand, particularly in systems with intermittent renewable energy sources.

  The installation of DG diversifies supply as well as demand, and increases the complexity of system balancing. As well as reducing demand when generation is not available, the use of DSM has evolved to include carbon considerations. Matching demand to renewable power production can improve

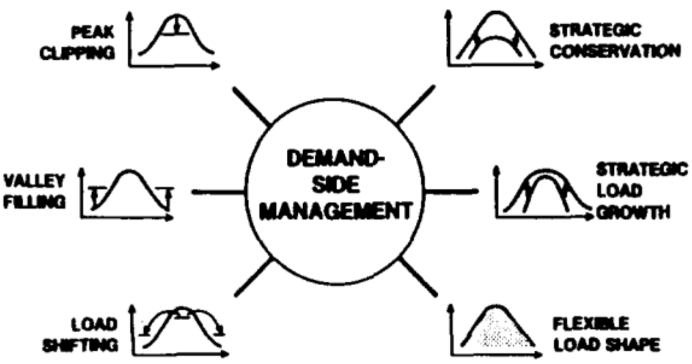

Figure 1.5: Load-shaping impacts of demand-side management. Each sub-graph shows a load peak and a new shape that can result from DSM. Figure reproduced from [34].

carbon efficiency; time-shifting of a load so intermittent supply can replace conventional fossil generation significantly reduces per-MW CO_2 emissions. DSM can be used to increase flexibility, such that intermittent sources are more fully exploited when available (including heat-demand-driven CHP as well as weather-dependent wind power), and less energy is consumed at other times. A direct benefit of this is that unfirm connections to the distribution network (i.e. where generators may be required to curtail output due to network conditions such as low demand or high production) can use DSM to reduce “spillage” of excess power.

- Improving efficiency of investment in transmission infrastructure.

  Currently, transmission assets typically only operate at 50% of capacity in order to accommodate the headroom required by demand peaks. As is the case with distribution, these peaks can be reduced through DSM, allowing a higher average utilisation.

It should be noted that DSM may have significant drawbacks, such as reducing load diversity. Load-shifting concentrates use into shorter time slots, so the proportion of devices in use at any one time increases after DSM interventions.

Customer engagement is vital to the effectiveness of DSM schemes, from acceptance of direct control of customer devices through to encouraging and maintaining behavioural change through market factors. OFGEM notes that its predicted scenario of increased use of electric heating (of space and water) and electric cars is highly compatible with DSM techniques[33].

1.5 Conclusion

The sections above describe the power systems domain, and the new technologies that must be integrated into the low voltage (LV) electricity network. The introduction of widespread distributed generation is being driven by concerns over energy security, climate change, and commercial considerations of costs and profits. These factors are also drivers towards Heat Pump (HP), Electric Vehicle (EV) and Demand-Side Management (DSM) adoption. These technologies will create problems and opportunities where network limits begin to impose constraints. Previous researchers have examined the impact of distributed generation on LV networks and developed control approaches to mitigate these effects for network operators. The proceeding chapters show that these issues have not been solved in their entirety and illustrate the significant open research problem.

1.6 Solution Overview

In this work, the importance and impact of location in LVDNs are examined with respect to the challenges of integration described above. New techniques are developed that can be used to identify the location of devices within electrical networks. The use of decentralised multi-agent systems in control applications is explored. The following new contributions are made to the field of distribution network control:

1. New electrical network models, data and tools.
2. New work in describing and identifying location.
3. Design of a novel, decentralised Multi-Agent System (MAS) incorporating new MAS techniques.
4. Design and build of a novel test rig to validate a location-based MAS approach for control.
5. Implementation and validation of a location-aware MAS control system for a voltage support task.

A case is built for location-aware systems through modelling the electrical network topology including the new technologies described above. New techniques are provided to detect or infer location by exploring the links between electrical topology and network theory and non-electrical forms of location. The multi-stakeholder context of the problem necessitates new multi-agent approaches, which are examined and exploited beyond the current state of the art. Finally, the argument for such a location-aware, decentralised MAS control system is tested and validated through implementation on a new test rig that incorporates real-time, online hardware-in-the-loop control with a dynamic electrical network simulation.

1.7 Structure of this Document

- The electricity distribution system is introduced in Chapter 1 and the context given for the work undertaken.
- In Chapter 2, the literature is reviewed on the integration of distributed generation and other technologies, control techniques, and coordination approaches, with a focus on multi-agent systems. This review is used to clarify the research problem. Conclusions are drawn for the ideal properties of a proposed new solution.
- In Chapter 3, electrical network theory is provided to explain the principles behind the research problem and to examine electrical aspects of location. Existing benchmarks are used to develop new electrical network models and create tools for examining, solving and testing solutions.
- Chapter 4 contains new work that develops ideas of location in electrical networks. New techniques are presented using geographical and electrical information to describe and derive network topology in a key step towards a location-aware system.
- The problem is examined in the context of the Multi-Agent System paradigm in Section 5.1. Agent theory is used to develop a decentralised approach, including the new concept of Intention Transfer, for a novel control system.

- The implementation of the new Multi-Agent System is described in Section 5.2. Proof-of-concept systems were used to demonstrate initial application of the various location and MAS components before development of a new test rig to allow on-line, real-time, hardware-in-the-loop testing of the system. This system incorporates the electrical and geographical aspects of location from Chapter 4, and operates on an electrical network described in Chapter 3.

- Results of the application of the control system using the real-time test facility are presented in Section 5.3; the new test facility is evaluated and the solution assessed in terms of the research problem statement.

- Finally, Chapter 6 contains a summary of the work, including discussion of its significance and wider context and the direction of future research in the field is outlined.

Chapter 2

Literature Review

This chapter examines integration of small-scale generation, electric vehicles and other new technologies into electrical distribution networks. By examining existing literature in this area, it introduces the concepts of Microgrids and the Small-Scale Energy Zone. It identifies how multi-agent system programming approaches can be used for coordination and control of LVDNs as well as highlighting existing examples of this approach. Finally, it uses this review of the state of LV network control technology to outline a significant and original research problem.

2.1 Low-Voltage Distribution Networks

2.1.1 Legacy LVDNs

As outlined in Section 1.2, *distribution* refers to the part of the network below 132kV and the LV section is at 230V single-phase (or 400V 3-phase) to 11kV. In the UK, the electricity system is composed of distribution networks that are connected by high-voltage *transmission* infrastructure.

The UK power network has been growing since the commissioning of the first modern three-phase distribution system attached to Neptune Bank Power Station in 1901. The longevity of power systems components was anticipated; according to the network planning literature of the 1980s in the UK,

> “The various items of equipment installed in power systems have long useful lives with some items remaining in service for 40 or 50 years. Thus, proposals... should not just cover present loads but be capable of meeting or being reinforced to meet future loads” [37].

Certainly, this has been borne out by experience — for example, as of 2012, the distribution network serving London and the South East still includes thousands of power transformers commissioned in the 1950-1960s and even earlier. Whilst the equipment lifespan may still support continued operation, the 60-year-old assumptions around network use and growth from initial installation are no longer appropriate. Whilst still allowing for future growth, “plans to meet future demands are generally based on the assumption that load patterns will not change significantly”[37]. With the recent and significant uptake of HPs, EVs and PV generation (as outlined in Sections 1.4 and 1.3.2), the growth models are out of date. Not only this, but the customer load profiles used to analyse the maximum power demand in design are no longer representative of realistic customer use patterns incorporating these technologies, due to the different characteristics of these new load (and generation) types in terms of magnitude, duration and time of use.

The long-term nature of electrical network planning for long-lasting assets means that as the electrical network changes significantly and rapidly in the coming years, there is a need to integrate DER into these sections of network that have existed — and will continue to exist — for decades. Consequently, as well as updated planning techniques for the future, there is a need for tools to allow these new technologies to fit within these legacy networks.

2.1.2 Properties of LVDNs

The characteristics of the LV segment are dependent upon a number of properties of the infrastructure (such as topology, transformer ratings and conductor types) as well as customer use patterns. Consequently, network characteristics vary widely with geographies of settlements and housing density, user types and environmental considerations.

LV networks can be designed in a way that reduces transformer requirements, improves security of supply and reduces cost. Possible topological arrangements are shown in Figure 2.1. The topology has an effect on the power flows through the system and the voltage profile along the wires. Various topologies may be in use at different points in the distribution network, but at the customer end open loops (with normally-open points depicted, Figure 2.1(d)), and the radial arrangement (Figure 2.1(e)) are common. In rural settings, long radial networks connected with overhead cables are the norm. In dense urban settings, underground cables are used. The urban LV network is predominantly still radial, but with some normally-open points: these switches isolate two halves of a loop under normal operation, but can be closed to connect them, particularly to facilitate maintenance[38] or to electrify part of the network where a conductor has been lost.

Less power is transferred over the conductors at the distribution level; the cost of thicker conductors (with lower resistance) is compared with the value of the power losses. The consequence is that in LVDNs, thinner conductors are used. In contrast to the transmission network, the resistance of the wires is significant compared to the inductance. Unlike in the transmission network, resistance cannot be neglected to simplify network analysis.

The supply voltage is divided between the feeder resistance and the load applied. Resistance is dependent upon conductor length. A longer conductor increases resistance; loading of the feeder increases

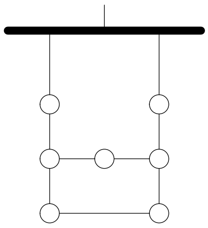

(a) Mesh network

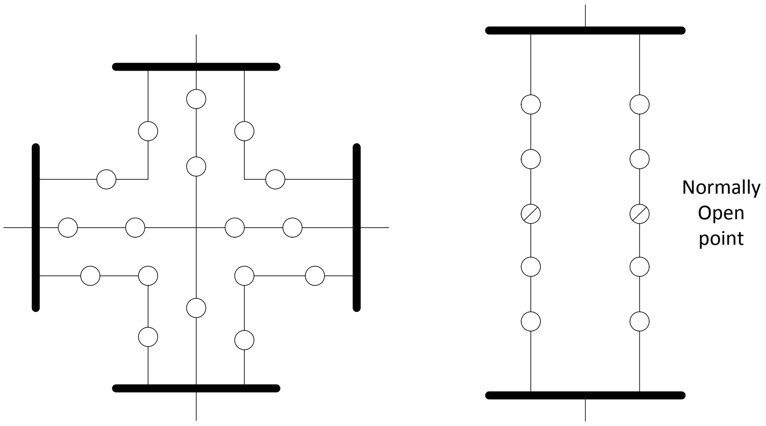

(b) Interconnected network

(c) Link arrangement

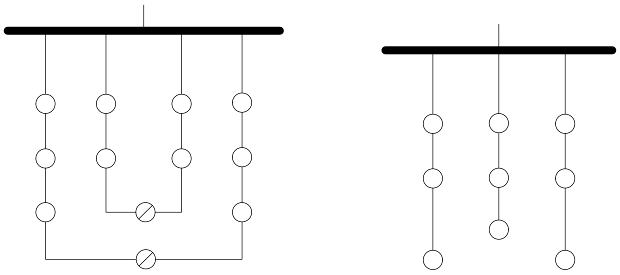

(d) Open loop

(e) Radial system

Figure 2.1: Network configurations
Types of network configuration, showing nodes connected in mesh, interconnect, link, open loop and radial systems [37]. Normally-open points are shown and may be closed to connect loops.

the current through the conductor. Since $V = IR$, there is a voltage drop across the conductor that increases with distance along the radial feeder and with feeder load. There is a difference between the nominal supply voltage and the actual voltage at customer equipment: the voltage drop across the conductor means that voltage at the remote (i.e. customer) end decreases with distance and load. UK law specifies that the voltage for the customer must not drop by more than 6% or rise by more than 10% of the nominal supply voltage[39]; equipment is designed to operate safely within this range, and deviation beyond it may cause damage. Consequently, the system must be designed to supply customers with power that remains within acceptable voltage limits.

The supply to individual domestic customers is typically single-phase. When customer loading patterns are not equally distributed across three phases, there will be some phase unbalance present in the network, which must be kept within acceptable limits. The Unbalance Factor can be calculated using the definition provided in [40]; the maximum permitted unbalance is 1.3% in the UK.

2.1.3 Current DER Integration Paradigm

As described under the regulatory framework in 1.3.1, currently, DG which outputs at less than 16A per phase may be connected to existing LV networks following the “fit-and-inform” principle in which there is no centralised planning of how or where DG is connected by consumers. Heat pumps and EVs do not require any such notification. The technical aspects of integration beyond this notification process and the state of the art in operating networks incorporating DER are detailed in Section 2.2.

2.2 Distributed Energy Resources

2.2.1 Impacts of DERs

The connection of DERs in distribution networks challenges many of the expectations regarding network properties in Section 2.1.2. In particular, the direction of power flow (usually from HV to LV), the voltage profile (decreasing along a radial feeder with distance from the transformer) and phase balance (with significant clustering effects) can no longer be assumed at any given point.

The impacts of DER have been explored extensively in the literature on DER integration. These include:

Voltage rise and drop

Attaching generators to the network causes the local voltage to rise. At a certain level, the voltage becomes unacceptably high. Similarly, for large loads such as HPs and EVs, the voltage reduces as the network becomes more heavily loaded.

Stability

Power electronic interfaces used to invert the supply from DG do not provide system inertia like conventional rotating machines; a higher proportion of power supply from SSEG reduces system stability.

Net power export

In conventional systems, the generator voltage is transformed up to the transmission level, then power flows down through successively lower voltage levels, from top to bottom as per Figure 1.1. With generation embedded in the distribution system, power may flow in the opposite direction to reach other loads if it is not consumed locally, with some currents reversed from the original design. Existing protection systems are not configured for these reverse power flows.

Unintentional islanding

An entirely self-supplied system may become islanded and continue to operate separately from the grid. The frequency will deviate from the remainder of the grid, causing stability issues for reconnection; frequency and phase must be rematched to avoid large transient currents and system oscillations. Safety is compromised: any on-line generators may still cause the system to be live, even if the grid supply is disconnected upstream in a way usually sufficient to ensure isolation. Maintenance may be difficult, as generators may need to be isolated individually.

Thermal overloading

Increased current from additional loading, reverse power flows and from phase unbalance causes increased heating of power systems components; care must be taken to ensure this does not exceed the existing design parameters.

Fault level

The inclusion of DER may change the current flows in the case of fault. This may render protection devices unable to detect faults, or increase the current that may be supplied past equipment ratings.

Unbalance

Clustering effects may occur where customers on the same phase acquire DERs in groups in a way that does not fit balanced load growth projections. The resulting voltage unbalance has a significant effect on phase current.

Loss of diversity

Similar items of equipment, and in particular with long duty cycles, reduce the coincidence factor (see Section 1.4) used to plan networks. The value of maximum power consumption per customer is sensitive to the duration of equipment use.

Harmonic injection

Power electronics that invert the supply from DG for grid connection superimpose harmonics on the Alternating Current (AC) network. These, and any Direct Current (DC) component, reduce power quality, increase losses and may damage other equipment.

Network losses

Generation local to consumption may reduce energy losses through reducing the distance over which power is distributed. However, when more energy is produced than locally consumed, losses are increased, particularly when SSEG operates at non-unity power factor. High-current loads (EVs, HPs) will increase I^2R losses.

The work by Trichakis[1] built on the impacts analysis by Lyons[41] to create a methodology to predict the maximum permissible volume of microgeneration on a given LV network without technical intervention. This was then expanded to include EVs[42]. The impact on voltage rise and regulation, phase unbalance, cable and transformer limits was dependent upon network configuration, necessitating an impact study on a per-substation basis — for example, despite the wider EU constraints on phase unbalance (%Voltage Unbalance Factor (VUF) of 2% versus 1.3% in the UK), the EU generic case study could only tolerate around half the asymmetrical SSEG volume compared to the UK generic network.

Whilst the impacts are configuration-specific, at times Trichakis argues that constraints on maximum SSEG volume within the UK system would vary by level, but not by the order in which they are encountered — although the results actually presented show some re-ordering of this list. In order of decreasing significance, these are: *a*) voltage unbalance (if SSEG is significantly clustered on a single phase); *b*) voltage rise; *c*) transformer thermal limits; *d*) voltage regulation; *e*) cable thermal limits; and *f*) network losses. Under the EU regulatory regime, however, voltage regulation and thermal overloading were more significant. These results informed their approach to create a control system targeted at voltage and thermal limits.

The impacts of EVs are primarily thermal overloading of transformers and voltage drop, with phase unbalance, cable ratings and harmonics not normally realistic problems. The threshold for network issues occurs where more than a third of customers use EVs[42], though Putrus *et al.* estimate it to be even lower[43].

A number of researchers have attempted similar research to analyse the impact of large amounts of DERs. Table 2.1 shows a summary of DER integration limits, extending the photovoltaics limits survey by Whitaker[44] to include EVs for comparable recent context[1]. These limits are for cases where no control or coordination is used. The table shows initial conservatism to maximum levels - 5% for PV and 10% for EVs in early stages of use, with increasing confidence closer to the present day- but still with voltage limits issues. Despite the increase in the maximum over time as understanding of impacts has improved, each of the authors places an upper limit, and it is clear that there is room for integration of PV to be facilitated and expanded through the use of active control techniques.

Approaches to improving integration of DER into distribution networks were explored in "Techniques for DER integration and active networks management" as part of the EU-wide SOLID-DER project[48]. The do-nothing approach may be valid for very low numbers of DER and in the short-term, but beyond monitoring (such as through G83 notification), the use of active control at the device and distribution network level is necessary.

Even after SOLID-DER, considerable obstacles remain in integration of DER. The 2014 review article

---

[1] Papdopoulos et al. also evaluate pre-2012 EV integration literature[45]. These studies align with the themes of thermal overloading and voltage drop, but also consider the acceptability of increased losses in distribution networks.

Source	Type	Maximum penetration	Cause of limit
Chalmers *et al.* (1985)	PV	5%	Ramping rates of conventional generation (centralised PV)
Jewell *et al.* (1988)	PV	15%	Reverse power swings during cloud transients (distributed PV)
Cyganski *et al.* (1989)	PV	(none specified)	Harmonics
EPRI (1990)	PV	>37%	No problems were caused by clouds, harmonics or fast transients at 37%
Barker *et al.* (2008)	PV	1.3%-36%	Unscheduled reverse power flows. Variation dependent upon geographical distribution of PV and is generation-mix-specific.
Asano *et al.* (1996)	PV	Equal to minimum feeder load	Voltage rise. Assumes no tap changers in MV/LV transformer banks.
Povlsen (2002), Kroposki and Vaughn(2003)	PV	<40%	Voltage regulation
Thompson and Infield (2007)	PV	33%	Voltage rise
Papadopoulos *et al.* (2012)[45]	EV	12.5%	Underground cable limits
Mu *et al.* (2014)[46]	EV	<25%	Voltage drop
Neaimeh *et al.* (2015)[47]	EV	60%	Voltage drop (but may be lower due to Voltage Unbalance

Table 2.1: Maximum DER penetration limits
Sources for maximum permissable penetration of DER on LVDNs before encountering network limits (with PV material from [44])

by Olivares *et al.*[49] enumerates some of the outstanding issues, which still include “development of new voltage and frequency control techniques to account for the increase in power-electronics-interfaced distributed generation”, and in particular “control mechanisms that exhibit a plug-and-play feature to allow for seamless integration over time”. In spite of approaches taken thus far in the literature towards solving these issues, there is still a need for methods that do not require extensive offline configuration, setup or maintenance to effectively lower barriers to entry for DER. The cost of both technology and effort in commissioning make plug-and-play functionality of particular importance.

A market approach is common in attempting to solve these constraints, such as the MAS EV charging current/voltage limits method presented in Weckx *et al.* [50].

2.2.2 Relevance of Location

Much of the literature discusses the importance of device location. The resistive nature of the LV network means that the power loss and voltage profile is dependent upon distance from the transformer. Voltage drop becomes a significant issue when a device is connected by a long line; rural connections in particular are vulnerable to voltage issues. Location can also refer to phase: nearby domestic customers may be connected to each other such that clustering effects become significant[42]. Alternatively, they may be on separate phases and consequently a customer experiencing power quality problems and voltage drop may derive no benefit from DG only a few metres away. Limits on the number of SSEG that may be installed in a passive network (due to the G83/1 system outlined in 1.3.1, for example) may not take account of a connection next to a transformer having a minimal effect on voltage and consequently unfairly limit access to the benefits of SSEG ownership and slow implementation of low-carbon technologies.

2.2.3 Active Approaches to DER Integration

Development of active approaches have continued well after SOLID-DER concluded. More recently, DNOs in the UK examined the integration of DER across distribution network scales in projects supported by the Low Carbon Networks Fund (LCNF); several operators undertook research in this field, and rather than network reinforcement, these projects analysed Smart Grid approaches to the issue. The HubNet consortium (see Section 2.2.5.1) summarised the outputs of the LCNF projects between 2009-2015[3]. It categorised the research effort into the following:

- Innovations for Network Operation
    - Storage
    - Flexible Demand
    - Generator Control
    - Network Configuration
    - Equipment for Active Regulation of Voltage
- Innovations for Network Visibility and Design

  - Real Time Thermal Ratings
  - Enhanced Network Monitoring
  - Enhanced Network Visualisation
  - Enhanced Understanding of Existing Demand
  - Enhanced Understanding of Low Carbon Technologys (LCTs)

Much of the research was aimed at the 11kV and above (especially for real-time thermal ratings and for large battery storage), but some of the work examined residential LCTs (i.e. small-scale DG, heat pumps and EVs). The analysis recommended new research into:

> "the geographical nature of flexible demand requirement".

This HubNet summary analysed the combined efforts of UK distribution network operators' progress in DER integration through the LCNF over more than 5 years with over £500m of research, and highlights a significant need for further work on small-scale DER integration with a focus on geographical aspects.

One of the key projects focused on residential LCTs was the Customer Led Network Revolution (CLNR)[32]. As part of CLNR, the approach in [51] outlines how a cluster of energy storage and demand-side response equipment could be used for voltage support, integrating customer equipment as essential network components. In comparison to approaches such as this that engage with DERs, some other projects within the cluster such as Customer Load Active System Services (CLASS)[52] examined active substation components to adjust voltage levels in order to improve network performance but without any integration with customer devices themselves. In the same paradigm, the network itself can be reconfigured, rather than adjusting DERs, such as in the system described by Capitanescu *et al.* where active network normally-open points can be used to dynamically adjust medium voltage (MV) network topology when required to assist in DG integration[53].

The move from simple active components to smart grids has developed significantly over time, and through a number of concepts and architectures.

Distributed Generation may be included in small, specially designed LV grids, which may or may not be islanded (isolated from a larger grid), with diverse generation technologies supplying a corresponding load. Whilst small grids were being designed and implemented earlier, Lasseter outlined the Microgrid in 2001 [54] and then formalised the concept in 2002 [55], identified its component parts and described a basic control mechanism.

In a Microgrid, a mix of sources and loads are controlled to provide power and achieve goals of reliability, efficiency and voltage regulation; to operate as a single dispatchable entity; and managed to function as either a grid-connected or islanded system. Consequently, it is of great relevance in the field of DER integration, as each item of equipment can be managed for both consumer function and for network support. Each element on the network can function autonomously: in its most basic form, each source has a controller that examines its local conditions and determines an appropriate operating point, but coordination across the system is also possible. Haziargyriou reviews the state of research efforts into Microgrids in [56] and also enumerates many demonstration projects worldwide.

The Small Scale Energy Zone (SSEZ) complements the Microgrid concept. Widespread DG, EVs and other technologies are incorporated over time into existing parts of the electricity distribution network without central planning by a DNO and then grouped together with a common control approach to form an active grid[57]. Whilst Microgrids and SSEZs are similar, their development is quite different; Microgrids are carefully designed to ensure correct network operation and matching of generation and loading after analysis of a particular scenario, whereas SSEZs grow organically and upon legacy systems as consumers and other stakeholders install new equipment.

The Smart Grid concept combines information technologies and control at and across any level in the electrical system. At the customer end, concerning integration of DG and EVs, it is thus highly related to Microgrids and SSEZs. Microgrids and SSEZs — as *de facto* Smart Grids — form a significant part of research activities into intelligent, active energy networks for DER integration.

2.2.4 SSEZ Approaches

Lyons described the SSEZ[41] and went on to create a physical model, creating the Durham Experimental Small Scale Energy Zone (ExSSEZ), shown in Fig. 2.2, and diagrammatically in Fig. 2.3, which were used for investigation of centralised control approaches[58].

Figure 2.2: Photograph of Durham Experimental Small-Scale Energy Zone facility

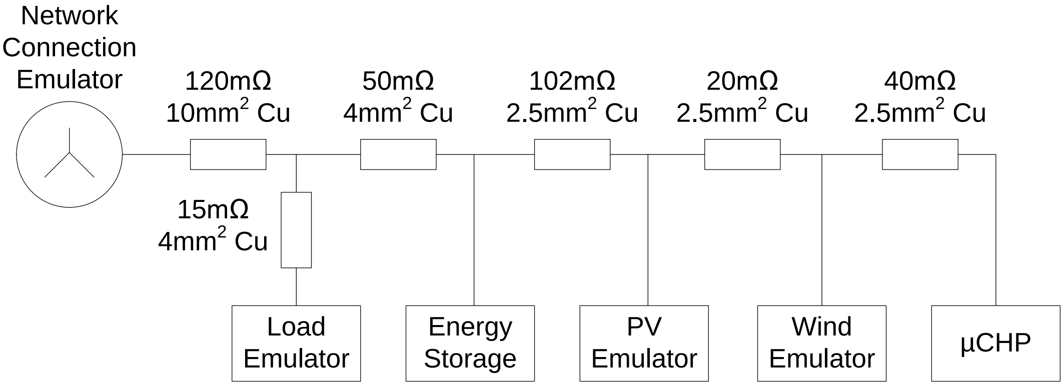

Figure 2.3: Durham ExSSEZ network diagram
Line model of the energy network of the Durham ExSSEZ LV network[1]

Based on the the impact analyses of Lyons and Trichakis outlined in 2.2.1, Trichakis identified the control requirements for an SSEZs to overcome the barriers to DER integration, and developed and implemented a partially-decentralised multi-agent system specifically for thermal and voltage limits[1]. Trichakis *et al.* then described the transformation from legacy, passive LV networks to multiple deployments of SSEZs, each zone capable of adapting to widespread integration of SSEG[57]. Cipcigan examined how many smaller generators could be coordinated using these SSEZs to provide a larger — and potentially more useful — dispatchable entity, with the benefit of offering ancillary services (such as voltage support — see [59]) as a Virtual Power Plant[60].

The control solution used by Trichakis has various advantages due to the strengths of the MAS paradigm. The properties and advantages of a MAS in this type of task are discussed in Section 2.3.2.3. However, the decentralised system actually implemented in [1] suffers from a fundamental problem: the control system relies on centralised information. A solution that removes this would be a considerable improvement.

Whilst covering the same research area, references to SSEZs have declined over time: more recently, the term “microgrid” has expanded from its original definition to include active approaches to grid-connected legacy networks.

2.2.5 Microgrid Approaches

Microgrid control, both in standalone operation and in grid-connected mode, has been the subject of extensive investigation. The “Trends in Microgrid Control” review[49] covers literature concerning DER integration as well as issues around control hierarchy, centralisation and decentralisation, markets, DER technologies and progress made in particular control techniques to date. The importance of the move to Smart Grids and the role of microgrids has resulted in the availability of international strategic research funding. Some of the larger research projects — and their relevance to this work — are outlined below.

2.2.5.1 UK Research

In the UK, major research in this field is funded by esearch Councils UK), primarily through Engineering and Physical Sciences Research Council (EPSRC).

HubNet is an RCUK (Research Councils UK) - funded consortium between seven universities within the SUPERGEN (SUstainable PowER GENeration, supply, transmission and storage) initiative within the RCUK Energy Programme, and is one of seven other energy hubs. Some of those other hubs are relevant here: SUPERGEN-HiDEF (formerly SUPERGEN-HDPS) uses the "cell" concept in its approach — this is similar technically to the SSEZ, but explicitly recognises roles of consumers, generators, regulators, network operators and other stakeholders as microcosms of existing electricity networks. SUPERGEN-FLEXNET has examined distributed voltage control through the use of power factor control[61]. However, the HubNet remit specifically includes smart grid infrastructure for electricity networks. It has been operating since 2011 and received follow-on funding in 2016. It cascades funding to projects within its remit and in particular has dedicated funding to distributed control.

The HubNet consortium has contributed significantly in this area, with over 117 publications to date. Two are book chapters, 63 are conference papers, and 53 are journal articles; 22 of these journal articles are on Smart Grids themselves (as opposed to, for example, power electronics or mega-grid technology), of which only one is on the topic of active control: in this approach, refrigerators are used for frequency control by adjusting demand[62]. The authors demonstrate that a random element is required to prevent overshoot from too many devices synchronising[2].

Of the remaining non-journal publications, the conference articles and consortium position papers are highly indicative of the state of the art. Its 2016 review[63] highlighted some research gaps by drawing on the literature and examining the state of current systems and regulatory frameworks.

A distributed agent-based multi-energy-vector control system was proposed by Arnold et al.[64]. They decomposed the energy systems into hubs which represent "the interface between the energy sources and transmission lines on the one hand and the power consumers on the other hand". These hubs could represent large demands or geographical areas. The distributed optimisation problem was solved iteratively. Control agents were associated with these hubs, with each agent responsible for its own (potentially geographical) sphere of control.

Whilst this still does not solve issues of predetermined control areas or iterative communication with large numbers of agents, this study in particular informed the approach of the consortium in tackling integrated energy systems, including decentralisation, and formed part of their research gap analysis [63] for future work.

---

[2] The issues of synchronisation were addressed for the voltage control problem solution in this thesis by using geographical distance to create sufficient disorder without compromising response.

Consequently, agent approaches have been used for modelling and control and in particular, for market-based control systems. Some work has examined whether there is benefit in using smart controllers at the distribution level for demand response[65] and demonstrated the effects of using DR in a specific community for control applications. Then, in applying agent-based approaches, Karangelos and Buffard use a learning agent-based controller to determine a consumer's participation in a demand response scheme[66]. Whilst this contributes to enabling more DER integration through allowing consumers to participate in the market, this approach makes no attempt to account for the distribution network operational constraints, and implementation of markets remains outside the scope of this work.

Even in 2014, the HubNet consortium still considered voltage control in the network and at the domestic meter to be an open issue[67]. However, the consortium identifies that smart technology in the domestic setting can play a role in achieving control.

Smart grid approaches for transformer equipment in LVDNs have examined the use of smart meters for state estimation[68]; however, in this case the state estimation relied on the power consumption being similar to the sum of all of the metered power usage plus some losses. This is not appropriate where smart meters are not fully deployed, or where there are unmetered users (e.g. for street lighting). There is significant progress yet to be made in leveraging information at the smart meter level in network operation.

The consortium will continue to operate until at least 2018. The Top-And-Tail consortium concluded in 2015 having examined the "last mile" of LVDNs and included examination of voltage regulation, but did not publish on active, agent-based control at this level. The latest HubNet Symposium included the "Distributed Intelligence for Network Operation and Control" theme, but as of the latest 2016 meeting, and despite the Integrated Energy Systems and the Distribution System Operations position papers[63, 67] showing the need for research in this area, LVDN operation (including voltage control and using their preferred agent-based approach) is still only examined in terms of markets, small numbers of agents and predetermined network topologies[69]. Overall, this indicates the research gap is still significant in terms of UK efforts in this area.

2.2.5.2 EU Research

After releasing a 10-year policy and strategy for sustainable energy in 2010[70], the European Union set out a timeline for a transition to low-carbon energy supply through the European Strategic Energy Technology Plan (SET-Plan). Alongside other topics, it includes the European Electricity Grid Initiative (EEGI)[71], which provides funding for smart grids research including demonstration facilities. DER integration is a theme under its research programme. As well as these European Industry Initiatives within the SET-Plan, other funding is available through a number of mechanisms. Currently, the European Commission supports large, international projects through the Horizon2020 research and innovation programme, to run 2014-2020. The Societal Challenges pillar includes Secure, Clean and Efficient Energy as a theme which incorporates Smart Grids. A great deal of work in microgrids was included in the the now-concluded Framework Programmes FP5-FP7, where active distribution network technology existed as research theme since cycle FP5.

Since then, considerable progress has been made. IRED[72] was a cluster of large-scale projects from FP5 and FP6 for improving distribution networks and integrating renewable energy sources. Table 2.2 lists the relevant projects and their abstracts, which are discussed below.

Table 2.2: Highlighted EU research projects within the IRED cluster

Project	Date	Abstract
Microgrids [73]	01/2003 - 12/2005	To investigate, develop and demonstrate the operation, control, protection, safety and telecommunication infrastructure of Micro-Grids and determine and quantify their economic benefits.
More Microgrids [74]	01/2006 - 12/2009	To increase of penetration of microgeneration in electrical networks through the exploitation and extension of the Microgrids concept, involving the investigation of alternative microgenerator control strategies and alternative network designs, development of new tools for multi-microgrids management operation and standardisation of technical and commercial protocols.
DISPOWER [75]	01/2002 - 12/2005	*"Distributed Generation with High Penetration of Renewable Energy Sources"* — To help to prepare the safe, reliable and high quality implementation of distributed generation into European grids focussing on the efficient integration of renewable energy sources.
DER-lab [76]	11/2005 - present	To support the sustainable integration of renewable energy sources (RES) and distributed energy resources (DER) in the electricity supply by developing common requirements, quality criteria, as well as proposing test and certification procedures concerning connection, safety, operation and communication of DER-components and systems. A major objective is to establish a durable European DER-Lab Network that will be a world player in this field.
FENIX [77]	01/2007 - 12/2009	To boost DER (Distributed Energy Resources) by maximizing their contribution to the electric power system, through aggregation into Large Scale Virtual Power Plants (LSVPP) and decentralized management.
SOLID-DER [78]	11/2005 - 10/2008	To tackle the barriers for further integration of DER, overcoming the lack of awareness and fragmentation in EU R&D results by consolidating all European DER research activities; to raise awareness of DER solutions and benefits in the new EU Member States, thereby addressing the specific issues and barriers faced here.

Whilst some of the projects focus on large-scale DG and transmission networks, others focus on integration of small generators. *Microgrids*[73] examined implementation of independent low-voltage grids, examining safety, communications, islanded operation and control and markets. Much of the literature in this field now uses the hierarchies and nomenclature set up in this project. It culminated in two key papers: one on multi-agent methods for microgrid control[79], and one presented at CIGRE summarising large-scale integration of DER through microgrids using a market-based centrally-controlled system[80].

More Microgrids[74] was the logical extension to the project and investigated issues of widespread deployment[81], such as standardisation, evolving distribution management structures and control strategies, particularly in terms of (de)centralisation. A strictly hierarchical system, such as the zone-profit-maximizing customer-cost-minimising central controller[82], utilised market bids from microgenerators. However, load-shedding for frequency response was devolved, where agents each established a ranked list of shedding responsibilies to determine operation (see Section 2.3.2.3). Field trials were of this system were conducted on a system implemented on the Greek island of Kythnos[83].

In another market-based approach, *FENIX*[77] used the Virtual Power Plant (VPP) as a vision for future power grids with extensive DG. It delivered a centrally-dispatchable VPP, used for the day-ahead market, 15-minute-ahead reserves and voltage control through setting reactive power at the SSEG[84]. Functionality was separated into the *technical* VPP, and the *commercial* VPP, examining market integration whilst facilitating system balancing functions for a system operator[85] by characterisation of a zone as a single entity.

DISPOWER[75] developed inverter technologies for grid integration, examining the problems of grid synchronisation, load balancing and limiting power output to prevent overloading. It also examined voltage rise mitigation in LV networks using load control (DSM). Voltage and real - reactive power were measured at strategic network points by a central controller. The central controller instructed customer loads to switch on to adjust demand via a radio communications link, reducing generation constraining and voltage rise[86]. In addition, the centralised load control system could be used to achieve the operational goal of zero (instantaneous) power export to the grid.

In particular, DER-lab[76], as a strategic institution facilitating standardisation for distributed energy resource technologies, outlived its sister projects. Through a network of European laboratories, it coordinated and implemented significant international collaborative studies.

With the IRED cluster concluded, EU funding continued to analyse DER integration through another 14 projects (INTEGRAL, IS-POWER, ADINE, CRISTAL, VSYNC, GROW-DERS, DESIRE, RELIANCE, CEERES, EU-DEEP, ADDRESS, MIRABEL, W2E and G4V) in the area in FP6-7. A full list of these projects and their abstracts is in Appendix A.1. Some were more targeted towards microgrids, decentralised and agent-based approaches; supporting ICT infrastructure research was also included. This infrastructure was recognised as a key requirement to support wide microgrid adoption, but now remains an open problem: *Web to Energy (W2E)*[87] was interested in the standardisation of communications infrastructure and data management as an enabling technology for smart grids, predominantly by extending existing control architectures. Aggregation of SSEG was included in its project objectives in including communication down to consumers at the metering level. IEC61850[88] is a communication standard intended for substation automation and its use in DER management and household appliance control was never anticipated, however, W2E examined extending IEC61850 for DSM, and provided an implementation of the standard for a VPP application by providing additional hardware between non-IEC61850 smart grid equipment (meters and sensors) and the control system[89]. The use of IEC61850 in DER integration has also been carried forward beyond W2E. A proof-of-concept DC microgrid controlled using this standard was developed in [90]. Later, some of the criticisms regarding the need for plug-and-play functionality were addressed in May 2015 when an adaptive architecture for microgrids using IEC61850 was proposed[91]. There are still barriers to its uptake in this configuration in the UK: the multi-stakeholder system means a microgrid controller hierarchy is inappropriate, and in a peer-to-peer agent-based configuration, IEC61850 is incompatible with the use of FIPA-SL (see Section 5.2.2.4 on MAS standards).

However, the main focus remained upon control strategies in distribution networks. In particular, of those 14 above, particularly relevant projects were INTEGRAL, ADDRESS and MIRABEL (formerly MIRACLE).

The *INTEGRAL* project sought to provide a model for the creation of an active network for DER integration. A major outcome was a framework for these active networks, divided into a series of issues and requirements: operationally, this would include *normal*, *critital*, and *emergency* operation modes, and approached via an aggregation approach[92]. Most of the solution requirements (scaleable, open, multi-actor, market-aligned) are echoed in Section 2.5. One aspect examined decentralised operation approaches for the three modes, although the final output of the project in this area only went as far as describing models and facilities for testing self-healing cells (through a microgrid approach) [93], rather than analysing decentralisation itself.

MIRABEL[94] (formerly *MIRACLE*) used market-based coordination methods[95] for demand-supply balancing. It used Experimental Microgrid at Centre for Renewable Energy Sources and Saving (CRES), and applied the VPP-like technology to create low-emissions zones through optimising energy sources on a carbon basis with a trial in Germany. The MIRABEL concept uses micro-requests for energy: a device will specify its energy requirement and how much flexibility (in terms of time-shifting) it has for this demand, and defines a way to describe flexibility[96]. The MIRABEL Energy Data Management System (EDMS) was devised in order to perform aggregation of these micro-requests at centralised trader nodes[97]. The project successfully demonstrated balancing for energy supply and demand, taking into account optimising for the availability of renewable energy, and optimising for price. This fits well with the energy trading structure in the UK: the demand and aggregator nodes map to consumers and their retailers. However, whilst the system accommodates the needs of a transmission operator for balancing purposes, the network constraints of the distribution network operator are not directly addressed in this approach[96] particularly where there may be geographical clustering effects from the aggregation process.

ADDRESS also approached DER integration through aggregation. Under the ADDRESS Technical Conceptual Architecture, each DER is operated via an Aggregator, and those Aggregators trade on markets for 24 discrete services, participating in a similar way to larger DG or conventional generation operators, retailers and traders[98]. Significantly, it highlights a need for a local intelligent service at the low-voltage susbstation that provides real-time topological knowledge of the network[99] although the project published no implemented solution for this. Intelligent solutions for network topology would fit the knowledge gap identified here.

Now the Framework Programme has concluded, microgrids and DER integration research continues through the Horizon2020 EU funding programme. The Low Carbon Energy research strand includes "LCE-02-2016 - Demonstration of smart grid, storage and system integration technologies with increasing share of renewables: distribution system". P2P-SmarTest will examine peer-to-peer (P2P) operation for integration of demand side flexibility (i.e. provision of ancillary services through DER)[100]; GOFLEX (Generalized Operational FLEXibility for Integrating Renewables in the Distribution Grid), starting in November 2016, will look at advanced aggregation and VPPs[101], alongside several other H2020 projects in the area. GOGLEX and P2P-SmartTest are included also in AppendixA.1). These will, however, primarily focus on market participation of small-scale DER operators. Both projects make explicit reference to enabling energy storage and the use of flexibility for distribution network operation and the use of agent and microgrid approaches, but are in their early phases.

The continuation of this significant funding stream underlines the importance of the microgrid paradigm in addressing issues of DERs in LVDNs. In particular, voltage control remains an open problem at low voltage[67]. With no current single, unified solution, it appears that microgrid and agent-based approaches seem likely to continue to be of interest to the European community for some time, and this work will provide a contribution in this space.

2.2.5.3 US Research

The US research into microgrids is not coordinated in a comparable fashion or scale to the EU programmes. The United States Department of Energy funds research programmes into distributed generation, primarily through the Office of Electricity Delivery and Energy Reliability. The California Energy Commission also commissions significant research into new and renewable technologies. Both of these organisations have been instrumental in the Consortium for Electric Reliability Technology Solutions (CERTS); this has contributed work in Microgrid research as a means to integrate DG. A major resource for this research is a full-scale distribution network facility comprising load and generation equipment up to a 13kV grid connection point[102]. The control approach is centralised, with "plug-and-play" capabilities as equipment is connected or removed. Microgrid disturbance response is achieved in milliseconds by autonomous local control, and a central Energy Manager individually dispatches microsources with a timescale of minutes, for either grid-connected or islanded operation. This corresponds with the conclusions reached by the EU research[103] that real-time centralised dispatching is inappropriate.

Subsequent to the Hatziargyriou review article[56] mentioned in Section 2.2.3, an energy bill in the US established "Smart Grids" research as a federal priority[104]. It created a task force to promote the use of digital information systems for control of the electric grid to improve reliability, security and efficiency. In particular, it provides for "Deployment and integration of distributed resources and generation, including renewable resources; [and] development and incorporation of demand response, demand-side resources, and energy-efficiency resources." Research in these areas has expanded rapidly, bolstered by economic stimulus funding into energy programmes.

Market participation is a common theme throughout research programmes[105]. Little explicit reference is made to agent-based control systems, although descriptions of control system functionality (communication, autonomy, learning capabilities) are common to the MAS approach.

2.2.5.4 Rest of World

Whilst the EU and US perform the majority of research into DER integration through large and collaborative initiatives, the challenges are still faced by other countries with a large uptake of DG. In Australia, the Commonwealth Scientific and Industrial Research Organisation (CSIRO) collaborates with other international research groups in multi-agent coordination of microgrids. In Japan, the New Energy and Industrial Technology Development Organization (NEDO) is a government body that has been involved in several microgrid projects including distribution network research for integration of very high levels of PV in existing networks[106], with a physical installation of 553 arrays in a single urban area; here, studies focused on how to avoid constraining generation to maximise energy yield. Private researchers particularly from Japanese technology companies have also examined microgrids comprising single buildings, and some cases to the scale of manufacturing sites; however, details of the research, methods and results are not publicly available for these corporate R&D programmes.

A number of international collaborations have been particularly significant in microgrids development, such as the €5million, 2014-2017 collaboration between Aalborg University, Tsinghua University, Kamstrup and the Shanghai Solar Energy Science & Technology Company on "Multi-energy Complementary Micro-grid Technology Research". Aalborg University has a microgrid research and demonstrator facility for AC and DC networks[107] in the Department of Energy Technology. Guerrero and the Power Electronic Systems group have published extensively on microgrid control and power electronics.

Much of the work published by this group considers islanded microgrids and DC microgrids, and the control strategies are predominantly hierarchical — such as in "Control Strategies for Islanded Microgrid using Enhances Hierarchical Control Structure with Multiple Current-Loop Damping Schemes"[108] — and is consequently not appropriate for the UK domestic context. The group has also used GPS technology in microgrid operation, albeit for its timing properties rather than geographical location, in common with many other power systems applications of GPS. In [109] the authors propose control of an islanded microgrid using GPS for synchronising timing for frequency management without requiring communications links.

However, some of the group's work is highly relevant on decentralisation and multi-agent systems for microgrid control. In 2016, the authors proposed in "Agent-Based Decentralized Control Method for Islanded Microgrids"[110] that a microgrid should be modelled as comprising an electrical network (with an electrical topology), with a control network (of a multi-agent system with associated communication topology) overlaid. This two-layer approach was then used to derive control laws for the island microgrid. Whilst identical topology for the electrical grid and communication network were used, the derivation could operate over a system where these two layers were asymmetric - albeit requiring complete information about the system in the control law planning stage. Further details of the group's work in multi-agent systems and decentralisation are discussed in Section 2.3.2.3.

2.2.6 Using Location in Integration

Some of the systems for DER integration use location in their solution (for example, Trichakis and Lyons in 2.2.4 or AGORA in 2.3.2.3). However, this review has only found existing work where network data is already known. In the context of attempts to decentralise the control system, then, there are no known solutions that use location or network data that is not provided centrally.

Some techniques have been devised to develop dynamic information about the network. For a fault detection system in the power systems domain, in He and Zhang[111], distributed measurements are used in conjunction with network topology information in order to add properties to a graph of transmission network nodes for the purpose of identifying the geographical location of network faults. Messaging is used to communicate local estimates of the graph properties through the network which are then reconstructed into an improved estimate of the global system. They coined the term *decentralised network inference* to describe this process. Again, however, this approach requires the network connections data in advance. A decentralised network inference method to derive topology information for distribution networks would obviate configuration effort.

Using a database of known electrical network information is problematic; data may be commercially sensitive or protected for security purposes. Coordination of multiple system operators is necessary, since the UK system is divided into regional monopolies. Data may not exist to a relevant detail level in all locations. Perhaps more significantly, such a database must be maintained and updated regularly. The ability of a controller to learn its network location without needing or building a full model of the entire system can avoid some of these problems.

2.3 Multi-Agent Systems in Power Systems

2.3.1 Applications Overview

Specific properties of agents are often highlighted within power systems applications: for example, social ability (communication, negotiation) is exploited for the coordination and aggregation of DERs into a VPP[112]. Ability to act autonomously improves agent-based transmission system secondary voltage control in both response time and reliability by removing total reliance on a single central controller, particularly in fault cases[113]. Whilst expanding the use of markets within power systems is sometimes perceived as an aim in itself, it can be a useful tool in solving other problems, such as in the market-based EV charge controllers for transformer current and voltage limits outlined in [50]. The participation of agents within markets is covered extensively in the literature both within and without the power systems domain.

Multi-agent techniques can be applied at all levels of power systems, from HV transmission networks down to individual appliances in houses or even aboard ships[114]. As well as applications in protection, network monitoring, and simulation, distributed control is a focus for MAS research[115]. In particular, the multi-agent system approach allows for development of a control system that takes into account the various stakeholders whilst maintaining the functional independence of all of its components. As a minimum, the following topics (including those in Sections 1.4 and 2.2) have all been addressed using agent functionality:

- thermal overloading of transformers
- voltage rise and drop
- voltage regulation
- voltage unbalance
- cable thermal limits
- frequency control
- islanding
- enabling DSM
- market participation
- maintenance planning
- topology reconfiguration.

The wide exploitation of specific and combined attributes of agents in these applications demonstrates that the multi-agent system approach has considerable merit in this field.

2.3.2 Decentralisation

2.3.2.1 Problems of Centralisation

In the liberalised electricity system described in Section 1.3.1, multiple stakeholders have different and sometimes conflicting control requirements. As a result, there is a lack of a clear control hierarchy for the resolution of network issues. The electricity market participants are not responsible for reliable network operation and the network operators are not responsible for operational decisions of generation equipment owners.

Centralised control has limitations including single points of failure and communication bottlenecks, as well as issues with system maintenance and scalability as components are added under G83/1. There is a lack of clear hierarchy since components are owned by different, self-interested entities. Market deregulation is recognised as a driver towards cooperation-based, distributed control[116] since there is not always a clear path from the entity responsible for a control goal to the entity responsible for its implementation. A decentralised approach can avoid hierarchy issues and is highly effective for local control, allowing local controllers to assume responsibility. However, more effort and coordination is required when dealing with adjusting or optimising global conditions. Studies in artificial intelligence have shown that combining markets with decentralised local controllers can yield a globally-effective system[117], and this result has clearly informed some approaches to microgrid control[103]. However, in the absence of an effective market (see Section 1.3.1) or even a clear definition of its participants, an alternative method of achieving global control through decentralisation is necessary. The advantages of decentralised control specifically outlined in [118] are *a*) scalability and openness; *b*) reliability and resilience; and *c*) communication efficiency, strongly suggesting that an LV control system should follow this approach.

2.3.2.2 Scalability

Scalability must be a basic facet of any system for distribution network control, so mechanisms for coordinating an appropriate multi-agent system must be suitable for large numbers of agents. The effectiveness of a group of agents in solving a problem reaches an (application-specific) threshold, above which the additional of more agents to a task impedes its solution[119], due to some resource constraint (such as communications bandwidth or physical space). Subsets of the agents can form ad-hoc teams for cooperation, removing the need to involve the entire agent society. A “dynamic, partial centralisation”[120] can create (temporary) teams and hierarchies for sub-problem or conflict resolution, providing a scalable and distributed approach. It can also accommodate the need to degrade gracefully with component failure by monitoring effectiveness and assigning new teams. The probability of some agent failing increases with agent society size and number of interactions, so a robust solution becomes increasingly important with larger systems, providing more weight to the arguments for a distributed solution.

2.3.2.3 Existing Approaches to Decentralisation

Communications can be used to improve microgrid operation, but are not essential in implementing distributed control techniques. The Distributed Fuzzy Load Controller[121] examines frequency in an islanded microgrid — a global property — to determine under/over-utilisation of microgeneration. Each load controller then autonomously calculates whether corrective action is necessary. The system incorporates a random element in the controller logic such that actions are fairly distributed between them over time.

One devolution example in the "More Microgrids" Kythnos project used a low-priority load-shedding DSM technique[83], triggered by low frequency. The central controller informs load agents of the total load to be shed. The loads announce their own consumption, and they all establish a ranked list of the others. Loads that were shed in the previous cycle are placed at the bottom of the list. An agent determines from its own position in the list whether or not it will shed load, and messages the central controller accordingly. If the sum of the shed loads is below the requirement, the central controller initiates another shedding cycle. Decentralisation in the Kythnos microgrid[103] refers to load or source controllers independently determining set-points based on bidding and negotiation through a central power market under normal operation: whilst agents retain autonomy of action, the system is still subject to the constraints of a single, central zone controller. Other examples of devolved control include the system to coordinate unit decommitment for maintenance described in [116].

Trichakis went some way to distributing control[1], but used a global database of measurement data, rather than giving each control agent a unique knowledge base. Reliance on global knowledge reduces an agent's ability to take action based on local knowledge. It also means that if an agent is unable to update a remote database for any reason, others will then act on incorrect information. Finally, the information representation method included network location data (phase location and a node reference number) which requires agent configuration and creates model maintenance issues as nodes are moved or changed, so the advantages of a plug-and-play system are not realised.

In existing power systems, power export from a zone may be limited by protection devices configured for unidirectional power flows, or by equipment ratings. Multi-agent systems are used to achieve operational goals of zero power export or as a coordinated exporting entity, so can be used to ensure these limits are not breached. Equipment upgrades could alleviate these limits but require reconfiguration of microgrid controllers. Consequently, there is a need to develop a system that allows for a changing electrical network. Static, human-configured, network-dependent designs are clearly inappropriate for the organic, unplanned SSEZ context. Unlike the configurations in [122] or [85], where the network constraints are determined at design by a DNO and included in a single centralised microgrid controller, the use of multiple Thermal Limits Agents (TLAs) in the scheme proposed by Trichakis [1] is highly suited to this purpose: a TLA sends commands to direct control agents (at loads or generators) when its local status indicates an alert or emergency state. Any number of constraint agents may be created, each taking effect at different system operating points. This implementation has considerable drawbacks, however. Firstly, the TLA is not located at the measurement equipment, but rather hosted at a DNO-operated facility with a global measurements database. Secondly, relevant agents or equipment are determined at the design stage and hard-coded into the system; no dynamic mechanism is proposed. Any form of change to load or generation would require considerable reconfiguration effort, effectively negating the positive aspects of scaling plug-and-play agents that impose independent network constraints. Some form of dynamic network characterisation, team generation or fuller decentralisation could improve this approach considerably.

One fully-decentralised, non-hierarchical control system is the “AEN”[123], an autonomous microgrid for integrating DG, enhanced by inter-device communication. The system comprises primary, secondary and tertiary control levels, which correspond with the priority of control goals in the distribution network. Primary control allows for operation even in the absence of communication in response to frequency and voltage and as a fail-safe. Secondary control then looks at voltages system-wide for power quality improvement. Finally, tertiary control performs economic optimisation. The lower two levels are decentralised by using a form of distributed averaging achieved through gossip communication. This entails repeated peer-to-peer interactions where agents exchange information and adjust their control set points accordingly, by “trading” small power output increments to keep total output constant. The economic optimum to be reached is the minimization of total cost of production. However, in the UK, the market is not complete: all power produced is purchased by the consumer at a fixed rate regardless of the cost of production. Since the objective of feed-in tariffs is to stimulate uptake of more expensive renewable technologies, it makes little sense to impose control based on economic efficiency. DSM is modelled as negative generation that will produce zero power when the unit cost reaches the correct incentive threshold. The binary nature of some loads is not appropriate for a proportional-integral control block for secondary control, and the discontinuity in the cost function is incompatible with the tertiary optimisation method; gossip-coordinated incremental power output change is not always possible. The framework for inter-device communication, AGORA[124], uses an overlay communications network such that devices on the same electrical network are strongly connected, with a gossip-step every three seconds. The experimental results show secondary and tertiary control response times of the order of tens of seconds. Whilst bandwidth and latency are unlikely to be an issue, this negates the benefits described above of fast response to avoid disconnection under G83/1. Most significantly, however, is that the power reliability control subsystem takes no account of network constraints such as transformer thermal limits, relying solely on voltage and frequency measurements to determine safe system operation.

The AGORA [124] system notes the importance of electrical network location of DERs in its control system. It uses location in order to determine connections between devices. However, it provides no mechanism for doing so: the information must still either be provided centrally, or derived somehow.

A notable collaboration between the US Department of Energy and CSIRO has resulted in a test facility employing GridAgents, a MAS with a central market and aggregation (like the VPP approach) where agents use a bulletin-board for exchanging information[125]. System planning algorithms of interest assume predetermined fixed groups of agents; researchers intend to examine algorithms for the formation of these planning groups including a topological approach similar to the overlay network of the AEN described above.

In [126], the authors describe a voltage control system in LV networks at 6kV. It uses power flow measurements to control Static Voltage Regulator (SVR) Static VAR Compensator (SVC) equipment where PV generation and EV use may cause large voltage fluctuations. The system combines centralised optimisation with resilience in the case where communications networks are unavailable. The system performs centralised optimisation which determines operating points and a model supplied to each control device for measurements across the system. When the communications network is unavailable, the control device uses its local measurements as inputs to the model, assuming a static load, to provide an approximation of the optimum with an error based on the deviation between the static load assumption and actual load. The model uses a local measurement of P and Q to extrapolate to the global state. Whilst a central server is required to set up the initial models, the local node uses its own model to estimate the state of the surrounding network for voltage control in a decentralised mode.

In “Droop-free Distributed Control for AC Microgrids”[127], the authors propose a system that provides cooperative multi-agent control. Significantly, even in 2016 the authors highlight the need for scalability, modularity, plug-and-play capability, and the need for the system to avoid the use of prior knowledge. The solution provides for sparsely-connected agents, each of which operates according to three layers to control voltage, reactive and real power at each node - and to account for exchange in grid-connected mode. The plug-and-play system also resolves where one of the DG units is removed and its agent goes off-line. The decentralisation is achieved via cooperation in communicating the reactive power supplied by each unit; each agent compares its set-point with those of its neighbours each time-step and adjusts itself in order to share the burden of network operational requirements for frequency and voltage. However, there is a significant move away from the decentralised approach for tertiary control as soon as stochastic sources (i.e. intermittent renewable generation) are included, and the system uses a single centralised tertiary controller these set-points and to allow power exchange in grid-connected operation. Whilst a single microgrid in this configuration would have plug-and-play functionality, the boundaries of each microgrid zone would need defining and a framework for establishing the (redundant) connectivity between the agents.

Each of these systems contributes in the area of DER integration, but in each case some significant aspects are missing. The next sections use the existing work to clarify the gaps in the literature that are yet to be fully developed and describe a solution that advances current methods of DER integration.

2.4 Research Problem Statement

The literature survey in the previous sections outlined research into DG integration and illustrated some of the attempts made thus far in addressing the concomitant LV network problems. Several systems have been designed to limit voltage rise and to address thermal overloading as part of a priority-based list of network control requirements. Some implementations incorporate DSM schemes for consumer loads and other equipment. Decentralisation of control and the absence of agent hierarchy can be used to increase the value of a system by lowering the barriers to its implementation: no maintenance of central controllers is needed; agents should require neither pre-programming with network topologies, nor design-time instructions about which controllers should cooperate. Some acknowledgement is made of the need for topology information, such as the intelligent real-time topology service specified in ADDRESS2011, (discussed in Section 2.2.5.2), but provided no implementation. Progress has been made in decentralisation, such as the gossip-based AEN (see Section 2.3.2.3), but there remains no system that simultaneously addresses the major technical problems of DG integration into a market-free, demand-driven legacy network context whilst also capturing the full benefits of a multi-agent solution.

Consequently, a scalable, network-independent, decentralised control system for operation of an evolving network with multiple system constraints remains an open research problem.

2.5 Problem Solution Requirements

Drawing on the survey of relevant literature, the idealised multi-agent-based control system to solve this problem would have the following characteristics:

1. Decentralised control

   The system should self-coordinate without the use of a central controller. Any aggregation must be distributed.

2. Multi-goal

   The system should allow for multiple and conflicting network constraints and control objectives. It should address the problems of voltage rise and thermal overloading as primary barriers to DG integration. Network constraints must be imposed dynamically to allow for these to change over time. The system must allow for the multi-stakeholder context of the liberalised electricity system.

3. Market-independent

   System coordination must be achieved without the use of supply or ancillary service markets.

4. Towards “zero” configuration

   The system should be capable of adapting to changing network equipment, topologies and participants, whether these occur regularly or perhaps once in the 40-year lifetime of a system. It must avoid the need for a DNO to initialize and maintain a controlled zone with electrical network configuration data. Control agents should require a minimum of setup to participate in the system.

5. Scalable

   The system should be capable of scaling from a few participants in grid-connected SSEZs up to a large number of agents connected across the distribution network.

6. Resilient

   The system should be able to tolerate failures of equipment, software and communications in a graceful fashion.

7. Standards-compliant

   The system must implement standard communication techniques in order to facilitate open access.

There is currently no system described in the literature that fulfills these criteria. A system with these combined characteristics constitutes a significant and original contribution in the field of DG integration and LV network control.

2.6 Summary

The uptake of DG, EV and DSM technology poses problems for the legacy installations of traditional, passive-controlled networks. Smart grids will not replace these in the short term, so a managed evolution to integrated networks is required. Decentralisation appears to be a very appropriate way to achieve this goal. Multi-agent systems provide a framework to solve a distributed control problem. Whilst some ground in this area has been covered, there are significant gaps in implementing such a system. Attempts by Trichakis and Lyons (see Section 2.2.4) have issues around the actual decentralisation by using prior knowledge of the network and a centralised storage system. Attempts using location (such as AGORA in Section 2.3.2.3) also suffer from the need to start with a network map. Other solutions may succesfully demonstrate forms of fully decentralised control but do not take into account device location, which is important for several network operational goals. Consequently, there is a need to explore fully decentralised control for DER integration that uses location without requiring a centralised database of known network data.

In conclusion, the solution to the research problem outlined in Section 2.4 would be a considerable development in the field of LV network control. This thesis will contribute new work in distributed network inference and voltage control using a multi-agent approach to fulfil these requirements and advance the state of the art of DER integration.

Chapter 3

Electrical Network Modelling

The electrical theory basis for understanding and modelling electrical networks is outlined in this chapter. Existing benchmark LV electrical networks are reviewed and used to create new models of legacy distribution networks incorporating widespread DER. These models were implemented using several software packages to provide electrical data and tools for the control solution presented in Chapter 5, as well as to illustrate the effects of DER in LV networks with a focus on the impact of device location.

3.1 Electrical Networks

In this section, the basic electrical principles are outlined for some of the phenomena caused by DERs in LVDN. The properties of appropriate models of these networks are discussed. The techniques used in extracting relevant results are presented, along with a discussion of software packages used to analyse complex networks.

3.1.1 Basic Network Theory

3.1.1.1 Voltage Drop

The electrical network delivers power to customer appliances through wires that have resistance and inductance, and this resistive property is relevant to the voltage supplied to the customer. It is this voltage that must be kept within the regulatory limits of no more than 6% drop or 10% rise of the nominal supply voltage[39]. The customer voltage depends upon both the properties of the distribution line and the size of customer load applied, as will be shown below.

3.1.1.1.1 Line impedance

To illustrate the voltage drop problem, a simple model is shown in Figure 3.1 showing a single load connected to an electrical generator. Resistor R_d and inductor L_d at the top of the diagram combine to

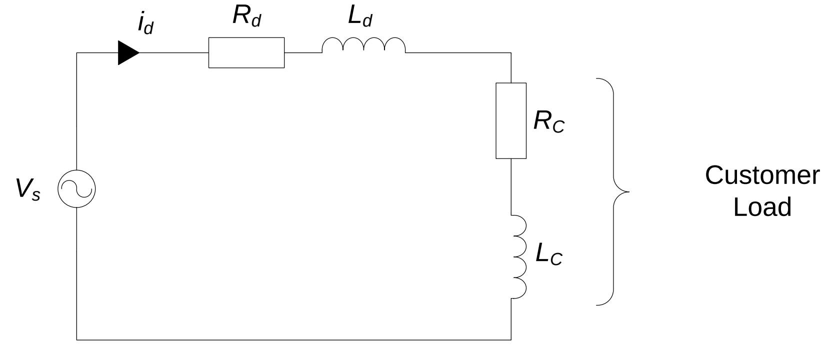

Figure 3.1: Electrical Network: simple line model
A very simple electrical network. The generator on the left represents a connection to the grid, with source voltage v_S. Current (i_d) flows through the distribution line to the customer load, right.

model the resistance and inductance of the distribution line. This represents a single phase distribution network for an AC system where a sinusoidal current i flows with peak magnitude I. The voltage drop across a resistor is given in Equation 3.1. An inductor opposes a change in current; the back Electromotive Force (EMF) produced across it is given in Equation 3.2

$$v_r = iR = I \sin \omega t \tag{3.1}$$

$$v_l = L\frac{di}{dt} = L\omega I \cos \omega t \tag{3.2}$$

where $\omega = 2\pi f$ for grid frequency f and $i = I \sin \omega t$. Note that the current lags the voltage by 90°.

A sinusoidal waveform can be projected onto a rotating vector in the complex plane, such as the voltage shown in Figure 3.2. Power system currents and voltages can be represented as vectors that rotate at the same rate: the angle between them is constant.

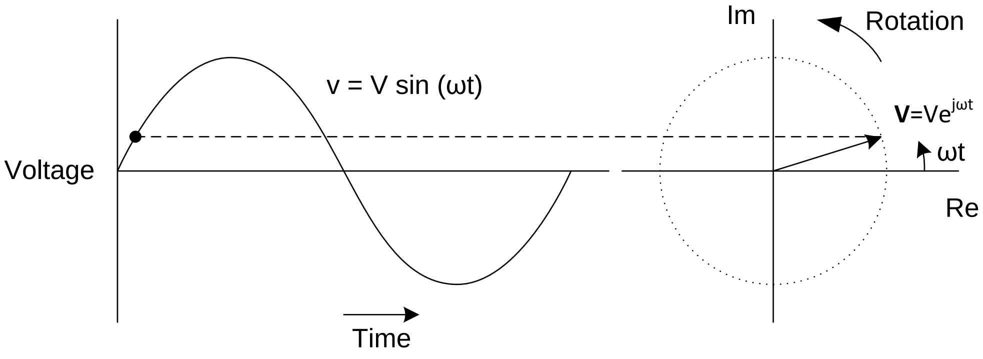

Figure 3.2: Rotating vector

The *reactance* is given by $X = \omega L$, and the *impedance* is the vector sum of the resistance and reactance to give the complex vector $\mathbf{Z} = R + jX$, where j is the unit vector along the imaginary axis of the phasor diagram. By the generalised form of Ohm's Law, potential difference is the product of the current and impedance vectors, $\mathbf{V} = \mathbf{IZ}$.

In a distribution line, the ratio of line resistance to reactance is high, i.e. $R \gg X$. Because of this, the *small angle approximation* is made: that is, the angle between the current and the voltage at both the source and the customer is approximately the same.

The current flows through the line to the customer load. This creates a voltage divider: the source voltage v_s is applied across the line and the load, with a potential difference across each. Then, using the small angle approximation, the voltage drop across the line is given by:

$$|V_d| = IR\cos\phi + IX\sin\phi \tag{3.3}$$

3.1.1.1.2 Customer loading

The total impedance in series in Figure 3.1 is $\mathbf{Z} + \mathbf{Z_C}$ (*i.e.* line + customer), and the current i_d is common to both resistances. Decreasing $\mathbf{Z_C}$ increases current through the cable, reducing the proportion of voltage dropped across the customer load and hence increasing the voltage drop across the distribution line.

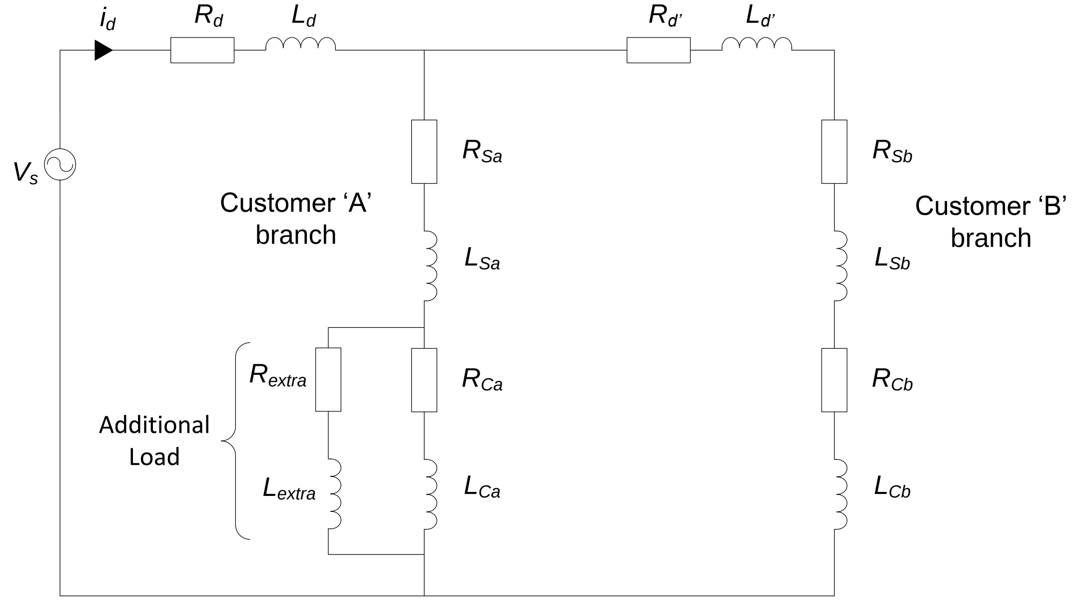

Figure 3.3: Electrical Network: two customers, extra loading
A two-customer model showing distribution lines, service cables and customer load. Additional loading has been added by Customer A.

A customer is connected to the distribution network via a service cable, adding resistance and inductance $R_S + L_S$. An expanded two-customer model is shown in Figure 3.3. By application of Kirchoff's Current Law, the current through the distribution line i_d must equal the sum of the current through the branches A and B.

Suppose that customer ‘A’ switches on some additional load. It is added as a parallel branch.The additional parallel load has the effect of reducing the equivalent impedance of the ‘A’ branch. Similarly, the equivalent impedance of ‘A’ and ‘B’ combined is reduced. Since the total impedance of line and loads in the circuit is decreased, the overall current i_d will increase.

The higher current causes an increased voltage drop across Z_d. The proportion of supply voltage applied across the customer loads is thus lower. With the ‘B’ branch unchanged, the voltage experienced by Customer B is reduced. Increased loading thus causes voltage drop locally and for other customers.

For the general case, in Figure 3.4, customers are connected at nodes along multiple distribution line segments, with indices $d_1...d_n$. Applying Kirchoff’s Voltage Law, the voltage drops in the mesh sum to the source voltage.

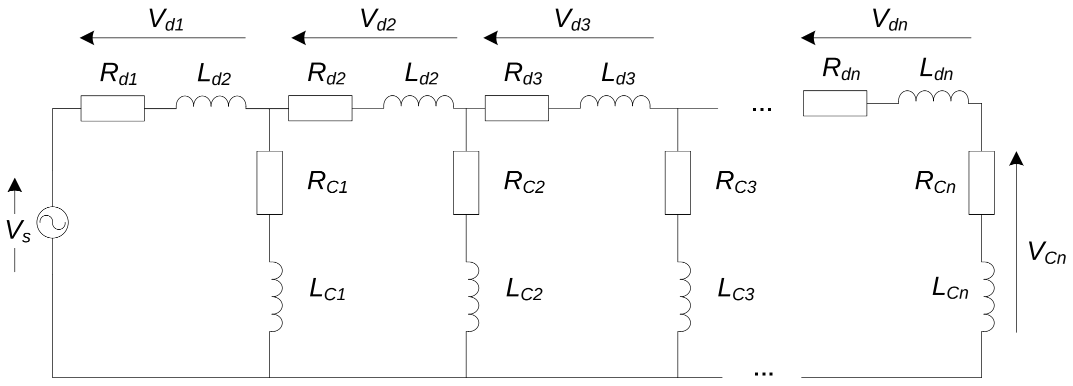

Figure 3.4: Electrical Network: general case
The simple network model, extended for multiple customers. Several distribution line segments are shown.

The voltage drops across those line impedances $Z_{d1}...Z_{dn}$ reduce the customer voltage at the corresponding connection node. Additional loading increases current through each upstream distribution line segment. For all customers, the effect of voltage drop caused by a customer load is increased with the electrical distance (i.e. resistance) from the source.

3.1.1.2 Analysis by Symmetric Components

The modelling above assumes a single-phase system for simplicity. The distribution network usually delivers single-phase power supplies to residential customers, with some exceptions for users with very high consumption. Figure 3.5 shows a three-phase system, with single-phase customer loads attached between one phase and a neutral point. The neutral point is common to the other phases.

The customer supply voltage waveforms are sine waves, with frequency f. In a *balanced* three-phase

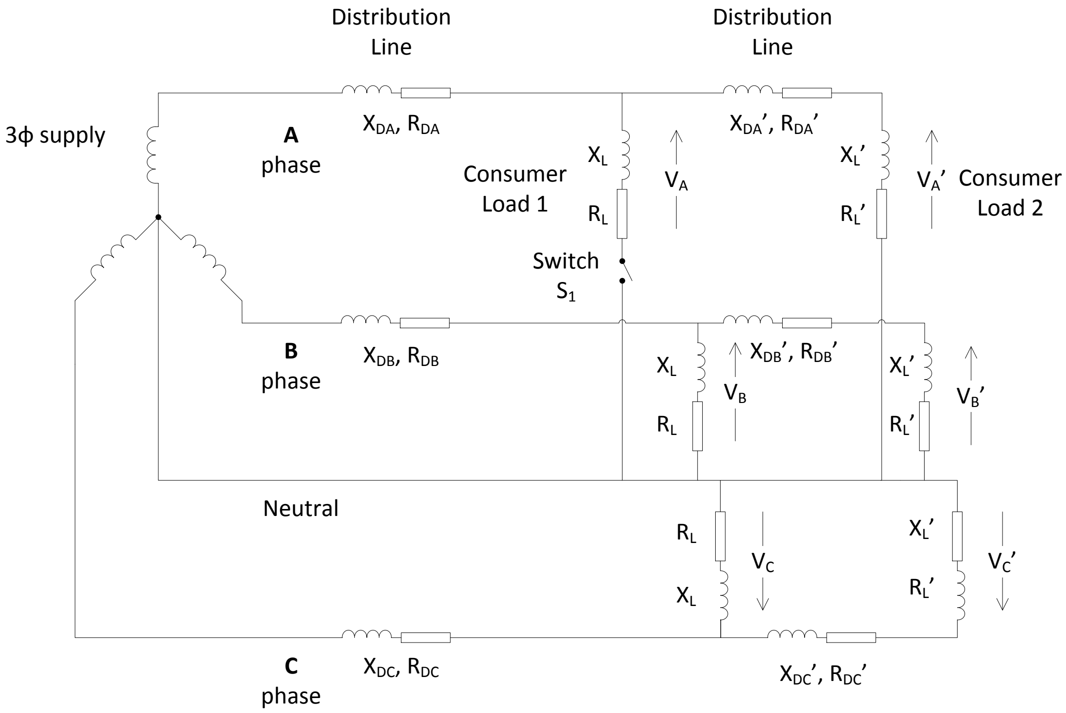

Figure 3.5: 3-phase, 4-wire network
The transformer, left, consists of three windings, connected to a common neutral point. Moving right, customer loads are connected between a single phase and a line that returns to this neutral point.

system the three phases exhibit the same magnitude and a phase shift of 120°, shown on the left in Figure 3.6. Similarly to Figure 3.2, by the application of Euler's formula, $e^{j\theta} = \cos\theta + j\sin\theta$, the sine waves can be projected onto real-imaginary axes, shown on the right, as vectors that rotate with angular velocity $\omega = 2\pi f$. Since the vectors rotate at the same rate, the plane may rotate to allow them to be considered as stationary phasors. In the balanced case, the resultant of adding the symmetric, equal-magnitude voltage phasors together is zero.

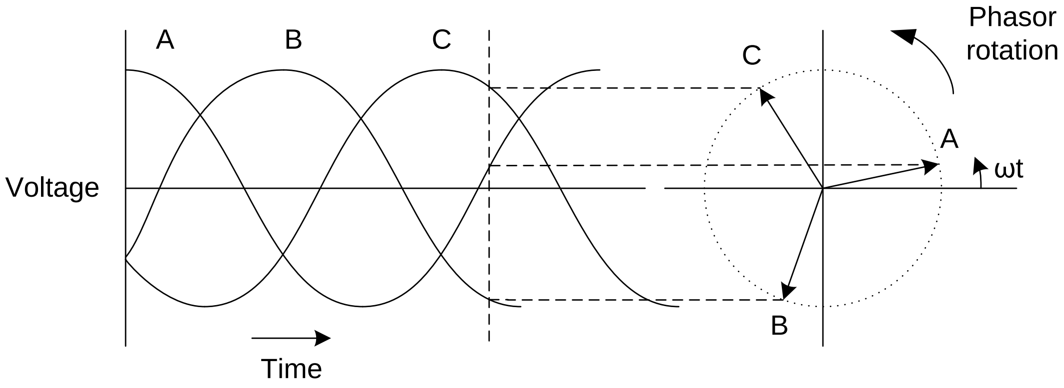

Figure 3.6: 3-phase to vector projection
The three sine waves (left) show the voltages of three network phases over time. On the right, a set of vectors are shown with 120° separation. Rotation of the vectors plots a circle, shown with a dotted line. The voltage changes with time can be projected onto the vector diagram: dashed lines show a snapshot at an instant in time. As time t increases, the vectors on the right rotate in concert through angle ωt.

In networks where residential customers are connected to single phases, large numbers of customers are distributed across the three phases such that the average demand across each is the same. However, at the LV end of the network, the diverse customer loads (i.e. impedances) attached to a particular feeder may not match perfectly at a given instant in time. Consequently, the current and voltage phasor magnitudes will not be the same and the phase shifts will vary. The result is that the sum of the phasors is not zero. Systems where the impedances do not match are *unbalanced*. The general case of an unbalanced, 4-wire network can be drawn and labelled as shown in Figure 3.7.

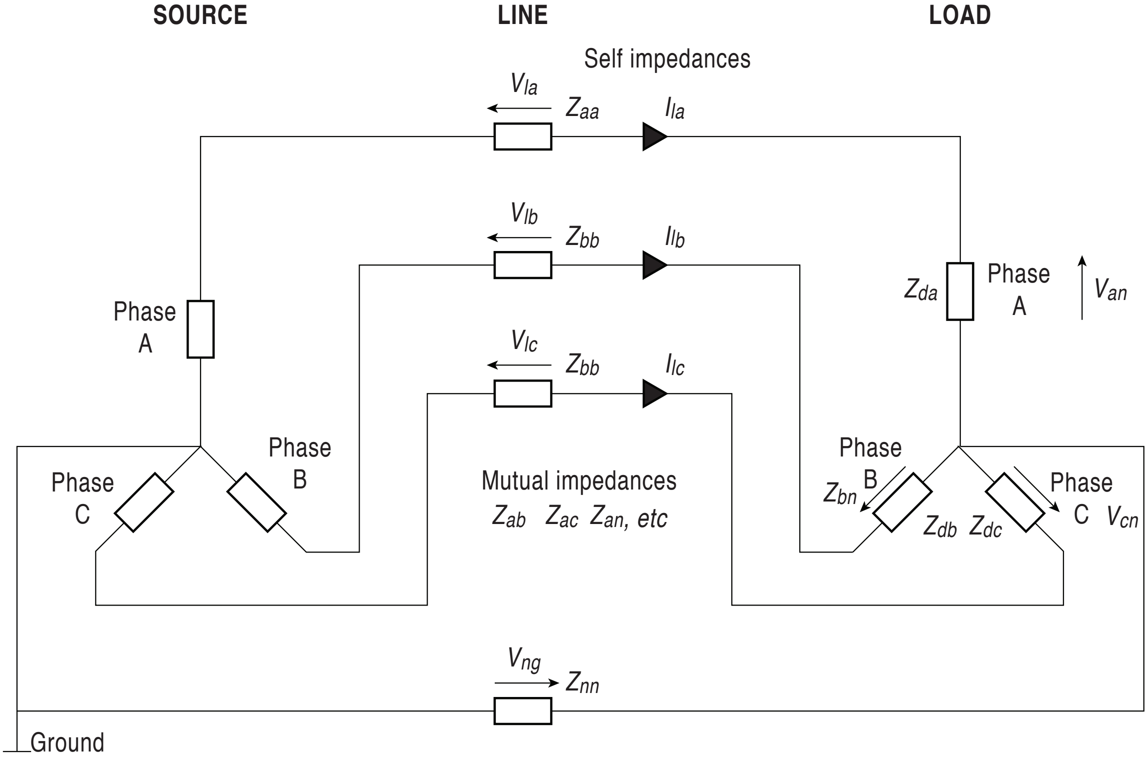

Figure 3.7: Three phase system
A three-phase, four-wire system showing source, distribution line and load.

Three voltages E_a, E_b, E_c can be represented as phasors, like A, B and C in Figure 3.6, each with their own magnitude and phase angle. The matrix of these voltages is thus:

$$E_{abc} = \begin{bmatrix} E_a \\ E_b \\ E_c \end{bmatrix} \tag{3.4}$$

To aid analysis of the system, these unbalanced system phasors can be described completely as the combination of symmetric components[128]: a superposition of three sets of symmetric phasors.

Three component phasors are denoted E_0, E_1, and E_2 (for zero, positive and negative *sequences* respectively):

$$E_{012} = \begin{bmatrix} E_0 \\ E_1 \\ E_2 \end{bmatrix} \tag{3.5}$$

The component phasors are rotated by multiples of 120° to create three sets of three symmetric phasors, according to the transformation matrix A:

$$A = \begin{bmatrix} 1 & 1 & 1 \\ 1 & a^2 & a \\ 1 & a & a^2 \end{bmatrix} \tag{3.6}$$

where $a = 1\angle 120^\circ$.

The sum of the transformed components, then, form a final asymmetric set that rotate in concert that can be used to describe the three electrical network phases at any point in time:

$$E_{abc} = AE_{012} \tag{3.7}$$

and similarly, to obtain the component phasors,

$$E_{012} = A^{-1}E_{abc} \tag{3.8}$$

3.1.1.3 Unbalance

In the LV network, the line impedance has $R \gg L$, and the self-impedances are large compared to the mutual impedances between the lines and line-neutral.

In the balanced case, the neutral would be at zero potential and carry no current back to the transformer, but in the unbalanced case, there is a nonzero resultant phasor.

A few definitions of voltage unbalance exist[129]; for the avoidance of doubt, here this unbalance is measured by the ratio of the negative sequence component to the positive sequence component:

$$\%VUF = \frac{V_2}{V_1} \tag{3.9}$$

and is limited to 1.3% in the UK[40].

3.1.1.4 Three-Phase Voltage Drop and Rise

Starting from the balanced case (Figure 3.8(a)), consider the case where a load is added to one phase: a parallel impedance branch is added, reducing the overall phase impedance. This causes additional current to flow. However, at this point the currents are no longer balanced, so $I_a + I_b + I_c \neq 0$, such that a current I_n flows in the neutral wire, and a voltage drop is present across Z_n. Unbalance in the network displaces the neutral wire voltage from ground[58].

The resulting voltage at the neutral is given by

$$V_N = \frac{\dfrac{E_a}{Z_{da} + Z_{la}} + \dfrac{E_b}{Z_{db} + Z_{lb}} + \dfrac{E_c}{Z_{dc} + Z_{lc}}}{\dfrac{1}{Z_{da} + Z_{la}} + \dfrac{1}{Z_{db} + Z_{lb}} + \dfrac{1}{Z_{dc} + Z_{lc}} + \dfrac{1}{Z_n}} \tag{3.10}$$

With additional load, the neutral is displaced from ground by voltage V_N, as shown in 3.8(b). This offset of the neutral must be accounted for in examining customer voltage levels. The resultant phase-neutral voltages at the customer for the unchanged phases B and C show a change both phase angle and magnitude compared to the balanced case.

The line, load and neutral-ground voltage phasors are dependent up on the power factor of the load as well as the line. Additional loading on phase A creates a neutral offset voltage, which can cause the phase-to-neutral voltage V_{CN} to rise —- one permutation of phasors is shown in Figure 3.8(b). With extra loading, the line voltage V_{LA} increases in length. An example of unbalance caused by reduced loading is shown in Figure 3.8(c), in which the length of V_{LA} is shorter than the other phases. Other variations in the remaining phasors will occur in different configurations of distribution networks: with an increased load of power factor close to 1, even the phase-to-ground voltage V_{CG} may rise, due to the increase of the magnitude and angle of V_{LC}, as shown in Figure 3.8(d).

For small variations at constant power factor, the neutral voltage phasor extends approximately linearly with added or reduced load, although at high levels the line properties become significant and this causes some noticeable rotation of the neutral voltage phasor.

To summarise, unbalanced load or generation on an individual phase causes voltage rises and drops relative to neutral, which are also experienced by customers on other phases. Since the distribution network is predominantly composed of single-phase customers with uncoordinated use, there will inevitably be unbalance, resulting in power flows in the neutral wire and neutral voltage offsets compared to ground, with their phasors dependent upon the nature (i.e. size and real/reactive components) of the customer load, the distribution line and the neutral wire impedances.

3.2 Network Modelling

The theory outlined above explains the basic principles involved in the network effects of increasing load or generation. To investigate the effects in more complex and realistic networks, some of the simplifications must be removed. However, additional complexity can render analysis by inspection somewhat unwieldy. With appropriate models, software packages can be used to automate these analyses. As well as better

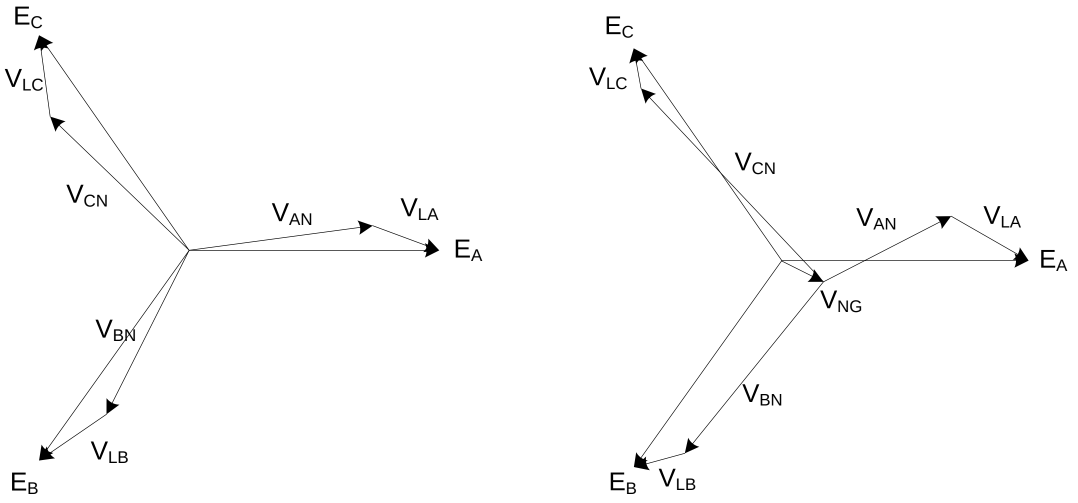

(a) Balanced case

(b) Unbalanced case — increased load on A. The neutral voltage offset is shown. As loading increases, the neutral offset phasor lengthens.

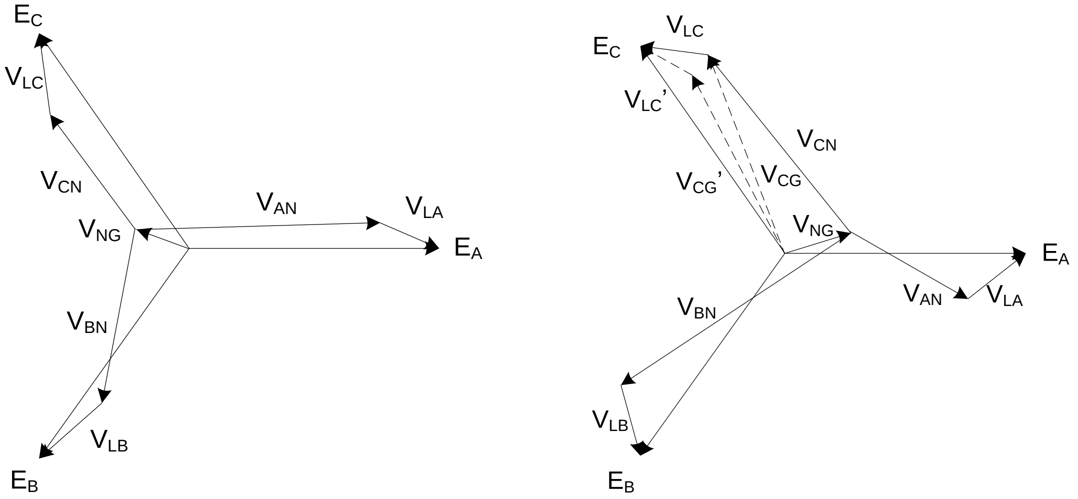

(c) Unbalanced case — reduced load on A. The neutral offset lengthens with further load reduction, but may also display rotation as the line properties become significant.

(d) Alternative unbalanced case with additional load on A, power factor = 1.0. The voltage drop across the line is higher as well as rotated further, increasing both V_{CN} and V_{GN}.

Figure 3.8: Phasor diagrams for an unbalanced system
Phasor diagrams for an unbalanced system, showing cases of extra load, reduced load and varying power factor. Lengths and angles are exaggerated for clarity.

understanding of a network, models can be used to examine the impacts and effectiveness of control systems imposed upon them. In this way, the model can provide a tool for both analysis and testing.

One network model was used provided a baseline throughout, but was implemented in three different software packages (PSCAD, IPSA and RSCAD), each for different purposes.

The PSCAD models were predominantly used to improve understanding of network behaviour with the highest levels of detail in various scenarios. PSCAD provided results of the unbalanced, single-phase, dynamic operation.

In contrast, the IPSA models were much simpler: only balanced, steady-state, 3-phase operation was supported, and the underlying network was simplified considerably. With these simplifications, the outputs were be used directly in a database of operating states to create a simulation to test a multi-agent system in software.

The RSCAD model was used to create a dynamic, unbalanced, single-phase real-time simulation of the electrical network that could be used in conjunction with hardware.

3.2.1 Model Properties

The challenges of DERs and network growth are present at the LV and MV level. This is the part of the electrical network below 132kV, as outlined in Section 1.2. In particular, the effects at the domestic level are of interest. Benchmark LV networks are described in the literature, and can be used to develop appropriate models.

3.3 Model Selection

The CIGRE EU Benchmark Low Voltage Microgrid[130] was adopted as a general-purpose LV study network by the EU Microgrids[73] project. It comprises three feeders: one residential, one commercial, and one industrial. The transformer and conductor properties are known, along with half-hourly loading characteristics. The benchmark microgrid includes proposed connection points and specifications for various DERs. However, for the UK residential context, the model is only partly appropriate: with six houses and two apartment blocks, the total of 26 domestic connections is low.

The UK Generic LVDN[131] was used instead, focusing more on residential connections: the feeder is intended only to represent domestic supply.

Ingram contends that “network impacts [may] be investigated on a relatively simple network representing a residential system which would give the most onerous condition for voltage regulation”[131]. This model does use the minimum standards for new service cable connections from UK DNO ENWL although legacy connections may not meet this specification. However, the design is highly symmetric to begin with, so analysis starts from a position of phase balance. In terms of customer connections, the model is relatively light: the DNO ENWL specifies a maximum of 200 customers per feeder. However, the 11kV/400V transformers are almost at their maximum loading (499.2kVA of maximum loading of 500kVA).

Some assumptions were made about the network in order to implement it here:

- No automatic voltage regulators present
- Perfectly transposed conductors, implying zero mutual inductance
- In common with the assumptions made in [132], additional grounding along the neutral wire is neglected.

The UK Generic model has no normally-open points (shown in Figure 2.1, for increasing system security). In some real-world cases, these are connected with solid neutral, i.e. only the three phases are normally open, and the neutral wire is connected.

3.3.1 Customers and Loading

The customer types of profile classes 1 and 2 as specified by ELEXON[19] are domestic consumers which may include those on an Economy 7 tariff. Averaged load profiles for these classes have maximum demands of 0.9kW and 1.8kW after diversity. The Average Daily Maximum Demand (ADMD) used in the UK Generic Network (UKGN) is 1.6kW across all customers.

In the context of exploring the future of distribution networks as they experience changing use patterns, a scenario was envisaged where the average daily maximum demand was assumed to have increased to 2.2kW per customer along one LV feeder. The minimum demand was maintained at 0.16kW.

The 1.6kW ADMD is a conservative estimate. The source network model data uses a ADMD of 1.3kVA per customer[131]. Measured data from the Netherlands smart meter programme confirms a current value of 1.2kW, but its examination of individual consumer loads gives further weight to the need to re-examine the standard method of calculation using a decades-old coincidence factor method. This is reconfirmed by recent results from the 2014 CLNR trials[133] of 8000 customers, which give ADMD results for heat pumps users of 1.69kW, and EV users of 1.79kW, up from a base load of 0.9kW — a 0.7–0.8kW increase per technology. This agrees with other results from the LCNF cluster that concluded there is a need to update ADMD figures[3]. In contrast, the ADMD in the model was increased by 0.8kW — the use of more than one of these technologies in a single household would significantly increase the average power and also increase the coincidence factor due to their long duty cycle. Consequently, the values in the model represent a conservative approach to future growth.

Furthermore, the distribution networks in the UK will have experienced growth in connections as houses are built and connected to existing networks since their initial planning. One segment at 230V has been extended by 10% (i.e. 3 customers per phase) to illustrate the effects of growth. This would require the transformer to operate beyond its rated capacity periodically, potentially leading to thermal issues.

These changes transform the generic network to include a heavily loaded feeder which will demonstrate the impacts of the combined challenges of electric vehicles and distributed generation on the UKs legacy networks.

3.4 New Tool Development

The networks were turned into tools for research through the use of software modelling packages.

PSCAD[134] provides an interface to the ElectroMagnetic Transients including DC (EMTDC) software package by Manitoba Hydro International. It uses the method described by Dommel[135]. Network components are converted to a number of connected equivalent networks, each composed of an impedance and a current source. A connection matrix of these equivalent networks is then used to create a set of (nonintegrable) differential equations. These are then computed iteratively over small timesteps using the trapezium rule to approximate the system response over time. The consequence is that an arbitrary, nonlinear network can be constructed, allowing for a 4-wire, unbalanced network to be simulated.

IPSA+[136], from TNEI, is used for static, balanced, 3-wire operation. It uses a Newton-Raphson method applied to a single-wire approach with the approximation that the phases are symmetric.

RSCAD[137] provides an interface to RTDS hardware, also from Manitoba Hydro. Similar to PSCAD/ EMTDC, it uses Dommel's method[135] in an ElectroMagnetic Transients Program (EMTP) approach, but with an implementation designed for parallel operation, where multiple processors perform the calculations simultaneously to provide real-time operation.

3.4.1 PSCAD Model

A validated PSCAD model of the UK Generic LVDN was completed by Trichakis[138]. This implementation was used as the basis from which modifications were made for PSCAD studies to better understand behaviour at more points in the network and to explore future conditions.

In common with the IPSA model above, the basic model was extended on one feeder from 96 to 134 customers. Some of the original model's simplifications were reduced: some load aggregation was removed. This additional complexity provided an additional, disaggregated feeder, to examine the links in network behaviour between devices at various points that were not previously visible.

SSEG was added, and variable levels of generation, phase clustering and uniformity could be selected. Figure 3.12 shows the household load model.

Figure 3.9 shows a high-level view of the PSCAD model. The overall model retains the basis from Ingram et al.[131], but the modifications discussed above have been made. Diagrams have been redrawn from the PSCAD interface for clarity.

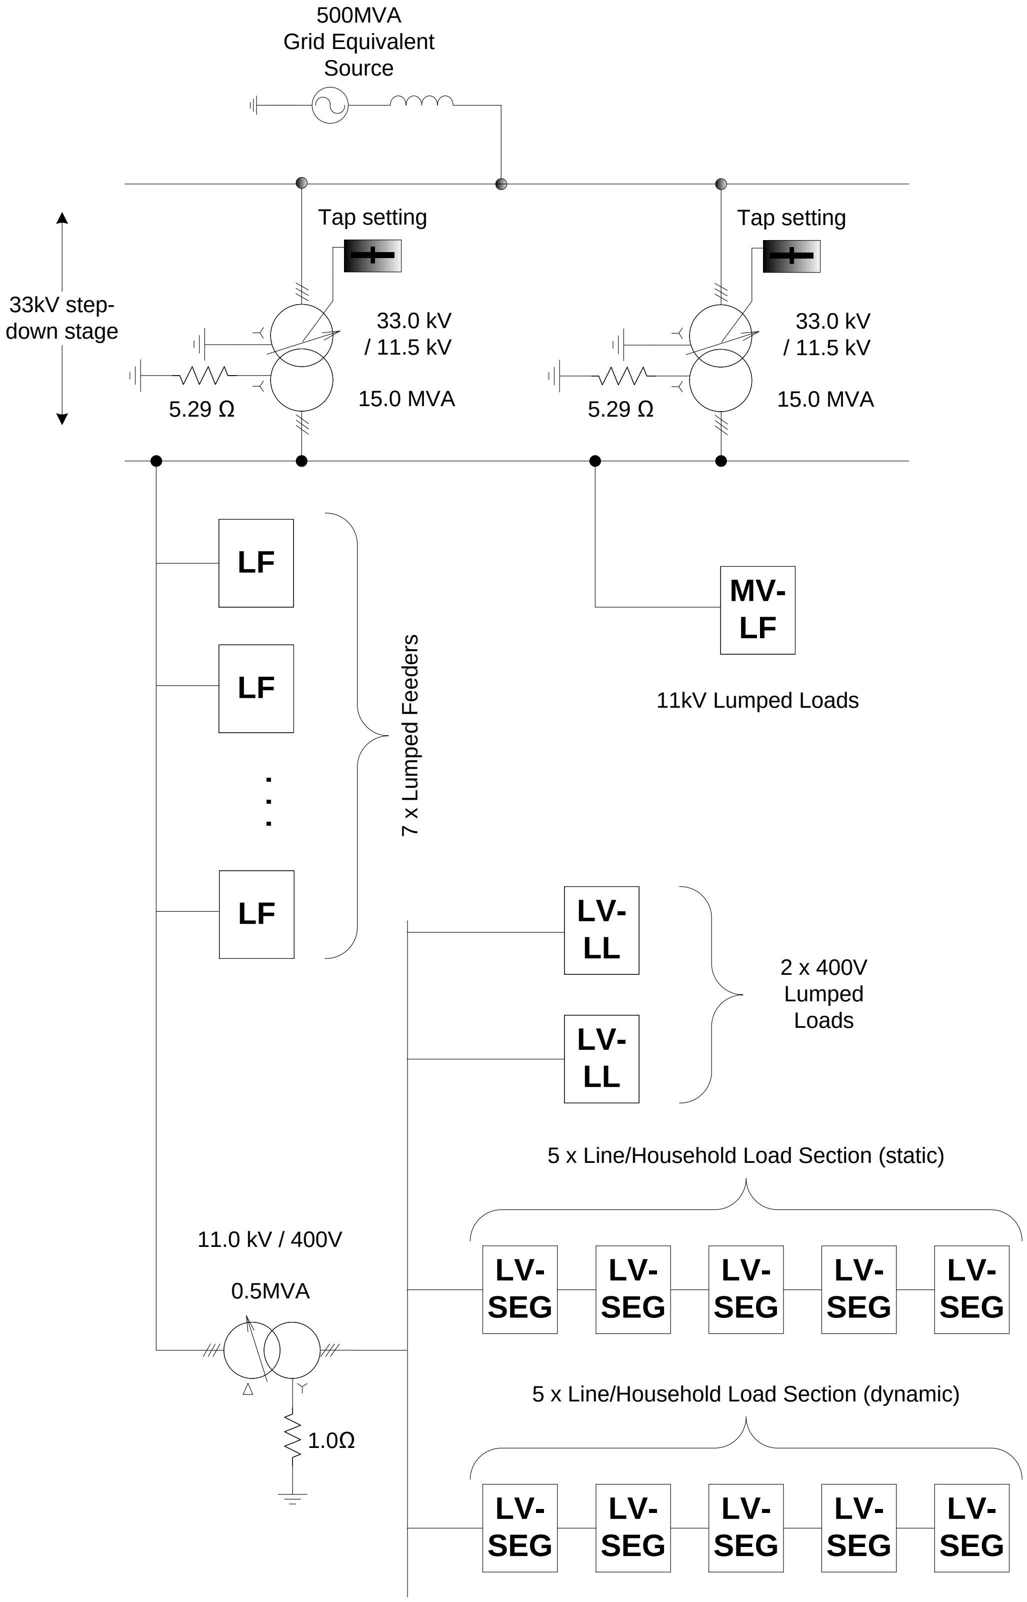

Figure 3.9: PSCAD model overview
The diagram shows a high-level view of the PSCAD model. Sections are expanded in Figures 3.10-3.12.

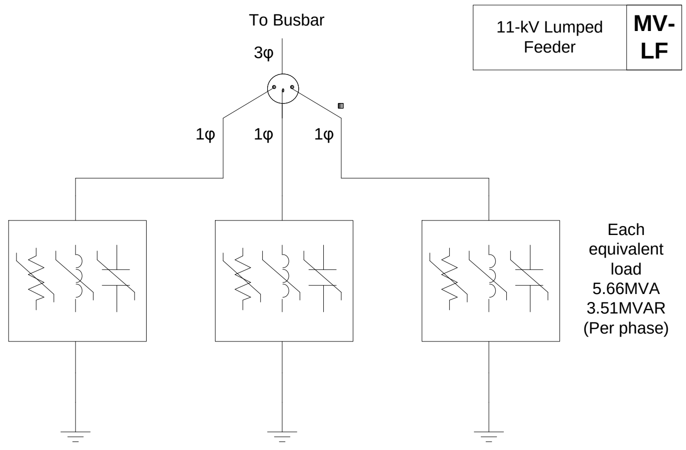

(a) Medium-voltage lumped feeder

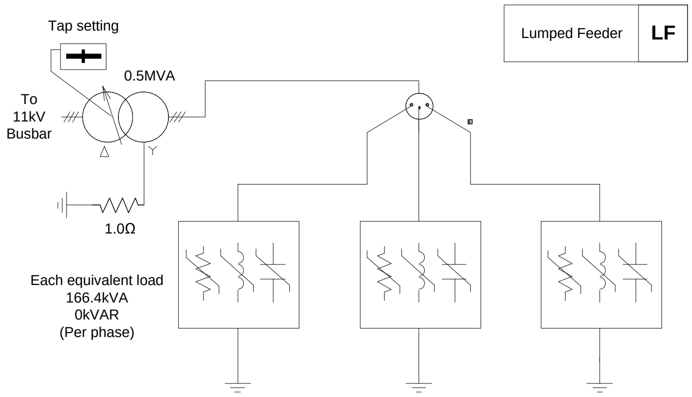

(b) Low-voltage lumped feeder

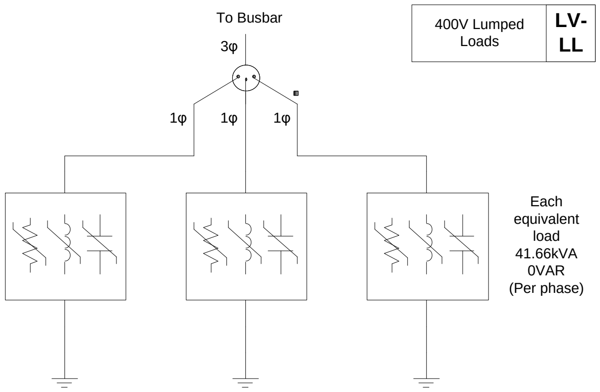

(c) Low-voltage lumped loads

Figure 3.10: PSCAD model: lumped items

These figures show the items in the model that remain lumped. Medium-voltage and low-voltage lumped feeders are represented simply in the UKGN[131]. In each case, the single-line representation of the 3-phase connection is expanded to show components on single phases.

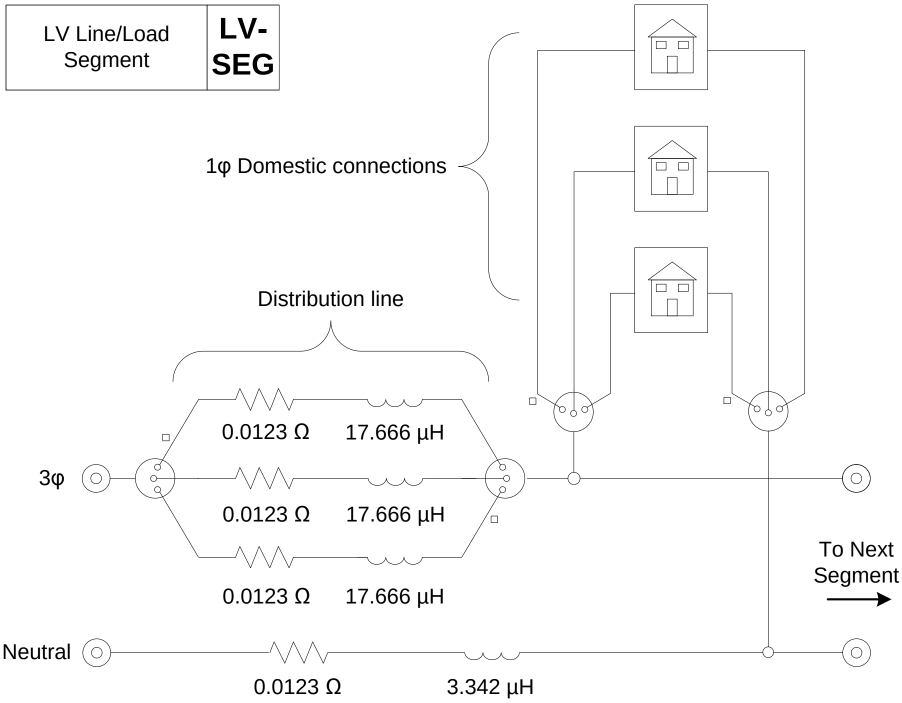

Figure 3.11: PSCAD model: single-phase connections
A segment of low-voltage network is shown. The single-line representation of the 3-phase distribution line (shown left) is split into single phases to show the impedances of each conductor of the distribution line. Clusters of houses are connected across single phases and then to a shared neutral point, shown in Figure 3.12. An LV feeder with disaggregated loads is composed of five similar segments.

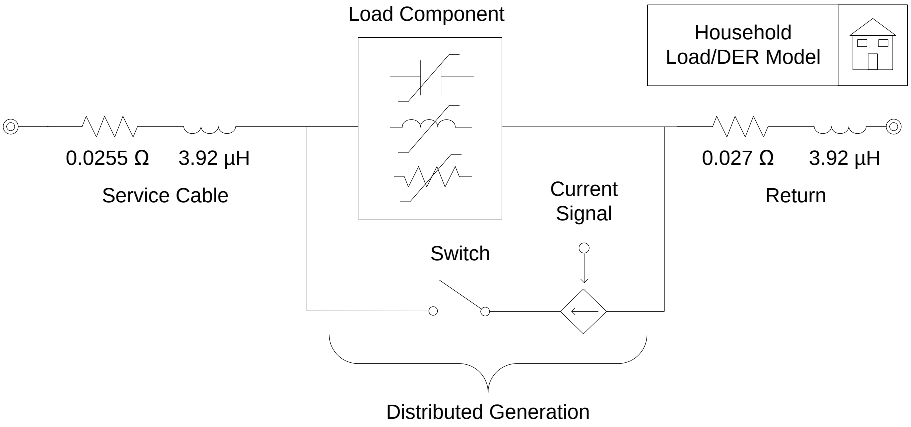

Figure 3.12: PSCAD model: household load and distributed generation
Small clusters of houses are connected to a single phase via a service cable and to a shared neutral. The distributed generation (where appropriate) is modelled as a current source, controlled via the current signal and enabled via the switch. The household consumption is is modelled with the variable load component. The load and generation controllers are not shown.

3.4.1.1 Modelling Results

PSCAD was used to model and simulate the network under normal operating conditions.

For future network growth scenarios, including clustering effects, EVs were modelled in a cluster at the end of the feeder, comprising a third of the total number of houses. This was replaced by EVs in a third of the houses within a single phase, and then extended to a third of the houses uniformly distributed throughout. A similar process was followed for the impact of embedded generators.

Qualitatively, the model behaved as expected: switching in additional load or generation caused behaviour in line with the theory described in Section 3.1.1.4. Clustering DERs on particular phases caused greater imbalance; devices at the end of the feeder exacerbated the voltage rise or drop. In particular, line-ground voltage rise was displayed by the system where a purely real load characteristic was used.

To illustrate likely impacts of individual items of equipment, a single EV charger (at 3kW) was added at the remote end.

At simulation time t=0.6s (to allow the system to stabilise), a 3kW EV charger was connected at the remote end of the feeder. Figure 3.13 shows the change in voltage profile at this point. The effect of the single connection is visible towards the busbar, but signal becomes very small: at the middle of the feeder, the visible change in voltage is 0.034%, and at the busbar this drops to only 0.0034%. At the device, the voltage dropped by 1.1%, or 2.41V. The effect of a single device connection was well within the +-3% D-code limit. Voltage rise and drop are visible on the other two phases, although considerably smaller than the effect on the loaded phase. The resulting unbalance causes current in the neutral wire, with a voltage phasor shown in Figure 3.14.

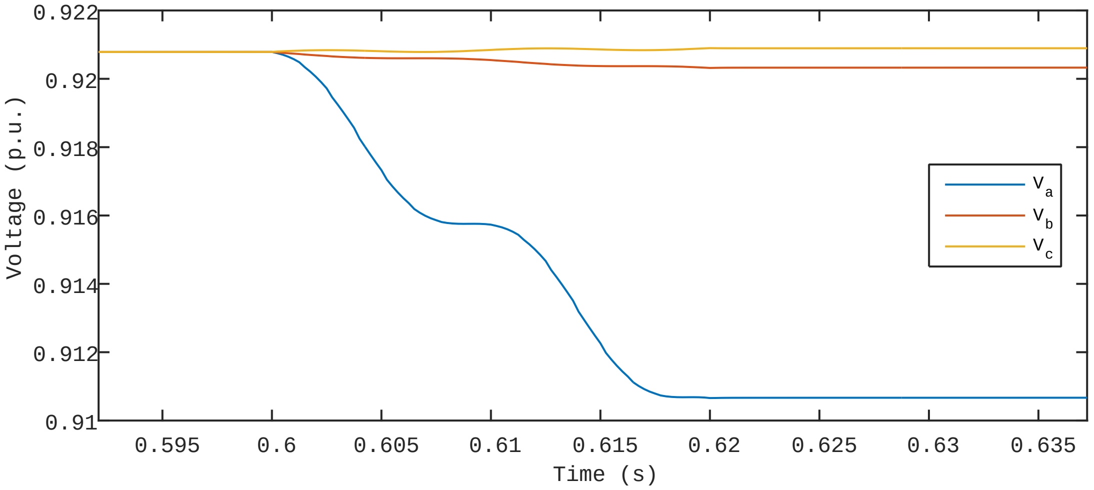

Figure 3.13: Effect of additional loading
Additional unbalanced loading on phase A causes changes in (per-unit) voltage. The lower line shows the voltage drop on phase A. Voltage drop is visible on phase B in the middle line and the top line shows a voltage rise on phase C.

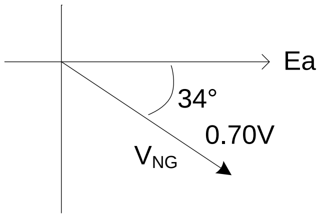

Figure 3.14: Neutral voltage phasor

Adding one additional 3kW load to the PSCAD model caused unbalance and a neutral-ground voltage V_{NG}, similar to those shown in Figure 3.8.

The outputs of this model have informed the development of the techniques for location derivation in Chapter 4.

3.4.2 IPSA+ Model

IPSA+[136] was used to recreate the UK Generic LVDN [131], as shown in Fig. 3.15. This includes a model for a 33kV grid connection and features a double transformer arrangement, connected to a busbar with aggregated loads, of which one is presented in more detail; this, in turn, is further decomposed to show a low-voltage transformer in front of clusters of individual customers.

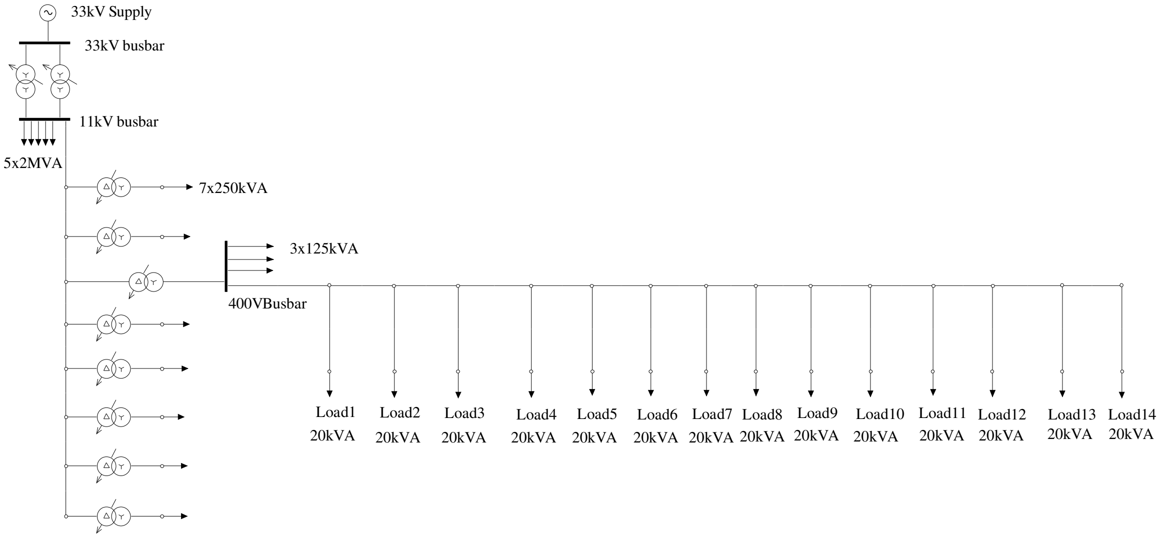

Figure 3.15: IPSA+ model of extended UK Generic LVDN

Along with the existing lumped-feeder simplifications in the benchmark specification, the customers were aggregated in groups of approximately nine and modelled as balanced three-phase loads.

Before additional load or generation was added to the reference case, the validation study specified in [131] was repeated to assess that the model behaved as expected. For each of the points on the network diagram shown in Figure 3.16, the voltage level was recorded (per-unit of nominal at that level) and is shown in Figure 3.17.

Three cases are used here: *a*) full load on 6 11kV feeders; *b*) full load on 3 11kV feeders, minimum load on 3 11kV feeders; and *c*) minimum load on 5 11kV feeders, full load on detailed feeder only. As

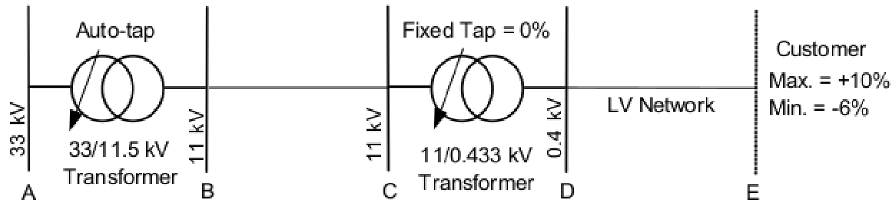

Figure 3.16: UKGN voltage regulation assessment locations
A simplified line drawing of the UK Generic Network, indicating the points A-E at which voltage regulation is assessed. Adapted from [131].

specified in the validation cases, the 33kV/11.5kV auto-tap transformer adjusted its position to keep the 11kV nominal output within the range 11.0-11.1kV (and can be seen in the tap change between case *a)* and *b)*). The source voltage (at Position A in Figure 3.16) was fixed at 1.0 p.u.; validation testing did not include varying the source voltage. The voltage regulation figures are shown in Table 3.1. The results indicate that the model behaves as expected and within voltage limits.

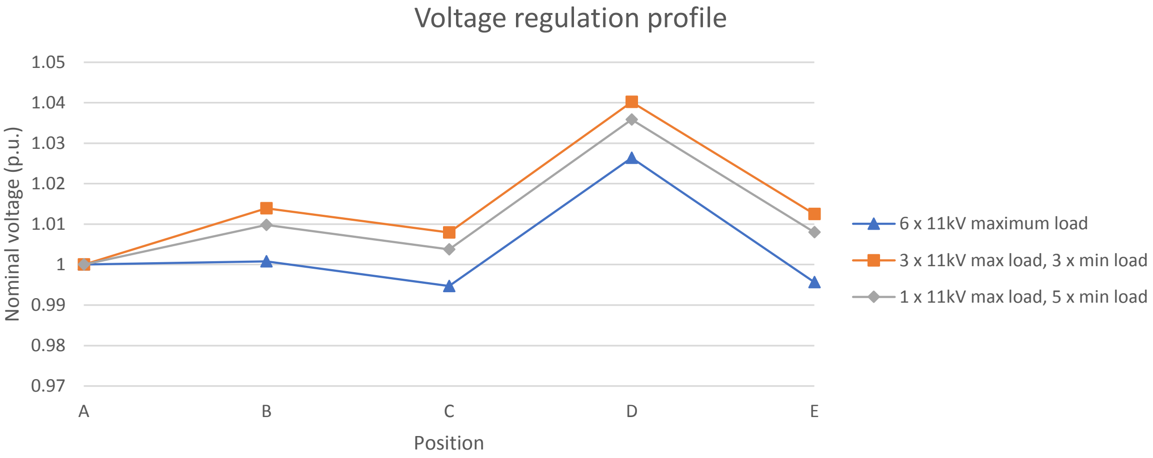

Figure 3.17: UKGN model validation
The validation study in [131] was repeated for three load cases for the pre-extension IPSA model. The figure indicates the voltage profiles at points A-E from Figure 3.16.

Validation Load Case	End load variation	Transformer tap setting
a) 6 11kV @ max, 0 @ min	3.1%	2%
b) 3 11kV @ max, 3 @ min	2.8%	3%
c) 1 11kV @ max, 5 @ min	2.8%	3%

Table 3.1: Voltage regulation and transformer tap settings for validation

To illustrate network growth scenarios, an extension to the basic model was made at the end of the detailed feeder: it was extended by 120m with the same 95mm^2 cable, and from 96 to 134 customers, which were aggregated into 14 groups, labelled Load1–14 on Fig. 3.15. The additional four aggregated loads and service cables were added at uniform distances along the extension. The ADMD was increased from 1.3kVA to 2.1kVA to reflect increased consumption associated with technologies such as EVs and HPs. This quite deliberately pushes the network beyond its original design. Operation beyond the current limits of the transformers may be possible for limited durations, and in a real-world case, at this point some level of network reinforcement or intervention might be necessary.

The software was used to compute the voltage levels at those 14 clustered load points with the system in balanced 3-phase steady-state operation, each steady state referred to below as an "operating point".

The model was scripted in Python in order to iterate through all the permutations of 14 aggregated loads, and later 6 aggregated SSEGs, switched either on or off to find the balanced steady steady state of the system at each operating point. For a load, an "off" state reduced the demand to a nonzero base load of the average minimum daily demand. This created a database of discrete steady-state voltage measurements for the network. This database was used for a computer simulation of the network state to a software load agent operating at the terminal of each aggregated load or generator through a Simulation Agent (SA) broadcasting the appropriate new voltage level at each location whenever the operating point changed. This formed a threshold for testing based on computer memory constraints. All the loads were fully controllable, and an SSEG penetration of up to 43% at 3kW per customer (as per the G83/1 regulations discussed in Chapter 2) could be simulated. This SSEG was modelled as a negative load, a common simplification which means no distinction is made between different sources such as CHP, micro-wind or PV. Overall, this allowed simple load profiles to be created to illustrate various scenarios.

Tap changers were set at 1.02p.u., 1.04p.u. and 1.07p.u. for the 11kV-0.433V transformer. The steady-state voltage at the test feeder varied according to the values shown in Table 3.2. Whilst the overvoltage case appears to be shown in all the scenarios, the undervoltage case only appears at lower tap settings. In these simulations, the other feeders experienced only light loading: the UK Generic Network assumes that the feeders will not experience maximum demand at the same time. The model showed undervoltage cases extending to include more than a third of customers in the 1.04p.u. case where this load factor is increased on other feeders, even before the case where adding additional demand was included for EVHP use. Table 3.2 also shows large values for voltage regulation (VR), and it is likely that DNO-set limits would be exceeded.

	Simulation results		
Tap setting	V_{min} (p.u.)	V_{max} (p.u.)	VR
1.02	0.93	1.11	9%
1.04	0.94	1.12	10%
1.07	0.98	1.14	8%

Table 3.2: Voltage ranges in simulations.

These new steady-state operating point databases were were published [139] and were incorporated into the proof-of-concept and cluster-based MAS testbeds described in Section 5.2.

3.4.3 RSCAD Model

RSCAD, like PSCAD, simplifies a network and represents it as an impedance matrix. Whilst RSCAD allows for construction of arbitrary networks, the complexity is limited by the the size of the matrix for which the hardware can compute the next state in real time. More complex networks can be simulated, but only if they incorporate components at a distance sufficient that they can be considered to be independent over a single time-step. A step time of 50μs is used, so a separation of 15km is required, which is not present in the UKGN. Consequently, several simplifications have to be made in comparison with the PSCAD model.

The overview of the RSCAD model is shown in Figure 3.18. The two-transformer arrangement at 33kV has been replaced with a single transformer at full power rating. Aggregated loads are balanced, three-phase loads, with the exception of the final 400V feeder. There is a limitation at the 400V level, since the neutral line for the (disaggregated) single phase would be common to all three phases at the customer end, but in this case is not shared with the other two phases.

Figure 3.19 shows the remote end of the feeder down to the household level. At the remote end, houses are modelled individually, where a house is at full or zero load. Houses usually retain a baseload (and average minimum of 160W is considered in the UKGN), but in comparison with the case of switching an individual household load of around 3kW, this is not considered significant. As such, at a more disaggregated level, switching events from the customer — by appliance use or by agent-based controller — are more accurately represented.

At each house, a node is shown, labelled $N_1 \dots N_{15}$. At each node, the voltage is recorded for each simulation timestep and output to the RTDS hardware analogue IO cards. Control outputs from control hardware are read by RTDS hardware digital IO cards at each timestep and applied as the state for each switch $BRK_1 \dots BRK_{15}$ which determines the connection state of that house's load.

Whilst the RSCAD model can output to screen to show real-time voltage levels, its main function is for real-time hardware integration into the MAS test-bed described in Section 5.2.5 in greater detail.

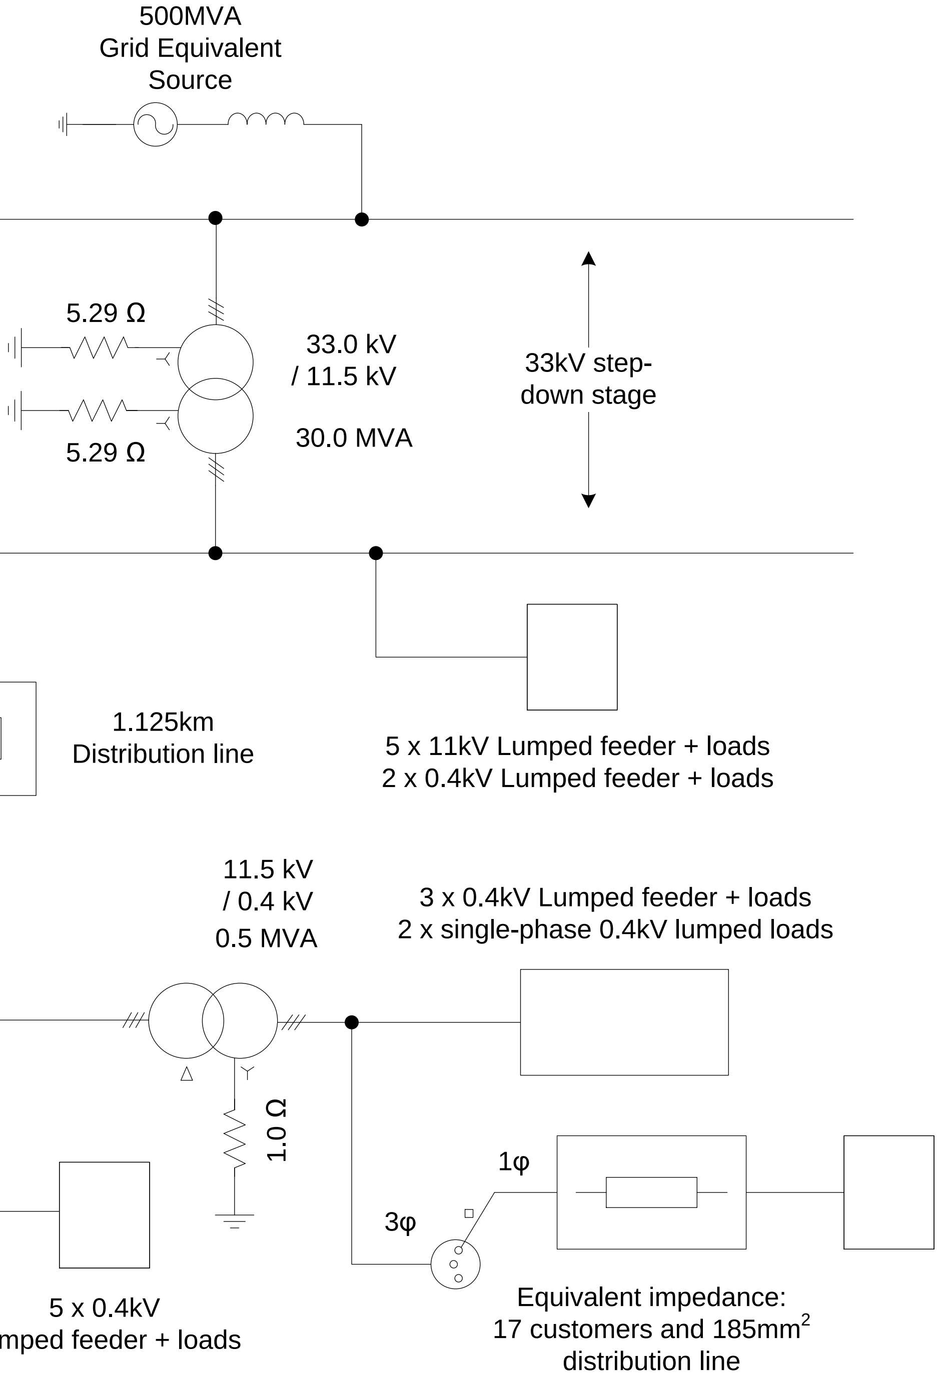

Figure 3.18: RSCAD model: Overview

High-level overview of the RSCAD model. Whilst the RSCAD implementation is based on the UKGN, some significant simplifications are made: the two-transformer arrangement is reduced, loading is much more aggregated, and some feeder properties are incorporated into the load models.

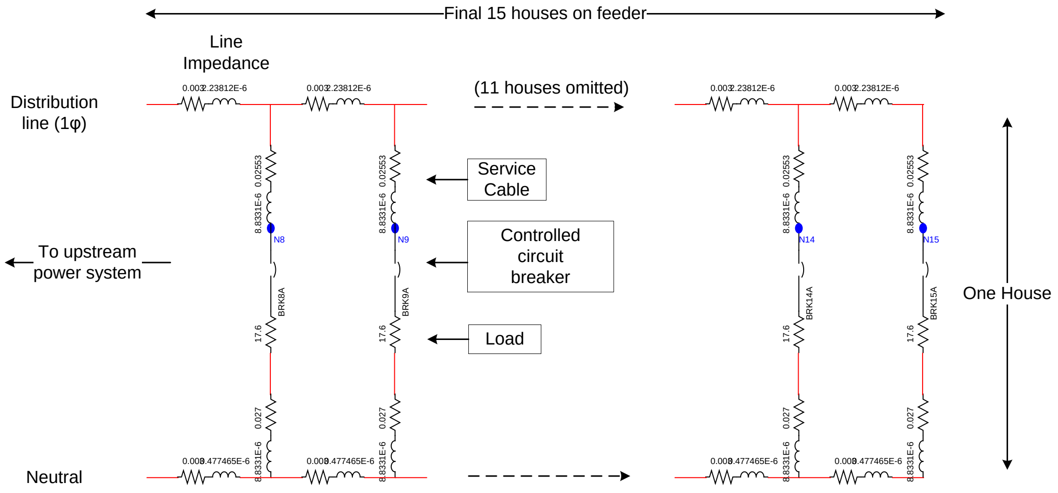

Figure 3.19: RSCAD model: Individual household connections
This section shows an extract of the RSCAD model for one phase, including household load, distribution line and service cables. The individual loads are modelled as static resistances with circuit breakers.

3.5 Conclusion

Relevant basic theory of electrical network behaviour was outlined in Section 3.1. The requirements for more complex network modelling were outlined in Section 3.2, and existing study networks and their qualities were discussed in Section 3.3. Software packages for modelling and simulation were discussed in Section 3.4 along with description of how tools were developed and adapted for the research task for three purposes: demonstrating complex network behaviour, allowing simple software simulation, and for real-time, dynamic, hardware-in-the-loop application.

The PSCAD model showed highly location-dependent effects, in line with the underlying theory; these properties are used in the development of location tools in Chapter 4; the IPSA+ and RSCAD models are then used directly as tools in the Multi-Agent Testbed described in Section 5.2.

Chapter 4

Location in Low-Voltage Distribution Networks

Concepts of location are developed in this chapter in terms of electrical location, geographical location and parallels between them. Latitude and longitude are combined with electrical measurements in an original technique to create a graph of proximate nodes. Electrical network theory and signal processing techniques are used to develop new methods for identifying electrical connections and applied on a dataset of voltage measurements collected from the real distribution network. These combined derived location data are used in constructing network models that are useful for control systems.

4.1 Location

4.1.1 Location in control

The specific impacts of Distributed Energy Resources (DERs) in Low-Voltage Distribution Networks (LVDNs) are explored in Chapter 3. As shown in the literature and from the results of the network modelling, there are several implications: voltage rise, reverse power flows and degraded power quality may result from large-scale integration of Distributed Generation (DG). EVs may cause network issues (primarily thermal overloading of transformers and voltage drop) where uptake reaches a third of customers, or even lower[43]. Connection phase affects unbalance and clusters of devices will impact more heavily than if they were uniformly distributed[132]. Thermal and voltage issues due to heavy loading are exacerbated by increasing the distance (and therefore line impedance) between equipment and transformer busbars: the LV network has resistive conductors, so power is wasted through ohmic heating of the distribution line.

Connection policies can be used to help prevent these problems. “Headrooom” is the spare capacity in the network for additional connections. Where headroom is available, A first-come-first-served approach would allow connection of power generation equipment up to the available system capacity, but could be considered unfair by customers paying the same service charge to DNOs but not permitted to connect SSEGs because of others installing equipment earlier.

Alternatively, controllers could be added to DERs to mitigate some of these issues — for example, in a feeder with 100 customers, a new PV array may be permitted to supply power back to the grid for up to one hundredth of the available headroom in the equipment. However, this may not be optimal, and result in a small handful of constrained generators operating alongside considerable unused headroom.

A control solution that does not take the wider network context into account will either operate more conservatively than required or risk the case of a locally acceptable operation causing unacceptable global impacts that are not assessed in any way by the control system. Consequently, a location-aware controller may be able to determine and operate within network-specific control requirements, or act in a more effective manner than a controller that does not have this information.

4.1.2 Describing location in LVDNs

There are numerous properties that can be considered as location factors in electrical networks, such as:

- nominal voltage level
- impedance between nodes
- connection phase
- position on radial feeder
- intermediate nodes

These properties begin to become analogous to geographical properties — impedance can be thought of as a kind of electrical distance; intermediate nodes are nearer neighbours. These properties are important: as shown in Chapter 3, loads connected at the same substation may contribute to a common power limit, but if two nodes are on different phases then their influence over each other's voltages may be limited.

These properties, then, could be represented for use in control. An appropriate representation must be developed to allow for their analysis, as conventional electrical network models only provide a partial solution. Furthermore, as discussed in Section 2.2.6, the availability of network data cannot be assumed, and may even be undesirable. In the absence of existing models or complete data, there is an open research question of how to construct a network model using locally available information that is useful for control.

4.1.3 Local and global information

A centralised control system could be pre-programmed with network topology including the location and properties of any DERs. This system could use some centrally-held global information to determine optimum settings to control individual devices. However, this has several shortcomings — in scalability and reliability[4], including single points of failure and communication bottlenecks, as well as issues with system maintenance and scalability as components are added under G83/1. A database of known network information is problematic: it may be commercially sensitive or protected for security purposes. Since the UK system is divided into regional monopolies, coordination of multiple system operators would be necessary. The multi-stakeholder context outlined in Section 1.3.1 makes it problematic to capture or describe the global context. Data may not exist to a relevant detail level in all locations. The database would have to be maintained and updated regularly.

The rationale for decentralised control is similar to the reasoning behind why an individual controller may be improved through self-configuration with location information and learning about its connections. Decentralisation and self-configuration (where the network topology is derived locally) fit a multi-stakeholder context where generation, load and network assets have different owners.

Consequently, a single owner for all the relevant data has significant disadvantages. Further, for a single owner for all the relevant data even to exist, there would need to be a change in the regulatory set-up of the power network.

Self-configuration means that user intervention is not required and there is no need for a DNO to release sensitive data or confidential information about network structure. This lowers the barriers to entry in LVDN control: the use of plug-and-play agents reduces the cost of implementation.

4.1.4 Electrical network and modelling

Figure 4.1 shows a conventional single-line diagram of a three-phase radial network. In comparison, Figure 4.2 shows an electrical network superimposed on a street map. It is generally indicative of the layout of the network – the phase connections are visible, as are the locations of various transformers. The properties of the power lines are not included, nor are the connections to individual houses. However, the layout can be used in constructing the model shown in Figure 4.1.

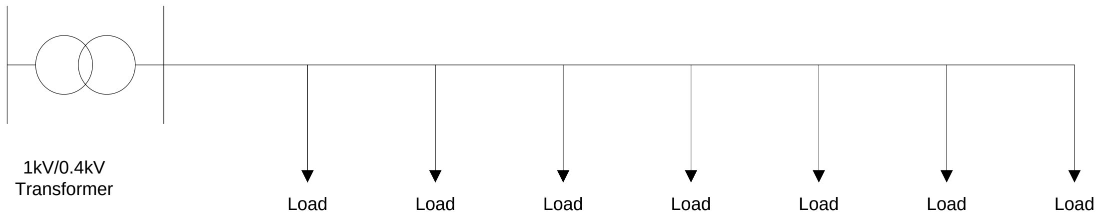

Figure 4.1: Radial feeder. The transformer (left) steps voltage down to 400V (or 240V single-phase). The single line running left to right represents a three phase distribution line; customers (usually) connect to a single phase of the supply.

Models inevitably include simplifications. The single line in Figure 4.1 represents all three phases; the loads are clusters of houses – this model only goes part of the way to describing some connections shown

Figure 4.2: Electrical network overlaid on a map
The image shows the electrical network overlaid on an outline street map of an urban area. Single, two- and three-phase lines are visible. The connections to individual properties are not marked. Figure reproduced from [140].

in the street map, and this sort of representation is widely used for static, balanced network studies.

There are other ways of representing distribution networks in a way that facilitates analysis. Kersting[140] details an effective network model that uses a geographical area, such as a housing estate, covering all the customers supplied by a transformer. The area is bounded by a trapezium. It is assumed to have uniform loading across customers, and the area thus has a "load density". Whilst the exact topology is not explicit in the model, the information is sufficient for estimating network behaviour (for approximating current draw and voltage drop) at any given point in the area and in aggregate.

The accuracy of such a model could be improved by the use of loading and impedance density functions for the given area — although when the assumption of balanced operation is removed, the advantages of the model simplicity are lost. Such a model is also difficult to construct without prior information about the characteristics of an area, including a method to allocate its geometry. Rather, by taking a graph-based approach, the nature of connection between network elements can be analysed: the system may be represented as a graph consisting of a number of customer nodes connected by edges with electrical network properties.

Depending upon the goals of network analysis, the underlying topology could be simplified even from Figure 4.1. A path structure (see Figure 4.3) might be appropriate for a voltage control task, since the impact of a given node is related to its distance in a resistive distribution network.

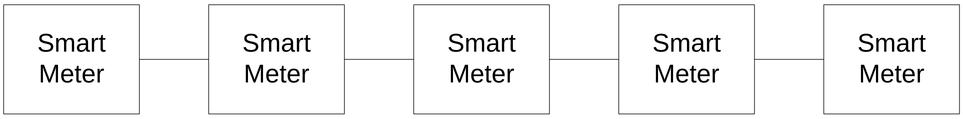

Figure 4.3: Simple path
Smart meter nodes connected as a path: they form a chain according to their position on a power line.

However, this simplification may go too far for a general-purpose model. It does not allow for information about phase, and branches in the network and clusters of houses (of the kind visible in Figure 4.2) are ignored.

To expand a little, some complexity can be included. Figure 4.4 gives a picture of clusters of houses connected to a transformer: Individual users are connected by service cables at a bus along a radial feeder (to points labelled as "dummy nodes" in Figure 4.4) – they are a point where the network branches but no active device exists at this point. This is a more complex graph: it uses a tree structure. Further detail can be added: for example, edges connecting the nodes may be given relevant properties such as a number to represent the impedance between them.

End (consumer) nodes connected to one phase are connected in terms of total load, as per Chapter 3, but separate to those on another phase for the purpose of line voltage. The structure of a graph used for control must then be relevant for its purpose.

Figure 4.5 gives a simple tree structure appropriate for thermal limits control. It neglects the charac-

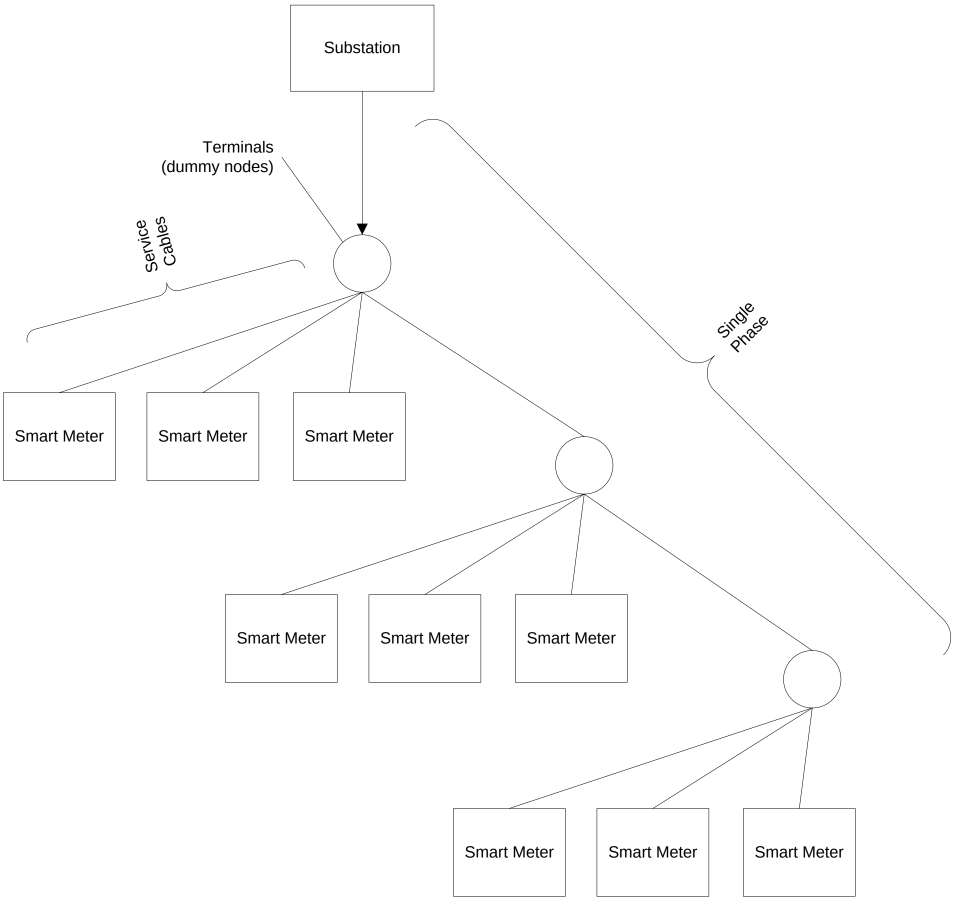

Figure 4.4: Tree view
Expanded tree view of distribution network. Smart meters are grouped around terminals; edges are more analogous to distribution lines.

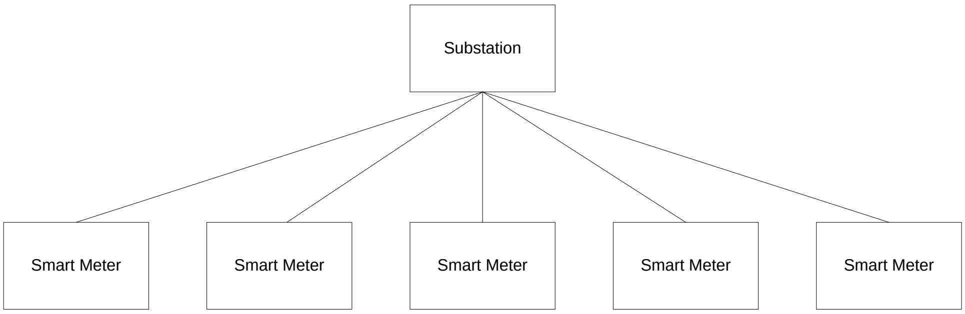

Figure 4.5: Simple tree
A smart meter at each customer is connected to the same transformer.

teristics of the distribution line. A number of customers (the nodes, represented by a smart meter) are connected (via the edges of the graph) to the common transformer node. Despite the graph not representing physical powerlines connecting devices, it may still be an adequate model for a thermal limits control task where all the nodes have an impact on total current draw.

As in both Kersting's area-based approximation[140] and the single-line simplification, the network model that is used by any individual controller need not necessarily adopt a topology that strictly follows that of the electrical connections. Rather, it is important that the model is useful in allowing the controller to assess its impact on the network. Instead of attempting create a model of all three phases, a hybrid of these simplifications outlined above would allow a controller to assess its link to other nodes for voltage and current. This has the advantage that the information held about the edges reflects the impact of the relevant node and can be updated as the result of controller action becomes known.

The final graph will thus combine a geographical approach with electrical assessment to form a model that will reflect electrical connection properties (rather than exact powerline layouts) to inform controller choices and allow for revision over time.

4.1.5 Types of Location

4.1.5.1 Electrical Topology Analogues

Impedance of a distribution line is proportional to its length, so for two connected devices the distance between them is analogous to the impedance between them. This premise forms the basis for the calculations in [140], which assumes a uniformly-distributed load in a residential setting to estimate voltage drop along a feeder.

Information for electrical network location information can be derived from several sources, such as an address within an IT network, Global Positioning System (GPS) coordinates, or by measuring local electrical network activity and comparing it with other locations in the system.

Adjacent houses may be connected to the same part of the electricity network. However, this is by no means guaranteed: houses may be connected to different phases, or even different feeders entirely.

Distant houses in rural areas may be connected over much larger distances than those in urban areas. Despite this, location data of various sorts offer an approximation to electrical network data. These can include geographical location, area derived from the mobile phone network, or through communications network properties.

As well as referring to location as the point of connection in an electrical network, location can also apply to geographical position or connection to an Information and Communication Technology (ICT) network. There may be overlap between these locations: for example, we may suppose that two geographically adjacent houses may be connected to the same electric distribution feeder. Location information can be derived from many sources: radio signals from GPS satellites give a precise geographical position; a postal code input upon installation provides an approximate area. Many smart meters are networked via mobile phone towers, which could be grouped by cell ID, or mapped to physical coverage area. These geographical properties can be used as an analogue to electrical location, although proximity is not necessarily a guarantee of a common connection. ICT network proximity (measured by latency or by common access point hardware) may indicate a common electrical connection. In particular, Power Line Communication (PLC) can be used directly to infer electrical connections between devices, since the distribution line itself is used as the medium for passing data.

Electrical network characteristics can be used to identify the nature of connections by using direct on-line measurement. Analysis of signals on the electrical network will show differences in connection phase, as well as identify where devices do not have any common connection.

4.1.5.2 Multi-Factor Location

Multiple sources of location information can be used together. Strong indicators — such as identifying common signals on the electrical network — can be combined with appropriately weighted weaker indicators, such as geographical coordinates or areas, to give both a location and a value for confidence in that assessment. A process whereby the agent learns about the effectiveness of its control actions allows a dynamic multi-factor location system to adjust its confidence in topology. Machine learning used in conjunction with assessment of the impact of each action would provide a degree of intelligence in control systems in a network.

Several layers could be used: control graphs of different types, such as Figure 4.5 or 4.3, could be combined with an overlay graph where each edge is allocated a weight that corresponds to the impact an action may have with respect to those nodes. Due to a reflexive process of adjusting confidences based on location data sources and observing responses to control decisions, the correctness of the initial tree becomes less relevant than knowledge of what its impacts are as the controller learns from experience.

This learning, combined approach is beyond the scope of this work, but the implementation of multiple assessments of location provides first steps towards an intelligent capability in this regard.

4.2 Location from Geographical Information

While latitude/longitude coordinates do not provide a complete picture of electrical connections, they can be used as a first filter to drastically reduce the search space of connected devices.

The Global Positioning System (GPS) allows a radio receiver to determine latitude/longitude coordinates accurate to within a few feet by the use of highly accurate time signals sent from a network of satellites. These can be used to determine the geographical distance between devices.

There is currently a lack of reliable or ubiquitous communication links between grid-connected devices, and the cost of deploying wired technology is high. The approach to the UK smart meter roll-out includes the use of the mobile telephone network[141], where each meter contains a SIM card to allow communication through GPRS, 3G or 4G connection. Under a standard cellular network, each device connects to a base station which covers a certain geographical area; the unique identifier assigned to each receiver can be used to identify devices within the same 'cell'. The ID number does not provide geographical location. Users may be nearby but in different cells; in areas of overlap the cell identity may change. The cell ID — along with signal strength data and other detected cells — can be used to provide a more precise location. There are companies partnered with the mobile phone network operators able to provide Geographical Information System (GIS)) data including base station ID, coverage areas and geographical location as a basis for a location service. Public implementations such as the W3C or Google Geolocation Application Program Interfaces (APIs) provide a direct location lookup using these data. For additional accuracy, other signals such as nearby WiFi networks may be included. Areas with high density of mobile cells and additional signals may provide accuracy to a few tens of metres; in extra-rural areas, this is reduced to the kilometre scale.

Postcode lookup provides a hybrid between geographical point and mobile-like cell structure. Postcodes are areas that cover clusters of nearby buildings, although some high-volume post users are assigned a unique postcode. Coverage areas are publicly available. The postcode of the device would be entered once during commissioning. Postcodes are usually static once assigned, but some areas have been restructured over time.

PLC offers one solution: devices superimpose data signals onto the electrical network. This provides an easy way to find nearby devices: they are connected to the same information network bus, and any connected to the same feeder will be able to read those signals. Additional equipment would be required to filter signals going up the distribution network, so as not to cause interference, and inter-feeder (or inter-phase) communications would require an addressing and bridging system. The system would need to be designed not to interfere with any existing PLC equipment.

Similarly, directional antennae may be used in conjunction with a radio communications network to identify the distance to neighbours sharing the link. Each agent maintains a list of neighbours whose signals it can detect and the signal strength and direction indicates its location. No absolute geographical position is known.

IP geo-location uses communications network connection information to determine location. At its most basic level, devices connected to the Internet are assigned an address which can be used to establish the Internet Service Provider (ISP), and provide a city of connection. Sometimes, more detailed analysis is possible, where network exchanges are known and additional data (such as packet round-trip time) is available, but the resolution and accuracy of this approach is insufficient for LV network connection information. It is possible, however, to assign addresses based on a known location; devices would then address a particular subnet to communicate with connected devices at the appropriate level. This unfortunately begs the question, since correctly assigning an address requires complete knowledge of the distribution network.

The GPS routing protocol requires each device to be connected to a tree of (virtual) routers, this increases points of failure, and creates an issue of maintenance and responsibility. Similar approaches using agents at dummy nodes are also excluded for the same reason.

For the purposes of this work, the acquisition of geographical data is simplified to using the coordinates from GPS to provide a straightforward geographical location that does not rely on additional (remote) equipment, or complex mutliparameter analysis.

4.2.1 Geographical tree creation

In the first instance, it is assumed that the distribution network is radial, has no spurs, and that the geographical distance increases with electrical distance, as in Figure 4.1. The aim is to create an undirected graph (Figure 4.6) of connected nodes on the same phase, where each node is a control agent at the customer.

Upon startup, a new agent registers with the agent platform and directories. It registers its own geographical coordinates, and requests those of its two nearest neighbours. This allows it to begin to construct a graph of nodes, shown in Figure 4.7, where the new device N' is situated between existing devices N_1 and N_2 on the network and the edge is the geographical distance. The graph is represented internally as a connection matrix. Self-connections are excluded and disconnections are represented as ∞. Each agent stores a partial representation of the full network topology.

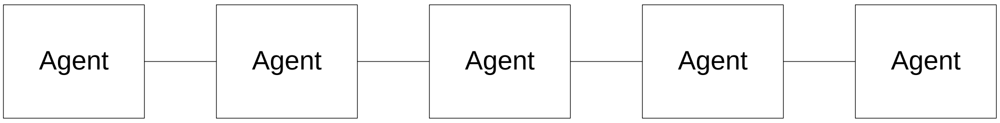

Figure 4.6: Agents are connected in a path. according to their position

Representing connections in this graph form already allows the agent to make deductions about the

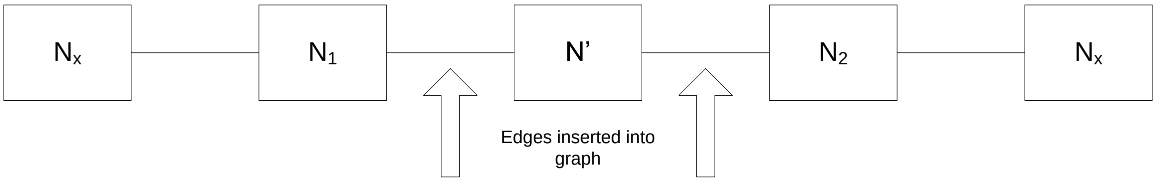

Figure 4.7: Creating a graph of connected customers. Here, a new customer is represented by a device N', and it is inserted into the graph G_1.

structure of the network: if the distances $N_1 - N'$ and $N_2 - N'$ are both less than $N_1 - N_2$, the device is between them (true in the 1D case below, requiring verification in the 2D case due to the possibility of complex geometry); otherwise, it is at the end of the graph. In the case where N' is not the most remote on the feeder, N_2 has knowledge of a node (e.g. N_3) on the feeder beyond N'; having found a connection, N' announces its existence to N_1 and N_2; if appropriate, N_2 responds with details of N_3.

If this is connection is found, N' is inserted into G_1 between N_1 and the next closest neighbour in G_1. This change is then propagated through the tree. If, however, the test shows the nodes are not connected to the same phase, N' selects the next-nearest neighbour that is not in G_1 and repeats the process. If it exhausts G_1, it has no neighbours on the same phase or feeder, and is the start of a new section of distribution network.

Figure 4.7 shows the updated graph $G1'$ after successfully identifying its location. The node N has been registered with the global database. The edge connecting N_1 and N_2 must be removed, and two connecting edges put in its place. Rather than the tree structure in Figure 4.5 and the more realistic Figure 4.4, this graph describes the simplified connection shown in Figure 4.3, and it is assumed that the order of two devices at a common terminal is not important.

4.2.2 2D Geometry

If a node fails to find a connection altogether, it forms the start of a new graph.

One potential shortcoming of this method is the potential for disjoint graphs to form where a connection existed. As well as in the complex geometries that may occur in an urban setting, this may also occur at a boundary such as an urban-rural transition (shown in Figure 4.8) where a high concentration of users drops off suddenly to long, sparse connections.

R_2 is the closest electrically connected node to a new node at R_1. If the node ends its attempts too early, then it will form an isolated tree, and nodes connecting between R_1 and R_2 will be isolated from one another despite their electrical connection.

Geographical distance does not necessarily increase with electrical distance. Consider the terraced housing in Figure 4.9. In this case, Load 1 is geographically closer to Load 12 than to Load 6. However, it is electrically more distant, so creating the graph in Figure 4.7 is not appropriate. In order to examine this, the node performs a walk.

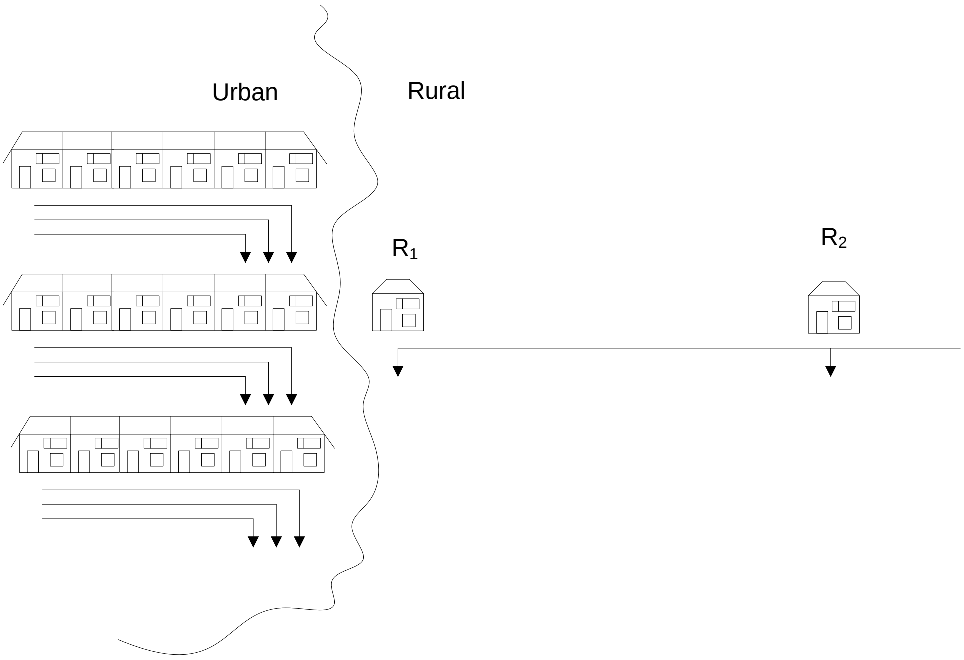

Figure 4.8: Urban-rural transition. In this exaggerated illustration, the node R_1 is close to a large number of electrically disconnected nodes, and may end up isolated from its neighbour R_2.

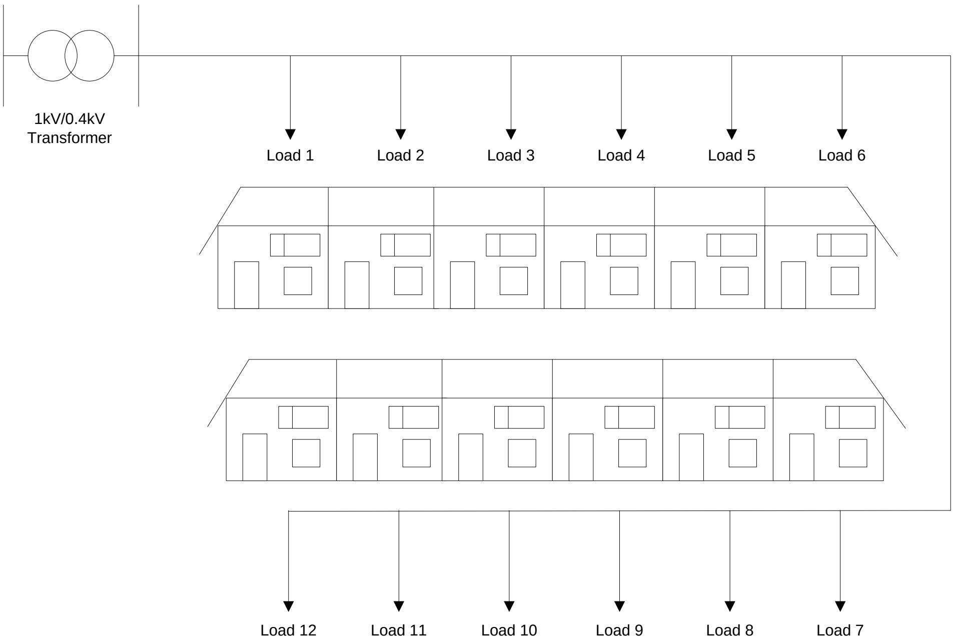

Figure 4.9: Geographical layout of network topology
The figure shows a hypothetical network with 2D geometry where a housing terrace has electrically distant but geographically proximate nodes.

4.2.2.1 Tree walk

Consider nodes at Loads 1, 3, 5, 7 and 9 in Figure 4.9. A new controller enters at Load 12, and it is closest geographically to the node at Load 1. N' receives the graph G_1. It walks the tree (potentially in two directions, if N_1 is not at the start or end of the tree). If, at any point, the next node is closer to N' than the previous one encountered on the walk, the network has the geographical property outlined above: here, distance from Load 12 increases for Loads 1, 3, and 5 — but decreases from Load 5 to Load 7, and again from Load 7 to Load 9, leading it to conclude it has this geometry. The node will be finally added at the minimum geographical distance.

4.2.3 Proof of concept

In order to test these initial approaches to tree creation, the UK Generic LVDN[131] was assumed to represent a linear UK residential street. The tool developed for testing in Section 3.4.2 was used for this purpose: arbitrary location co-ordinates were assigned to each load. Figure 4.10 shows the schematic of the proof-of-concept network: an extract of the remote end of the UK Generic LVDN is shown; the dashed arrows to the left indicate the upstream network beyond the 11kV busbar. The loads (indicated with arrows and labelled) are connected along a hypothetical linear housing terrace, depicted below. Each load point is allocated arbitrary (x,y) geographical coordinates in 3D space. An additional load, indicated by the dashed line, will be connected for illustration. A controller at each load will determine its own local connection graph.

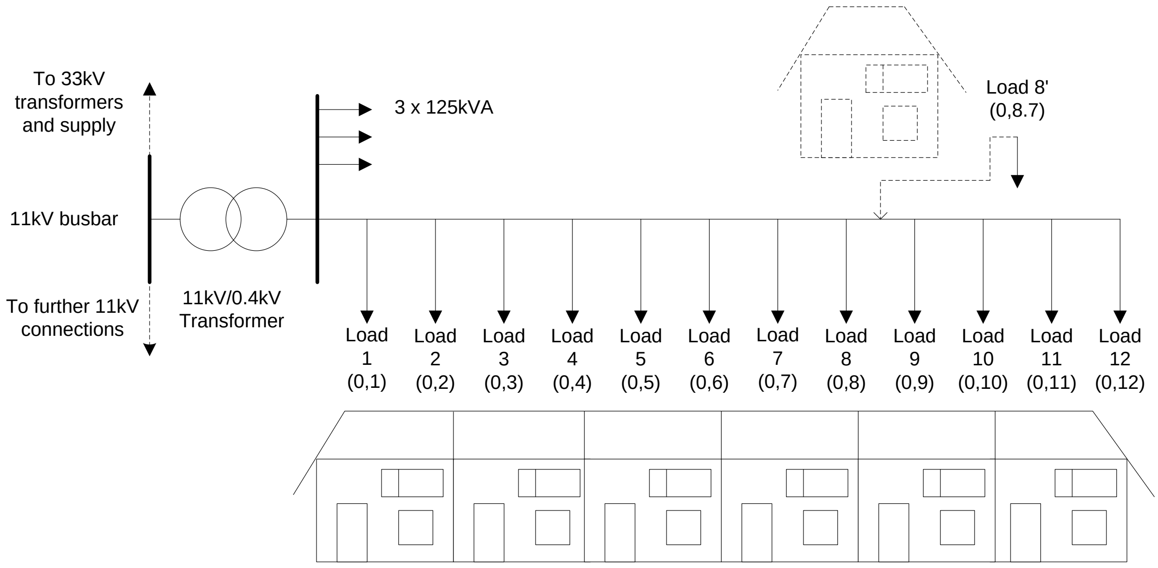

Figure 4.10: Proof-of-concept geographical graph
The figure shows the remote end of the UK Generic Network[131], with loads at coordinates (x,y) represeenting a hypothetical linear housing terrace. An additional load, indicated at the house with dashed line, will be connected.

A Multi-Agent System (MAS) was used to represent the load nodes[1]. Each agent supported the algorithm developed above in Section 4.2.1. The size of the graph representation held by each agent was arbitrarity limited: for a full implementation, a trade-off is required between minimising graph size for internal storage, processing and planning, whilst keeping a reasonable representation of its sphere of influence[2]. For the purpose of the proof-of-concept, each node would attempt to build a graph of three nodes only, to demonstrate the cases in Section 4.2.1. This forced the graph to be limited and without representing the complete network, whilst allowing each node knowledge of its surrounding peers.

The load nodes each joined the network in sequence 1-12. To illustrate, the output of the end node L12 is given in Figure 4.11, edited for readability.

```
Load Agent 12 L12
Current location 0 12
[...]
finished all setup
[Find neighbours]
Received neighbour list (# 0 L12; # 1 L11; # 2 L10 [...])
[Measure distances]
[Connection matrix]
[   ][L12, L11, L10]
[L12][inf, 100, inf]
[L11][100, inf, 100]
[L10][inf, 100, inf]
```

Figure 4.11: Edited debug output from proof-of-concept graph agent

The agent starts; takes a list of nearby agents, and draws up the connection matrix. Node 12 is at the extreme right of the network. Figure 4.12 shows the simple graph from the connection matrix output. Similar graphs were calculated by each node L1...L11 and verified against the original connection of the test feeder in Figure 4.10.

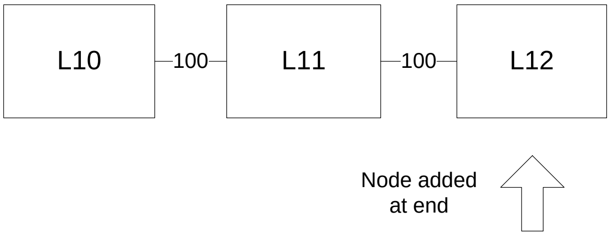

Figure 4.12: Local graph for Node 12

Next, to examine node insertion, the additional load shown in Figure 4.10 was added at (0, 8.7). Edited debug output is shown in Figure 4.13. This shows Load 8.7 successfully inserted into the graph between L8 and L9. The graph from the connection matrix is shown in Figure 4.14.

At this point, an update would be propagated to the surrounding nodes via a whisper protocol[3],

---

[1] For reference, the technical set-up for the hardware platform and multi-agent framework are described in detail in Section 5.2.3

[2] In Section 4.3.7, it is argued that computational intensity of processing the full graph of N nodes increases with N^2.

[3] Simple, unicast message propagation, also commonly referred to as *gossip* or *telephone*.

```
Load Agent L8.7
Current location 0 8.7
[...]
finished all setup
[Find neighbours]
[Nearby neighbour list] (# 0 L8.7; # 1 L9; # 2 L8 [...])
[Measure distances]
[Connection matrix]
[    ][L8.7 L9 L8 ]
[L8.7][inf, 33, 66 ]
[L9  ][33, inf, inf]
[L8  ][66, inf, inf]
```

Figure 4.13: Edited debug output from node insert agent

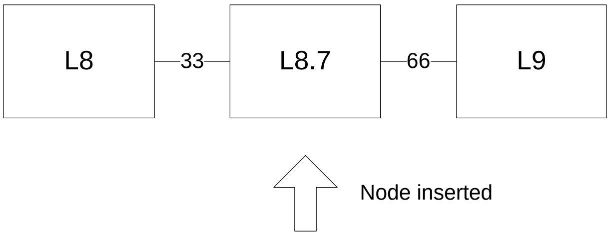

Figure 4.14: Local graph for node L8.7 showing node insertion

such that the other graphs would be updated after the joining event, but this was not performed for proof-of-concept.

Similar graphs were calculated for each node L1...L11 and verified against the original connection in Figure 4.10. Thus, each agent L in the network has constructed a graph G_L that represents its geographical conditions constrained by local knowledge. These geographically-derived graphs will require verification of electrical connection, which will be discussed in Section 4.3.

4.3 Location from electrical signals

4.3.1 Overview

Local electrical measurements of voltage and current can be used by a node to infer electrical location information.

Local current measurements can be used to determine whether events on power systems are caused by loads at the customer. These events can be characterised as *upstream* or *downstream*[142]. These can be used for determining the equivalent impedance of the grid at the node. However, this does not directly give information about the impact of nearby nodes and requires combination with other information to characterise the nature of the connection between them.

Another approach is to examine local voltages. Consider the electrical network consisting of an AC generator, an RL distribution line and a load impedance, shown in Figure 4.15.

When we consider two users connected to a common electrical network separated by a power line

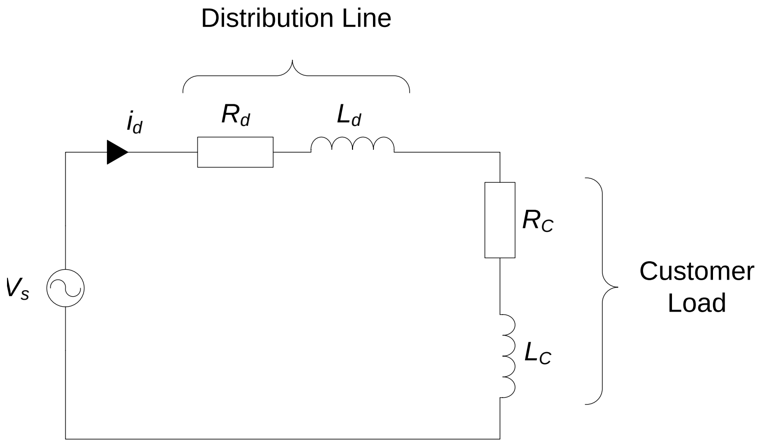

(a) Single customer

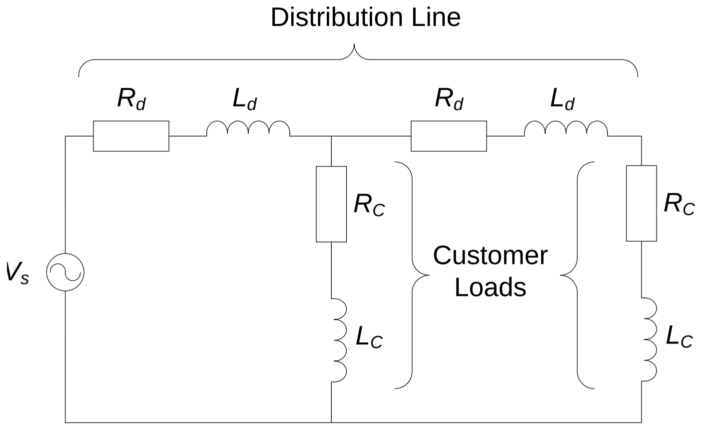

(b) Two customers

Figure 4.15: Simplified power system model
The two diagrams show the model for a single customer and for two customers connected in a radial network.

(Figure 4.15(b)), there will be a difference in the line-neutral voltage at those two nodes. The voltage for a household consumer is permitted to vary from a nominal value. In the UK, this is 220V, ±10%[39]. In the traditional distribution network, the voltage will decrease over the line with increasing line length from the transformer.

The voltage along the line may be altered by the voltage output of the transformer. The windings of the transformer will have a number of 'taps' that allow connection to some numbers of turns that deviate from the nominal ratio by a small amount. The transformer will usually be configured with an output voltage a little above the nominal value so that at the remote end it will still be above the lower permitted voltage limits after accounting for the voltage drop over the line. A change in voltage at a higher voltage level will change the voltage lower down the system, since the transformers merely operate on a ratio of one voltage level to another. At higher voltages, the winding tap of the transformer may be changed whilst it is still energised; these are not frequently changed at the low voltage end of the system, although this may change as On-Load Tap Changers (OLTCs) become cheaper and viable as these active network technologies are applied deeper into distribution networks.

A change in load will change the current through the line. Because of the resistance of the line, increasing loading will decrease voltage. Conversely, there will also be a rise in voltage towards any generator (e.g. PV panel) connected to the system.

As previously described, changes in transformer taps, load or generation will cause changes in the magnitude of the voltage that will be experienced by two electrically connected nodes, with the change including a phase shift dictated by the impedance of the line between them.

Household consumers use the electricity network by connecting to a single phase. Areas are split so that the aggregated consumption patterns give, on average, balanced power use across the phases. However, in reality they are not perfectly balanced.

A radial feeder experiences voltage drop across its length. By comparing voltage levels with surrounding nodes, it may be feasible to create an ordered list of nodes with increasing distance from the transformer. This requires a simplifying assumption that neglects the contribution of the impedance of the service cable to the property; uneven loading; and different lengths of cable. Depending on the application, however, strict ordering may not be necessary: the size of the error caused by this assumption may determine whether a poorly ordered list is a sufficient approximation of the reality to remain useful for control. In any event, however, the voltage rise caused by DG unfortunately rules this out as a reliable technique as networks transition to smart grids incorporating DERs.

Instantaneous measurements, then, do not provide sufficient location information. Observations over time allow a node to observe the behaviour of the network from its location. Communication with other nodes then allows for these behaviours to be compared, and for network topology information to be inferred. It is useful to consider how location data that might be derived from observations over different timescales, and the challenges that arise from these respective approaches.

4.3.2 Signal processing techniques

In considering how local measurements taken over time may be used to infer location, there are existing signal processing tools that can be adapted to this application. The following section gives an overview of signal cross-correlation, hypothesis testing and filtering; the subsequent sections then discuss the available signals and how these techniques may be applied.

4.3.2.1 Cross-correlation

The cross-correlation of two signals f and g gives a measure of the similarity between them.

$$(f \star g)(\tau) \overset{\text{def}}{=} \int_{-\infty}^{\infty} f^*(t)\ g(t+\tau)\,dt \tag{4.1}$$

That is, cross-correlation at time τ is the sum of the product of each data point pair $(f(t), g(t+\tau))$ of the two waveforms.

For a system using sampled (i.e. discrete) data, normalised to the number of samples,

$$(f \star g)[m] \overset{\text{def}}{=} \frac{1}{N} \sum_{n=0}^{N-1} f[n]\ g[n+m] \tag{4.2}$$

Equation 4.2 allows finite-length, discrete signals that share sampling intervals to be examined: the result is a new sequence that is the combined length of series f and g. Whilst 4.1 is centred around time $t = 0$, the resulting series of 4.2 begins at sample $m = 0$.

Where the waveform is most similar, the product is maximised. A time delay between the two signals becomes apparent, as the maximum will be at some value that is not $\tau = 0$. For sampled data, this means that a maximum at $m = n$ indicates there is no phase shift between the two signals.

4.3.2.2 Hypothesis testing

4.3.2.2.1 Choosing from alternatives

In considering the application of cross-correlation to establish the existence of a connection between nodes and their phase allocations, there are several alternatives that are possible. The following cases are possible and of use in the control problem for voltage:

- nodes not connected to the same feeder.
- nodes connected to the same feeder, but different phases
- nodes connected to the same feeder on the same phase

Other location information may be useful - multiple feeders may emerge from a single transformer, and the information might be used in control for thermal limits at that transformer. For voltage limits control, however, these are the appropriate cases to include in the electrical topology map since the modelling in Chapter 3 shows the voltage linkage from customer use is minor across multiple feeders.

Hypothesis testing provides a method for choosing between a number of options given a test statistic, based on the probability of the individual observation occurring for each alternative.

4.3.2.2.2 Probabilistic Approach

The probabilities of the cases occurring may be known in advance, *a priori.* The probability, P, of the hypothesis H_n occurring is

$$P(H_n) \tag{4.3}$$

The probability, P, of the hypothesis H_n being correct, given a particular observation x, is denoted:

$$P(H_n|x) \tag{4.4}$$

These probabilities are after measurement of x, and hence *a posteriori.* A simple decision rule is to choose H_0 if

$$P(H_0|x) > P(H_1|x) \tag{4.5}$$

That is,

$$\frac{P(H_0|x)}{P(H_1|x)} > 1 \tag{4.6}$$

The probability of a particular observation can be expressed as a Probability Density Function (PDF), *a priori* in $p(x)$ and *a posteriori* in $p_n(x)$ for each case H_n. Then,

$$P(H_0|x) = \frac{p_0(x)P(H_0)}{p(x)} \tag{4.7}$$

This can then be substituted back into the test in Equation 4.6 to give:

$$\frac{p_0(x)P(H_0)}{p_1(x)[1 - P(H_0)]} > 1 \tag{4.8}$$

where $P(H_1) = 1 - P(H_0)$. Rearranging, choose H_1 if

$$\frac{p_1(x)}{p_0(x)} \geq \frac{P(H_0)}{1 - P(H_0)} \tag{4.9}$$

The ratio on the left of Equation 4.9 is the *likelihood ratio.*

4.3.2.2.3 Error rates and *a priori* information

With two overlapping PDFs, it is obvious that an observation in the overlap region could have occurred due to either case. It is useful to define the likelihood of an incorrect hypothesis being chosen. Consider the decision, D_1, to choose H_1. Assuming that the case was, in fact, H_0, the probability the decision is incorrect is

$$P(D_1|H_0) \tag{4.10}$$

and is the probability of 'false alarm' (i.e. a Type I error). Similarly,

$$P(D_0|H_1) \tag{4.11}$$

is the probability of a false negative (i.e. a Type II error).

The *a priori* probabilities of the cases may not be known, although the conditional PDFs may be. It can be shown that the error rates can be used to define the threshold for a decision rule. The Neyman-Pearson criterion can be used to define the threshold to maximise the probability of detection for a given probability of false alarm. Consider the region R_1 in which H_1 is chosen; then the probability of false alarm is the area of the PDF for H_0 in that region:

$$P(D_1|H_0) = \int_{R_1} p_0(x)dx \tag{4.12}$$

The solution for the lower bound of R_1 is then the threshold for the likelihood ratio test (Equation 4.9), and may be obtained without the *a priori* probabilities.

4.3.2.3 Filtering

The components that make up the power system include customer equipment and power system equipment. Customer equipment might include kettles or power showers switched on for a few minutes at a time, or lower-power equipment such as a refrigerator with a known duty cycle. Network equipment includes components such as transformers with tap-changers used for voltage control. Filtering allows frequency components of signals to be isolated or removed. Filters can be used to identify patterns in voltage measurements that describe particular equipment, or longer patterns of change that are characteristic of a particular area.

4.3.3 Micro timescale: milliseconds

This section examines the millisecond-timescale impacts of connecting loads to the network, and how these are related to location in the electrical network.

When a large, simple load is connected, it causes a drop in the voltage magnitude at the load. The voltage observed at the terminal will exhibit a change in the sine waveform amplitude, shown in Figure 4.16. A second customer connected to the network will experience a similar change in voltage magnitude when the first load is connected.

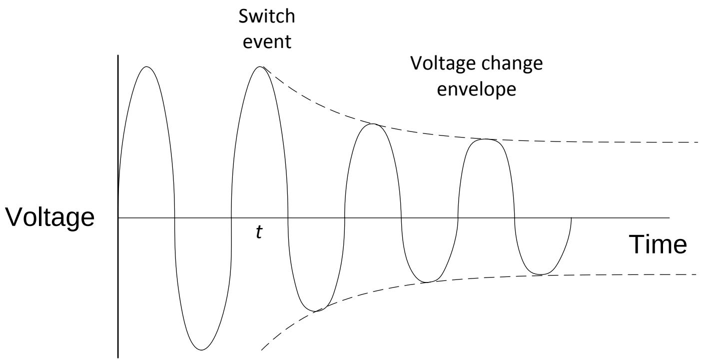

Figure 4.16: Voltage drop during switching event
The figure shows a switching event at time t at peak voltage. The voltage decreases along the envelope indicated by a dashed line.

An observer at the customer terminals might be able to use this to verify the existence of a connection between those customers: if there is a voltage magnitude change at exactly the same point on the waveform, a common connection between them can be inferred. However, if this is not observed, it might be reasonable to conclude the customers have different supplies.

In the three-phase power system, customers may be supplied by different phases. Single-phase loading causes unbalance, which can be observed on the other phases as a decrease or increase in voltage. The change happens at the point in time when the load is connected, but observations across different phases will show this change at a different point on the waveform due to the 120° phase shift.

Similarly, this could be used to identify connection. As shown in Chapter 3, two connections on the same phase will observe the magnitude change at the same time and same point on the waveform; connections on different phases will observe a small (positive or negative, depending upon phase and reactive power) magnitude change at the same time caused by unbalance.

In some cases, the envelope may describe a fast change, but the characteristics of the line and loads may be such that the voltage changes over the course of several cycles, as in Figure 4.16. Whilst attempts could be made to identify phase based on the point on the waveform at which change begins, this is challenging due to the potential variation in equipment, noise, event timing and so forth, especially considering the low magnitude of the voltage variation in question. Figure 4.17 shows a series of 5 cycles around a switching event of a 1kW load - but pinpointing the exact moment of connection certainly cannot be performed by inspection.

The results from chapter 3 show the voltage change from additional loading to be no more than a few percent, and well below the +10, -6 % voltage rise/drop limits. There are many sources of noise in

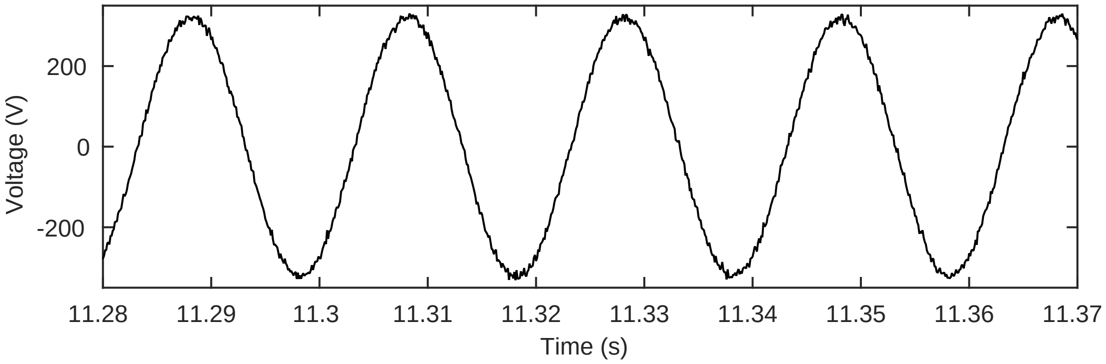

Figure 4.17: Network voltage during switching event
The figure shows 6 cycles around a voltage switching event. Noise is visible at the wave peaks. The moment of switching is not clearly visible.

the power system, including the characteristics of various loads as well as other customers connecting or disconnecting equipment, all contributing to the noise in the voltage sine wave observed at each point. Figure 4.18 shows a very typical peak of a sine wave recorded from the mains at 10kHz. A 1% (2.3V) change in RMS voltage corresponds to an amplitude change of approximately 3V - but the figure shows noise in excess of this value.

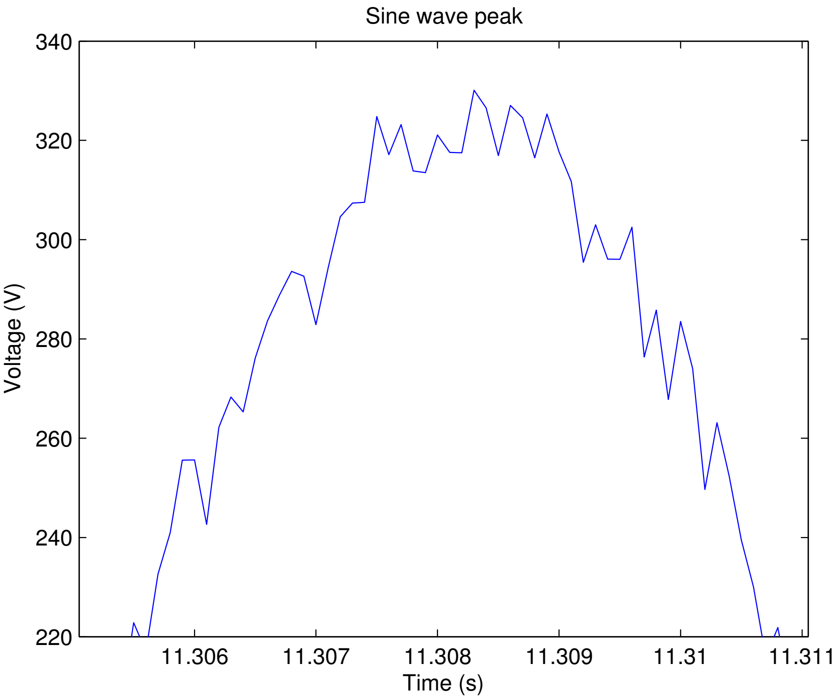

Figure 4.18: Captured real network voltage
The plot shows the absolute voltage measured on the real power network over the peak of a cycle. Noise and other frequencies are visible superimposed on the base sine wave.

However, the RMS voltage measurements do show a clear change, as shown in Figure 4.19.

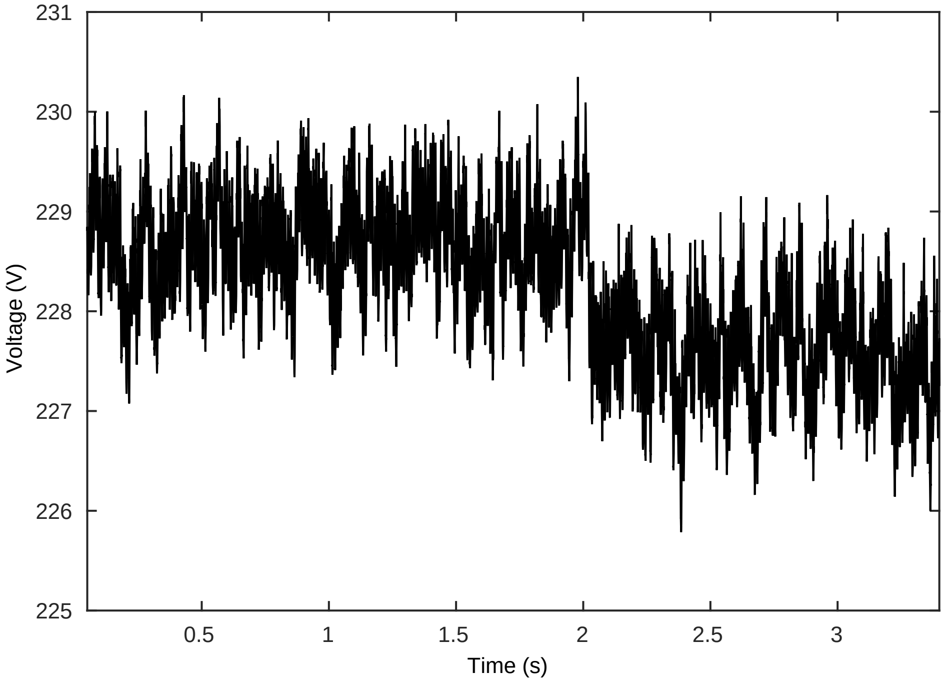

Figure 4.19: RMS voltage change during switching event

By comparing measured voltage with a modelled voltage change, the moment of switching could be found to within a small number of samples by performing the cross-correlation to determine the phase shift between the reference and the measurements.

Using the RMS voltage simplifies the voltage model for comparison, because it removes the significance of the point along the waveform at which switching occurred. The change in voltage results in a ramp in the RMS voltage. The modelled RMS ramp from a fast 3% change is shown in Figure 4.20.

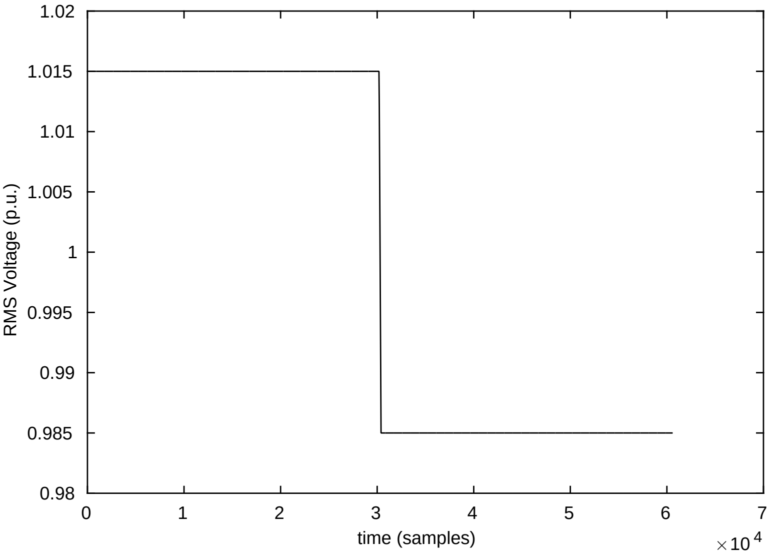

Figure 4.20: RMS Voltage ramp model for a 3V step

A test case was set up to assess this approach with a network layout according to Figure 4.21. A switching event was monitored at multiple points on the network. The results of the cross-correlation of the model and the measured data is shown in Figure 4.22. The data show the signal at the switching point is aligned at sample n=95080.

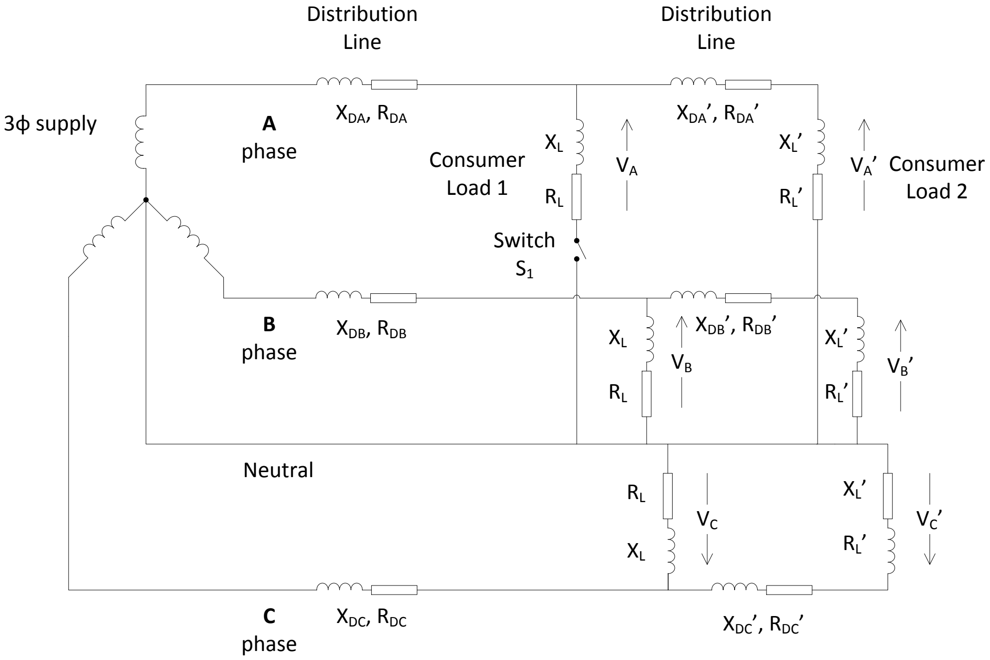

Figure 4.21: 3-phase network for switching analysis
The network setup in the figure was implemented in the Smart Grid Laboratory. A switching event was triggered, placing demand at Customer Load 1, causing the voltages in the network to change. These were recorded at V_A, V_A', V_B, V_B', V_C and V_C'.

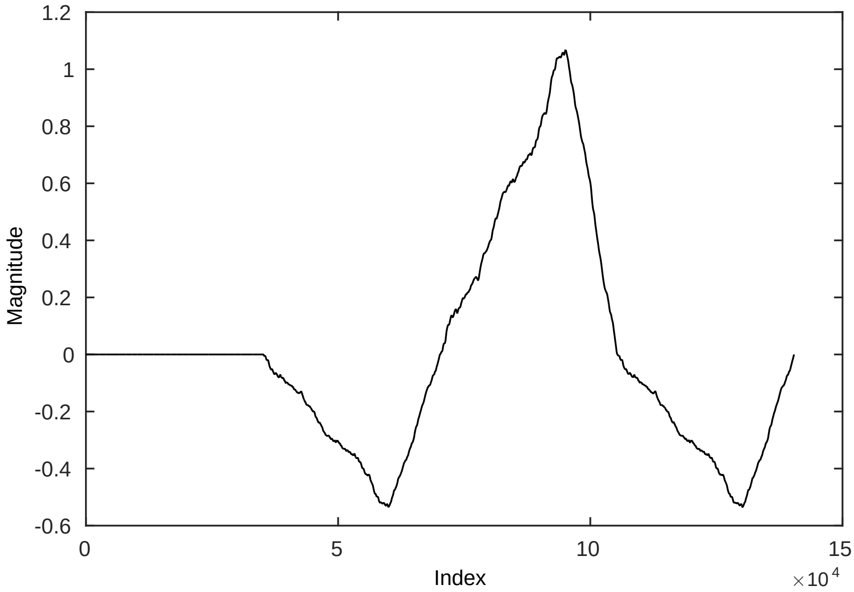

Figure 4.22: Cross-correlation of RMS change and recorded data
The figure shows cross-correlation of the 3V RMS ramp model and a measured signal containing switching.

Figure 4.23 shows the results of the cross-correlation performed for multiple points on the network. Two closely aligned signals are shown for V_A and V_A', with the maximum obtained 80 samples apart. Signal V_B, however, shows a little negative correlation, and shows some lag. Interpreting the results with any confidence was problematic given the characteristics of different loads, the level of noise in the system and the potential configurations of equipment in a real system.

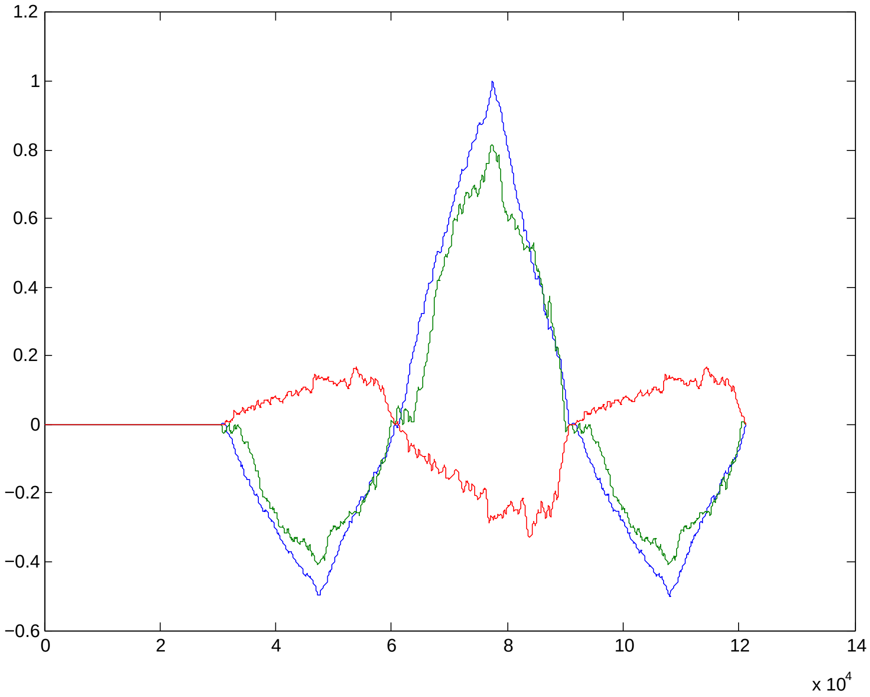

Figure 4.23: Cross-correlation results across network
The normalised cross-correlations are given for V_A, V_A' and V_B. The correlation for V_B is not strongly positive, and its peak exhibits lag. The other measured voltages are omitted for clarity.

When considering how to implement such a technique, further practicalities become apparent: it is necessary for the sampling to be synchronised to between the observers. Whilst GPS receivers can be used to provide timestamps accurate to microsecond-level, this functionality is not commonly implemented in commodity hardware. Consequently, voltage must be examined at the macro level. The next section will explore how the network behaves over the course of seconds, minutes or hours and use these properties for analysis.

4.3.4 Macro timescale events: seconds and minutes

On the macro scale, it is useful to examine the RMS voltage and its behaviour over time. The voltage changes over the course of a day, caused by a variety of power systems components including automatic voltage regulation systems as well as customer loads connecting and disconnecting, with fluctuations lasting from under a second to cyclic daily phenomena. With the power system model described in Figure 4.15, there will be local variations common to connected nodes — it is proposed that these variations be used as a fingerprinting method to determine the connection between customers.

In the electrical network as described in 4.1.4, customers close to each other will experience a similar voltage to each other; disconnected customers will have different profiles. There will be some difference in the former case — the loads and distribution lines will distort this signal.

The cross-correlation of two RMS voltage signals over a common 24-hour time period will be used to assess the potential connection between nodes on the electrical network. The magnitude and the peak of the cross-correlation between the voltage signals measured by customers will give an indication of the similarity. A high peak at the midpoint suggests a very similar signal aligned to the same timebase, and indicates a common customer connection.

Cross-correlation was used in [143] in a small residential network where the data on allocation of phase connection to a household was unavailable. The study sampled voltage at properties on the network, and performed the cross-correlation with a known three-phase reference location. The resulting series with the highest maximum was selected as a method to determine the phase of the house. However, in the context of this research, a three-phase reference location is not available. The study also made no assessment of the phase shift indicated by the location of the maximum. Further, the case where nodes may not be connected is not examined. .

Voltage variations will occur as a result of daily, predictable consumption patterns. Control equipment operates to adjust the voltage in an area to maintain performance within acceptable limits – but the timing of operation of these items will vary from area to area. Filters can be used to attempt to improve the cross-correlation results: the use of a low-pass filter increases the visibility of daily variation and slower power system control responses. In particular, the cross-correlation between areas under the influence of different control equipment will exhibit a phase lag as their controllers intervene at different times. This will be shown by a variation in the index of the peak away from the midpoint. Similarly, a highpass filter can be used to exclude consumption patterns and control equipment, such that the magnitude of the peak of the cross-correlation will be more reflective of local variation.

4.3.5 Data collection

The idealised models described in Chapter 3 contain neither noise nor power system events such as transformer tap changes. It is useful to use logged voltage measurements from a real power network to provide a variety of events comprising different customer equipment operating on a noisy connection.

OpenEnergyMonitor.org energy monitors were used to log the voltage at 10s (0.1Hz) resolution on the legacy network Durham, UK, with the setup as per Figure 4.24.

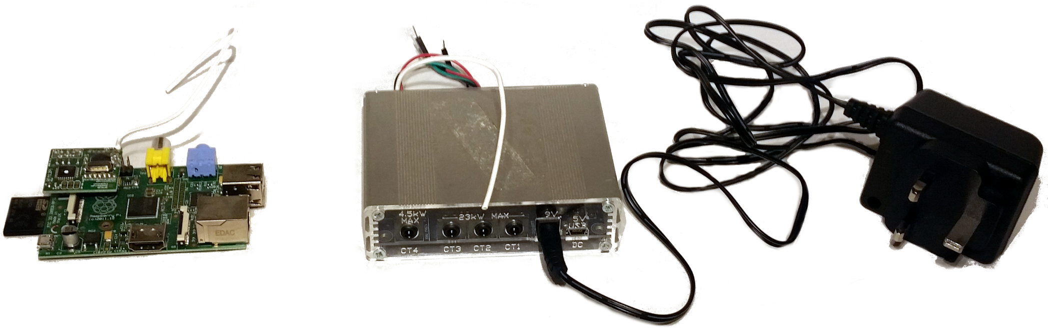

Figure 4.24: Datalogger setup
Datalogger setup. The image shows the data hub (Raspberry Pi, left) with radio component visible; OpenEnergyMonitor datalogger (centre) and plug-in voltage transformer (right).

Over the course of a three-week period, the voltage was logged in 4 properties, A-D. Three were adjacent to each other (A,B,C), with two on the same phase (A,B) and one on a different phase (C). This was verified by measuring the voltages on a common timebase. The verification of the 120° phase difference in (A,C) can be seen in Figure 4.25. The fourth property (D) was in same city but on a different part of the network, approximately 2.6km away.

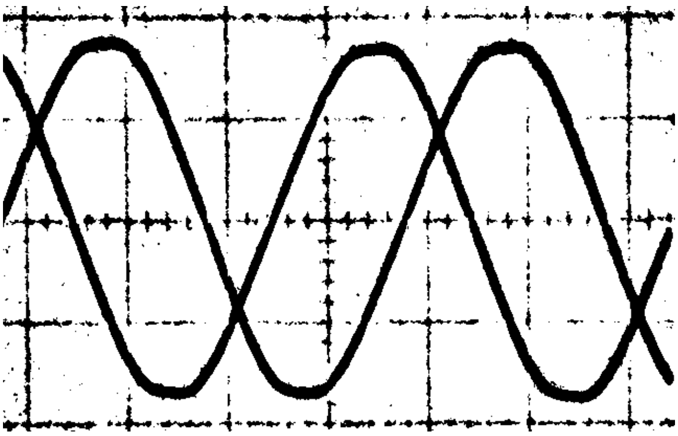

Figure 4.25: Identifying phase for datalogging
An oscilloscope was used to observe the voltage waveforms at properties A and C simultaneously to verify the nature of the electrical connection between them; this figure shows the 120° phase difference between them.

The connections are shown in Table 4.1, and these connections can be represented graphically as shown in Figure 4.26.

	A	B	C	D
A	-	phase	feeder	disconnected
B	phase	-	feeder	disconnected
C	feeder	feeder	-	disconnected
D	disconnected	disconnected	disconnected	-

Table 4.1: Electrical conections between distribution network nodes

Dataloggers were placed with customers A-D on different electrical connections: 'phase' denotes two properties connected to the same phase and feeder, 'feeder' denotes two properties on a different phase on the same feeder, and 'disconnected' is used where properties do not share a secondary substation.

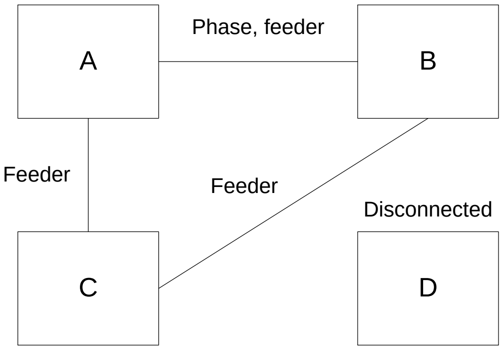

Figure 4.26: Node connection graph

The graph shows the customer nodes A-D from a real distribtution network; the edge properties denote the nature of the connections between them.

The data were collected and processed to create 21 sets of matched 24-hour time series. Figures 4.27(a) and 4.27(b) show RMS voltage profiles for weekdays and weekends that are typical of the recorded data. Before further processing, a small number (less than 0.1%) of spurious 0V readings were removed from the dataset. Missing datapoints were interpolated linearly to ensure timescale-consistent sample signals.

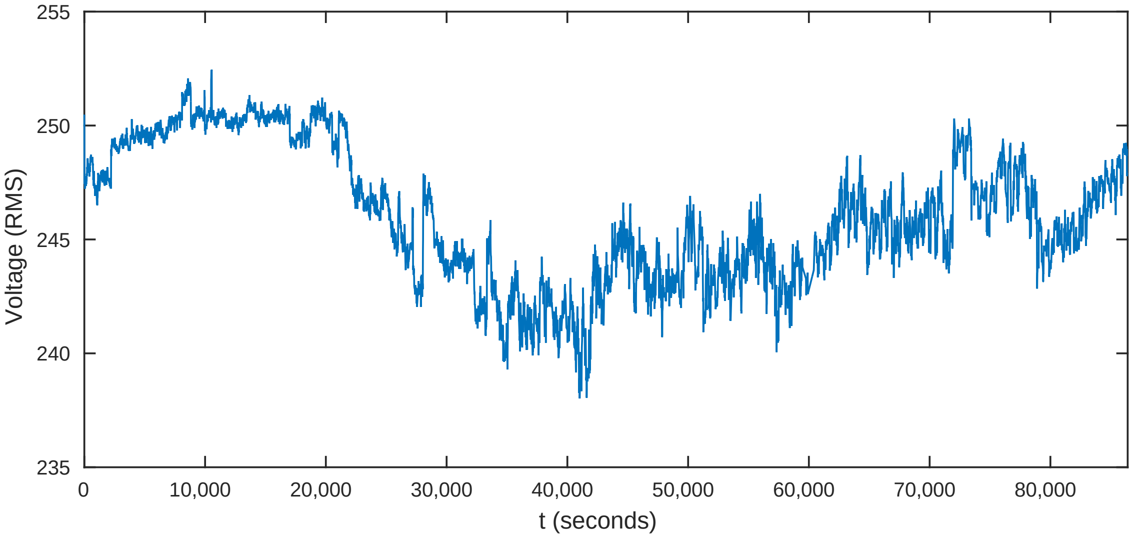

(a) Weekday

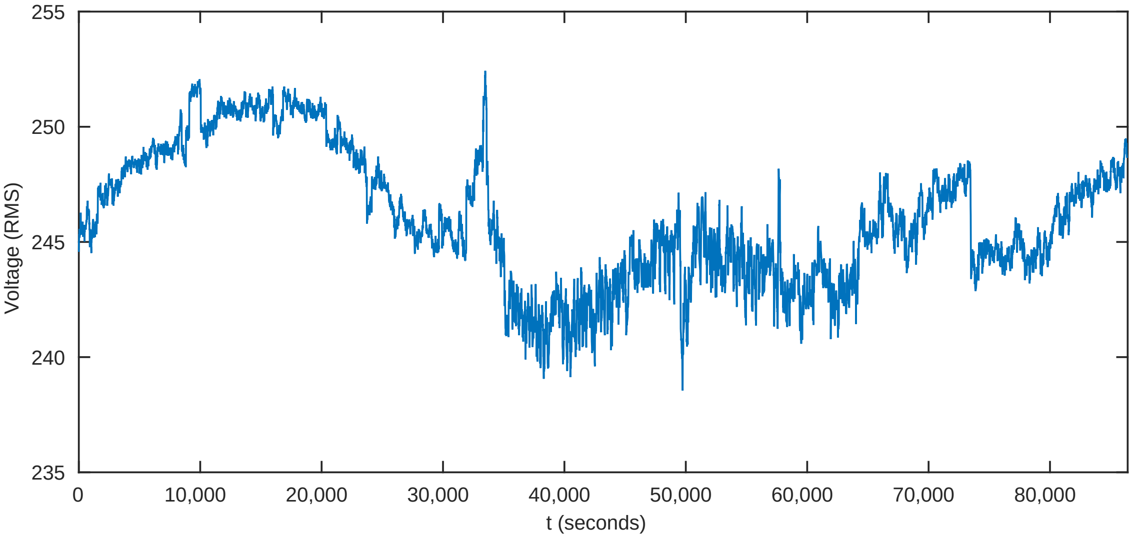

(b) Weekend

Figure 4.27: Weekday and weekend voltage profiles

Consumers typically have different weekday and weekend power usage. The RMS voltage at the customer varies over the course of 24 hours: the graphs show two typical 24-hour (86400 seconds) voltage profiles in a dwelling, showing weekday and weekend network behaviour. A power system event is visible in Figure 4.27(b), where a significant step change can be seen at approximately t=75,000, potentially as a result of a transformer tap change operation.

4.3.6 Application and results

4.3.6.1 Filtering

Low- and high-pass filters were used to improve the ability to discriminate between connection types by isolating different kinds of variation, as per 4.3.4. The filters were designed in MATLAB[144] with the properties in Table 4.2. These filters were then applied to the signals before cross-correlation.

Type	Filter	Pass frequency	Stop frequency
High-pass	Equiripple FIR	2×10^{-4}Hz	0.5×10^{-4}Hz
Low-pass	Equiripple FIR	0.5×10^{-4}Hz	2×10^{-4}Hz

Table 4.2: Filters for event isolation

The high-pass Finite Impulse Response (FIR) filter removes frequencies with period greater than 1.5 hours, so highlights (local) short-term variation; the low-pass FIR filter emphasises events longer than 1.5 hours, more likely to be characteristic of the power system.

4.3.6.2 Cross-correlation

To allow for analysis of different node pairs and remove the effect of individual voltage levels across the power system, the offset in RMS voltage (approximately 230V) was removed from each time series, and a normalised cross-corelation was used in order to compare peak maxima. By doing so, the effect of absolute local voltage (which will depend upon factors such as the service cable properties) and of node proximity to a stiff grid is reduced. The variations in voltage will be smaller, but are normalised.

The cross-correlation between the properties was performed in MATLAB[144] for each recorded 24-hour period for signal pairs (A,B), (A,C) and (C,D) to examine the phase-connected, feeder-connected, and disconnected cases. The three plots in Figure 4.28 show the (unfiltered) results for each pair.

4.3.6.2.1 Peak Maxima

The first, Figure 4.28(a), shows a clear peak in the middle with a magnitude very close to 1. From this, the voltage fluctuations in properties A and B are very similar. The peak in the middle shows that no time-shifting is necessary for the maximum similarity between the signals — since they are recorded simultaneously, this is to be expected.

The second, Figure 4.28(b), shows a similar spike, although with a lower peak magnitude. The third, Figure 4.28(c), shows the disconnected case. There is still a higher value around the centre, although the shoulders around the peak are wider than in the connected case, giving a weaker similarity than in Figure 4.28(b).

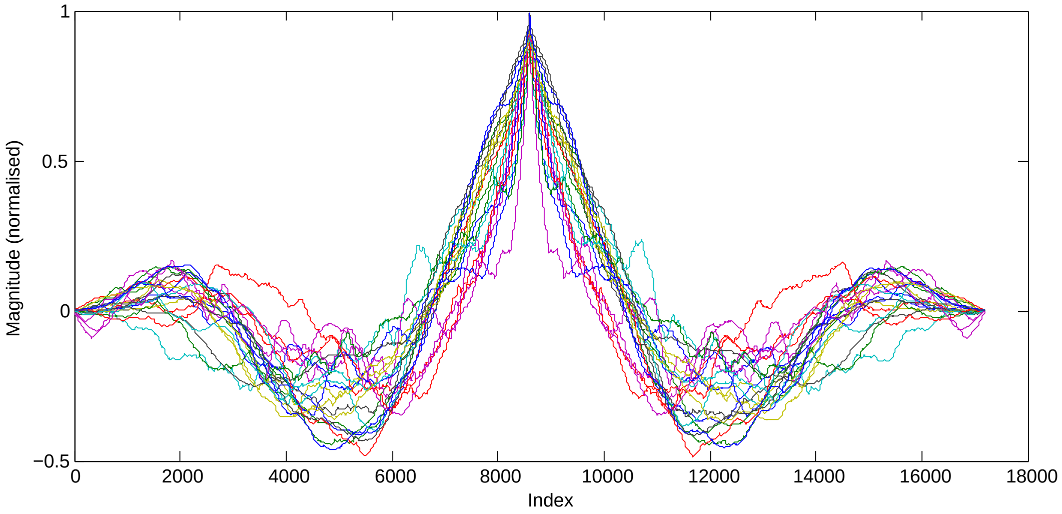

(a) (A,B)

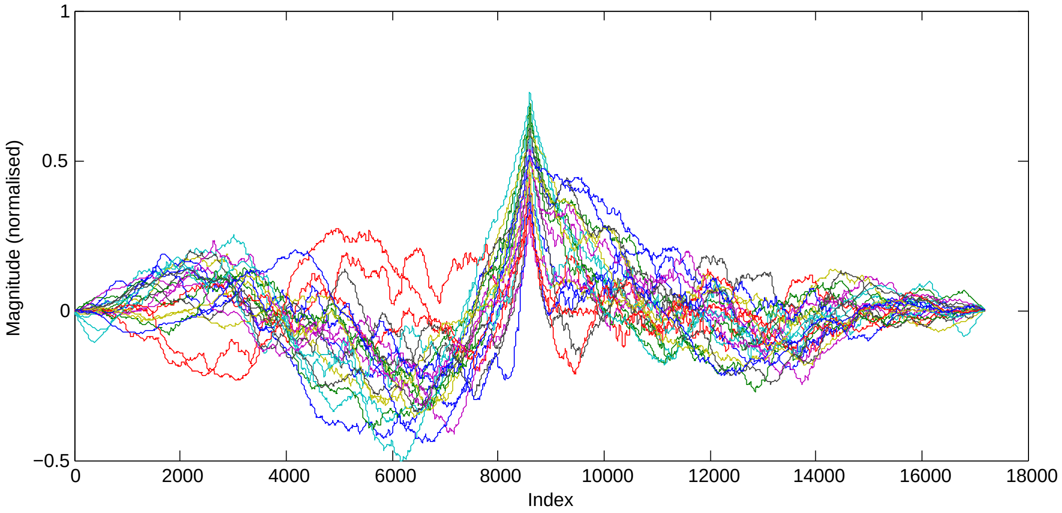

(b) (A,C)

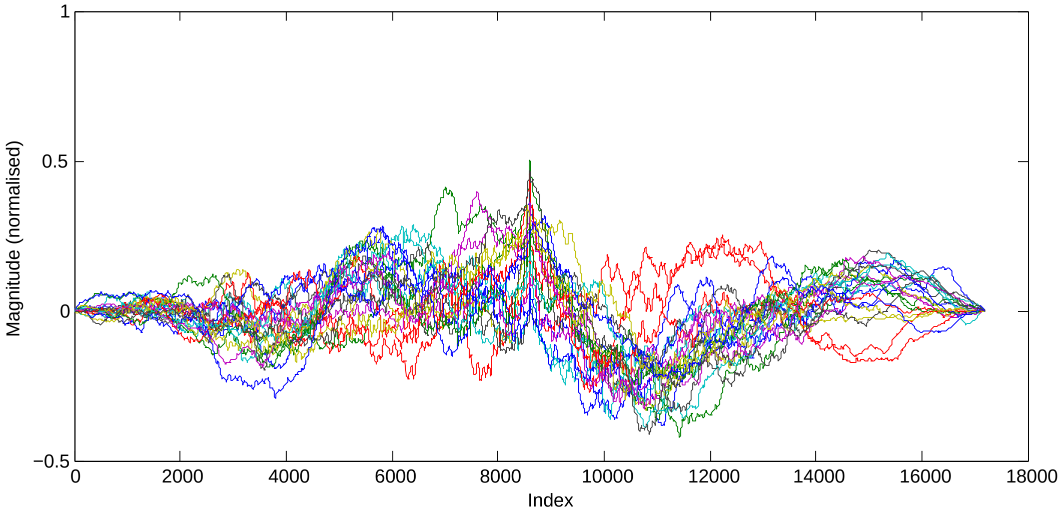

(c) (C,D)

Figure 4.28: Cross-correlation of voltage profiles
The figures show the normalised cross-correlation of 24-hour time series for different pairs of consumers.

To visualise the useful components, then, the peak magnitudes and indices may illustrated in a box plot. The central box is bounded by the upper and lower quartiles of the data, with the mean indicated inside. The range of the data is shown by a dashed line. The box gives an indication of the shape of the distribution including its skewness.

The box plot in Figure 4.29 contains the distribution of daily peak magnitudes for each time series and each signal pair. There is very little variation for (A,B) - the range is less than 0.1%. The mean peak magnitude is considerably lower for the other two cases. The connected case is clearly identifiable.

The results of the cross-correlation after applying the filters are shown in Figure 4.30. The distribution for the disconnected case shows higher magnitude as upstream power system characteristics become more significant.

The high-pass filter results, shown in Figure 4.31, reduce the peak of the feeder-connected (A,C) case. This shows how the similarities are dominated by phase-balanced system behaviour rather than local consumption — which is in line with the modelling in Chapter 3, where single-phase loading produces relatively small interaction with other phases, and potentially contributing to negative correlation. The change, then, shows sufficient balanced low-frequency behaviour to discriminate between the cases where the nodes are connected to a common feeder and the disconnected case.

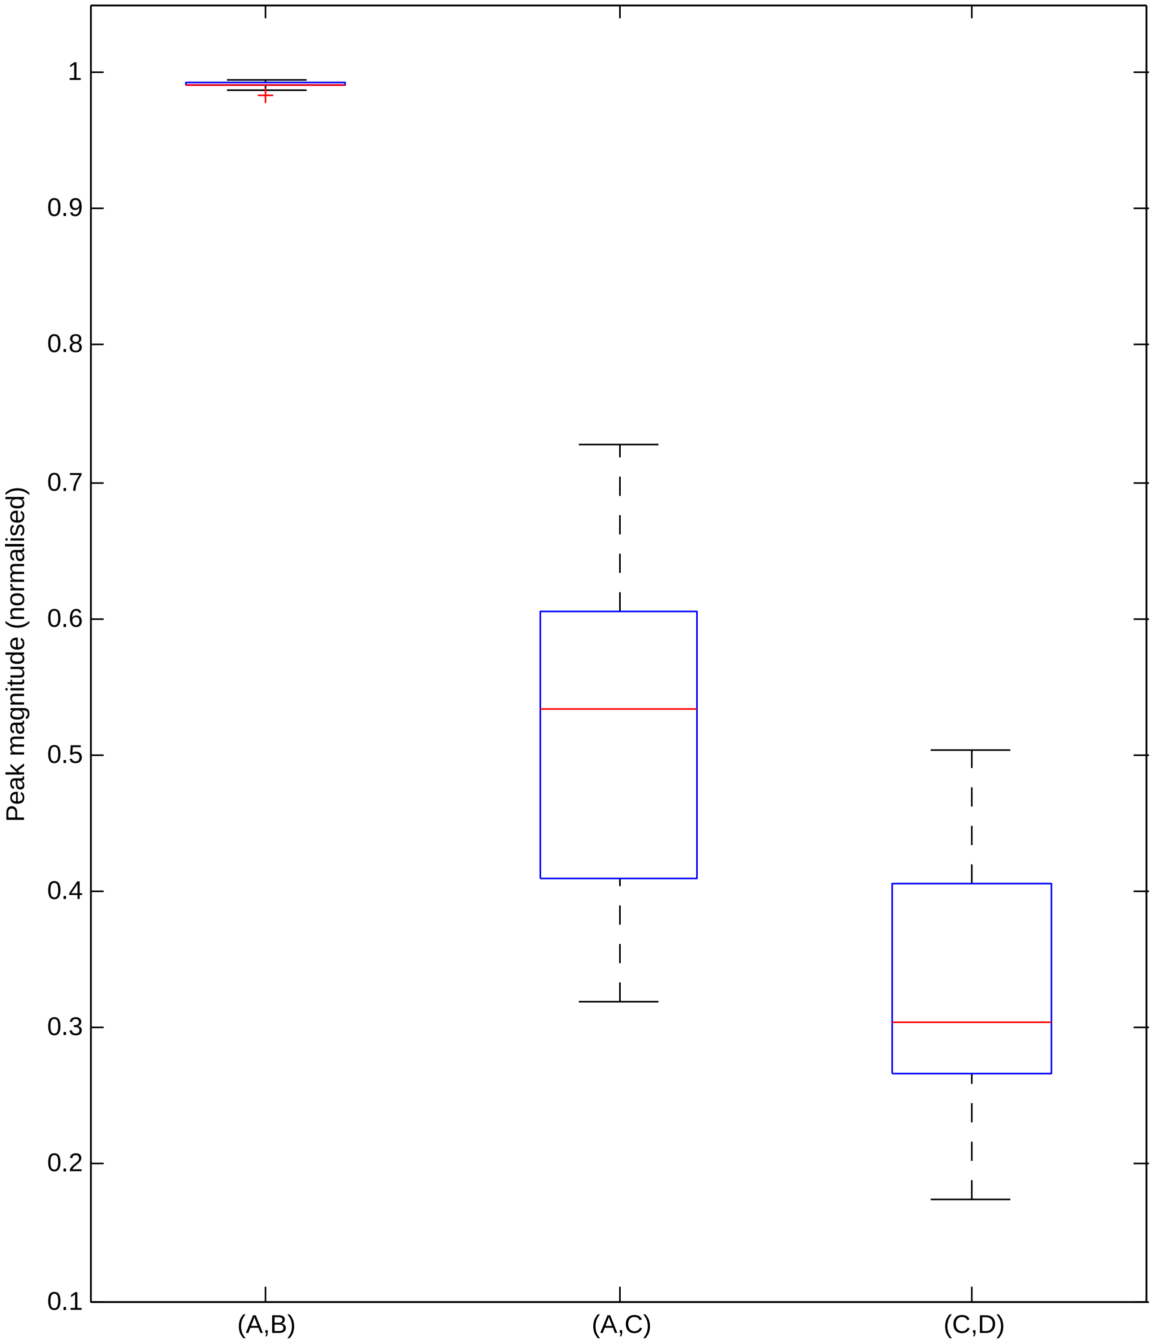

Figure 4.29: Cross-correlation peak magnitude
The box plot contains the distribution of peak magnitudes for each signal pair.

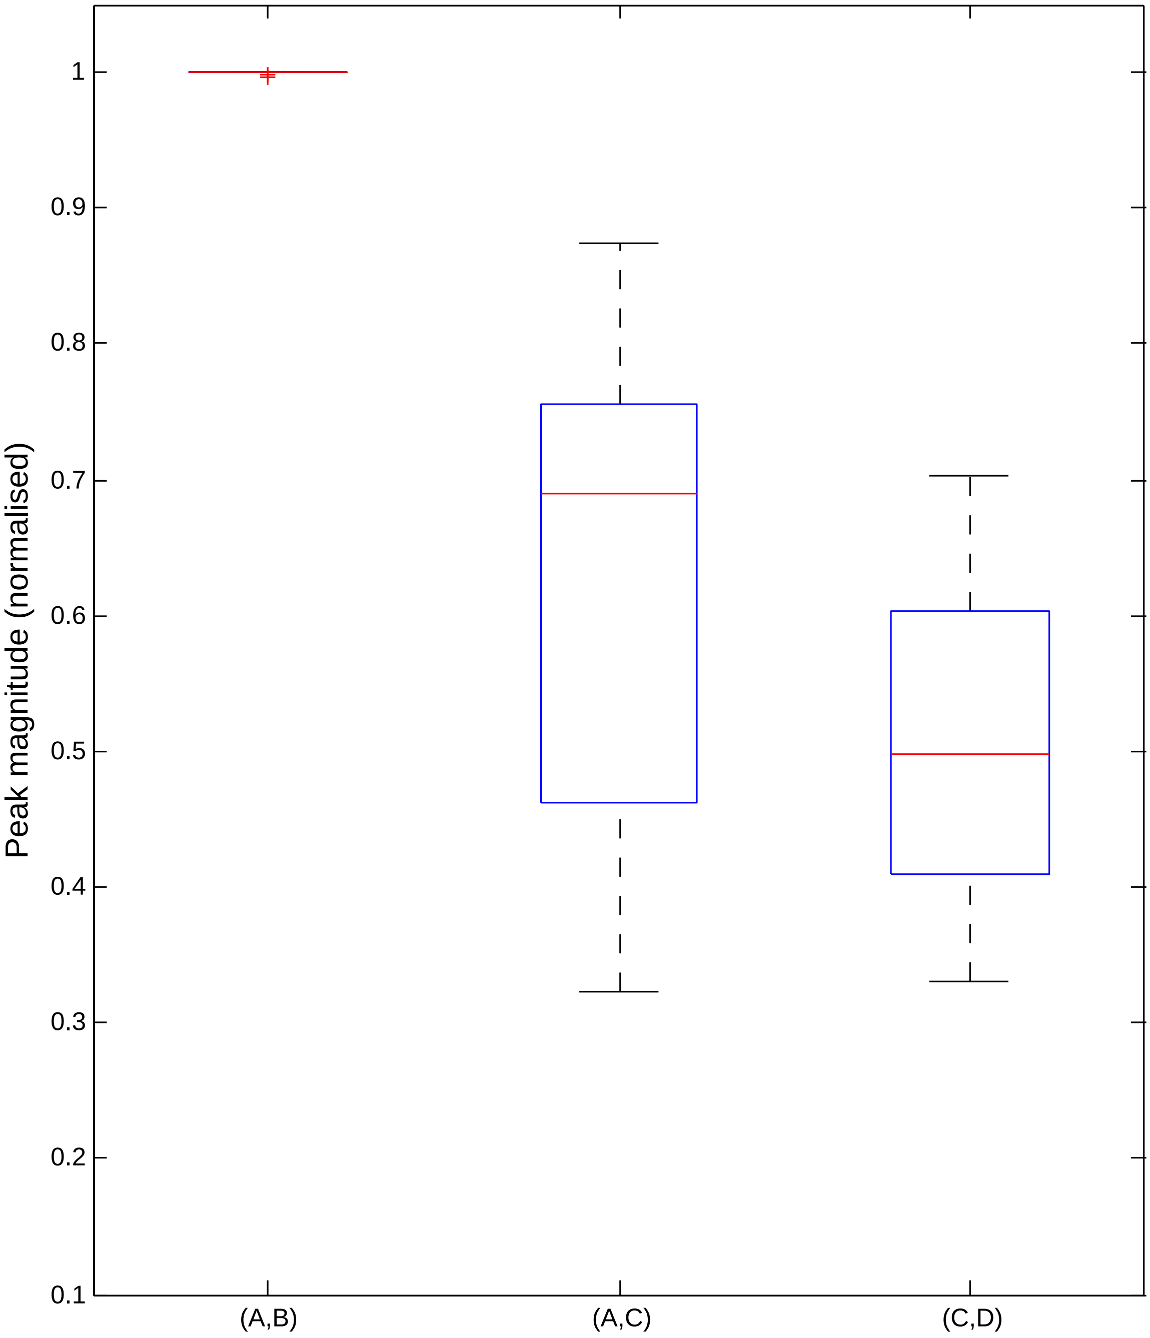

Figure 4.30: Cross-correlation peak magnitude, after low-pass filter
The box plot contains the distribution of peak magnitudes for each signal pair after a low-pass filter has removed high-frequency components

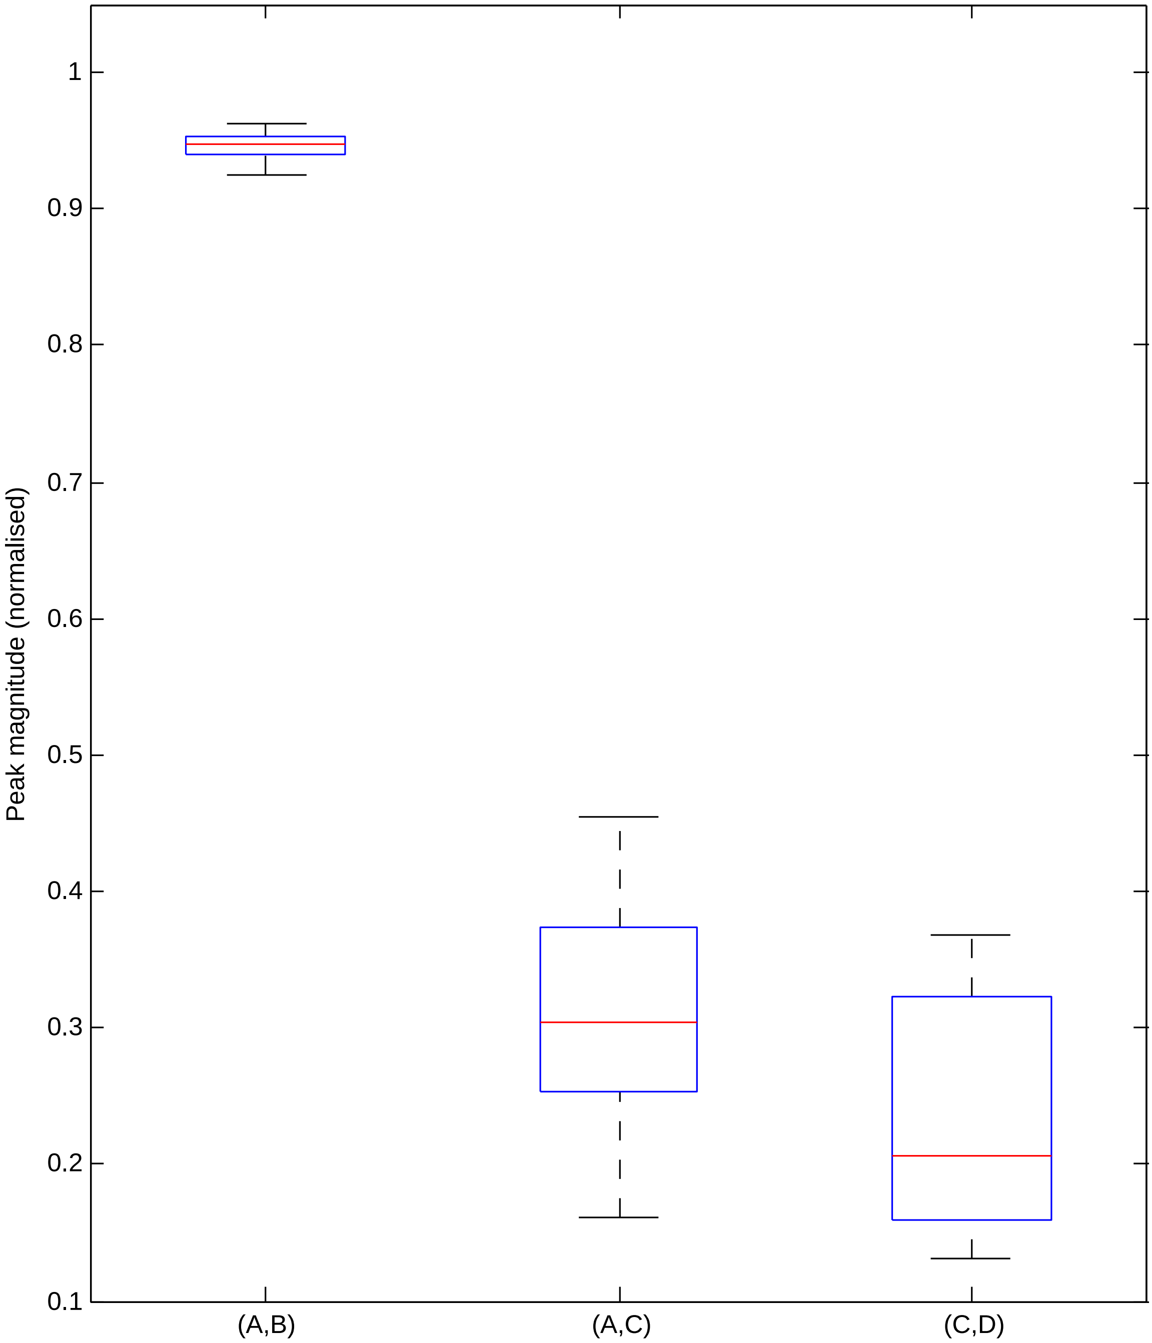

Figure 4.31: Cross-correlation peak magnitude, after high-pass filter
The box plot contains the distribution of peak magnitudes for each signal pair after a high-pass filter has removed low-frequency components

4.3.6.2.2 Peak Index

The peaks shown in Figure 4.28 do not all match perfectly with the midpoint. The further away the peak is from the centre (point 8641), the worse the similarity between them, suggesting a lower similarity between the corresponding signals. The distribution of the peak in each case is shown in Figure 4.32.

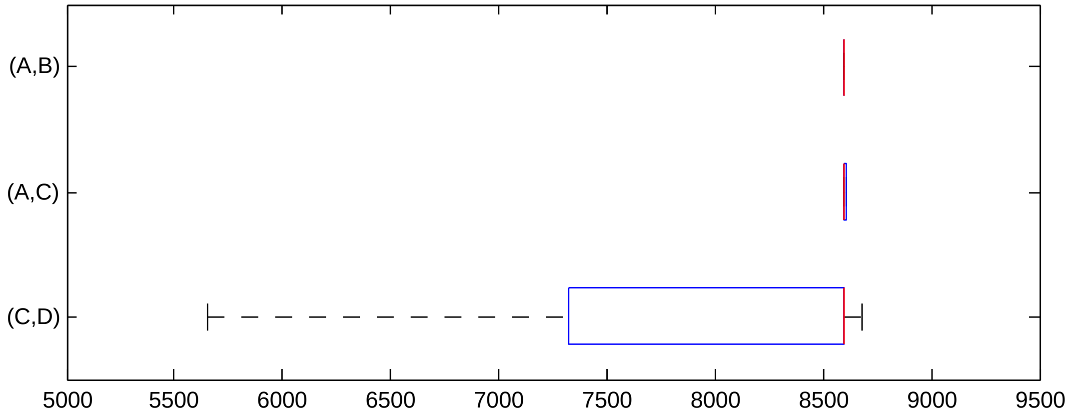

(a) No filtering

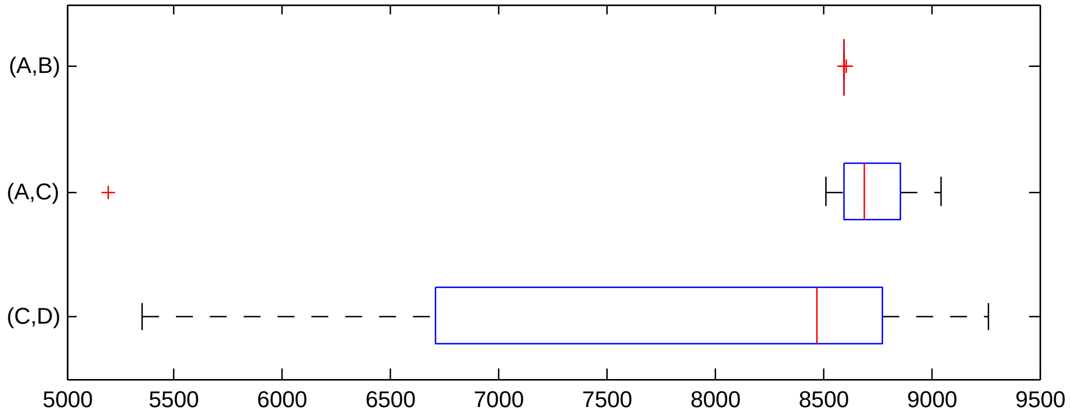

(b) Low-pass filter

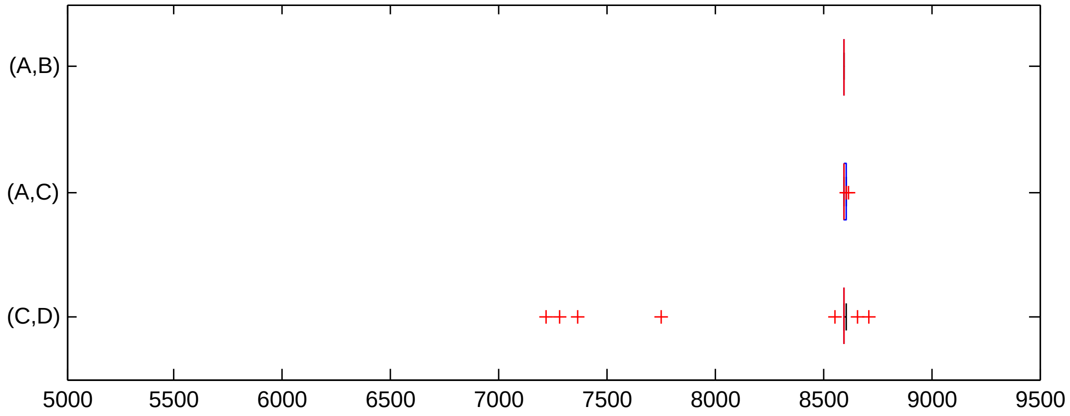

(c) High-pass filter

Figure 4.32: Index of cross-correlation peaks
The x-axis shows a scale of a section around the midpoint of the cross-correlations. The boxes show the distribution of the peaks for the signal pairs. Deviation from the midpoint indicates lead or lag.

The position of the peak in the phase-connected (A,B) cases are at 8641, i.e. the midpoint, but in (A,C) the values are distributed - although with the mean still the midpoint; and in (C,D) showing potentially several minutes of delay from the others.

In particular, the distribution for (C,D) in Figure 4.32(b) shows shows D leading C by intervals of up to 1800 samples, or 30 minutes. With the low-pass filter applied, this is a clear demonstration of the combination of differing network loading patterns and timing of voltage control equipment.

For disconnected equipment, lead and lag are equally likely, so the important factor is the likelihood of deviation from the midpoint. This particular distribution, then, shows the difference in the network but is only illustrative of the two network sections in question and does not reflect a general case. Rather, this may be better represented by a high probability of nonzero shift.

Based on the descriptions of the effects of electrical network structure and equipment on cross-correlation in Section 4.3.4 and the observations in Section 4.3.6.2, descriptions of the cases may be constructed; these are shown in Table 4.3.

Connection	Peak magnitude	Peak Index	High-Pass Filter effect
Feeder connected, phase-connected	very high	at midpoint	limited change
Feeder-connected, phase-disconnected	medium	around midpoint	reduced peak
Disconnected	low	(large deviations)	limited change

Table 4.3: Connection identification matrix

These qualitative descriptions will require quantitative determination in order for them to be used.

4.3.6.3 Hypothesis Testing

Thresholds must be determined for the descriptions in Table 4.3. In order to decide which category a result falls into, the hypothesis testing described in 4.3.2.2 can be applied. Further work will be required to provide confidence in a generalised set of thresholds. However, the peak magnitudes of the feeder-disconnected, phase-disconnected case and the fully disconnected case show some overlap in the data gathered, so the decision threshold is developed for this case.

4.3.6.3.1 Parameters

The value of the maximum of the cross-correlation function described above is a random variable. Similarly, the index of the maximum of the cross-correlation is a random variable.

In the case of deciding between “nodes connected to the same phase of a common feeder” and “nodes connected to a different phase of a common feeder”, the *a priori* probabilities can be assumed to be $\frac{1}{3}$ and $\frac{2}{3}$ respectively, since there are three possible configurations: connected to the same phase, a leading phase or a lagging phase. However, the relative probability of the disconnected case is not known in advance.

4.3.6.3.2 Peak Magnitude

Some skew is visible in the box plots in Figures 4.29-4.31, but the peak magnitudes will be modelled as normally distributed.

It is assumed at this point that the peak magnitude results are broadly indicative of the three cases - however, in the absence of further work regarding the sensitivity of the peak magnitude to network properties, the results for the filtered cases are more similar. The exact thresholds will vary; they determine and describe the "sphere of influence" of both customer action and network components. Consequently, the unfiltered data (Figure 4.29) is used, showing the distinction between the disconnected and the connected case.

To fit a normal distribution to the (unfiltered) data, then, the following properties for are calculated:

Connection	Mean μ	Standard deviation σ
(A,C) - feeder-connected	0.520	0.124
(C,D) - disconnected	0.327	0.096

Table 4.4: Distribution properties

Fitting the normal distribution with these parameters provides the PDFs shown in Figure 4.33.

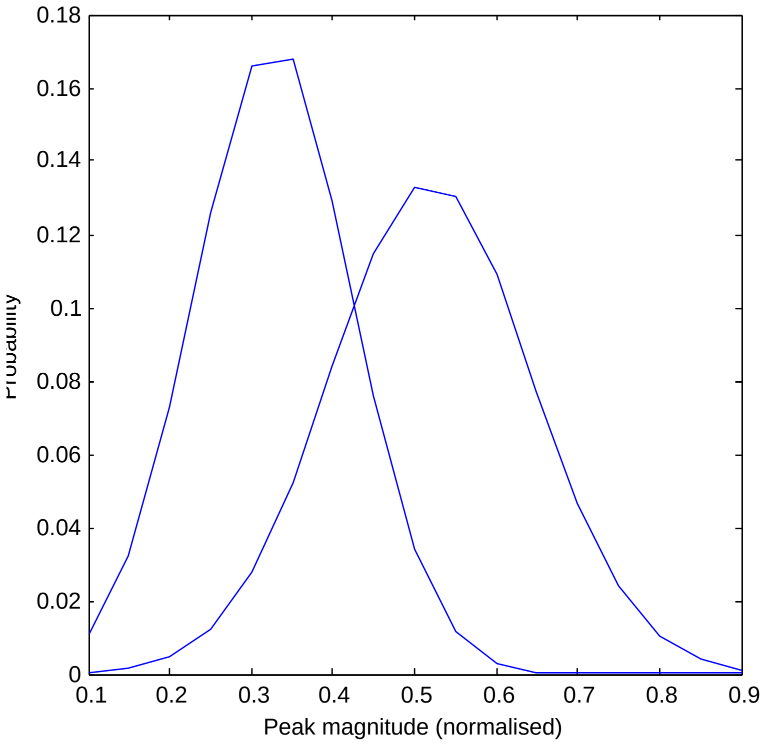

Figure 4.33: Connection PDFs
The figure shows the probability density functions of the disconnected (left) and feeder-connected (right) cases.

The hypotheses in this case are:

$$\begin{aligned} &H_0\text{: Nodes are disconnected} \\ &H_1\text{: Nodes are connected to common feeder} \end{aligned} \tag{4.13}$$

There is a cost associated with an incorrect decision: a false detection will result in controller action that will not achieve its objectives; a false negative will exclude the controller from taking action where it could provide support. The cost of excess action, however, is more significant. Rather than attempting to use cost functions, this preference is reflected in the choice of a low false alarm rate of 0.05.

Following the definition of a gaussian distribution, the *a posteriori* probabilities are:

$$p(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{\frac{-(x-\mu)^2}{2\sigma^2}} \tag{4.14}$$

The decision threshold for a given false alarm rate, then, is the boundary of the region described by Equation 4.12. Consequently,

$$P(D_1|H_0) = 0.05 = \int_{R_1} p_0(x)dx = \int_{\gamma}^{\infty} \frac{1}{\sqrt{2\pi}\sigma} e^{\frac{-(x-\mu)^2}{2\sigma^2}} dx \tag{4.15}$$

Substituting μ and σ from 4.4,

$$\gamma = 0.485 \tag{4.16}$$

For any peak magnitude measurement $x \leq \gamma$, the disconnected case will be assumed.

4.3.7 Proof of Concept

With the macro-timescale properties described qualitatively in 4.3.4, the quantitative parameters were derived using real data in Section 4.3.6. In order to verify the approach, the parameters can be recalculated for the connections of nodes A-D in Section 4.3.5 to establish repeatability.

Verification must be performed on data not used in the original derivation. An additional 24-hour dataset was collected from locations A-D with the method decribed in Section 4.3.5. These four signals are shown in Figure 4.34.

The methods described in Sections 4.3.4 - 4.3.6 were applied. The cross-correlation between signals was calculated using MATLAB for each pair in A-D (i.e. 12 pairs, excluding same-node pairs). The normalised peak magnitude and peak index were recorded. These are shown in the matrix in Table 4.5.

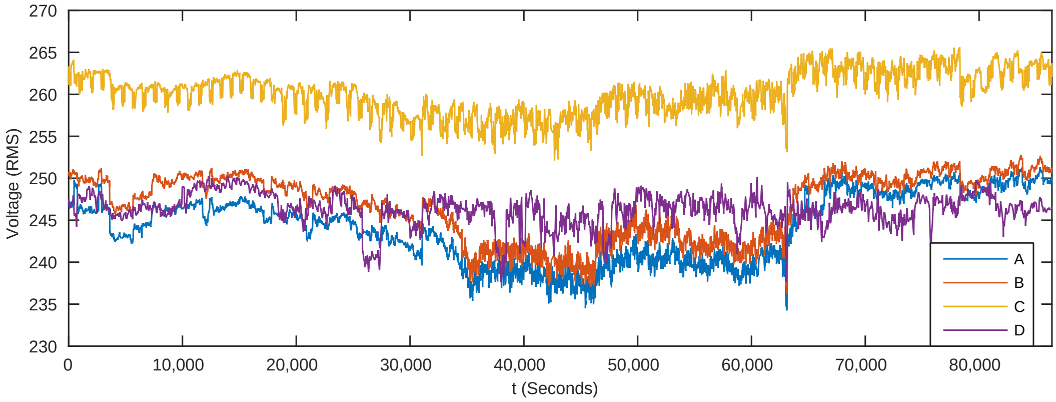

Figure 4.34: 24h recorded time series
The figure shows the verification dataset: RMS voltage measured at locations A-D recorded at 10s intervals to a common 24-hour timebase.

By comparing the results with the thresholds developed in Section 4.3.6, an adjacency matrix may be formed, shown in Table 4.6. In this case, the peak index occurs at the mid-point for the D cases; it still clearly falls into the disconnected category. The adjacency matrix is then used to create the full connection graph, as shown in Figure 4.35

	A	B	C	D
A	-	0.97, 8641	0.78, 8665	0.33, 8641
B	0.97, 8641	-	0.74, 8665	0.36, 8642
C	0.78, 8617	0.74, 8617	-	0.33, 8641
D	0.33, 8641	0.36, 8640	0.33, 8641	-

Table 4.5: Proof-of-concept: trial results
The dataset shows the normalised peak magnitude (0-1) and the peak index (1-17281) of the cross-correlation of recorded node voltages

	A	B	C	D
A	-	phase	feeder	disconnected
B	phase	-	feeder	disconnected
C	feeder	feeder	-	disconnected
D	disconnected	disconnected	disconnected	-

Table 4.6: Proof-of-concept: adjacency matrix result
The results are translated into connection properties through the developed thresholds and given here for each non-self pair. Note that this figure indicates the same connections indicated in Figure 4.26.

When compared with the known connection properties of nodes A-D as shown in Table 4.1 and Figure 4.26, these results confirm successful use of the electrical identification techniques developed.

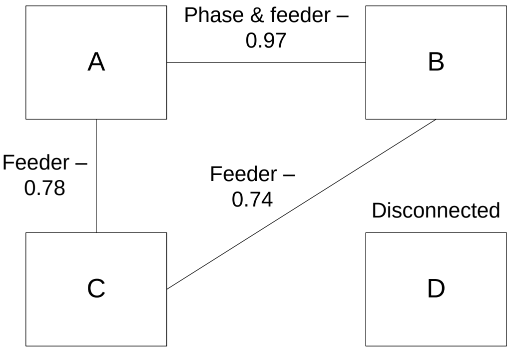

Figure 4.35: Final graph of electrical connections
The figure shows the connections indicated by the results of the proof-of-concept verification. The properties of the connections are indicated by the text by the connecting edges.

These properties can be used to construct a path for voltage control (see Figure 4.6) and a star (or tree) for current control (see Figure 4.5). A node will calculate the cross-correlation between itself and potential neighbours. Beginning from node A, then, a path and a star can be constructed from these properties, as shown in Figure 4.36: Figure 4.36(a) shows the path and Figure 4.36(b) shows the star. Importantly, the edge between B and C is not visible in Figure 4.36(b): no assessment is made of the connection between the two nodes by A, since as the number N of neighbours grows, the number of cross-correlations to be calculated increases with N^2 and rapidly become a significant burden for calculation. However, the node's own internal representation still carries information about which nodes are connected to it.

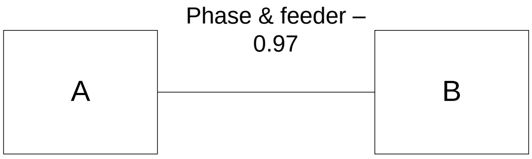

(a) Voltage control: path

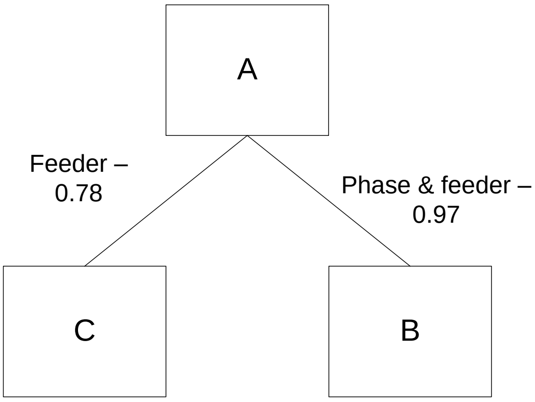

(b) Current control: star

Figure 4.36: Electrically-derived graphs for Node A

This technique, then, has been shown to correctly identify the location properties (as shown in Figure

4.35) and to be used in graph structures for internal representation of connection properties (as shown in Figure 4.36). The geographical and electrical approaches must be combined: Section 4.4 provides a summary of the complete approach.

4.4 Final location consolidation

A final graph can constructed by combining all these data: Firstly, the geographical approach selects likely candidates. The measured signals are then analysed: a confirmation of the phase-connected case for voltage control can be obtained with a lower threshold for peak magnitude of 0.9. The node is thus incorporated into the path graph with the structure from the geographical tree creation process (see Section 4.2.1. The feeder-connected attribute can be used to create a star graph suitable for current control: if the phase connection is not shown, further analysis is performed: as shown above, the criteria for inclusion as sharing a feeder are both the index of the peak occurring near the midpoint of the cross-correlation, and a peak magnitude above $\gamma = 0.485$; otherwise, it is added to a list of disconnected nodes. The node is thus incorporated into the graph, and the edge weighted with the 3-tuple (distance, connection type, peak magnitude) to distinguish the overlay of voltage/current graphs and the connection quality.

For the purposes of this work, the initial approach the voltage control task will only necessitate the phase-connected (i.e. path) graph.

4.5 Conclusion

The importance of location was discussed in Section 4.1. New concepts of location in the electrical network are discussed and location analogues and sources of location data are explored. A geographical approach was outlined to derive location information, and a new method was presented using distances and tree walks for using this information in creating a network graph suitable for use by a control system. The properties of the electrical network were explored as a source for location data; signals were captured from the live power system and then used to establish the real-world validity of a novel, non-intrusive method of confirming the existence and nature of connection between points on a network. The model can consequently be used in both voltage and current control, and is derived from locally available information. The combined new geographical approach, electrical analysis and graphing method will be taken forward as key functionality in the implementation of a location-learning multi-agent control system to be described in Chapter 5.

Chapter 5

Multi-Agent Systems for LVDN Control

This chapter contains the implementation of location-based control for Low-Voltage Distribution Networks (LVDNs) using a novel real-time test rig.

The impacts of Distributed Energy Resources (DERs) on LVDNs are described in the literature (see Section 2.2.1) and Chapter 3 outlined models incorporating them to demonstrate these effects. This chapter details a Multi-Agent System (MAS) approach as a means to tackling those technical challenges.

The design and testing of this system required new facilities. The development of multi-agent testbeds is described in Section 5.2 and voltage control agents were applied that use the techniques developed in previous chapters to demonstrate the capabilities of those new facilities.

A review of MAS theory is presented at the start of this chapter, where decentralisation is examined as both challenge and tool in electrical network control. It introduces concepts of multi-agent systems and cognitive models of agents. Existing paradigms of (non-) hierarchical societies of agents are described through the lens of distributed control systems. Existing and new approaches are described for group problem solving through communication. Formal descriptions of agent attitudes are used to describe and critique *intention* and the new technique of Intention Transfer is presented as a decentralised opt-in collaboration tool.

In Section 5.2.2, the novel control agents developed to address the research problem are outlined along with the components of the MAS software platform used. Three original MAS test bed hardware implementations are described: a proof-of-concept system, a cluster supercomputer implementation, and a unique electrical network/control software/hardware-in-loop platform.

The results of the system operating on the network models are given in Section 5.3. The performance of the location-based control approach is examined over the course of the development of the project.

Finally, the novel approach against the problem statement is evaluated alongside the features of the new test facilities developed to demonstrate it.

5.1 Multi-Agent Systems

5.1.1 Agent Concept

Agents are widely described as *autonomous, proactive, reactive* (and, often, *social*) *entities*. They might be software programs, humans or hardware articles with these properties. Agents reason (at varying levels of complexity) and make decisions for themselves based on their own knowledge. They will actively seek to achieve or maintain a particular state or to solve a problem. They might change operational strategies to do so if necessary. Agents may adjust and adapt to external changes, creating a closed-loop feedback approach to affecting their environment. It is often expedient for agents to communicate directly in order to share information, negotiate or achieve a better solution through collaboration. In some cases, agents are capable of learning, or exhibiting intelligence.

Multiple autonomous, reactive, proactive and social software entities acting on the same domain can be referred to as a Multi-Agent System (MAS).

In software, MAS is a programming approach that uses independent software entities to perform actions and communicate to achieve some functionality. Multi-agent systems are distinct from monolithic (single-agent) programming, in which one piece of software has complete control over all of the system inputs and outputs, and the client-server model where resources are centralised and accessed in a hierarchical fashion.

Multi-Agent Systems have been used to create robust, resilient and effective systems in many industries, such as manufacturing and telecommunications. In some applications, agents are used for modelling rather than acting upon a physical system, for example by representing human behaviours in order to establish emergent system-wide phenomena. There are platforms and toolkits for agent software development in different applications, standards for information exchange, and proposals and specifications specific to power engineering.

In the power systems domain, agents could be control routines, monitoring devices or a proxy for a consumer; in this way, the multi-agent system becomes analogous to the electricity system itself, with each element represented in some fashion by a connected agent. Communication of information about the state of the distribution network, independent decisionmaking, and negotiation allows deduction of rational, relevant control actions.

5.1.2 Multi-Agent Systems in Power Systems

Multi-Agent Systems have been applied in power systems from generation, transmission, distribution, supply and domestic control. The review by McArthur et al. [115] discusses distribution network applications in protection, network monitoring, and simulation, and highlight that distributed control is a particular focus for MAS research[115].

In particular, the following topics have all been addressed using agent functionality:

- thermal overloading of transformers
- voltage rise and drop
- voltage regulation
- voltage unbalance
- cable thermal limits
- frequency control
- islanding
- enabling DSM
- market participation
- maintenance planning
- topology reconfiguration

The properties of agents are evidently appropriate for the power systems domain. Distributed control approaches are of interest: purely autonomous agents can operate to create an emergent control system, such as in [145]. However, currently underexploited is the social aspect of agents in control, at least in part because of the timing requirements for different control goals[1] as well as the difficulties of reliable communication. As an approach to the research problem, then, the social and communicative aspects of multi-agent systems may be exploited to improve LV network behaviour.

In constructing a multi-agent system, it is useful to understand how an agent can make rational decisions: to this end, there are several cognitive models of an agent that can be used to create rules and programmes that result in an entity exhibiting behaviour of the type described above. The following section describes appropriate modelling, which is used later in the chapter to develop reasoning and communication approaches.

5.1.3 Distributed, Decentralised, Devolved Control

Control could be considered to be distributed either by a *distributed algorithm* (i.e. performance of a shared algorithm by multiple participants), or by changing hierarchies in decisionmaking and action.

Hierarchy manifests in the degree to which an agent will obey another, and the autonomy or initiative an agent possesses to achieve an objective. A hybrid of the two is devolution, where some centrally determined requirement or goal in a hierarchical organisation of agents is achieved by passing the responsibility for determining a solution and its implementation to lower-level, locally-relevant agents. Often, however, any combination of these ideas is simply called “decentralisation”, being any system that does not use some form of service-oriented or client-server model.

Different systems of organising agents will result in varying hierarchy and autonomy.

5.1.4 Communication

The MAS paradigm involves independent (software) entities which operate autonomously in order to achieve individual goals. However, individual- or system-level goals may be hard or impossible to achieve where only individual actions are available. Agents may work together through communication in order to achieve these goals. The literature is approached from the perspective of the electrical power systems domain, with the aim of identifying techniques that may be useful and transferable to distribution network control.

5.1.5 Cognitive Models of Agents

It is useful to consider the internal structure of a software agent to assist understanding of its operation and to facilitate implementation. From the outside, agents are considered to be “black boxes”. Agents are usually unaware of the internal structure of other agents and how this relates to their behaviour. A “cognitive model” is a description of an agent’s subsystems that a programmer can use to design how it responds to stimulus in order to simplify its construction. There are several ways to describe the “mind” of an agent — existing models include BDI, CLARION, DUAL and others. Much like a programming language, the choice of analysis model is partly an implementation choice with some advantages to each. For reasons of familiarity, the Beliefs, Desires, Intentions (BDI) model is used here.

The BDI cognitive model[146] provides a way to compartmentalise agent reasoning systems. A complex agent is broken down into a knowledge base (beliefs), a set of goals the agent may wish to achieve (desires), and, having decided to act, a commitment to a set of actions the agent may take to actually achieve them (intentions).

Separation in this way allows an agent to perform reasoning about each of these discrete components and make decisions about them, such as whether it is feasible to take a certain action, or whether it believes a goal has been achieved. This facilitates the internal decision-making process.

In the electrical domain, then, an EV controller agent could be modelled in this way: with a set of beliefs (e.g. the current battery status), desires (e.g. ‘a full EV battery’ or‘’minimise the cost of charging) and intentions (e.g. ‘charge the EV battery’). This informs the construction of the various functions (sensing, evaluation, and action routines) that make up the overall controller.

Functions such as communication interpretation and goal management can be placed alongside the BDI model to allow the agent to select its behaviour and perform social functions. An agent may possess conflicting desires, but must not rationally attempt to enact incompatible intentions simultaneously. Intentions are implemented at the lowest level by choosing an appropriate *plan* (sometimes known as a *recipe*) from a library of functions which could perform a physical or communicative act. Beliefs only reflect the agent's (incomplete) knowledge about the world it inhabits, and may be incorrect or out of date. The terms *goal* and *desire* are used interchangeably here.

Cohen and Levesque [147] describe a domain of possible worlds, each connected by sequence of non-overlapping events which occur along some temporal index. An agent (with incomplete information) may inhabit a world compatible with its perceptions. The agent can make a transition to a compatible world (including the current world) at some other point on the temporal index, linked by its beliefs and goals; the latter are expressed as *choices*. Between two consecutive points on the temporal index, an event may occur which changes the state of the world that an agent perceives that it inhabits. This event may also alter the accessibility of other worlds.

Cohen and Levesque use a modal logic representation to describe agents; agent beliefs; goals, intentions and other mental attitudes; information and communication; truth propositions; and how choices and actions affect the real world. Modal logic notation is used throughout this thesis.

In a communicative agent society, multiple agents could operate upon their environments simultaneously to achieve their goals. Some similar goals might be held by multiple agents — in this context, this could be an overall control scheme. The following sections discuss forms of interaction between agents to create group actions.

5.1.6 Agent Collaboration

Agent interaction can be used to achieve a goal, whether locally decided or globally designed. Agents may cooperate and work together to achieve a common goal. The notion of "joint action" (and correspondingly, joint *intention*, see Section 5.1.10.2) involves two or more agents operating simultaneously towards a common goal, but more importantly incorporates an element of belief that other agents are also acting to achieve the same ends[148]. This requires explicit communication of these beliefs. Other concepts in cooperative action include the formation of interdependent *plans*, where agents share a common solution and members execute the relevant components[149].

However, the application of a common goal and its representation in a devolved agent-based system is problematic. An agent may be willing to cooperate, but be unable to do so due to some operational constraint. Or, it may cooperate, but know that doing so is to the detriment of the requesting agent. This can be a particular issue where agents only perceive their immediate environment but can form beliefs about the rest of the world through communication. Other approaches, such as competition or market-driven negotiation may be redundant or irrelevant in the context.

The following section explores systems of organisation as a means to achieve joint action in multi-agent societies.

5.1.7 Organisational Paradigms

Agent societal structure has a significant impact on the effectiveness of groups of agents at solving particular tasks[150]. These structures have strengths and weaknesses and there are differences in the way these groups are formed. Some structures could promote altruistic or competitive cooperation; others may restrict agent interactions dependent upon location or function. Table 5.1.7 gives a summary of the key types, strengths and weaknesses the various structures. Whilst *coalition* has a specific meaning described below, it is more generally a group that agrees to attempt a common goal; other structures are effectively special cases of coalitions.

5.1.7.1 Society

The finite set of agents N from Definition 1 is a *society*; however, society has properties not expressed in this model of a game. As well as a set of agents, society incorporates ideas of conventions for interactions. In the society, agents must act according to these social laws; a society may enforce these laws with penalties such as fines or exclusion. These can be used to create a particular behaviour, for example in a microgrid where a customer may not draw more than 5kW on pain of having supply cut off. Social laws and norms (for example, electricity unit trading in pounds sterling) form part of the environment the agents exist in. However, these laws require representation or internalisation within agents if they are to exploit cooperation or coordination[151]. *Hierarchy*, *distribution* and *autonomy* are properties of an agent society that are partly resultant from social laws and norms, but also derive from the choice of organisational paradigm used to achieve a task.

5.1.7.2 Coalition

Any subset S of a society N can be a *coalition.* A coalition is a set of agents formed to solve a specific goal. They are formed dynamically and are only relevant whilst the goal is maintained. Coalitions may be formed whether agents are cooperative or self-interested. Self-interested rational agents will only join a coalition if the resulting payout is individually fair, that is, no worse than if the agent acted alone. Whilst a coalition is a flat structure, there may be hierarchies within it; coalitions may be nested, or one agent may have some centralising or coordinating role. A leading agent may be decided by an election between the agents before action, or determined by the nature of the goal to be fulfilled. Horling and Lesser give the example of a deadline negotiated by a coalition leader, rather than by its component agents[150]. Coalitions have similarities with *federations* described below, but federations have a distinct and persistent facilitator role. Coalitions formed from a society must have common communication and

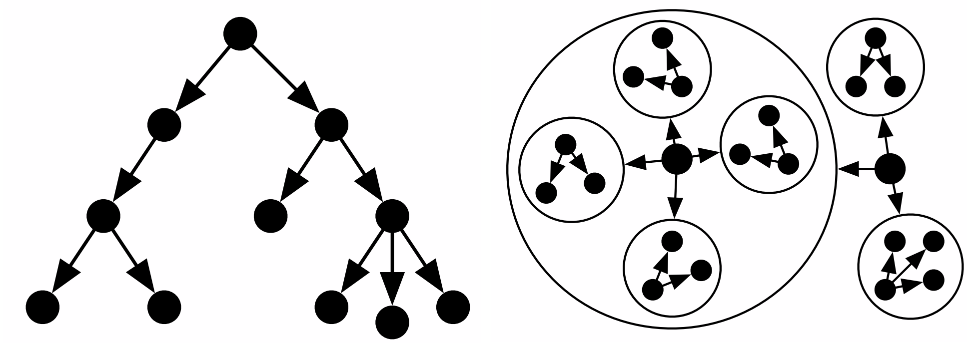

(a) Hierarchy (b) Holarchy

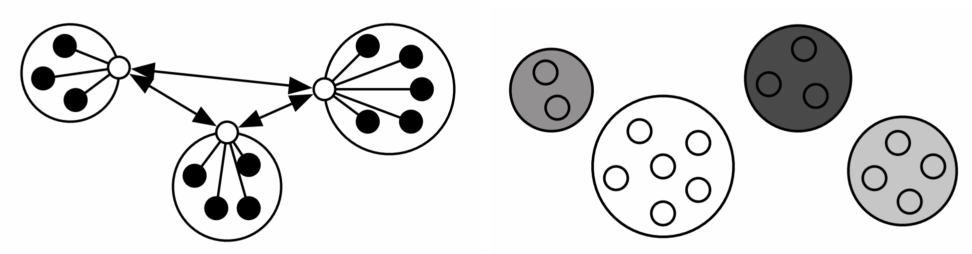

(c) Federation (d) Congregation

Figure 5.1: Organisational Structures
Four key organisational Structures (reproduced from [150]) incorporating aspects of hierarchy, delegation, independence and collaboration.

Paradigm	Key characteristic	Benefits	Drawbacks
Hierarchy	Decomposition	Maps to many common domains; may scale well	Potentially brittle; can lead to bottlenecks or delays
Holarchy	Decomposition with autonomy	Exploits autonomy of functional units	Must organise holons; lack of predicatable performance
Coalition	Dynamic, goal-directed	Exploits strength in numbers	Short-term benefits may not outweigh organisation construction costs
Team	Group-level cohesion	Addresses larger grained problems; task-centric	Increased communication
Congregation	Long-lived, utility directed	Facilitates agent discovery	Sets may be overly restrictive
Society	Open system	Public services; well-defined conventions	Potentially complex; agents may require additional society-related capabilities
Federation	Middle agents	Matchmaking, brokering, translation services; facilitates dynamic agent pool	Intermediaries become bottlenecks
Market	Competition through pricing	Good at allocation increased utility through centralisation; increased fairness through bidding	Potential for collusion, malicious behaviour; allocation decision complexity can be high
Matrix	Multiple managers	Resource sharing; multiply influenced agents	Potential for conflicts; need for increased agent sophistication
Compound	Concurrent organisations	Exploits benefits of several organisational styles	Increased sophistication; drawbacks of several organisational styles

Table 5.1: Organisationsal Paradigms
Comparing the qualities of various organisation paradigms (reproduced from [150]).

negotiation capabilities, whereas federations may be used to interface between incompatible groups.

5.1.7.3 Hierarchy

In a hierarchy, agents are arranged in the structure in Fig. 5.1(a) which shows a directed graph where instructions from higher levels are passed to nodes (agents) at lower levels. These relationships are typically determined *a priori* and are related to the functional role of heterogeneous agents. Whilst instructions from higher levels are obeyed by those lower down, there is often scope for agents responsible for implementation to act autonomously to determine *how* the instructions should be fulfilled from its array of capabilities.

Hierarchies can be used to improve scalability by dividing agents into smaller groups, each with a controller. No single agent is directly responsible for too many others. However, this architecture can be brittle: nodes become single points of failure.

Data flows back up the graph to assist the decisionmaking agents. There is a trade-off to be made about which data are passed up to higher levels; more data can result in better decisions, but as data is aggregated up progressive levels, the communication requirements may rapidly become a burden.

A role-based approach can be used to explicitly represent hierarchical relationships, as well as social constraints imposed on agent actions[152]. Barbuceanu rejects the necessity of mutual goals in the context of multiple subordinates working on subtasks whilst ignorant of a supertask and criticises the joint intention model as overly restrictive in the communication required to abandon a commitment; however, he does not consider the subtask as the joint goal held with a supervisor. In this highly-decomposed hierarchical structure, the agreements comprise two participants only. Goals and plans to act are created dynamically by inference from communication and the predetermined obligations, permissions and interdictions inherited from a defined social role.

In other contexts, however, the obligation to act derived from hierarchy results in a rational plan to act, adopt joint intention or otherwise collaborate: the penalty for not obeying hierarchy may be explicitly represented in a utility function, or an agent's reasoning functions may include "hierarchy rationality" as well as economic rationality. As with Distributed Problem Solving (DPS), individual economic functions may not be relevant.

5.1.7.4 Congregation

The *congregation* grouping characteristic — Fig. 5.1(d) — is modelled after human interactions. The congregation is a group that persists in time (unlike a coalition) to perform a function. The function itself is one important facet (such as common religious practice, after which it is named), although other aspects — such as the location of the members, are also significant. Even virtual congregations may have a form of location, such as the campaigning group 38 Degrees which organises online, often in subsets of its membership. In the power systems domain, microgrids tend not to use this structure: whilst microgrids may include clustering of participants, even in a single forum (i.e. a market interface), the activity is usually for trading between individuals, rather than achieving a common goal. Congregation examples might be more readily found in Virtual Power Plants (VPPs) where a number of participants combine efforts to provide a particular service, presented externally as a single effort through an aggregator.

The persistance and commonality of a congregation allows its members to improve performance of the task of identifying collaborators by reducing the search space to suitable partners. If partner selection has cost, this may increase utility by selecting counterparts with a higher likelihood of more successful cooperation. Agents in congregations maintain individual rationality, rather than optimising a group payoff[153]. In this implementation, congregations are formed over successive rounds of a game and all members of a congregation are treated as a coalition to receive a payoff. In other implementations where congregations do not directly play a game, coalitions may be formed from congregation members. Congregations are self-selecting: if agents do not achieve any benefit, they may infer that a different congregation is more suited to them; an agent may increase its own utility but also that of the congregation it leaves by establishing that their respective goals or methods are incompatible. Clans are a complementary concept to congregations, but have a notion of trust and require agents to be invited to join, rather than using self-selection[150].

One property of the congregating process in [153] is the use of labelers, who assign identifying information to a locus (commmon feature) that connects congregators. (Potential) Congregators must share a common semantic understanding of these labels. These labels could be information about function or membership, and congregators rank these labels in order or preference. Congregators estimate the utility they might derive from membership of a congregation and then choose the locus that maximises this payoff. Labels need not actually include any information about a congregation; congregators may learn which are more preferable through trial-and-error, rather than identifying common ground.

The task-*satisficing* coalition formation method in [154] draws on concepts similar to congregations. However, in this case agents are not explicitly part of a congregation, but rather each agent maintains an internal representation of neighbours as a group of those it believes are able to assist in a particular task. To solve a multi-agent problem, they choose an initial coalition based on a ranking of agent capabilities, past performance and knowledge of the environment. A negotiation phase selects willing agents from those identified as able. If the formation of this coalition is unsuccessful, it will attempt to form another coalition with lower-ranked agents. Whilst it is not guaranteed that a successful coalition will be formed, it is reasonable to start with a coalition of agents that have previously cooperated successfully. The agents are also capable of learning satisficing negotiation strategies that allow fast, albeit non-optimal, cooperation by re-attempting successful proposals instead of trying to maximise payoff.

5.1.8 The Cooperation Process

Wooldridge and Jennings[155] provide a simple definition of cooperation in a formalisation of Cooperative Problem Solving (CPS): “Cooperative problem solving (CPS) occurs when a group of logically decentralized agents choose to work together to achieve a common goal”. The stages involved in this are as follows:

- Recognition: The CPS process begins when some agent recognises the potential for cooperative action; this recognition may come about because an agent has a goal that it is unable to achieve in isolation, or, more generally, because the agent prefers assistance.
- Group formation: During this stage, the agent that recognised the potential for cooperative action in the previous stage solicits assistance. If this stage is successful, then it will end with a group having a joint commitment to collective action.
- Plan formation: During this stage, the agents attempt to negotiate a joint plan that they believe will achieve the desired goal.
- Group action: During this stage, the newly agreed plan of joint action is executed by the agents, which maintain a close-knit relationship throughout; this relationship is defined by an agreed social convention, which every agent follows.

Though game theory aspects are discussed in more depth later, it is useful define groups and results as follows[1]:

Definition 1 *Coalitional Game: The game is a pair,* (N, V) *where*

- N *is a finite set of agents*
- S *is the members of a coalition,* $S \subseteq N$
- $v : 2^N \mapsto \mathbb{R}$ *is a payoff to each coalition* S*, also called the characteristic function*

A key assumption of many analyses and strategies, particularly those identified in Section 5.1.9, is that the payoff function is known in advance for agents to evaluate.

Wooldridge and Jennings[155] point out a simplification in their formalisation of the *recognition* task, that an agent knows in advance the identity of the group it believes could cooperate to achieve a goal. However, this results in a statement of the (mental) conditions that allow for cooperation:

---

[1] Other representations are available, such as weighted majority[156].

$$
\begin{aligned}
(\text{PfC } i\ \phi) \overset{\text{def}}{=} & (\text{Goal } i\ \phi) \wedge \exists g \cdot (\text{Bel } i (\text{J-Can } g\ \phi)) \wedge \\
& \neg(\text{Can } i\ \phi\ ) \vee \\
& (\text{Bel } i\ \forall \alpha \cdot (\text{Agt } \alpha\ i) \wedge (\text{Achieves } \alpha\ \phi) \Rightarrow (\text{Goal } i\ (\text{Doesn't } \alpha\ )))
\end{aligned}
$$

That is to say, agent i has the Potential for Cooperation (PfC) to achieve goal ϕ iff:

- it holds the goal ϕ
- there exists a group g, and i believes that g has the capability to achieve ϕ
- that i cannot achieve ϕ, or, that it has a goal not to perform an action that would otherwise achieve ϕ (perhaps on grounds of cost, or other conflicting objective)

Given that an agent has recognised that there is potential for some form of action a group may take, the actual method of putting together that group may vary considerably: the social, organisational or economic contexts the agents exist in have a considerable impact in the subsequent team formation task. The *team formation* or *coalition formation* step in cooperation described above is a large topic in itself. Section 5.1.7.2 describes coalitions and the formation problem, but there are other organisational structures that can be used; in contrast to a coalition formed dynamically in response to a problem, alternative structures are possible, and may differ in terms of approach, autonomy and persistence. Putting together a group to solve a problem requires communication. Foundation for Intelligent Physical Agents (FIPA) (Foundation for Intelligent Physical Agents) standardised the language FIPA Semantic Language (FIPA-SL)[157] and set of communicative acts[158], following the semantics of the Beliefs-Desires Intentions agent architecture[147]. FIPA-SL can be used for representing many types of information and is applicable in many of the situations where communication is necessary, not just in establishing a group of agents. *Speech acts* are the communication of information that are in themselves actions; concepts such as *propose* or *agree* are speech acts, and these can be both rational actions and message contents to be exchanged. There is a need for interaction protocols to give meaning to these messages from the context of an exchange. FIPA has also determined the contract-net protocol[159] as a standardised process of interactions in task allocation with negotiation, but stops short of describing the process for group formation; The formation problem is structure-dependent, but there are approaches to generalising the problem[160] through *cooperation protocols.*

Agents may be fully acquiescent to requests of others, acting in the best interests of the wider group. This is referred to as DPS, Cooperative Problem Solving (CPS), or Cooperative MAS[161]. In DPS, agents cooperate in order to increase the overall outcome of the system, and are not concerned with personal payoffs. There is a considerable body of work on the topic of DPS, and Multi-Agent Cooperation and Collaboration has been a national conference theme since 1991 at the American Association for Artificial Intelligence (AAAI). However, cooperation can still be achieved where agents are self-interested. As Shoham and Leyton-Brown point out, cooperation does not mean a loss of self-direction:

> It does not mean that... each agent is agreeable and will follow arbitrary instructions. Rather, it means that the basic modeling unit is the group rather than the individual agent.[162]

A game-theoretic approach incorporating some payoff mechanism can produce a system-wide emerging result. Whilst this may not be directed by a central decision, a well-designed mechanism may implicitly direct the system behaviour without sacrificing agent autonomy, providing a balance between self-interested agents and a fully altruistic DPS implementation. In this context, we may design markets where agents bid to supply or purchase a resource or service. We may describe the system in terms of payoff for the result of certain interactions, and engineer these such that successful strategies will incorporate a cooperative element.

Where we have a superadditive environment, any addition of an agent to a coalition does no harm:

Definition 2 *Superadditive game:* $\forall S, T \subset N$, *if* $S \cap T = \emptyset, \implies v(S \cup T) \geq v(S) + v(T)$

However, it this is often not the case in real-world application. Communication and complexity both increase with the number of agents, creating a cost to add each new agent that may not be outweighed by the corresponding benefit. There may also be task interactions or conflicting goals specific to the domain. For these or other reasons, the payoff function $v(S)$ may be dependent upon the members of the coalition, rather than the action they may produce. The *grand coalition* is the coalition of all the agents in a finite set of players; in the case of a superadditive game, the grand coalition is an optimum, providing the maximum payoff in the game. However, non-superadditive environments do not guarantee this optimum, and there must be a means to decide which agents should act together.

Having identified potential for cooperation and formed some group of agents to work together, the agents must then actually perform the tasks required. In contrast to apparently coordinated but emergent behaviour of autonomous actors, Levesque et al. argue that for joint action, agents acting together must have an awareness of the mental state of the other agents in the group[148]. This is a stronger requirement than outlined above that agents know there is the possibility for cooperation[155]. Following this definition, they give an outline of the agent mental attitudes required to define joint action. A mutual belief must be held about a proposition, allowing a mutual goal to be held. Building on "Intention is Choice with Commitment"[147], we can generalise to a group of agents with joint commitment to the goal: they may have joint intention.

Definition 3 *Joint intention is a joint commitment to perform an action while mutually believing throughout the action that the agents are doing it.*

Further, agents are required to inform the other participants of the joint goal if they come to believe it has been achieved or is impossible. On the one hand, the requirement to inform the other agents upon discovery of success or failure makes for a more strongly held attitude, but on the other it creates a problem: attempts are not abandoned based on uncertainty about other agents, but on whether the task is achievable. There should be a mechanism to identify when another agent has stopped participating for any reason (whether as a result of a defection in self-interest, or communication failure or otherwise). Agents are not committed to particular individual actions that may be irrelevant due to inaction on the part of other agents. It leaves the actual execution of actions to later work. The overall joint problem may be broken down into subtasks which may have a precedence order or require a specific agent to undertake, particularly in a society of heterogeneous agents. Some further coordination may be required; but the formulation above means that all of the agents remain mutually committed to the goal, even where they are currently inactive and dependent upon some other event.

5.1.8.1 Recognition

An agent may identify a need for cooperation through analysis of its beliefs, desires and intentions (or corresponding rationalisation from any other cognitive model). This forms part of the means-end reasoning step in performing actions: it must establish that another agent may be capable of an action that would contribute to objective achievement[151]. This reasoning step is implementation-specific and hence beyond the scope of this document. However, the simple determination that cooperation is possible is insufficient, and these agents must be identified.

Communication may form a part of the recognition process. In the coordination approach presented by Durfee and Montgomery, agents broadcast high-level information about their behaviour[163]. In the coordination problem, agents may adjust their behaviour to resolve conflicts by exchanging information about their actvities arranged around 6 dimensions (*who, what, when, where, how* and *why*). This exchange can also facilitate recognition of potential positive interactions as an agent discovers alignments or opportunities in these behaviours. At a lower level of detail, intention transfer[164] is a recognition technique that allows an agent to adopt another's commitment as its own when a group-level goal is identified. A common theme to the agent discovery processes is information broadcast, such as in a Call for Proposals announcement[158], but there are techniques that do not require this communication function.

A representation of properties of other agents may permit reasoning to identify their ability to assist in some particular task. To use this approach, an agent must maintain a list of agents and their capabilities. Facilitation agents can be used to streamline the process of identifying collaborators. A "yellow pages" agent maintains a list of agents and their capabilities. Agents register with this service when joining a society; an agent may interrogate this list to find those with the capabilities relevant to a particular task, meaning relevant data is communicated when necessary, rather than broadcast in advance. The initial coalition generation approach in [154] uses a learning technique to store previous successful interactions to rank potential collaborators for a new task.

Multiple agents may be in a position to collaborate, and possibly with different outcomes or payoff functions. Decisions on how (or even *if*) to optimise the system form a part of the group formation task, discussed in Section 5.1.8.2.

5.1.8.2 Group Formation

Once a group of collaborators (with a relevant organisational structure) has been identified, a process is required to consolidate the group and commit them to solving the task.

5.1.8.2.1 Group consolidation

Whichever structure is used, the coordination approach should incorporate a dynamic element that allows for the maintenance and failure that is a feature of real-life systems[165]. As well as the formation team function, failure tolerance and self-diagnosis should extend throughout the life of the collaboration effort.

Ketcphel enumerates desired characteristics of a method to form coalitions[166]:

- Stability: that agents will not defect to another coalition, usually by satisfying individual rationality and maximising payoff.
- Efficiency: The method should have low communication and computational cost.
- Decentralisation: Agents should participate directly to increase resilience and scalability, and to reduce bottlenecks.
- Symmetry: agents should have similar computational requirements. Addition of members should not disrupt the system.

Ketchpel's algorithm consists of an iterative process to join pairs of entities at a time. If we consider a task that can only be solved by a coalition of many agents, it may not be individually rational for an agent to join a small unsuccessful coalition in the first iterations. This method may thus fail to produce a useful outcome where the domain does not satisfy the superadditivity assumption. However, the process of communication for partner identification, calculation of utility, negotiation and unification into a single group still hold for other approaches to the group formation problem.

Formation of a coalition is computationally complex. Enumerating the possible coalitions to find an optimum solution is an NP-complete problem, increasing in difficulty with the size of an agent society.

Shehory and Kraus [161] present a method of solving a problem that accounts for task precedence and the capabilities of heterogeneous agents, but acknowledge that scalability is an issue. Computing optima requires full knowledge of tasks, agent capabilities and payoffs, which may not be feasible. However, communication requirements for negotiation and formation may also pose a problem. Sandholm and Lesser describe how *bounded rational* agents can make a decision on coalition choice by trading off solution quality for computational complexity[167]. Alternatively, past performance can be used to select preferred collaborators: in the iterative process in [154], an agent maintains a list of likely collaborators, proposes it, and if necessary tries again with those ranked lower. Directly solving a system optimum is not always necessary; in a market approach, agents will join a coalition of service buyers and sellers where a price is acceptable to them. Repeated interactions with agents joining, leaving or forming coalitions should cause payoffs to approach some equilibrium.

5.1.8.2.2 Group agreement

In the DPS approach, the negotiation phase consists simply of asking an agent to cooperate followed by acceptance or rejection. However, where market or other considerations apply, there may be negotiation or bargaining. The game theory aspects of what offers are made, whether they will be accepted or even adhered to will be discussed in Section 5.1.9.

The calculation phase involves each agent determining the value of a coalition to itself; this will inform offers it makes or is willing to accept in the negotiation phase. An agent may demand a particular payoff to agree to participate. However, a well-defined payoff function is not always available, and different agents may have different estimates of the value of the coalition. Ketchpel proposes that one agent is elected as the coalition manager through an auction, and takes on the risk. The manager pays its partner a fixed amount, but receives any surplus or deficit of utility[168]. In common with the algorithm in [166], large coalitions are built up from smaller ones over an iterative process, and in each instance the managerial relationship is maintained such that the non-managing agents always receive the agreed payment.

Bargaining consists of a cycle of offers, counteroffers and finally acceptance or termination when no more offers are made. This might be at the point where an offer has been made, accepted and other parties no longer try to disrupt that coalition[156]. The bargaining process is of more relevance when forming coalitions competing to maximise payoff to members; however, it offers a mechanism to resolve the case where two groups discover and attempt to solve the same task.

5.1.8.3 Dialogue

Burmeister, Haddadi and Sundermeyer proposed a generic cooperation protocol[160], arguing that a dialogue approach communicates more information through context than the act or contents of communication itself. Agents follow branching steps through a conversation as messages are exchanged. Messages themselves comprise type, descriptor and content organised in a particular syntax; purpose is derived from the position in the protocol tree. The authors refer to dialogue types as *cooperation patterns*, such as contracting, bargaining or persuading, in contrast to the message types such as *inform, accept* or *reject*.

FIPA Contract-Net[159] is a contracting pattern. It specifies how a task can be assigned to other agents without the process of establishing the joint commitment described in Section 5.1.8. Larger tasks are decomposed into well-defined subtasks. A Call for Proposals (CfP) is issued, and other agents bid for the opportunity to perform the task; the initiating agent then awards the tasks to the agents, and a contract is formed between those agents. Large tasks are thus formed of many two-party cooperative efforts where one pays the other to execute the subtask, rather than a combined effort where all parties are aware of a higher goal.

Joint commitment requires establishing joint beliefs. The commitment step requires negotiation on the part of the agents in order for them rationally to commit to action, whether to confirm agent acceptance an offer of coalition or to propose a specific payoff distribution. This can create a need for an intermediate pre-commitment step[151]. After group selection, a negotiating position is proposed. Once this is accepted, the joint commitment is formed. At each step it is rational for an agent to communicate some information about its beliefs or bargaining position. At the point where negotiation concludes, the agent believes that the other agent is either already committed, or will eventually become committed. These rational communication steps form the negotiation protocol.

5.1.8.4 Group Action

5.1.8.4.1 Task allocation

Before task allocation is determined, the joint ability (J-Can) component of Potential for Cooperation may not be satisfied. A task to achieve a goal can be decomposed into subtasks, and these tasks contracted-out to other agents. Design effort may be required in setting up a specific task decomposition and communication — an agent may not know what the subtasks are.

5.1.9 Market Mechanisms in Cooperation

Agent interactions for cooperation can be a result of an effective market-driven environment. Agents are typically modelled as rational economic actors, which allows the application of game theory to explain or design particular outcomes. However, there are factors regarding fairness and properties of the context that will determine just how applicable this model may be. This section introduces game theory aspects of cooperation and discusses some of those issues.

5.1.9.1 Payoff Distribution

In terms of coalitional games,

> "The central question in coalitional game theory is the division of the payoff to the grand coalition among the agents[162]".

If we consider division of the value $v(S)$ of a coalition as a necessary function to facilitate collaboration, there is a requirement for some payment mechanism.

Definition 4 *Transferrable Utility Assumption: The payoff to a coalition may be freely divided and distributed to its members.*

Utility has no fixed scale or unit, but money has a convenient linear value which can be used to assess an actor's utility. Payoff can usually be represented by money. However, payoff is not always indivisible. In money transactions, this might only be an issue with, for example, a distributed system where small computing tasks are outsourced to other agents at a cost of fractions of a penny. Payoff is more difficult to transfer in non-money transactions where a good is only available in whole units, or only the first instance has value. In a electric network, a payoff could be permits for a number of wind turbines to supply power; however, a coalition owning a virtual wind farm cannot allocate fractions of permits to its participants, and there is no point in allocating two permits to a single turbine.

Inter-agent transfers of utility are called *sidepayments*. This mechanism causes Definition 4 to hold. These payments may be made in advance of a coalition forming: agent A pays agent B to form a coalition $S_1\{A, B\}$ rather than B's preference $S_2\{B, C\}$. This secures a higher payoff for A by compensating the loss of utility to B. The end result is indistinguishable from an alternative distribution of the payoff $v(S_1)$ between A and B. The inverse of Definition 4 is nontransferable utility, where $v(S)$ is a pre-defined vector containing payoffs to individual agents.

With transferrable utility or individual payoffs, the grand coalition may not always be formed; agents may prefer to form other coalitions that maximise personal payoffs. Consider (N, v) with two successful coalitions: $S_1 = \{A, B, C, D\}$ and $S_2 = \{A, E\}$. If v is split equally between coalition members, there is an incentive for agents $\{A, E\}$ to form S_2. In this coalitional game, the task has still been performed and the payoff achieved, but the payoff only benefits two agents in N. Given that v is dependent upon S, this may also be at the cost of lower group utility. Whether this is *fair* in a given agent society is problematic as there are many ways of defining fairness. A Pareto optimal coalition is found if there is not an alternative in which *all* the agents may achieve a better (or at least equal) payoff. Whilst a group may achieve a higher overall payoff, it is necessarily at the expense of one of its members.

Definition 5 *Core: the set of Pareto optimal coalitions.*

The core satisfies group, coalitional and individual rationality. However, the core may be empty where no set of coalitions satisfies all three. Coreless games require a mechanism to resolve these conflicts, to avoid a repeating cycle of proposal, counterproposal and defection[156].

5.1.10 Decentralisation

The idea of Intention Transfer was developed as a novel method to facilitate agent participation in event resolution[164]. With Intention Transfer, an agent announces its intentions to an agent society, and others may "opt-in" to perform non-explicit joint action outside of existing conventions for cooperative problem solving. This permits agents operating with incomplete information to assist others where it is appropriate, and also to plan accordingly.

5.1.10.1 Model of Agent Existence

In the Cohen and Levesque possible-worlds model described in Section 5.1.5, an agent may access (that is, inhabit) a world consistent with its beliefs about that next world (through operator B). Further, Cohen and Levesque introduce the operator G: a world may access through G all worlds in which the statement ϕ is true. This G is referred to as a "goal". An achievement goal may be represented as choosing a world in which ϕ is possible, but currently false; this conveniently drops the goal once ϕ becomes true. Sadek prefers *choice* (operator C, being distinct from an achievement goal), where an agent *chooses* the next world in line with its desires[169]. An agent may only choose C (or G) that are compatible with B; if, in the next world, the agent believes ϕ, it may not simultaneously choose a world in which$\neg\ \phi$ is true. This *realism constraint* is defined as $C \subseteq B$. The violation of this definition of the realism constraint is possible, although normally irrational (e.g. "Choose a world in which ϕ is true but you believe it not to be"). The constraint is better represented as C(i,ϕ)$\Rightarrow \neg$ B(i,$\neg\ \phi$). A combination of realism-constrained choices and beliefs prevents an agent from choices which are impossible through reducing the set of worlds to those which are compatible. At the next point in time, the agent will inhabit some world in the compatible subset made from a combination of *realistic* choices (constrained by beliefs); *free* choices, where an agent selects a preference for some property it has no current knowledge of; and open choices from its indifference to any other properties. Practically, however, the idea of world choice is rather intangible in its application to real software agents; FIPA refers to C, a desire [158], and does not include it in its operators of B, I, and the *persistent goal PG* as choice without an attached commitment or plan.

5.1.10.2 Intention

Goals and choices do not convey any necessity or compulsion on the part of the agent to act to attempt to bring about a transition to the desired world. In order to solve this problem, Intention I is defined in "Intention is Choice with Commitment"[147] as the combination of C operator and a persistent compulsion to act towards achieving a goal. Without this made explicit, an agent may never actually take steps towards fulfilling a goal, or may abandon its attempts too soon. The agent only drops an intention once it believes it is unfeasible to achieve or the new target state has been brought about. In software implementation under the BDI model described above, an active intention will result in routines being selected from a library of available plans and executed until the agent believes the task is finished, irrelevant or is unfeasible. Cohen & Levesque[147] illustrate two kinds of intention: INTEND_1, which concerns actions (e.g. "intend to switch on a load"), and INTEND_2, which concerns a state of the world (e.g. "intend to take action so that it is true that the load has been switched on"). An agent may have executed all applicable plans to try to achieve the result of its intention; however, this may have been insufficient. It may then be rational for an agent to abandon its intention, reasoning that because it has no more plans to achieve a still-unrealised final state, its intention is unfeasible. It may then reassume it as a Persistent Goal. In this implementation, however, the agent retains the intention. For simplicity, it is assumed (although not stated explicitly) that an agent's goals are achievable. It could be argued that responding to requests and informing others of its intention *is* a plan for its achievement, although it may no longer be directly responsible for its effectuation.

This is not a problem, however: Cohen and Levesque's INTEND_2 permits an intention to be held even alongside the belief that another agent might achieve the same end. Whilst not implemented here, Intention Transfer does not preclude the possibility of an agent dropping its intention once another has adopted it if it has a belief that another agent will achieve the same result.

The two forms of intention share a similar representation: INTEND_1 is of the form

$$\text{(I Agent (done VUp))}$$

where I is the intention operator, *Agent* is the agent that has committed to some action, and *VUp* is the generic action of directly raising the voltage. The expression

$$\text{(done VUp)}$$

is the statement that action *VUp* has been performed.

INTEND_2 can be illustrated with the voltage comparison proposition *voltageabovelowerbound*:

$$\text{(I Agent (B Agent voltageabovelowerbound))}$$

which means that the agent commits to action such that it should eventually come to believe that the proposition is true, where B is the belief operator.

5.1.10.3 Intention Transfer

The application of a common goal and its representation in an agent-based system is problematic in conventional power systems hierarchies with direct instructions or requests. An agent may be willing to cooperate, but be unable to do so due to some operational constraint. Or, it may cooperate, but know that doing so is to the detriment of the requesting agent. This can be a particular issue where agents only perceive their immediate environment, but can form beliefs about the rest of the world through communication. Other factors, such as negotiation to cooperate (perhaps using market factors) and settlement on a common plan [149] may be redundant or irrelevant. A social agent may wish to communicate its intentions to other agents in the system. With Intention Transfer, an agent announces its intentions to an agent society, and others may "opt-in" to this intention to perform non-explicit joint action outside of the conventions for cooperative problem solving. This permits agents operating with incomplete information to assist others where it is appropriate, and also to plan accordingly. To illustrate this, voltage drop is considered as one of the control system goals. The opt-in basis without a negotiation or acknowledgement of commitment means there is no compulsion on the part of the receiving agent to adopt an intention. However, if an agent's logic dictates that it is willing and able to assist, the agent receiving such an intention may then adopt it as

$$(\text{I Agent (B RemoteAgent voltageabovelowerbound)})$$

which would be held until informed of success; thus, the intention has been transferred.

INTEND_2 is used for this application as a scalable and future-proof way to describe a property to be changed. Agents with plans with the rational effect of achieving (or working towards) the same desired world state (rather than a commonly-named similar action) would be able to participate, even as new device types are added to the electrical network.

There are key distinctions from an agent performing a direct REQUEST for assistance (such as in [113]): Firstly, the initiating agent is not required to specify which agent it wishes to perform some action. Secondly, a receiving agent may conclude that others may be better situated to assist, based on knowledge unavailable to the initiating agent. Alternatively, an intention may be adopted as a persistent but non-planned goal that may become available in the future, rather than flatly refusing cooperation.

A more general communication of agent intention can be useful; other agents can reason about the intentions of others to determine their own actions, particularly in regard to forward planning. However, this specifically examines Intention Transfer as an approach for decentralised system control where agents opt-in to achieve a common goal.

With a method for opting-in to action described, one further challenge is to decide which of the willing agents should assist. Mechanisms for doing so include multi-agent system negotiation, elections or auctions, centralised allocation, technical assessment or random chance. Without some form of intervention here, an entire agent society might attempt to switch simultaneously to solve a short-term minor problem at a single node, with potentially serious adverse consequences. In a real system this would be an unreasonable response, so there must be limitations on the extent of action.

5.1.10.4 Applying Intention Transfer

The application of this novel approach requires a multi-agent test platform, and adaptation to the power system domain. It must also address the redundant-action issue outlined above. In this implementation, an agent calculates a short delay before acting, based on a function of the difference between its own voltage and that of an agent announcing a problem. As previously noted in Chapter 4, in more complex systems voltage may not necessarily indicate network location - however, this added an important plug-and-play facet to agents: without any configuration or prior knowlege, the closer it is to a problem, the more effective it is likely to be in its assistance.

With this modification, this multi-agent systems algorithm can be tested on the power systems voltage control case used for illustration above. The next section describes novel testbed systems that were developed to combine power systems with a multi-agent platform; finally, results for trial of Intention Transfer in power system control are presented in Section 5.3.3.

5.2 Test Systems

5.2.1 Overview

This section documents the design and creation of three novel test systems that integrate power systems and multi-agent systems for control system trials.

Initial testing and development used a single PC to develop a MAS with basic communicating agents without any sophisticated underlying electrical model. It was apparent early on that this single computer was insufficiently powerful for the task of emulating the behaviour of embedded controllers continuously monitoring network conditions for more than a few agent instances. In addition, development toward a modelling approach incorporating real-time simulation, multiple physical nodes and hardware-in-the-loop was intended to create a realistic testing environment.

The three systems that were developed were *a*) a proof-of-concept MAS/distribution network simulation system; *b*) a cluster supercomputer implementation; and *c*) a unique electrical network/hardware-in-the-loop/software platform. The multi-agent system approach remained common across the three systems, and is presented here first.

5.2.2 Common Multi-Agent System

5.2.2.1 Overall Platform

The common platform comprises the non-hardware elements of the environment in which the software agents exist. This includes the choice of messaging system and network protocols, inter-agent language and ontology, and societal structure.

Agents can be implemented in a number of ways to operate successfully within this environment; however, other common choices were made in their development: the control agents (and facilitation agents in 5.2.3) were developed using the same programming language and with the same agent development tool. These agents were also developed with a common approach to the internal logic systems - the BDI agent cognitive architecture.

5.2.2.2 Agent Cognitive Architecture

Cognitive architecture — whether BDI, CLARION, DUAL or any other — is an implementation choice. Agents are unaware of the internal structure of others; division of an agent into subsystems serves to simplify the construction of an agent to the programmer. For reasons of familiarity, the basic Belief-Desire-Intention model is used.

This BDI model is often implemented in slightly different ways. One approach is to use layers of management, coordination and execution, where overall goals are selected, a general approach taken, and then plans within that general approach reasoned upon and chosen based on their perceived feasibility. Another is to use an interpreter which provides the facility for reasoning through its connection with the agent's belief base, as PRS [170] does. A developer can choose any architecture or psychological model: the existence of rigid rulesets to determine agent actions, or highly specific applications, or a need for flexibility of response may inform this choice. After implementation, an agent functions as a "black box" to all others on the system.

BDI has been implemented in several frameworks which have been designed from different standpoints to facilitate agent development, building on theory of knowledge representation and cognitive models. [171] presents an introduction to several of these frameworks, in which their respective developers outline their approach and rationale.

Several toolkits are available for implementing the BDI cognitive architecture, of which several are examined in [171]. JASON uses a language called AgentSpeak(L), which is specifically designed as to incorporate representation of BDI attributes, inter-agent communication and internal reasoning mechanisms. JACK uses BDI representation internally for high-level beliefs, plans and intentions, but emphasises flexibility in communication, platform and programming language support as vital in the multi-agent approach. The JADE framework (discussed in Section 5.2.2.7) does not implement BDI directly, but rather provides a scheduler that allows a developer to implement goal, intention and plan management through addition and removal of behaviours, whilst providing tools to facilitate compliance with FIPA standards. Use of existing frameworks is useful to save development time and facilitate interoperability.

5.2.2.3 Agent Society

The discussion in the sections above indicates that holarchy is an appropriate societal structure for a localised control goal; however, in this implementation the opt-in cooperation needs even less hierarchy and is akin to a congregation.

5.2.2.4 Multi-Agent System Standards

5.2.2.4.1 Interoperability

Foundation for Intelligent Physical Agents (FIPA) was set up in 1996 with the intention of forming standards to facilitate interoperation between agents and multi-agent systems. Their standards define FIPA Agent Communication Language (FIPA-ACL), sequences of communication for particular interaction types, and content languages for representing and communicating knowledge within those messages and interactions. This includes FIPA Semantic Language (FIPA-SL).

5.2.2.4.2 Communication

The Communicative Act Library Specification [158] formalises FIPA-compliant discourse types, e.g. *INFORM, REQUEST, PROPOSE, AGREE*; it also describes the basis for representing information in message structure through FIPA-ACL and in content through FIPA-SL. The specification of FIPA-SL [157] implements the first-order modal logic used in this chapter. The state of agents and various properties of the world can be queried and exchanged by combining these two standards.

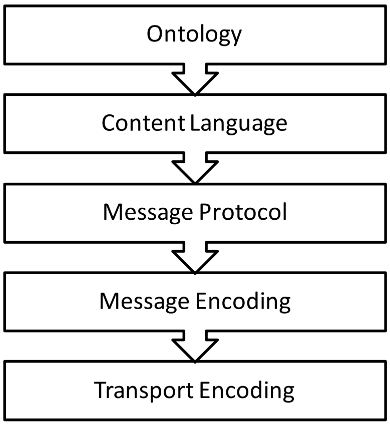

Figure 5.2: Information encoding structure

Figure 5.2 shows layers of information encoding for agent data exchange. The encoding of information is independent of the structures (or ontology) used to describe the domain.

The use of XML-based encoding of data for communication was explored as an alternative to the standard String encoding used in FIPA-SL, but the latter had a smaller packet size and some parsing functions already available through the Java Agent DEvelopment framework (JADE) platform.

5.2.2.5 Language and Logic

FIPA-SL is based from work built on by Cohen and Levesque[147] for the basic agent BDI structure, which used a “possible-worlds” basis and a first-order modal language (later, the meaning of agent *intention*, a mental attitude of an agent that may be communicated, was revised and clarified in [169]). FIPA-SL allows for communication of descriptions of agent beliefs, desires and intentions, as well as logical propositions about these. Definitions in the ontology are terms that can be included in dialogue; these complex messages allow agents to rationalise about the mental attitude of other agents, as well as the physical reality of qualitative or quantitative data.

5.2.2.6 Ontology

Ontology is a study of *being*: specifically, description of concepts and relationships between them. *An* ontology is used in this context to refer to a specification of descriptions within a domain. A common set of data structures in the form of an ontology is used to allow all agents on the network to communicate using a shared language. An ontology can be used to create a subsumption and composition (*is-a* and *has-a*) directed acyclic graph linking all the components and information relevant to concepts within the domain. The relationships can be used to allow an agent to perform reasoning about properties of the world.

An ontology describes how a concept is represented as combinations of basic data types (such as *string* and *integer*) and other concepts. These can be instantiated with data from the physical world to represent reality. Consequently, terms in the defined ontology can be inserted into agent discourse in order to allow reasoning and discussion of real entities, concepts and properties.

Ontology, as description of *being*, is inherently domain-specific. There are some relevant standards for describing items within electrical distribution networks, such as the data models within IEC 61850 for Electrical Substation Automation design. However, these do not cleanly map onto all of the relevant entities and concepts. The ontology developed by Trichakis[1] goes some way to describing the problem domain. However, the approach relied on a centralised approach with known network locations, which introduces configuration and maintenance barriers to using agents as an expedient to a plug-and-play system. Other design decisions, such as the use of “customerVoltage”, make little sense when considering flexibility — such as monitoring voltage at DNO assets, e.g. transformers. There is a need to communicate data that is not easily represented in this system.

Building on the basic ideas, however, a relevant ontology was developed, using the Protégé ontology development tool[172]. The ontology defines actions, predicates and concepts specific to the SSEZ domain: concepts, such as ‘generator’ or ‘load’ can have fixed attributes (e.g. ‘rating’ or ‘state’). Predicates allow comparison of concepts and the perceived world (e.g. ‘generator voltagelevel 1.02pu’). In addition to electrical network components, concepts of geography (such as GPS coordinates) were included. The ontology shared by all the agents on a system is a common structure of items that could be represented in several ways internally; this ontology is then used as the dictionary that forms inter-agent discourse in FIPA-SL.

5.2.2.7 Development Platform

Java Agent DEvelopment framework (JADE)[173] is a set of widely-used open-source Java libraries which contain tools to assist implementation of a multi-agent system; it provides a FIPA-compliant agent platform, with an Agent Management System (AMS) responsible for agent addressing, and a Directory Facilitator (DF) which contains a register of services offered by each agent [171]. It is hardware- and network-independent, and inter-agent messaging is handled transparently to the developer at a high level regardless of host platform or network medium. It has a low memory footprint (100kB) for use in embedded systems. A skeleton JADE agent has a task scheduler and the ability to register with the agent platform and perform messaging. Custom Java functions which extend a “behaviour” can then be added to and removed from the scheduler to achieve agent functionality. JADE is widely used in several industries, and has been proposed and adopted in power systems control applications (e.g. [125], [174]). As well as the advantages outlined above, it has an active developer community, acceptance in existing production and mission-critical applications (particularly in telephony), and provides a well-documented fault-tolerant platform.

An add-on to JADE was developed in 2005 which adds an interpretive engine using FIPA-SL in order to provide functionality shown in Figure 5.2.2.7. The central interpretive engine more closely follows the theoretical model in [175]. The add-on ceased to be maintained by its developers toward the end of 2009, and was an incomplete implementation. However, this software provided useful insight into agent structures, and the message parsing functions of this add-on were integrated and adapted to facilitate the use of FIPA-SL in this application.

5.2.2.8 Control Agents

Based on the common platform building blocks described in Sections 5.2.2.2 - 5.2.2.7, two novel and similar control agents were written: one for controllable loads (Load Agent (LA)), and one for microgenerators (Generator Agent (GA)).

They were programmed in Java as a result of the decision to use the JADE platform.

Following the choice of a BDI cognitive architecture, agent functionality was broken down into libraries

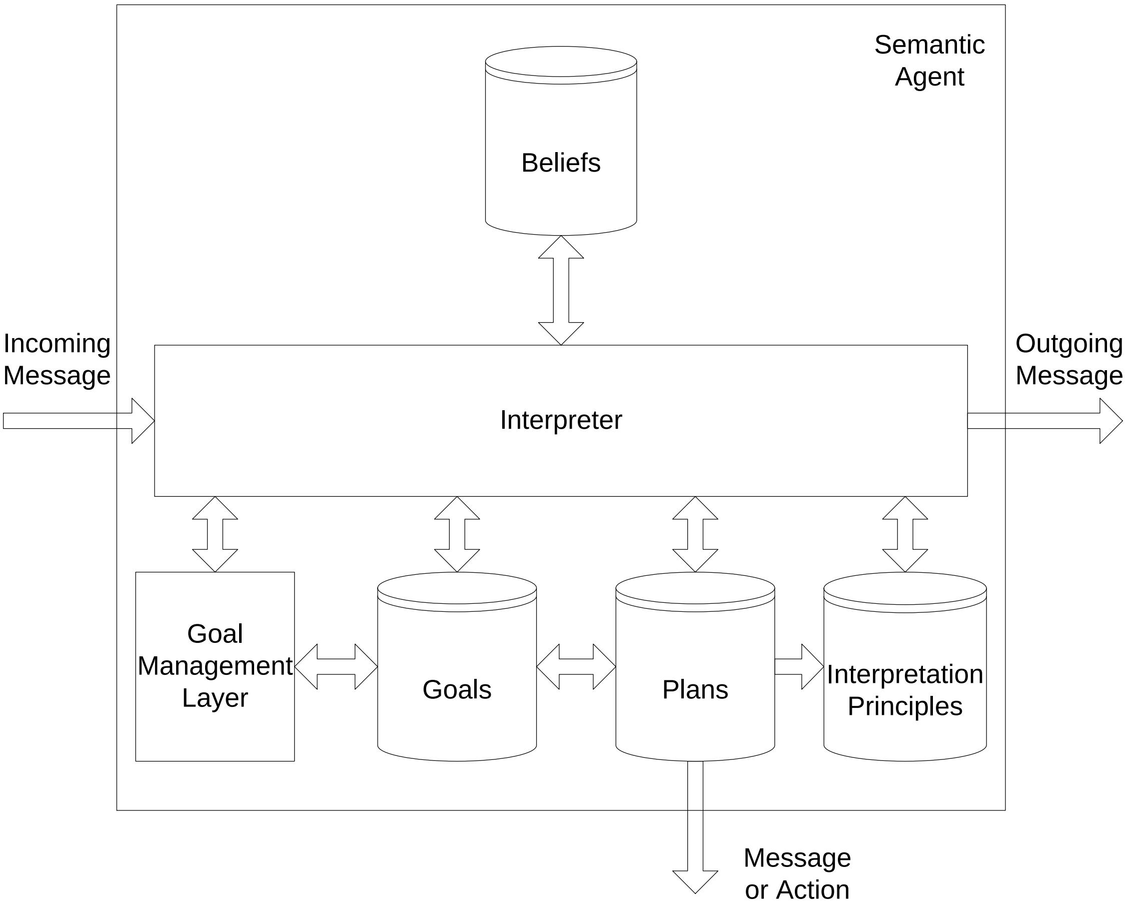

Figure 5.3: Semantic-enabled BDI architecture
An interpretive layer is used to rationalise about each of the components of a semantic agent; the diagram shows the connections between layers and libraries linked to an interpreting engine.

of goals, intentions and plans. Figure 5.4 shows the basic agent structure used, in common with the theory outlined above; Figure 5.5 shows the specific planning, deliberative and acting functions for a Load Agent separated into different layers, and the links between each.

The JADE platform contains a scheduler that executes code related to Behaviour objects in turn from a list. The first objects added to the scheduler are a goal management object and an object to maintain the agent's senses — the former tests achievability, priority and relevance of goals and can add and remove them accordingly; the latter updates the agent's perceptions of its environment.

Functions may be added to and removed from the scheduler by behaviours in the layer above as well as resulting from their natural completion. The bidirectional arrows highlight the communication of information about function success, feasibility and current state to an adjacent layer.

When action needs to be taken in order to achieve a goal, an appropriate intention is selected from the Intention library and added to the Scheduler's list.

The Intention objects are added to commit the agent to action; this include strategies for plan selection. Whilst the goals and intentions are widely applicable, the plans may be hardware-specific. Hardware and

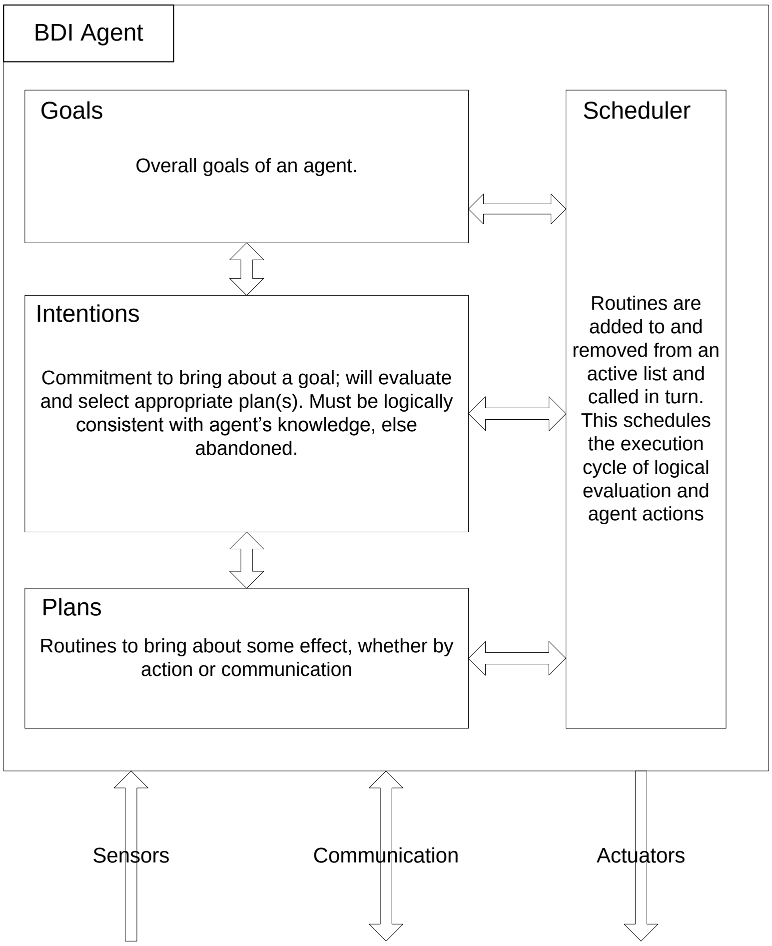

Figure 5.4: BDI agent subroutines architecture description

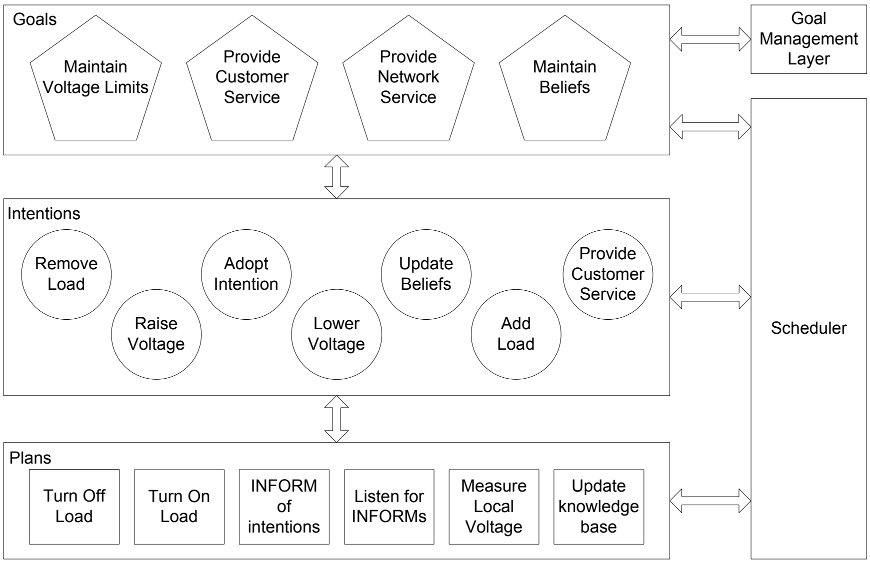

Figure 5.5: Load Agent Structure
BDI agent subroutines for a Load Agent showing goal, intention and plan libraries

low-level implementation is discussed in Section 5.2.5.2.

Without a specific interpreter provided by the platform, the behaviour objects were written to functions to allow them to be evaluated in some respects such as if they are achievable, or if they conflict with others. For example, Goal objects consist of a function to evaluate the possibility of their achievement, since a rational agent must drop unachievable goals. With Plan objects, for example, the switchLoadOff plan cannot be executed if the load is already switched off; its isRunnable() function may be tested in order for an Intention object to evaluate its suitability in a given case.

When a contingency occurs, the agent uses the JADE middleware to encode and broadcast the intention to all the other agents on the system using the Agent Communication Language (ACL) INFORM performative. An extra step is included in the software implementation: an agent is responsible for ensuring the accuracy of its internal belief base, so a message (QUERY-IF) is sent to the originating agent to request confirmation that the intention is still held before the receiving agent acts to assist.

The format and basic functions of the LA and GA are similar, although the agent internal logic was adjusted to give behaviour relevant to the controllled equipment during a voltage contingency. The agents were given a simple schedule for demand or generation, which could either be high- or low-priority. These schedules were intended to illustrate demand and production in a repeatable way and to allow for forcing of contingencies, rather than using load profiles or probabilistic models.

5.2.3 Proof-of-Concept System

5.2.3.1 Introduction

5.2.3.2 Hardware

The software was adapted to run on multiple computers, providing indication of system performance when tested with a real communications medium.

This multi-node test system comprised a network of 15 PCs, shown in Figure 5.6. They ranged in specification and connection, from a 32-bit Intel 2GHz processor with IEEE 802.11g WiFi connection at 24Mbps, up to a dual-core 64-bit 2.4GHz AMD Opteron; Each of the control nodes were connected with IEEE 802.3ab gigabit ethernet; the main platform node was connected to this router via Wireless (802.11g at 54Mb/s).

The Simulation Agent (described in Section 5.2.3.4) and Agent Management System were hosted on the node connected by wireless network. The remaining 14 nodes each hosted a Load Agent and were on the wired portion of the network.

5.2.3.3 Electrical Network

The electrical network, based on the UK Generic Network discussed in Chapter 3, was modelled in IPSA. A python script was devised to iterate across the feasible combinations of equipment states to generate

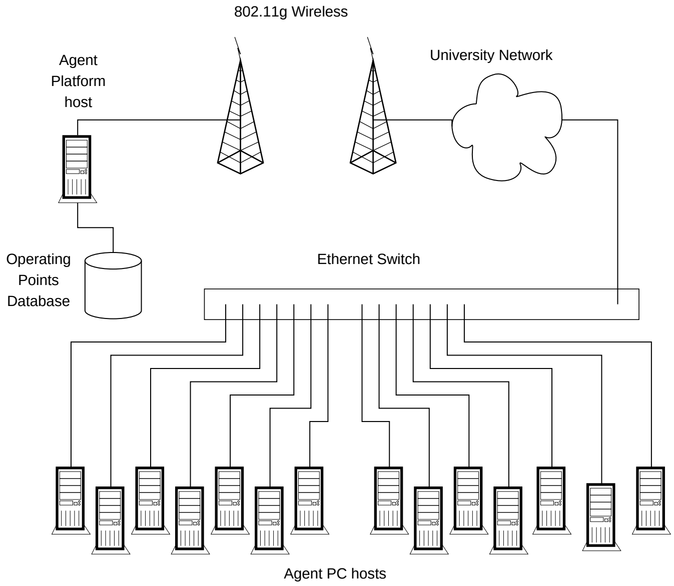

Figure 5.6: Multi-node networked PC testbed

an exhaustive set of static network operating points and corresponding equipment voltages. These data were put into a database. The voltage-monitoring components of the Load Agents were provided with measurements from this database.

5.2.3.4 Software

The Simulation Agent hosted the operating points database and supplied the Load Agents with a voltage level via the standard JADE messaging subsystem, rather than a real AC waveform to be read by a physical voltmeter.

5.2.4 High-Performance Cluster

5.2.4.1 Introduction

In the proof-of-concept system, only 14 load agents could be operated due to memory constraints. The agents were ported to a computer cluster with many identical nodes. This proved more flexible for testing and initialising agents through the use of scripting files, and expanded the simulation beyond 14 nodes. This system is described in this section.

5.2.4.2 Hardware

Durham University High Performance Computing Service has a supercomputer known as the Hamilton Cluster, which consists of many networked multi-core Intel Xeon E5520 2.26 GHz processors that can run programs in parallel within a common environment. In this case, each agent (including the SA) was assigned to an individual core by the Cluster's scheduling engine. When the cluster was commissioned, up to 1800 cores could be made available, and this has subsequently been expanded to 1950+ cores. Supercomputer time on this scale was not available for this project, but in principle could allow simulations that could cover several low-voltage feeders simultaneously.

5.2.4.3 Electrical Network

The operating points database was re-generated to include distributed generation. 14 load agents and 6 generator agents (i.e. a total of 20) agents were trialled, using the same IPSA+ model.

In practice, the operating points database used an inefficient data structure; when combined with 4GB RAM available per core, this limited the size of tests size to fewer than 30 agents. However, this was considered sufficient to examine the underlying principles on the understanding that a fast and lean database system could be applied in future to allow for testing at scale.

5.2.4.4 Software

5.2.4.4.1 Control agents

The agent structure remained substantially similar to those in the proof-of-concept system; the generator and load agents were similar except that a generator should not disconnect in an undervoltage case.

5.2.4.4.2 Startup tools

To execute agents across multiple processors, a wrapper was written in C using MPI[2] to request and access a number of cores, and then to start a relevant agent on each available processor.

5.2.5 Real-time Hardware-in-Loop Platform

5.2.5.1 Introduction

Figure 5.7 shows the overall system outline, which includes the RSCAD model and Real-Time Digital Simulator (RTDS); the Beagleboard target platform and associated Input/Output (I/O) hardware; network router; and agent platform host/Dynamic Host Configuration Protocol (DHCP) server.

5.2.5.2 Hardware

5.2.5.2.1 Embedded Controllers - Overview

A range of embedded controller target hardware was considered. In order for the test facility to be both flexible and relevant, a widely available general-purpose single-board computer was chosen with the idea that off-the-shelf components could be integrated and tested, reducing costs of development and large-scale deployment.

Figure 5.8 shows the arrangement of the embedded controller hardware. Each controller was connected to the network switch (bottom right). Attached to each controller is a daughterboard (not pictured) with an onboard Analogue-to-Digital Converter (ADC) which provides access to the RTDS voltage signals.

---

[2] Message Passing Interface, a multi-purpose tool for communication, and synchronising and scheduling parallel code on cluster computers

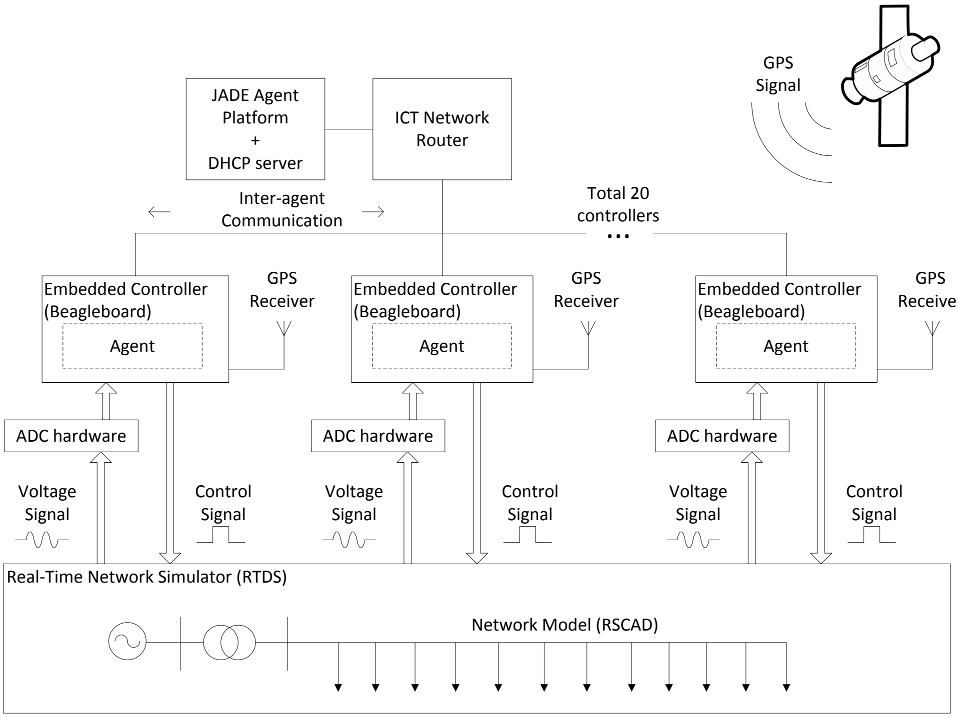

Figure 5.7: Multi-Agent Hardware-In-Loop Test Platform

Analogue and digital signals are routed through an interface box (bottom left) that also handles voltage conversion for the control signals and contains a transistor array to provide current to the optoisolated inputs on the RTDS.

5.2.5.2.2 Embedded controller mainboard

The embedded controller is the Beagleboard-xM revision C, which is a small (approximately 90mm × 90mm) single-board general-purpose computer. The Beagleboard uses a 1GHz ARM Cortex-8 processor, with 512MB RAM. The board also hosts USB and ethernet adapters. An array of 20 beagleboards were mounted to a frame with associated network cabling and power supply.

At the time of purchase, the Beagleboard-xM cost approximately £100. However, there has since been considerable expansion in the single-board computer market — the Beagleboard Black (an updated model) is now available for £44; the Raspberry Pi 2 is now available (albeit with a 0.9GHz processor) for £26 in single units. Whilst the use of general-purpose computing was previously perceived as too expensive for this application, the move to such low-cost hardware makes a future of ubiquitous location-aware controllers a realistic possibility.

5.2.5.2.3 GPS receiver

The system used a GlobalSat BU-353 USB GPS recievers.

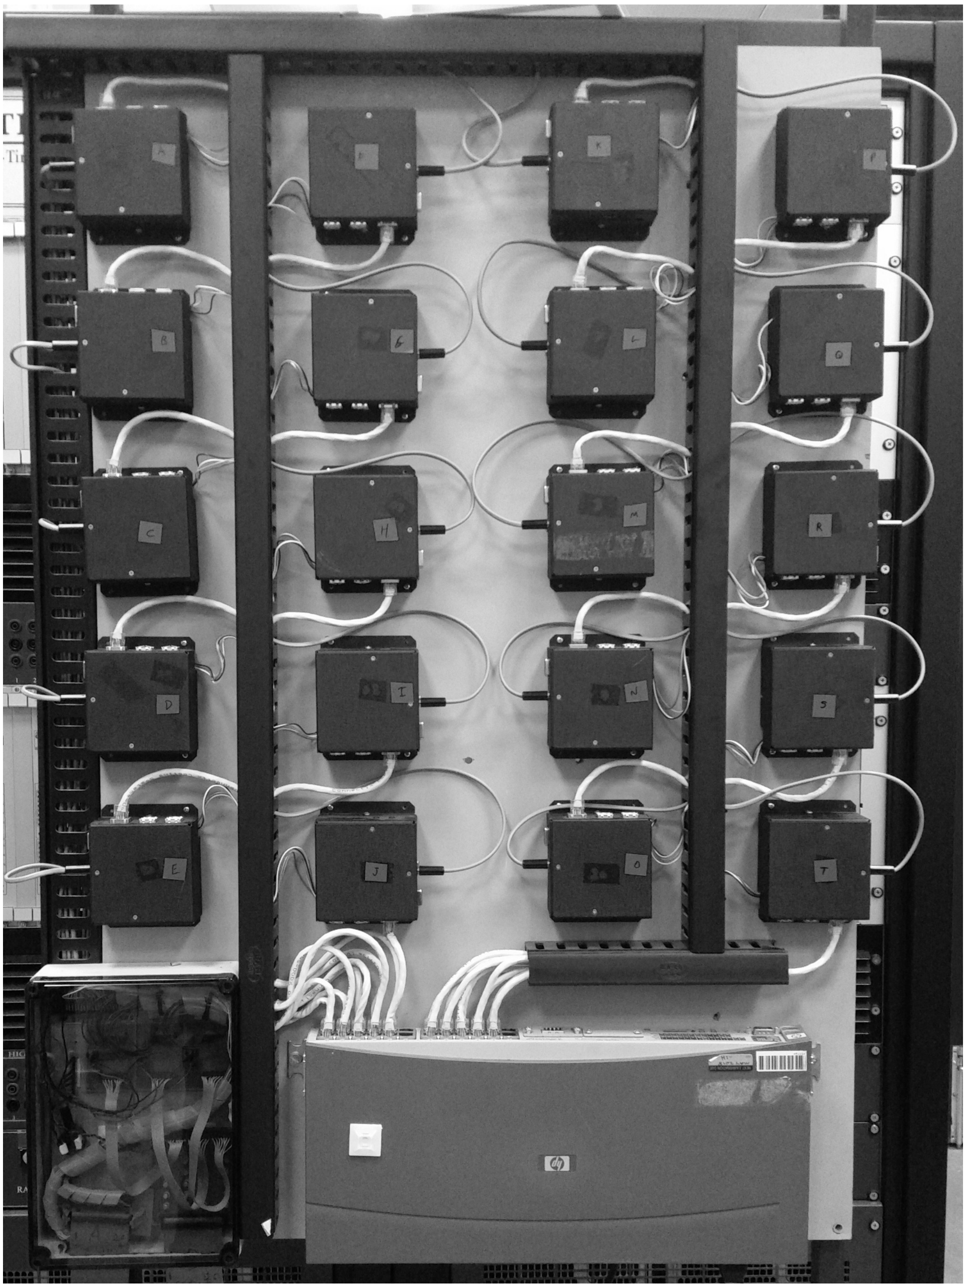

Figure 5.8: Multi-Agent Hardware-In-Loop Test Platform - Controller array
The array of 20 beagleboard controllers were mounted to a frame. Also shown: network switch, test rig interface box featuring voltage conversion/transistor array for RTDS input.

5.2.5.2.4 Signals I/O hardware

The General Purpose Input/Output (GPIO) channels used for control signals from the beagleboard to the RTDS required more current than could be supplied, so were passed through a transistor array.

The Beagleboard's onboard ADC proved to be highly nonlinear and not stable over time and thus unsuitable for this application. An external 14-bit ADC was used: the AD7899, by Analog Devices, was selected. The ADC has an accuracy of ±2 Least Significant Bits (LSBs), and the maximum sample

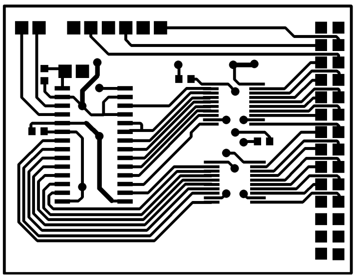

(a) Circuit board - top

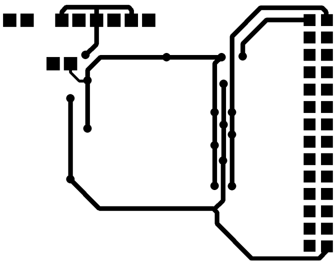

(b) Circuit board - bottom

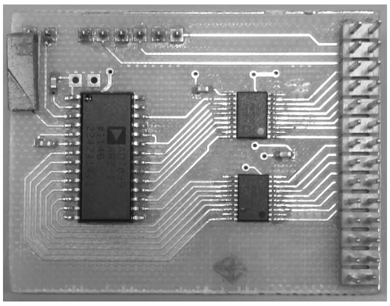

(c) Daughterboard- assembled

Figure 5.9: I/O daughterboard design
The obverse and reverse of the daugterboard are shown in (a) and (b); the AD7899 ADC and the ADG33 voltage converter pair are shown assembled in (c).

frequency of 400kHz was well above that supplied by the RTDS. A pair of ADG3300 (with a maximum data rate of 50Mbps, well in excess of the 400kHz sample rate) were used to convert voltage levels from the ADC outputs to the beagleboard inputs. A daughterboard was designed and fabricated to interface the ADC with the beagleboard, shown in Figure 5.9.

Whilst the data collection in Section 4.3.5 used an off-the-shelf plug-in transformer, for the purpose of laboratory testing a break-out box was constructed to allow the system to be plugged into a domestic mains socket or to the RTDS without the need recalibration, using a voltage transducer with a range of -10V to +10V to represent a real deflection of -500V to +500V. The output range of the RTDS was set to be the same, meaning the controller hardware could be connected to either domestic sockets or to the RTDS hardware. The break-out box is shown in Figure 5.10.

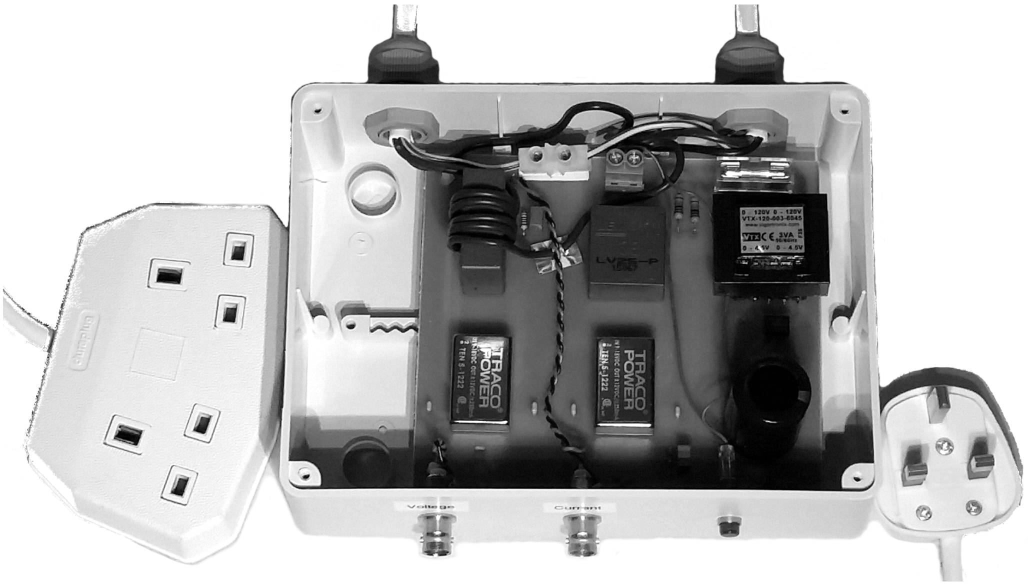

Figure 5.10: 240V Voltage transducer
The break-out box for measuring voltage. The domestic 3-pin plug and sockets are either side of the enclosure, with voltage and current signal outputs for direct measurement at the controllers. The LV25-P voltage transducer is shown as the middle component of the top row, which is accurate to within 1% and linear to within 0.2%.

5.2.5.2.5 Agent platform and router

The JADE agent platform was hosted on a laptop computer. A Hewlett Packard 2524 network switch provided IEEE 802.3u 100Mb/s ethernet between each of the embedded controllers and the agent platform.

5.2.5.2.6 RTDS

The RTDS is a dedicated electrical network simlulator manufactured by RTDS Technologies. It uses 1.7GHz Freescale MC7448 RISC processors, which have a dedicated vector execution unit making them highly suitable for solving electrical networks using the method described in [135]. The Real-Time Operating System (OS) manages the execution, which allows for a guaranteed 50μS timestep. This corresponds to a system simulated to a 20kHz precision.

5.2.5.3 Electrical Network

An RSCAD model, described in Chapter 3, is compiled into matrices and code to be executed by the RTDS.

5.2.5.4 Software

5.2.5.4.1 DHCP server, agent platform and location

The JADE agent platform was hosted on a computer connected to the router. This computer also hosted a DHCP server to allocate network addresses to the controller. By communicating with this machine, control agents could register with the DHCP server, join the agent platform and register their GPS coordinates; this allowed them to search for their nearest neighbours by geographical distance. A facilitator agent was written to provide this functionality.

5.2.5.4.2 Embedded controller platform

Linux was used as the OS for the beagleboards. A 2.6-family kernel was customised and compiled, and the remaining system based off the Ångström distribution[3] to allow a flexible platform for a wide range of software and hardware compatibility. Software written for the target could be easily ported to other hardware.

5.2.5.4.3 GPS

The gpsd daemon[4] and gps4java[5] were used to translate from the GPS hardware output to a format the agents could read.

5.2.5.4.4 ADC system

A driver was written in C++ for the Linux platform to allow the agent to interface with the ADC. The driver activates a clock on the beagleboard, and routes the signal to the begin-conversion pin on the ADC. An interrupt-driven routine copies the conversion result to a buffer in beagleboard memory, guaranteeing the required sample rate. Because the beagleboard processor executes the steps to copy the conversion result into memory, the sample rate was set to 5kHz as a compromise between agent code time and sampling system time. The buffer could then be accessed through the driver.

Two versions were implemented: one to allow the instantaneous voltage to measured, and the other to return a buffer for RMS voltage calculation. The former was used to generate short-time datasets to compute frequency components via Fast Fourier Transform (FFT) and cross-correlation in section 4.3.3, and then the latter for logging measurements for longer durations more akin to a high-resolution smart meter for 4.3.4.

---

[3] http://www.angstrom-distribution.org

[4] http://www.catb.org/gpsd

[5] http://taimos.github.io/GPSd4Java

5.3 Results

5.3.1 Introduction

The three novel multi-agent test systems were each created for a specific purpose: firstly, a proof-of-concept system using multiple desktop PC nodes; secondly, a supercomputer system with a large number of processor cores; and thirdly, a novel test rig comprising networked embedded controller target boards and electrical network simulation hardware.

The proof-of-concept system was used to trial a multi-agent control system that used a simple congregation of load controller agents and Intention Transfer to solve an undervoltage problem.

The cluster computer was used for a facility to scale up to test more agents. It was trialled with a combination of load and generator controllers to solve an undervoltage problem.

The real-time, hardware-in-loop test rig was designed to use the novel location derivation and multi-agent techniques in voltage control, implemented on physical embedded controllers on a real-time simulated arbitrary dynamic electrical distribution system.

All three of these test systems used the same multi-agent framework, JADE, as described in Section 5.2.2.7.

5.3.2 JADE Framework

The JADE-provided Agent Management System was run in fault-tolerant mode: in the event that the main instance of the agent platform stopped responding (whether through software fault or network outage), one of several redundant platform instances took over. JADE has a facility to provide this fault-tolerance through the use of federated, communicating facilitator nodes. In this way, the platform itself forms a distributed system, since these nodes could be hosted in multiple locations by several pairties. The platform did not fail in any of the testing for the duration of the project. As in [1], a straightforward N-1 test of resilience was performed: the system demonstrated successful recovery when one of the management system nodes was deliberately terminated.

5.3.3 Intention Transfer

5.3.3.1 Test outline

Initial testing was performed using the proof-of-concept system to demonstrate Intention Transfer with a voltage drop problem[164].

In line with the technical challenges of integrating DER outlined in Section 2.2.1, control agents were developed to tackle deviations from voltage limits, caused by too much consumption (undervoltage) or too much generation (overvoltage), with limits set by the UK regulations of acceptable fluctuation from nominal voltage. In the UK, voltage must not drop by more than 6% or rise by more than 10% of the nominal supply voltage[39].

Control agents were developed for a multi-agent system that used a simple congregation. These were implemented on the proof-of-concept test system.

5.3.3.2 Future Smart Grid Scenario

Whilst there is currently no relevant regulatory regime to enable household-scale DSM across the UK context, it is assumed that in a future smart grid there is a willingness on the part of householders to allow some flexibility in their energy consumption, provided there are appropriate control tools. As outlined in Section 2.4, a future scenario of a particpatory, collaborative smart grid is used, such that the agents do not require direct market mechanisms to act. Load, storage and generation can all be resources in this future, so load and control agents were programmed as flexible entities with an interest in their surrounding network.

This hypothetical future creates an application a of location-aware agent-based control system where a key goal is voltage support. Other work considers regulations and acceptability of control decisions to various stakeholders in LVDNs, so for the case of developing location in control, the DSM schemes applied were kept simple.

5.3.3.3 Control Agents

Initial testing used load agents with a basic functionality described in Section 5.2.2.8. The DSM scheme for consumer use did not implement a real-world consumer load profile, but rather the agents were given a simple schedule for demand or generation, which could either be high- or low-priority. These schedules were intended to illustrate demand and production in a repeatable way and to allow for forcing of contingencies, rather than using load profiles or probabilistic models. For proof-of-concept functionality, only the load agents were used in testing the undervoltage case. The control agents used absolute voltage magnitudes to determine their action priority.

The simulation was run for a set of use conditions which would ordinarily result in a combination of 1–4 agents experiencing voltage excursions, occurring sequentially, and separated by a short pause. 50% of the customers were randomly assigned to be able to provide assistance through DSM in any given combination of loads.

5.3.3.4 Simulation Agent

The Simulation Agent hosted the electrical network model and supplied the Load Agents with voltage information via standard JADE messaging. The update process typically took 84ms, although it ranged from 40ms to 172ms, depending upon factors including number of agents, message queues and presence of a sniffer used to examine network traffic. To contextualise this, 84ms is approximately 4 cycles at

50Hz.

5.3.3.5 Intention Transfer Dialogue Trial

An extract of communication recorded by JADE between the agents is shown reconstructed in Figure 5.11.

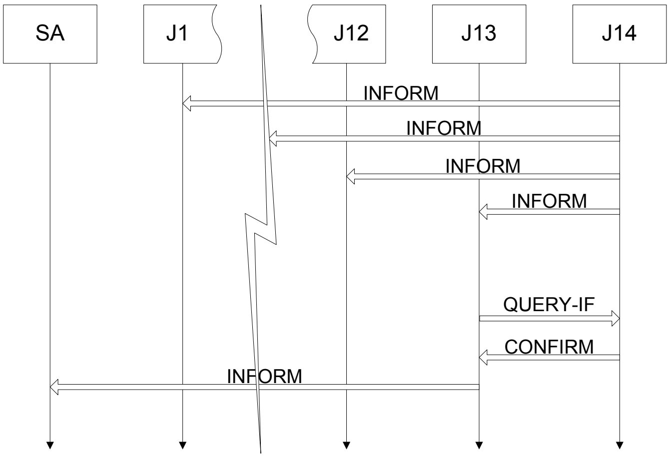

Figure 5.11: Communications extract

In this dialogue (simplified here for clarity), the Load Agent J14 (located at Load14, Figure 3.15) broadcasts its intention with an INFORM message:

(implies
(B J14@pwrsvr (LocalVoltage J14 0.939))
(I J14 (B J14 lvgtll)))

The “implies” operator is used to connect the belief that the local voltage is under 0.94 p.u. with the intention to bring about the truth of the proposition *lvgtll* (an abbreviation of “Local Voltage Greater Than Lower Limit”). This helps other agents to determine priority of this control effort. J13 (correspondingly located at Load13) accepts and adopts this intention. It waits a few milliseconds before responding, and then sends J14 the following QUERY-IF:

(B J14 lvgtll))

Since the voltage at J14 is still below the statutory limit, the *lvgtll* proposition is not true. J14 returns a DISCONFIRM message with the same content. J13 reasons that since it still believes the proposition is false, it takes action to remedy the voltage excursion, and sends a message to the Simulation Agent that it has switched off. Agents J10 and J11 also assist before the voltage at J14 can be updated to be above 0.94 p.u. Agents which subsequently query the truth of *lvgtll* are sent a CONFIRM; because the end condition of intention has been fulfilled, the intention is dropped and no further action is taken. A few minutes after this event has occurred, J13 refreshes its belief base (not shown) by performing the same QUERY-IF; J14 confirms the proposition has become true, and J13 returns to normal operation.

In a single operation to resolve a voltage event, 25 messages were generated between the load agents within 1500ms. ACL and FIPA-SL running on the JADE messaging system has a fairly low data density; content ranged from 82B –3kB, with approximately 500B of additional metadata per message. This is indicative of a minimum required connection speed of 120kb/s. Whilst this was by no means a problem for either test system, consideration needs to be given to network bandwidth and latency in future smart grids. Implementation of multicasting for the voltage contingency announcement would reduce bandwidth requirements, and is now supported in JADE[176]. Whilst message size could be reduced, data structures are constrained by the FIPA standards which trade large message sizes for flexibility, clarity and openness.

One problem that arose during development was that all low-priority loads would switch off simultaneously in order to assist network operation. This was mitigated by the use of a nearest-neighbour analysis: agents calculated a delay before action, arbitrarily based on local proximity to a belief of the remote voltage. Agents linked closer proximity on the network with better ability to assist. This provided a partial solution.

The number of DSM switching events was dependent upon the delay, a function of ΔV^2 , and agents in close proximity opted-in too rapidly for the SA to respond. By tuning this function, fewer switching events occurred. Voltage proximity is only a simplistic method to assess ability to assist, due to the complexities of LVDNs incorporating DER.

5.3.4 Combined Load and Generation

As a baseline, the system was trialled where no agents were able to assist. This yielded a baseline result where generators disconnected themselves under the G83/1 regulations in the overvoltage case, and no loads were shed in the undervoltage case.

The simulation was run for a set of use conditions which would ordinarily result in a combination of 1–4 agents experiencing voltage excursions, occurring sequentially, and separated by a short pause. 50% of the agents were randomly assigned to be able to provide assistance through DSM in any given combination of loads. As designed, the combination of high demand and generation with an extended feeder resulted in the network simulation database supplying under- and overvoltages.

Figure 5.12 shows the change in voltage at the end of the feeder during the first voltage limits excursion for one combination of DSM availability for the load-only case. In this typical case, the voltage at the end of the feeder (at Load14, Figure 3.15) dropped to 0.939 p.u. This lasted 360ms. The agents at Load10, Load11 and Load13 assisted and the voltage was raised in three steps to 0.972 p.u. The excursion in which the voltage at Load12, Load13 and Load14 was below limits lasted 220ms where the agents at Load9, Load10 and Load11 assisted to raise the voltage to 0.963 p.u. Agents refreshed

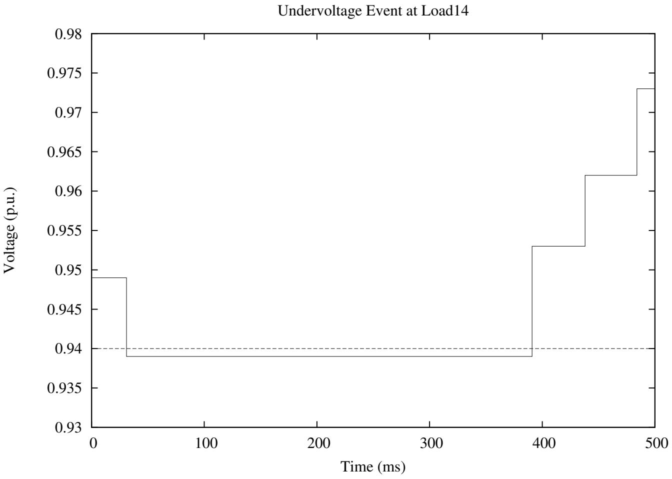

Figure 5.12: Voltage Excursion at Load14

their belief base every 5 minutes as a means of dropping an intention which is no longer relevant; this parameter would require tuning and would reflect the characteristics of physical plant. This resulted in two new network behaviours which were dependent upon the random DSM assignments. In some trials, subsequent undervoltages were avoided entirely because agents continued to assist for a short time period after the contingency was solved. Some trials exhibited an additional (but less severe) undervoltage when an assisting agent returned to its original state. Further tests with reduced load diversity showed an increased response time as more remote agents assisted. As an indication of the ability of the system to control DG, the G83/1 threshold of 1500ms was never exceeded if sufficient cooperation was available. As diversity reduced and fewer loads could assist, this situation became more probable. A maximum of 30% of customers with the ability to switch off was sufficient to guarantee voltage restoration; normally, fewer were required. The scheme operated in a similar fashion when GAs were included, with the system reducing the number of generator disconnections.

Without the control approach, the 4 LV network nodes at the remote end of the feeder (representing 38 customers) would have experienced breaches of the statutory voltage limits under the test scheme. With the distributed multi-agent control system, some breaches were avoided entirely, no breach exceeded 1500ms, and was typically resolved with 3 DSM switchoff events.

Over the course of development a significant issue was discovered in the JADE platform: in the default implementation, the execution of some agent behaviour could be prevented by the existence of other recurring behaviours. The Scheduler class in the JADE libraries determines execution of behaviours (plans, intentions and any other routines). However, it does not guarantee execution. If the final behaviour in the execution list (`readyBehaviours`) adds a new behaviour, that new behaviour is *not* executed in the current cycle around the list. A patch was written for the the JADE 3.7 libraries to fix two serious bugs in the behaviour scheduler. Routines to demonstrate the errors and relevant patches were circulated to the open-source community and are included in Appendices B.1 to B.3.

5.3.5 Real-Time, Hardware-in-Loop Testing

5.3.5.1 System Outline

The overall setup of the real-time, hardware-in-loop Embedded Controller test platform were previously outlined in Section 5.2.5, and shown diagrammatically as per Figure 5.7.

The UK Generic LVDN-based test network, implemented in RSCAD in Section 3.4.3 was integrated into the RTDS, and is shown simplified in Figure 5.13(a). The embedded controllers were connected to the simulated electrical network at the points indicated in Figure 5.13(b). The physical analogue outputs of the RTDS supplied the voltage waveforms of the simulated electrical network at the load points indicated to agents (labelled J1-J14) hosted on the embedded controller test rig.

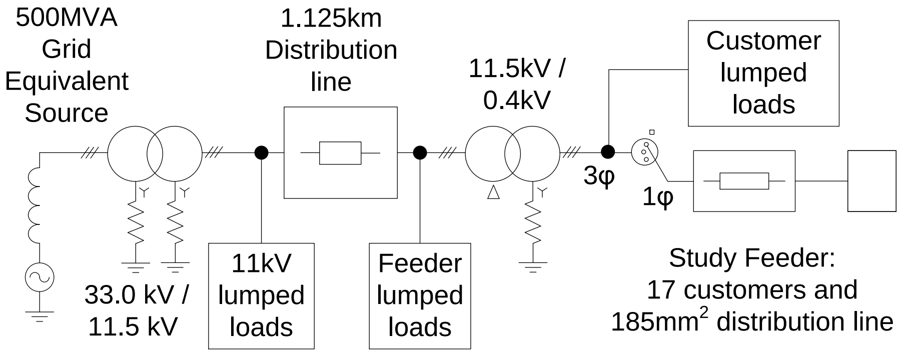

(a) Real-time testing: network overview

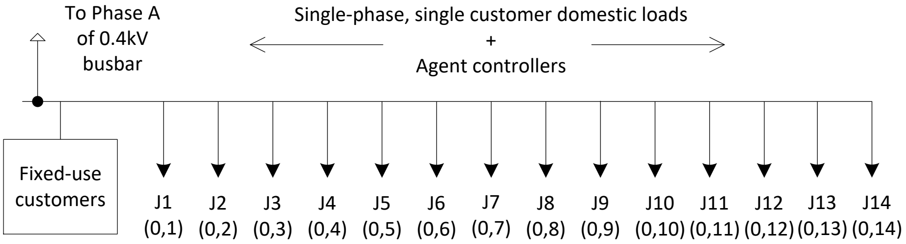

(b) Real-time testing: study feeder

Figure 5.13: Real-time network overview and agent controller connections
The UK Generic LVDN, implemented in RSCAD and shown simplified in a) was integrated into the RTDS. Subfigure b) shows the expanded view of the study feeder remote end (i.e. domestic connections).

The individual sub-systems were tested before a final trial of the complete location-aware system for a voltage control scenario.

5.3.5.2 GPS System

Initial testing of the GPS subsystem indicated successful location acquisition, as shown in Figure 5.14. In principle, the embedded controllers could be removed from the test rig and placed anywhere on the electricity network, given an appropriate network connection. However, in the lab environment, the proximity of the devices to each other meant that the margin of error in GPS measurement overlapped. The receivers were placed sequentially in a line approximately North-South. Figure 5.15 shows the reported arrangement of the controllers, in which the order was not maintained. Subsequent to confirming that the Beagleboards could collect and use real GPS data, the capability to supply arbitrary coordinates to the controllers was used to specify the relevant test geometry.

```
Load Agent G1
Current location 54.7671439 -1.569591591
Initialised array
finished all setup
```

Figure 5.14: Debug output from GPS-enabled agent

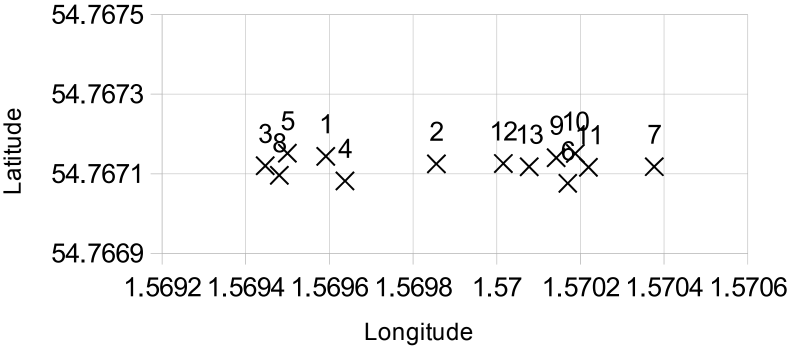

Figure 5.15: GPS coordinate overlaps of controllers in lab
The figure shows the reported locations of controllers 1-13. The lower numbers are clustered to the left, higher to the right, but are not in the correct order due to the margin of error of the GPS receiver.

5.3.5.3 Deriving Location

The full graph representation of the agents on the UK Generic LVDN test feeder (as implemented in the RSCAD model) is shown in Figure 5.16. For consistency, the agents used supplied coordinates after the successful GPS testing in Section 5.3.5.2. The edge weightings are omitted for clarity.

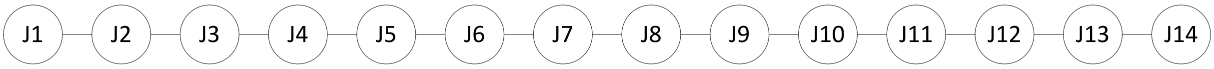

Figure 5.16: "Complete information" graph
The full graph of the control agents connected at the feeder remote end.

The GPS-enabled agents registered their location and constructed the connection matrix as an initial action before beginning to monitor local voltage conditions. Figure 5.17 shows Agent J14 starting up, identifying its nearest neighbours ('J13' and 'J12') and proceeding to monitor voltage. Some of the output has been omitted for clarity.

As in Section 4.2.3, each of the agents undergoes the same process to create a sub-graph for its own internal representation. These use the path structure for voltage control on this single-phase system; the current control tree graph was not calculated in this step. The agents store this as an adjacency matrix

```
Load Agent J14
finished all setup
Register
14.0;0.0  [longitude/latitude]
mystring: (Nearest_X 10 (agent-identifier :name J14@pwrsvr...) (GPSLocation :Longitude
"0.0" :Latitude "14.0") (sequence))
Sent
Incoming Location Response: ((= (sequence (Neighbour :NeighbourAID (agent-identifier
:name J14@pwrsvr :NeighbourLoc (GPSLocation :Longitude "0.0" :Latitude "14.0"))
(Neighbour :NeighbourAID (agent-identifier :name J13@pwrsvr :NeighbourLoc (GPSLocation
:Longitude "0.0" :Latitude "13.0")) (Neighbour :NeighbourAID (agent-identifier :name
J12@pwrsvr :NeighbourLoc (GPSLocation :Longitude "0.0" :Latitude "12.0" ... ))
FirstMeasure 111.19492664455872
...
i=0
j=0
nnm[i][j] = Infinity
j=1
nnm[i][j] = 111.19492664455872
i=1
j=0
nnm[i][j] = 111.19492664455872
j=1
nnm[i][j] = Infinity
...
volts: 1.024702
...
```

$$\begin{array}{c|ccc|} & J_{14} & J_{13} & J_{12} \\ J_{14} & \infty & 111.2 & \infty \\ J_{13} & 111.2 & \infty & 111.2 \\ J_{12} & \infty & 111.2 & \infty \end{array}$$

Figure 5.17: GPS-enabled agent

GPS-enabled agent beginning to construct local graph. The figure shows the initial debug output, showing registration of its coordinates and communication; it constructs the connection matrix for voltage control (at the bottom) and proceeds to voltage-monitoring.

as shown in Figure 5.17. None of the agents store a complete representation. Three of the sub-graphs calculated in the trial are shown in Figures 5.18(a)-5.18(c).

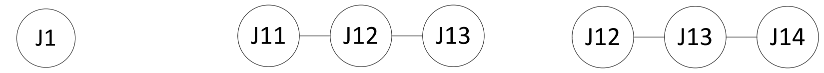

(a) Sub-graph: start node J1

(b) Sub-graph: middle node

(c) Sub-graph: end node J14

Figure 5.18: Final graphs
The graph of the entire feeder (see Figure 5.16) is represented via a number of sub-graphs by the agents. Subfigures a-c show a subset of these for the near-busbar, middle and remote end cases. Edge weights are omitted for clarity.

Agent J1 entered the system first of all. With no other agents to map, and no whisper protocol yet implemented to perform graph updates (as per Section 4.14), it is the only member of its graph. Figure 5.18(b) shows a sub-graph from the middle of the feeder. Figure 5.18(c) shows the graph for the agent J14 at the remote end: when the agent calculates the distance between J12-J14, the result (222.4) is longer than the other two, and consequently no edge is inserted. The two nearest neighbours thus form a path to the left. These demonstrate the successful application of the geographical algorithm developed in Section 4.2.1 on the embedded hardware platform.

5.3.5.4 Voltage Acquisition

The break-out box, voltage acquisition daugterboard, and driver software on the embedded controllers were successfully integrated into the agent design. Figure 5.19 shows the per-unit voltage output of the the UK Generic LVDN RTDS simulation during normal operation (i.e. no under/overvoltage events present and no control scheme applied) over 2 minutes, as measured by a GPS-enabled control agent with integrated voltage acquisition hardware and bespoke driver. System noise in the region of 0.1% is visible.

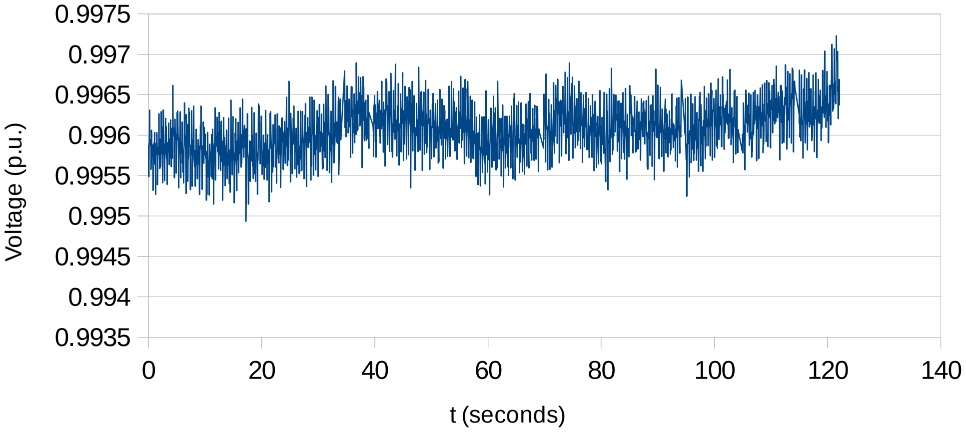

Figure 5.19: 2-minute voltage measurement from embedded controller
A 2-minute voltage profile of RTDS-simulated network voltage from the UK Generic LVDN, as measured at one of the embedded controllers.

With the full integrated electrical network simulation and hardware platform, the MAS algorithms for location derivation and the control task could be implemented.

5.3.5.5 Final Location-Aware, MAS Hardware-in-Loop Voltage Control Demonstration

Up to this point, then, the following components of the MAS control system have been independently implemented:

- Electrical network modelling (in Section 3.4)
- MAS platform
- Intention Transfer
- Voltage control using loads and generators
- Hardware platform for physical agents
- Location derivation

In this section, these components were combined in order to apply location-aware agents for voltage control. A successful demonstration would show the complete process from agent registration and location derivation, to voltage monitoring, to control action, and through to resolution of the network event.

The demand profile of the customers on the feeder was designed in order to create heavy loading (as described in Section 3.3.1) whilst remaining realistic. This loading would create a voltage dip outside of regulatory limits; control agents would demonstrate restoration to within limits compared to the no-control case.

With the overall setup as described in Section 5.3.5.1, the set of load control agents J1-J14 were initiated and the agent actions logged.

To examine the behaviour of the system, the interactions between the agent J14 and J13 and the platform are shown here: the full process diagram is shown in 5.20. The first actions after initiation are shown on the left: agent registration and graph building. The output graph is shown in Figure 5.21.

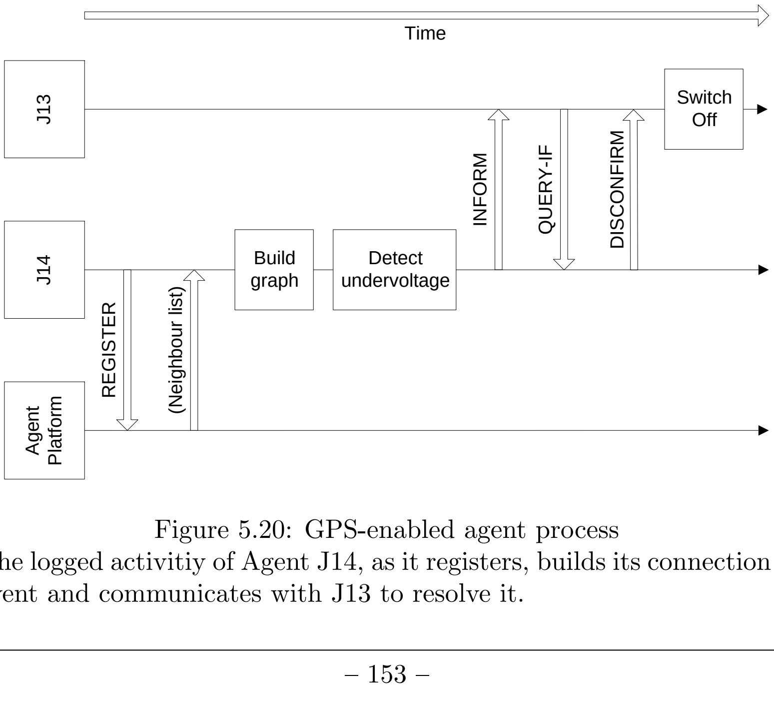

Figure 5.20: GPS-enabled agent process
This figure shows the logged activitiy of Agent J14, as it registers, builds its connection graph, experiences an undervoltage event and communicates with J13 to resolve it.

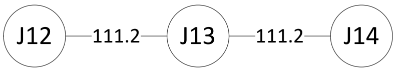

Figure 5.21: End node internal graph
Agent J14 builds its connection matrix upon startup, represented here as a weighted undirected graph.

The agents on the feeder measured the real-time analogue voltage signals. The voltages recorded by agent J14 are shown in 5.22. At a short time after the graph creation process, the model exceeded operational limits: the statutory lower voltage limit (-6% from nominal, i.e. 0.94p.u.) is shown with a loose-dashed line. For agent J14, the voltage excursion below the limit begins at t=517ms.

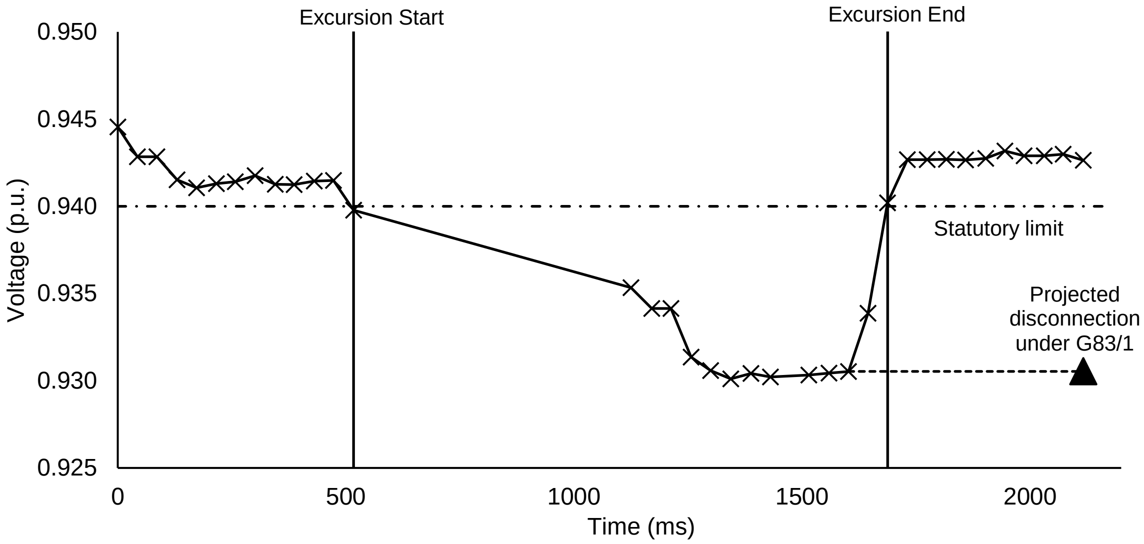

Figure 5.22: GPS-enabled agent voltage excursion measurements
A 2200ms per-unit voltage capture by agent J14. The statutory lower voltage limit (0.94p.u.) is shown with a loose-dashed line. The voltage experienced by the agent dips below acceptable limits between the excursion start and end vertical bars at t=517ms and t=1688ms. The voltage for no control action is projected with a dashed line to t=2117ms.

Shortly after t=517ms, agent J14 contacted J13, the agent directly adjacent on the feeder, to request assistance, as per the dialog visible in Figure 5.20. J13 was well-placed to assist, since its impact on voltage is higher than those closer to the busbar. It determined to assist; it reconfirmed the necessity of action (to reduce excessive control actions from voltage transients), then delayed its EV charging for a time, reducing network load. This process occurred for 6 agents requesting assistance from adjacent agents almost simultaneously (not shown), and resulted in 4 controller switching operations. There was less communication in this case compared to the proof-of-concept system: the agent only messaged with its connected agents, rather than the entire pool.

At t=1688ms, the voltage was restored to within limits: a duration of 1150ms from the original excursion. This restoration time falls within the range of 220-1500ms of the combined system shown in

Section 5.3.4, and is within the 1600ms requirement of the regulation G83/1[9]. Without the controller switching events, the system would have continued in the undervoltage condition: the lower dashed line on the right of Figure 5.22 shows system voltage projected for no controller action until 1600ms after the initial undervoltage, where disconnection of SSEG becomes mandatory,

In the solution by Trichakis [1] (as discussed in Section 2.3.2.3), where the full network geometry is stored in a preconfigured centralised database, a voltage excursion developed between t=0.5s – t=1.0s, and was resolved by t=1.015s; (i.e. in the range 15 – 515ms).

Consequently, these results show a successful demonstration of the location-aware decentralised MAS avoiding disconnection under G83/1, and with the new solution providing comparable response times and with increased decentralisation over the Trichakis implementation through a dynamic network topology creation process.

5.3.6 Discussion

In terms of the control system itself, the gaps in the literature highlighted in the earlier chapters can be revisited in this context.

Over the course of the work, strengths, limitations and weaknesses were found for the system and each of the test facilities. One key issue is around fairness and switching operations, particularly since the DSM scheme is not sophisticated.

Each iteration of testbed contributed to the development of the control system, and the testing facility itself should also be evaluated.

5.3.6.1 Research Challenges

The system was developed to take account of each of the following strands of the research challenge outlined in Section 2.4:

- Decentralised control

  The agents form congregations and take individual control actions. No centralised controller is present. The rudimentary control techniques proved useful for maintaining voltage within limits.

- Multi-goal

  The multi-stakeholder context included voltage as a control goal without sacrificing customer equipment use priorities. Whilst the location could be useful in tackling voltage unbalance, as of writing this functionality was not implemented.

- Market-independent

  The future scenario included a regulatory setup in which the agents were collaborative. Markets were not used. However, some flexibility is lost since agents cannot be offered compensation to incentivise action

- Towards "zero" configuration

  The only configuration required by the control agents was their use profile, which would either be set up in the ordinary course of equipment use (such as setting heating times for a HP) or would emerge from their function such as PV output. Upon startup, the agents automatically derived location and used that information in a control system without additional user intervention.

- Scalable

  Local congregations of agents with connection graphs can be used to create groups of appropriate size, but the impacts of larger numbers of agents were not quantified. The supercomputer implementation was developed to allow testing at scale to be performed for future systems.

- Resilient

  The agent platform ran in a fault-tolerant mode. The opt-in basis for agent action operated over the top of an existing disconnection regulation regime, allowing the system to fall back onto predefined safety standards in the case of outage or component failure.

- Standards-compliant

  Relevant multi-agent standards were applied throughout development and all the system components were open-source, ensuring that interoperability (including third-party developments) can be guaranteed into the future.

5.3.6.2 Fairness and Equity

An agent delayed assistance by waiting for a time specified by a function of the difference between its own voltage and that of an agent announcing a problem. This worked for the simple linear feeder in this simulation, but this technique is flawed for real-world cases. Firstly, it takes no account of three-phase operation, which could contribute to unbalance problems or just result in ineffective DSM switching events. These should be avoided, since any switching operation can be seen as a reduction in supply quality. Secondly, it does not adjust well to branching networks, where a node may be close to another in voltage but not in network location, reducing the effectiveness of its assistance. Agents which are aware of network location are able to make better decisions; whilst the voltage-difference method is a good start, better information could improve network operation and improve agent plug-and-play capability. Finally, it means that agents at the remote end of the feeder assist disproportionately more: whilst this may be a more technically effective solution, it is not an equitable one.

Although the delay function was effective in producing an improved voltage profile, it requires tuning.

For an agent close to a voltage problem, the differences in delay to assistance is not significant compared to the inter-agent message latency; as a result, several agents may assist in cases where one would be sufficient. This tuning is case-specific, however. It would require extensive testing to attempt to specify a generally-optimised catch-all-cases function.

5.3.6.3 Test Facility

Initially, ubiquitous control was considered to be unrealistic on grounds of cost; the use of inexpensive, general-purpose hardware now gives scope for a variety of agent applications where direct control would be prohibitively expensive but small-scale actions can be aggregated into providing real value.

The specific application developed was for domestic control on LVDNs, but the embedded controller platform has wider relevance: the facility can host control agents written for any item of equipment represented in the RTDS, whether at transmission or distribution level, and read from and feed back into the RTDS for closed-loop operation. Separately from this work, the RTDS has been linked to a PAS 2000 4-quadrant power amplifier, allowing items of equipment (at the 400V level) to interact with a physically simulated network, such that the multi-agent system can incorporate both real and simulated equipment in closed-loop operation, increasing the value of the facility yet further.

Consequently, the test rigs are a new resource for testing communicating energy resource controllers in the MAS paradigm. The implementation of a combination of real-time simulation with hardware-in-the-loop controllers and a networked MAS is a significant development as a smart grid test resource.

5.4 Conclusions

In the first part of this chapter, the multi-agent system paradigm was examined, in particular when considering the power systems domain. After considering various societal structures, techniques in group interaction were explored. The novel technique of Intention Transfer was presented as a tool in coordinating a nonhierarchical, decentralised agent control system. It was argued that the INTEND_2 form is more flexible and appropriate for describing group action. These approaches were combined to produce the design of new control agents — software entities that can be applied across multiple control goals and which use their communicative and social abilities within a society in order to assist it in solving a control problem.

A novel MAS was developed to control SSEZ entities in a fully decentralised fashion. The design was implemented and tested on several iterations of new multi-agent test rigs.

Intention Transfer proved to be a successful original technique to facilitate opt-in-based operation[164]. The use of network inference through location identification proved to be a useful new tool in control[177].

The agents were plug-and-play capable, requiring no technical configuration; this aspect is valued as a way to lower barriers to implementation and warrants further investigation. The MAS achieved an improvement in network operation from the existing G83/1 regulations in terms of customer voltage levels and reduced contingency durations, and improved the availability of SSEG through the reduction of mandatory disconnections. This facilitates integration of controllable loads and high SSEG penetrations; it adds value for DG owners and may delay network reinforcement for DNOs.

The final embedded controller platform was demonstrated to be an effective new test facility for multi-agent systems and will be a useful resource in the field of smart grid develoment.

Chapter 6

Conclusion

This thesis began with an introduction to the challenges of integrating DER and the motivation and context for research in the area. The state of the art and its shortcomings were examined by reviewing the literature in Chapter 2, and this was used to devise requirements for the solution to the problem in Section 2.4. In Chapter 3, theory was outlined to explain electrical phenomena, indicating the significance of network topology. The link between topology and ideas of location were developed in Chapter 4 as well as techniques for deriving topology. Decentralised control was examined in Chapter 5 through the use of Multi-Agent System techniques. Test facilities were created to study the results of the control system devised to solve the Research Problem (in Section 2.4) in the context of the importance of location in network operation. The development of these facilities was also described in Chapter 5, and the implementation of location-aware agents was documented and evaluated.

In this chapter, the progress made in each of the preceding chapters is reviewed. It includes discussion of the wider implementation of control systems using location and their use in the move from legacy, passive LVDNs towards a truly smart grid. Finally, this thesis concludes with a summary of its outcomes.

6.1 Research Context

The integration of DERs is discussed in Chapter 1 in terms of the increase of renewable generation, electric vehicles and heat pumps at the domestic level. It describes the legacy network of the UK and its regulatory framework in which which these new technologies must operate. The key technical constraints are explained in Chapter 2. Of these, voltage rise and drop and unbalance are specifically highlighted, for two reasons: they are the top two on the list of electrical barriers to higher penetrations of DG, and because the G83/1 regulations specify that DG must disconnect in undervoltage conditions where they could instead be used to support the network. A solution to voltage rise and drop would facilitate wider SSEG uptake and can add value across stakeholders: economically, a producer would continue to supply

— and therefore sell — electricity, and a network operator would be able to defer network reinforcement costs; environmental benefits occur from maximising the potential output of low-carbon generation. The voltage unbalance issue does not have the same disconnection requirements for DG owners under G83/1, but are a wider issue of power quality that could result in customer equipment damage if left unchecked. The voltage rise and drop issues are linked to voltage unbalance electrically as explained in Chapter 3.

The context and review combine to givea summary statement of the problem:

> “...there remains no system that simultaneously addresses the major technical problems of DG integration into a market-free, demand-driven legacy network context whilst also capturing the full benefits of a multi-agent solution. Consequently, a scalable, network-independent, decentralised control system for operation of an evolving network with multiple system constraints remains an open research problem. ”

The set of criteria for evaluating the solution have been identified as:

- Decentralised control
- Multi-goal — operating across multiple technical constraints
- Market-independent
- Towards “zero” configuration — for “plug-and-play” operation
- Scalable
- Resilient
- Standards-compliant

The scoping of this work was co-developed with industry partner E.ON through the EPSRC CASE (Cooperative Awards in Science and Technology) mechanism. This approach highlighted some of these regulatory issues and specific technical challenges. It also helped guide some implementation decisions. Some of these were related to system structure (for example, the examination of the multi-stakeholder context and the decision to investigate a non-market-based approach), whereas others were related to low-level implementation (such as the use of a passive monitoring approach to detect electrical connections, rather than an active approach such as a controller placing a signal onto the electrical network for others to detect). These considerations evolving from industry feedback were useful to provide some focus areas over and above those provided by the gaps identified in the literature.

6.2 Location in LVDNs

The theory underpinning the voltage rise/voltage drop phenomenon and voltage unbalance was introduced in Chapter 3. The relationship between network impedance and voltage changes caused by DER connection was explained using this theory. Since the impedance increases with length of distribution line, the distance between equipment is related to the voltage drop between them. The effects of phase unbalance in three-phase networks were also explored. This analysis demonstrated that distribution network topology is important for understanding and solving the constraints on DER integration.

The UK Generic Network in the existing literature was examined as a benchmark electrical network and was used to create new models of legacy distribution networks incorporating widespread DER, which were implemented as research tools for this work using several software packages. The resulting models demonstrated the problems outlined in Chapter 2 in line with the theory.

IPSA+ modelling is valid for static, balanced operation. The new model was used for aggregated groups of customers and was scripted to provide an exhaustive database of operating points which could be used as in software simulation.

In part due to the limitations of balanced analysis, the model was also implemented in RSCAD — this implementation included modelling down to individual customers. This model could be translated to the real-time digital simulator for use in the test rig described in Chapter 5.

The combination of theory and models provided demonstration of technical issues related to location and tools for analysing and testing solutions.

6.3 Location in Control

An order of priority for barriers to DER is described in the literature review. From a control perspective, this starts with voltage drop and rise and phase unbalance. There are constraints on these phenomena imposed by regulation. Control of the network was considered in the context of maintaining operation within these regulations in non-fault conditions.

The electrical properties are related to the control aims. Ideas of location are developed in Chapter 4 by drawing parallels between electrical properties — especially topology — and geographical position, such as the analogy between impedance and distance.

Initial attempts to identify connection types by using instantaneous voltage measurements showed limited success, so electrical connection was determined by measuring voltage over a 24-hour period and comparing it through cross-correlation; test data were collected from the real distribution network in multiple locations at 10s resolution through household energy monitors to demonstrate applicability in smart meter-like hardware and were used to establish criteria in determining connection type.

Connection between entities was examined through graph representation of electrical networks. Finally, a method was devised to create a graph of proximate nodes by finding geographical location and confirming electrical connection; the nodes are connected by edges weighted with a (distance, connection type, correlation peak magnitude) 3-tuple suitable for voltage and current control purposes.

6.4 Multi-Agent Systems

The multi-agent system paradigm was explored conceptually to establish its applicability for the research problem. It was argued in Chapter 5 that the MAS paradigm provides a framework for decentralisation, dependent upon societal structure, and highlighted the use of congregations in this application. Cooperation and collaboration were examined, including game-theoretic aspects, although the use of markets was excluded.

The agent INTEND attitude can be expressed in two forms: commitment to a specific action, or commitment to a state having been achieved. The argument was made that the latter form is more flexible and appropriate. Formal descriptions of INTEND_2 were used to develop a novel opt-in system for voltage control. In particular, the agents in this system use their communicative and social abilities within a society in order to assist in solving a control problem.

6.5 Application

The prospect of ubiquitous computing makes the use of agents in diverse electrical equipment ever more likely. With the advent of low-cost small-form-factor general-purpose computers, agents may be more easily deployed.

The multi-agent system was combined with location techniques to create a solution to the research problem. Three test facilities were created in the development of the system: a proof-of-concept computer network demonstrating Intention Transfer, a cluster supercomputer using load and generator agents on a larger electrical network, and an embedded controller platform featuring real-time electrical network simulation interfaced with controller hardware with real signal I/O and including GPS capability.

In the course of test facility development, patches were developed for the industry-standard JADE platform. Hardware modules were designed and manufactured to interface the embedded controllers with the RTDS I/O signals. The testbed developed for the MAS trials provided a unique and vital resource.

The system demonstrated successful recovery from voltage excursions using an opt-in demand-side management approach. The system used a form of location in agent configuration in each incarnation, and ultimately used geographical coordinates to construct a connection graph which was then used to determine messaging and control action to resolve a voltage excursion within the G83/1 time limits, improving the economic and environmental benefits of the connected microgeneration and delaying the need for network reinforcement.

6.6 Future Work

A method for confirming the presence of an electrical connection is set out in 4.3. Progress was made towards implementing this method using the embedded controller test rig, but has not yet been fully realised. Consequently, the immediate next step for this work would be a demonstration of the integration of both the GPS and electrically-derived location information in the same agent, incorporating the overlay of connection information for combined voltage/current control. A structured investigation of the effects of network geometry in test scenarios would then give more confidence in the specific algorithm for deriving network graphs.

The embedded controller platform currently exists on an ICT network that is isolated from the Hamilton Cluster. Connecting the two networks would provide significant additional value: as well as agents hosted on the RTDS-connected embedded controllers, an agent society of thousands could feasibly be tested — even connecting agent-based human behavioural models to equipment controller models to electrical network models end-to-end. Some electrical signals could be supplied to agents on the cluster (through an abstraction layer similar to the Simulation Agent used in this work, rather than hardware-in-the-loop). This would create an at-scale testbed across the full spectrum of the power system.

Machine learning is briefly discussed in Section 4.1.5.2. Use of confidence factors in different location measures, as well as evaluating the effectiveness of previous control interventions, can allow an agent to improve its graph over time and future decisions. This rationalising process creates intelligence in control — a further project implementing Artificial Intelligence (AI) techniques is therefore the next step to a network of agents incorporating automatic configuration, location-based, adaptive and social decision-making in a truly smart grid.

6.7 Wider implications of research

This study was focused on the area of small-scale DER at household voltages, but the relevance to industry of the central themes of this work go wider than this specific application. At any level within electrical networks, there are control issues where network topology is significant; the wider issue addressed by this work is that that deriving location and using it in a multi-agent system in the power domain is feasible and useful.

The multi-agent system testing explored in this work was not feasible before the construction of the embedded controller test rig. Again, such a system has wider application: the RTDS allows for simulation of arbitrary electrical networks (including at the transmission level as well as for LVDNs), the platform can be configured for any geographical arrangement of equipment, with an industry-standard MAS platform and general-purpose, low-cost controller hardware. This facility represents considerable progress towards a fuller exploitation of agent capabilities throughout electrical power systems.

6.8 Summary

This work argues that location is an important aspect for the control and integration of distributed energy resources. It identifies various sources of location data and develops new geographical and electrical techniques for deriving network topology using GPS and 24-hour voltage logs. It examines the multi-agent system paradigm and societal structures as an approach to a multi-stakeholder domain, and concludes that the decentralised context is best served by a nonhierarchical approach, in which this non-market-based implementation indicates the congregation structure. Through formal description of the agent attitude INTEND_2, the novel technique of Intention Transfer was applied to an agent congregation to provide an opt-in, collaborative system. Test facilities for multi-agent systems were developed and culminated in a new embedded controller test platform that integrated a real-time dynamic electrical network simulator to provide a full-feedback system with control hardware in-the-loop. Finally, a multi-agent control system was developed and implemented that used location data in the process of providing demand-side response to a voltage excursion, with the goals of improving power quality, reducing generator disconnections and deferring network reinforcement. The resulting communicating and self-organising energy agent community, as demonstrated on a unique hardware-in-the-loop platform, provides an application model and test facility to inspire agent-based, location-aware smart grid applications across the power systems domain.

References

[1] P. Trichakis, “Multi agent systems for the active management of electrical distribution networks,” Ph.D. dissertation, University of Durham, 2009.

[2] DTI, “Network management systems for active distribution networks - a feasibility study,” Department for Trade and Industry, Tech. Rep., 2004.

[3] D. Frame, K. Bell, and S. McArthur, “A review and synthesis of the outcomes from low carbon networks fund projects,” HubNet, White paper, 2016.

[4] HM Government, “The Electricity Act,” HMSO, Primary Legislation, 1989.

[5] ——, “The Energy Act,” HMSO, Primary Legislation, 2008.

[6] ——, “The Energy Act,” HMSO, Primary Legislation, 2013.

[7] T. T. Onifade, “Hybrid renewable energy support policy in the power sector: The contracts for difference and capacity market case study,” *Energy Policy*, vol. 95, pp. 390 – 401, 2016.

[8] ENA, *The Distribution Code*, [online] http://www.energynetworks.info/the-distribution-code, Std. 12, 2010.

[9] EA, *Recommendations For The Connection Of Small-Scale Embedded Generators (Up To 16A Per Phase) In Parallel With Public Low-Voltage Distribution Networks*, The Electricity Association Engineering Recommendation G83/1, 2003.

[10] P. C. Taylor and J. Bialek. (2010, November) Smart grids: The broken value chain. Durham Energy Institute / Department for Energy and Climate Change. [Online]. Available: www.dur.ac.uk/resources/dei/DEI_DECCWorkshopnotes_3-11-2010v4nm.pdf

[11] R. Howard and Z. Bengherbi, “Power 2.0: Building a smarter, greener, cheaper electricity system,” Policy Exchange, White paper, 2016.

[12] DECC, “Implementing electricity market reform (EMR): Finalised policy positions for implementation of EMR,” DECC, White paper, 2014.

[13] HM Government, “The feed-in tariffs (specified maximum capacity and functions) order,” HMSO, Statutory Instrument, 2010.

[14] DECC, “Solar photovoltaics deployment,” Department for Energy and Climate Change, National Statistics Publication, April 2016.

[15] OFGEM, “Feed-in tariff generation and export payment rate table,” OFGEM, White Paper, April 2016.

[16] National Grid, “Enhanced frequency response market information report,” National Grid, Market Information Report, 2016.

[17] HM Government, “New plug-in car grant levels from March 2016,” Office for Low Emission Vehicles (OLEV), Guidance document, 2016.

[18] ——. (2016, April) Plug-in car and van grants. Office for Low Emission Vehicles. [Online]. Available: https://www.gov.uk/plug-in-car-van-grants

[19] ELEXON. (2013) Load profiles and their use in electricity settlement. [online]. [Online]. Available: https://www.elexon.co.uk/wp-content/uploads/2013/11/load_profiles_v2.0_cgi.pdf

[20] G. Strbac, “Demand side management: Benefits and challenges,” *Energy Policy*, vol. 36, pp. 4419–4426, 2008.

[21] DTI, “Our energy future - creating a low carbon economy,” Department for Trade and Industry, White paper, February 2003.

[22] DECC, “Digest of UK energy statistics,” Department of Energy and Climate Change, National Statistics Publication, 2010.

[23] HM Government, “Planning our electric future: a white paper for secure, affordable and lowcarbon electricity,” Department of Energy and Climate Change, White paper, 2011.

[24] DECC, “Monthly central feed-in tariff register statistics,” Department for Energy and Climate Change, National Statistics Publication, 2014.

[25] P. Wilson, B. Betts, R. Bolton, A. Davison, P. Lucy, G. Norris, J. Sheppard, and H. Windebank, “Powering the future,” Parsons Brinckerhoff, Tech. Rep., December 2009.

[26] DECC, “The UK low carbon transition plan,” Department of Energy and Climate Change, National Strategy Document, July 2009.

[27] W. Kempton and J. Tomic, “Vehicle-to-grid power implementation: From stabilizing the grid to supporting large-scale renewable energy,” *Journal of Power Sources*, vol. 144, no. 1, pp. 280–294, 2005.

[28] DECC, “Digest of UK energy statistics,” Department of Energy and Climate Change, National Statistics Publication, July 2014.

[29] D. Hart-Davis. (2016) A note on variations in UK/GB grid electricity CO_2 intensity with time. [Online]. Available: http://www.earth.org.uk/note-on-UK-grid-CO2-intensity-variations.html

[30] G. Strbac, C. K. Gan, M. Aunedi, V. Stanojevic, P. Djapic, J. Dejvises, P. Mancarella, A. Hawkes, and D. Pudjianto, “Benefits of advanced smart metering for demand response based control of distribution networks,” Energy Networks Association, Tech. Rep., 2010.

[31] J. Hong, N. J. Kelly, I. Richardson, and M. Thomson, “Assessing heat pumps as flexible load,” *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, vol. 227, no. 1, pp. 30–42, 2012.

[32] CLNR. (2014) Customer led network revolution. Project Website. [Online]. Available: http://www.networkrevolution.co.uk

[33] OFGEM, “Demand side response discussion paper,” OFGEM, Discussion Paper 82/10, July 2010.

[34] C. Gellings, “The concept of demand-side management for electric utilities,” *Proceedings of the IEEE*, vol. 73, no. 10, pp. 1468 – 1470, 1985.

[35] R. Shaw, M. Attree, T. Jackson, and M. Kay, “The value of reducing distribution losses by domestic load-shifting: a network perspective,” *Energy Policy*, vol. 37, pp. 3159–3167, 2009.

[36] G. Strbac, N. Jenkins, and T. Green, “Future network technologies,” DTI, Tech. Rep., 2006.

[37] E. Lakervi and E. J. Holmes, *Electricity distribution network design.* Peter Peregrinus, 1989.

[38] E. Lakervi and E. J. Holmes, Eds., *Electricity Distribution Network Design*, 2nd ed. Peter Peregrinus, 2003.

[39] HM Government, “The electricity safety, quality and continuity regulations,” HMSO, Statutory Instrument, 2002.

[40] *Planning limits for voltage unbalance in the UK for 132kV and below*, Energy Networks Association Std. Engineering Recommendation P29, 1990.

[41] P. Lyons, P. C. Taylor, L. M. Cipcigan, P. Trichakis, and A. Wilson, “Small scale energy zones and the impacts of high concentrations of small scale embedded generators,” in *Universities Power Engineering Conference*, 2006.

[42] S. Slater, M. Dolman, P. C. Taylor, P. Trichakis, and J. Shine, “Strategies for the uptake of electric vehicles and associated infrastructure implications,” Element Energy Limited, Tech. Rep., October 2009.

[43] G. A. Putrus, P. Suwanapingkarl, D. Johnston, E. C. Bentley, and M. Narayana, “Impact of electric vehicles on power distribution networks,” in *5th IEEE Vehicle Power and Propulsion Conference*, September 2009.

[44] C. Whitaker, J. Newmiller, M. Ropp, and B. Norris, “Distributed photovoltaic systems design and technology requirements,” Sandia National Laboratories, Tech. Rep., 2008.

[45] P. Papadopoulos, S. Skarvelis-Kazakos, I. Grau, L. M. Cipcigan, and N. Jenkins, “Electric vehicles’ impact on British distribution networks,” *IET Electrical Systems in Transportation*, vol. 2, no. 3, pp. 91–102, September 2012.

[46] Y. Mu, J. Wu, N. Jenkins, H. Jia, and C. Wang, “A spatialtemporal model for grid impact analysis of plug-in electric vehicles,” *Applied Energy*, vol. 114, pp. 456 – 465, 2014. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S030626191300826X

[47] M. Neaimeh, R. Wardle, A. M. Jenkins, J. Yi, G. Hill, P. F. Lyons, Y. Hbner, P. T. Blythe, and P. C. Taylor, “A probabilistic approach to combining smart meter and electric vehicle charging data to investigate distribution network impacts,” *Applied Energy*, vol. 157, pp. 688 – 698, 2015.

[48] B. Blažič and I. Papič, “Techniques for DER integration and active networks management,” SOLID-DER, Tech. Rep., 2008.

[49] D. Olivares, A. Mehrizi-Sani, A. Etemadi, C. Canizares, R. Iravani, M. Kazerani, A. Hajimiragha, O. Gomis-Bellmunt, M. Saeedifard, R. Palma-Behnke, G. Jimenez-Estevez, and N. Hatziargyriou, “Trends in microgrid control,” *IEEE Transactions on Smart Grid*, vol. 5, no. 4, pp. 1905–1919, July 2014.

[50] S. Weckx, R. D’Hulst, B. Claessens, and J. Driesensam, “Multiagent charging of electric vehicles respecting distribution transformer loading and voltage limits,” *IEEE Transactions on Smart Grid*, vol. 5, no. 6, pp. 2857–2867, Nov 2014.

[51] J. Yi, P. Wang, P. C. Taylor, P. J. Davison, P. F. Lyons, D. Liang, S. Brown, and D. Roberts, "Distribution network voltage control using energy storage and demand side response," in *2012 3rd IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe)*, Oct 2012, pp. 1–8.

[52] CLASS. (2013) Customer load active system services. Project Website. [Online]. Available: http://www.enwl.co.uk/class/

[53] F. Capitanescu, L. F. Ochoa, H. Margossian, and N. D. Hatziargyriou, "Assessing the potential of network reconfiguration to improve distributed generation hosting capacity in active distribution systems," *IEEE Transactions on Power Systems*, vol. 30, no. 1, pp. 346–356, Jan 2015.

[54] R. H. Lasseter, "Microgrids [distributed energy generation]," in *IEEE Power Engineering Society Winter Meeting*, vol. 1, 2001, pp. 146–159.

[55] ——, "Microgrids," in *IEEE Power Engineering Society Winter Meeting*, vol. 1, 2002, pp. 305–308.

[56] N. Hatziargyriou, H. Asano, R. Iravani, and C. Marnay, "Microgrids: an overview of ongoing research, development and demonstration projects," *IEEE Power and Energy Magazine*, vol. 5, no. 4, pp. 78–94, 2007.

[57] P. Trichakis, P. C. Taylor, P. Lyons, and R. Hair, "Transforming low voltage networks into small-scale energy zones," in *Proceedings of the Institution of Civil Engineers: Energy*, vol. 162, February 2009, pp. 37–46.

[58] P. Lyons, "Experimental investigation and evaluation of future active distribution networks," Ph.D. dissertation, University of Durham, 2009.

[59] "Ancillary service provision from distributed generation," ILEX Energy Consulting, DTI, Tech. Rep. 04/1738, 2004.

[60] L. M. Cipcigan and P. C. Taylor, "A generic model of a virtual power station consisting of small scale energy zones," in *CIRED*, 2007.

[61] T. Sansawatt, L. Ochoa, and G. Harrison, "Integrating distributed generation using decentralised voltage regulation," in *Power and Energy Society General Meeting, 2010 IEEE*, July 2010, pp. 1–6.

[62] D. Angeli and P. A. Kountouriotis, "A stochastic approach to "dynamic-demand" refrigerator control," *IEEE Transactions on Control Systems Technology*, vol. 20, no. 3, pp. 581–592, May 2012.

[63] M. Abeysekera, J. Wu, and N. Jenkins, "Integrated energy systems: An overview of benefits, analysis methods, research gaps and opportunities," SUPERGEN HubNet, Position paper 13, 2016.

[64] M. Arnold, R. R. Negenborn, G. Andersson, and B. De Schutter, "Distributed control applied to combined electricity and natural gas infrastructures," in *Proceedings of the First International Conference on Infrastructure Systems: Building Networks for a Brighter Future (INFRA)*. IEEE, 2009.

[65] A. Fazeli, E. Christopher, C. M. Johnson, M. Gillott, and M. Sumner, "Investigating the effects of dynamic demand side management within intelligent smart energy communities of future decentralized power system," in *Innovative Smart Grid Technologies (ISGT Europe), 2011 2nd IEEE PES International Conference and Exhibition on*, Dec 2011, pp. 1–8.

[66] E. Karangelos and F. Bouffard, "Integrating demand response into agent-based models of electricity markets," in *Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference on*, Oct 2012, pp. 1308–1315.

[67] T. Green, N. Jenkins *et al.*, "Issues for distribution system operation at 2030," SUPERGEN Hub-Net, Position Paper 3, 2014.

[68] Q. Chen, D. Kaleshi, and Z. Fan, "Inferring low voltage transformer state using only smart metering data," in *Innovative Smart Grid Technologies Europe (ISGT EUROPE)*. IEEE PES, Oct 2013, pp. 1–5.

[69] P. Taylor, "SUPERGEN HubNet and symposium agenda," private communication, 2016.

[70] "Energy 2020: a strategy for competitive, sustainable and secure energy," European Commission, Tech. Rep. Com(2010) 639, 2010.

[71] "The European electricity grid initiative - roadmap 2010-18," European Commission, Tech. Rep., 2010.

[72] IRED. (2011) IRED: Integration of renewable sources and distributed generation into the European electricity grid. EU Project. [Online]. Available: www.ired-cluster.org

[73] MICROGRIDS. (2011) Microgrids. EU Project. [Online]. Available: http://www.microgrids.eu

[74] MORE MICROGRIDS. (2011) More Microgrids. EU Project. [Online]. Available: http://www.microgrids.eu

[75] DISPOWER. (2011) DISPOWER: Distributed generation with high penetration of renewable energy sources. EU Project. [Online]. Available: http://www.dispower.org

[76] DER-LAB. (2011) DERlab: European distributed energy resources laboratories. EU Project. [Online]. Available: http://www.der-lab.net

[77] FENIX. (2011) FENIX: Flexible electricity network to integrate the expected energy revolution. EU Project. [Online]. Available: http://www.fenix-project.org

[78] SOLID-DER. (2011) SOLID-DER: Coordination action to consolidate RTD activities for large-scale integration of DER into the European electricity market. EU Project. [Online]. Available: http://www.solid-der.org

[79] J. Oyarzabal, J. Jimeno, A. Engler, C. Hardt, and J. Ruela, "Agent based micro grid management system," in *Proceedings of the International Conference on Future Power Systems*, November 2005.

[80] N. Hatziargyriou, N. Jenkins, G. Strbac, J. A. P. Lopes, J. Ruela, A. Engler, J. Oyarzabal, G. Kariniotakis, and A. Amorim, "Microgrids - large scale integration of microgeneration to low voltage grids," in *CIGRE*, 2006.

[81] N. Hatziargyriou, "Advanced architectures and control concepts for more microgrids: Executive summary report final results," EU More Microgrids Consortium, Tech. Rep., December 2009.

[82] A. G. Tsikalakis and N. D. Hatziagyriou, "Centralized control for optimising microgrids operation," *IEEE Transactions on Energy Conversion*, vol. 23, no. 1, pp. 241–248, March 2008.

[83] S. Chatzivasiliadis, N. Hatziargyriou, and A. Dimeas, "Development of an agent based intelligent control system for microgrids," in *IEEE Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century*, July 2008.

[84] J. Marti, O. Abarrategui, and A. Gonzalez, "Real experience with active DER in distribution networks," in *Electricity Distribution - Part 1, 2009. CIRED 2009. 20th International Conference and Exhibition on*, 2009, pp. 1 –4.

[85] D. Pudjianto, C. Ramsay, and G. Strbac, "Virtual power plant and system integration of distributed energy resources," *IET Renewable Power Generation*, vol. 1, no. 1, pp. 10 –16, March 2007.

[86] J. Mutale, "Final report: Development of a distributed intelligent load controller (DILC)," DISPOWER, Workpackage Deliverable 1.4, November 2004.

[87] W2E. (2011) Web 2 Energy: The intelligent energy network of the future. EU Project Website. http://www.web2energy.com/.

[88] *Communication networks and systems in substations.*, International Electrotechnical Commission (IEC) Std. IEC61 850.

[89] B. Fenn, O. Hopp, M. Ahner, B. M. Buchholz, V. Buehner, A. Do, N. He, M. Marschollek, W. Wagner, and Z. A. Styczynski, "Advanced technologies of demand side integration by VPPs and through smart metering in households experiences from a lighthouse project," in *CIGRE*, 2012.

[90] S. Sučić, J. G. Havelka, and T. Dragičević, “A device-level service-oriented middleware platform for self-manageable DC microgrid applications utilizing semantic-enabled distributed energy resources,” *International Journal of Electrical Power & Energy Systems*, vol. 54, pp. 576 – 588, 2014.

[91] W. Deng, W. Pei, Z. Shen, Z. Zhao, and H. Qu, “Adaptive micro-grid operation based on iec 61850,” *Energies*, vol. 8, no. 5, pp. 4455–4475, May 2015.

[92] G. Peppink, R. Kamphuis, K. Kok, A. Dimeas, E. Karfopoulos, N. Hatziargyriou, N. Hadjsaid, R. Caire, R. Gustavsson, J. M. Salas, H. Niesing, J. van der Velde, L. Tena, F. Bliek, M. Eijgelaar, L. Hamilton, and H. Akkermans, *INTEGRAL: ICT-Platform Based Distributed Control in Electricity Grids with a Large Share of Distributed Energy Resources and Renewable Energy Sources.* Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, ch. 23, pp. 215–224. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-19322-4_23

[93] N. Hadjsaid, R. Caire, and B. Raison, “Decentralized operating modes for electrical distribution systems with distributed energy resources,” in *2009 IEEE Power Energy Society General Meeting*, July 2009, pp. 1–4.

[94] MIRACLE. (2011) MIRACLE: Micro-request-based aggregation, forecasting and scheduling of energy demand, supply and distribution. EU Project. [Online]. Available: http://www.miracle-project.eu

[95] H. Berthold, M. Boehm, C. Nychtis, R. Rantzau, J. Verhoosel, M. Konsman, Z. Marinsek, and B. Filipic, “Micro-request-based aggregation, forecasting and scheduling of energy demand, supply and distribution : Yearly public report for the first year,” SAP AG, Tech. Rep., December 2010.

[96] *MIRABEL-ONE: Initial draft of the MIRABEL Standard*, MIRABEL Std.

[97] M. Boehm, L. Dannecker, A. Doms, E. Dovgan, B. Filipic, U. Fischer, W. Lehner, T. B. Pedersen, Y. Pitarch, L. Siksnys, and T. Tusar, “Data management in the MIRABEL smart grid system,” in *Proceedings of the 1st Workshop on Energy Data Management (EnDM 2012, 15th International Conference on Extending Database Technology*, Berlin, Germany, 2012.

[98] R. Belhomme, M. Sebastian, A. Diop, M. Entem, F. Bouffard, G. Valtorta, A. D. Simone, R. Cerero, C. Yuen, S. Karkkainen, and W. Fritz, “ADDRESS technical and commercial conceptual architectures,” ADDRESS, Tech. Rep. Deliverable 1.1, 2009.

[99] ADDRESS. (2011) ADDRESS: Active distribution networks with full integration of demand and distributed energy resources. EU Project Website. [Online]. Available: http://www.addressfp7.org

[100] P2P-SmartTest. (2016) P2P-SmartTest: Peer-to-Peer Smart energy distribution networks. EU Project. [Online]. Available: http://cordis.europa.eu/project/rcn/194444_en.html

[101] GOFLEX. (2016) GOFLEX: Generalized Operational FLEXibility for integrating renewables in the distribution grid. EU Project. [Online]. Available: http://cordis.europa.eu/project/rcn/206239_en.html

[102] J. Eto, R. Lasseter, B. Schenkman, J. Stevens, D. Klapp, H. Volkommer, E. Linton, H. Hurtado, and J. Roy, "Overview of the CERTS microgrid laboratory test bed," in *Integration of Wide-Scale Renewable Resources Into the Power Delivery System, 2009 CIGRE/IEEE PES Joint Symposium*, July 2009.

[103] A. Dimeas and N. Hatziargyriou, "Operation of a multiagent system for microgrid control," *IEEE Transactions on Power Systems*, vol. 20, no. 3, pp. 1447 – 1455, aug. 2005.

[104] USC, "Energy independence and security act of 2007," United States Congress, Public Law, December 2007.

[105] X. Mamo, "Update of the profiling and mapping of intelligent grid R&D programs," EPRI, Tech. Rep., 2008.

[106] NEDO. (2006) Demonstrative project on grid-interconnection of clustered photovoltaic power generation systems. NEDO. [Online]. Available: http://www.nedo.go.jp/english/publications/brochures/pdf/ota-project_nedo.pdf

[107] L. Meng, A. Luna, E. R. Daz, B. Sun, T. Dragicevic, M. Savaghebi, J. C. Vasquez, J. M. Guerrero, M. Graells, and F. Andrade, "Flexible system integration and advanced hierarchical control architectures in the microgrid research laboratory of Aalborg University," *IEEE Transactions on Industry Applications*, vol. 52, no. 2, pp. 1736–1749, March 2016.

[108] Y. Han, P. Shen, X. Zhao, and J. Guerrero, "Control strategies for islanded microgrid using enhanced hierarchical control structure with multiple current-loop damping schemes," *IEEE Transactions on Smart Grid*, vol. PP, no. 99, 2016.

[109] M. S. Golsorkhi, D. D. C. Lu, and J. M. Guerrero, "A GPS-based decentralized control method for islanded microgrids," *IEEE Transactions on Power Electronics*, vol. 32, no. 2, pp. 1615–1625, Feb 2017.

[110] Q. Li, F. Chen, M. Chen, J. Guerrero, and D. Abbott, "Agent-based decentralized control method for islanded microgrids," *IEEE Transactions on Smart Grid*, vol. 7, no. 2, pp. 637–649, 2016.

[111] M. He and J. Zhang, "A dependency graph approach for fault detection and localization towards secure smart grid," *Smart Grid, IEEE Transactions on*, vol. 2, no. 2, pp. 342–351, June 2011.

[112] A. Dimeas and N. Hatziargyriou, "Agent based control of virtual power plants," in *International Conference on Intelligent Systems Applications to Power Systems*, November 2007.

[113] H. F. Wang, "Multi-agent co-ordination for the secondary voltage control in power-system contingencies," in *IEE Proceedings - Generation, Transmission and Distribution*, vol. 148, 2001, pp. 61–66.

[114] A. Al-Hinaii, "Multi-agent control and operation of electric power distribution systems," Ph.D. dissertation, West Virginia University, 2005.

[115] S. McArthur, E. Davidson, V. Catterson, A. Dimeas, N. Hatziargyriou, F. Ponci, and T. Funabashi, "Multi-agent systems for power engineering applications part i: Concepts, approaches, and technical challenges," *IEEE Transactions on Power Systems*, vol. 22, no. 4, pp. 1743 –1752, November 2007.

[116] T. Nagata, H. Fujita, and H. Sasaki, "Decentralized approach to normal operations for power system network," in *Proceedings of the 13th International Conference on Intelligent Systems Application to Power Systems*, November 2005.

[117] F. Ygge and H. Akkermans, "Decentralized markets versus central control: A comparative study," *Journal of Artificial Intelligence Research*, vol. 11, pp. 301–333, 1999.

[118] P. Trichakis, P. C. Taylor, G. Coates, and L. M. Cipcigan, "Distributed control approach for small-scale energy zones," in *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, vol. 222 (2), 2008, pp. 137–147.

[119] A. Rosenfeld, G. A. Kaminka, and S. Kraus, *Coordination of large-scale multiagent systems.* New York: Springer, 2006, ch. A study of scalability properties in robotic teams, pp. 28–51.

[120] P. Scerri, R. Vincent, and R. Mailler, *Coordination of Large-Scale Multiagent Systems.* New York: Springer, 2006, ch. Comparing Three Approaches to Large-Scale Coordination, pp. 53–71.

[121] P. C. Taylor and N. Jenkins, "The development and application of distributed fuzzy load control to an autonomous wind diesel system," in *Proceedings of the EWEA Wind Power for the 21st Century Conference*, September 2000, pp. 290–293.

[122] N. Hatziargyriou, A. Dimeas, A. Tsikalakis, J. Lopes, G. Karniotakis, and J. Oyarzabal, "Management of microgrids in market environment," in *International Conference on Future Power Systems*, November 2005.

[123] K. De Brabandere, K. Vanthournout, J. Driesen, G. Deconinck, and R. Belmans, "Control of microgrids," in *Power Engineering Society General Meeting, 2007. IEEE*, 2007, pp. 1 –7.

[124] K. Vanthournout, K. D. Brabandere, E. Haesen, J. V. den Keybus, G. Deconinck, and R. Belmans, "Agora: distributed tertiary control of distributed resources," in *Proceedings of the 15th Power systems Computation Conference*, Liege, Belgium, August 2005.

[125] G. James, D. Cohen, R. Dodier, G. Platt, and D. Palmer, "A deployed multi-agent framework for distributed energy applications," in *Proceedings of the fifth international joint conference on Autonomous Agents and Multiagent systems*, 2006, pp. 676–678.

[126] S. Omi, M. Kumagai, and K. Yamane, "Super decentralized control on the basis of centrally optimized control model for distribution voltage regulation," in *CIRED*, 2013.

[127] V. Nasirian, Q. Shafiee, J. Guerrero, F. L. Lewis, and A. Davoudi, "Droop-free distributed control for AC microgrids," *IEEE Transactions on Power Electronics*, vol. 31, no. 2, pp. 1600 – 1617, 2 2016.

[128] C. F. Wagner and R. D. Evans, *Symmetrical Components.* McGraw-Hill, 1933.

[129] M. H. J. Bollen, "Definitions of voltage unbalance," *IEEE Power Engineering Review*, pp. 49–50, 2002.

[130] S. Papathanassiou, N. D. Hatziagyriou, and K. Strunz, "A benchmark low voltage microgrid network," in *CIGRE*, April 2005.

[131] S. Ingram, S. Probert, and K. Jackson, "The impact of small scale embedded generation on the operating parameters of distribution networks," PB Power and DTI, Tech. Rep., 2003.

[132] P. Trichakis, P. C. Taylor, R. Hair, and P. Lyons, "Predicting the technical impacts of high levels of small scale embedded generators on low voltage networks," *IET Renewable Power Generation*, vol. 2, no. 4, pp. 249–262, 2008.

[133] C. Thompson, P. Foster, E. Lodge, and D. Miller, "Customer led network revolution: A guide to the load & generation profile datasets," Northern Powergrid, Tech. Rep. CLNR-L071, August 2014.

[134] Manitoba HVDC. PSCAD 4.2.1. DVD-ROM. Manitoba HVDC Research Centre Inc.

[135] H. W. Dommel, "Digital computer solution of electromagnetic transients in single- and multiphase networks," *IEEE Transactions on Power Apparatus and Systems*, vol. 88, pp. 388–399, 1969.

[136] TNEI. (2010) IPSA+ 1.6.5. CD-ROM. TNEI Services Ltd.

[137] RTDS Technologies Inc. RSCAD. DVD-ROM. RTDS Technologies Inc.

[138] P. Trichakis, P. C. Taylor, L. M. Cipcigan, P. F. Lyons, R. Hair, , and T. Ma, “An investigation of voltage unbalance in low voltage distribution networks with high levels of SSEG,” in *International Universities Power Engineering Conference*, 2006.

[139] H. Hutchinson. (2017) UKGN IPSA simulation - location awareness in multi-agent control of distributed energy resources - Hutchinson - 2017.zip. Online. Durham University. [Online]. Available: doi:10.15128/r247429913c

[140] W. H. Kersting, *Distribution System Modelling and Analysis.* CRC Press, 2002.

[141] E. Taylor, “Smart meter wide area network DCC assurance strategy,” Data and Communications Company, Tech. Rep., 2016.

[142] K. O. H. Pedersen, A. H. Nielsen, and N. K. Poulsen, “Short-circuit impedance measurement,” *IEE Proceedings - Generation, Transmission and Distribution*, vol. 150, pp. 169–174, March 2003.

[143] H. Pezeshki and P. J. Wolfs, “Consumer phase identification in a three phase unbalanced LV distribution network,” in *IEEE Power Engineering Society Innovative Smart Grid Technologies Europe*, Berlin, October 2012.

[144] MathWorks. MATLAB. Mass., USA. [Online]. Available: http://uk.mathworks.com

[145] K. Pandiaraj, P. C. Taylor, N. Jenkins, and C. Robb, “Distributed load control of autonomous renewable energy systems,” *IEEE Transactions on Energy Conversion*, vol. 16, pp. 14–19, 2001.

[146] A. S. Rao and M. P. Georgeff, *Readings in Agents.* Morgan Kaufmann Publishers, Inc., 1998, ch. Modeling rational agents within a BDI-Architecture, pp. 317–328.

[147] P. R. Cohen and H. J. Levesque, “Intention is choice with commitment,” *Artificial Intelligence*, vol. 42, no. 2-3, pp. 213–261, 1990.

[148] H. J. Levesque, P. R. Cohen, and J. H. T. Nunes, “On acting together,” in *Proceedings of the National conference on Artificial intelligence.* AAAI, 1990, pp. 94–99.

[149] N. R. Jennings, “Controlling cooperative problem solving in industrial multi-agent systems using joint intentions,” *Artificial Intelligence*, vol. 75, pp. 195–240, 1995.

[150] B. Horling and V. Lesser, “A survey of multi-agent organizational paradigms,” *The Knowledge Engineering Review*, vol. 19, no. 4, pp. 281–316, 2005.

[151] A. Haddadi, *Communication and Cooperation in Agent Systems*, ser. Lecture Notes in Artificial Intelligence, J. G. Carbonell and J. Siekmann, Eds. Berlin: Springer, 1996, no. 1056.

[152] M. Barbuceanu, "Coordinating agents by role based social constraints and conversation plans," in *Proceedings of the National Conference on Artificial Intelligence.* AAAI, 1997, pp. 16–21.

[153] C. H. Brooks, E. H. Durfee, and A. Armstrong, "An introduction to congregating in multiagent systems," in *Proceedings of Fourth International Conference on MultiAgent Systems*, 2000, pp. 79–86.

[154] L.-K. Soh, C. Tsatsoulis, and H. Sevay, *Distributed Sensor Networks: a Multiagent Perspective.* Kluwer Academic Publishers, 2003, ch. A Satisficing, Negotiated and Learning Coalition Formation Architecture, pp. 109–138.

[155] M. J. Wooldridge and N. R. Jennings, "Formalizing the cooperative problem solving process," in *Proceedings of the Thirteenth International Workshop on Distributed Artificial Intelligence*, 1994.

[156] J. P. Kahan and A. Rapoport, *Theories of Coalition Formation.* New Jersey: Lawrence Erlbaum Associates, 1984.

[157] FIPA, *FIPA SL Content Language Specification*, FIPA Standard 00 008, Rev. I, December 2002.

[158] ——, *FIPA Communicative Act Library Specification*, FIPA Standard 00 037, Rev. J, December 2002.

[159] ——, *FIPA Contract Net Interaction Protocol*, FIPA Standard 00 029, Rev. H, 2002.

[160] B. Burmeister, A. Haddadi, and K. Sundermeyer, *From Reaction to Cognition: 5th European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, ser. Lecture Notes in Computer Science. Springer, 1993, no. 957, ch. Generic, configurable, cooperation protocols for multi-agent systems, pp. 155–171.

[161] O. Shehory and S. Kraus, "Methods for task allocation via agent coalition formation," *Artificial Intelligence*, vol. 101, pp. 165–200, 1998.

[162] Y. Shoham and K. Leyton-Brown, *Multiagent Systems: Algorithmic, Game-Theoretic and Logical Foundations.* Cambridge University Press, 2009.

[163] E. H. Durfee and T. A. Montgomery, "A hierarchical protocol for coordinating multiagent behaviors," in *Proceedings of the Eighth National Conference on Artificial Intelligence.* AAAI, 1990, pp. 86–93.

[164] G. Hutchinson, P. C. Taylor, G. Coates, and R. Hair, "A multi-agent system for decentralised control of low voltage distribution networks," in *Universities Power Engineering Conference (UPEC)*, 2010.

[165] M. Tambe, "Agent architectures for flexible, practical teamwork," in *Proceedings of the National Conference on Artificial Intelligence.* AAAI, 1997, pp. 22–28.

[166] S. P. Ketchpel, *From Reaction to Cognition: 5th European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, ser. Lecture Notes in Computer Science. Springer, 1993, no. 957, ch. Coalition Formation among Autonomous Agents, pp. 73–88.

[167] T. W. Sandholm and V. R. Lesser, "Coalition formation among bounded rational agents," in *Proceedings of the fourteenth international joint conference on Artificial intelligence - Volume 1.* San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1995, pp. 662–669.

[168] S. Ketchpel, "Forming coalitions in the face of uncertain rewards," in *Proceedings of the twelfth national conference on Artificial intelligence.* AAAI, 1994, pp. 414–419.

[169] M. D. Sadek, "A study in the logic of intention," in *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning*, vol. 1, Cambridge, MA, 1992, pp. 462–473.

[170] M. P. Georgeff and F. F. Ingrand, "Decision-making in an embedded reasoning system," in *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, vol. 2, August 1989, pp. 972–978.

[171] R. H. Bordini, M. Dastani, J. Dix, and A. E. F. Seghrouchni, Eds., *Multi-Agent Programming: Languages, Tools and Applications.* Springer, 2005.

[172] Stanford Center for Biomedical Informatics Research. Protégé. [Online]. Stanford University. http://protege.stanford.edu/.

[173] TI. Java Agent DEvelopment framework (JADE) 3.7. [Online]. Available: http://jade.tilab.com

[174] A. L. Dimeas and N. D. Hatziargyriou, "A MAS architecture for microgrids control," in *Proceedings of the 13th International Conference on Intelligent Systems Application to Power Systems.*, 2005.

[175] A. S. Rao and M. P. Georgeff, "BDI agents: From theory to practice," in *Proceedings of the First International Conference on Multi-Agent Systems*, June 1995.

[176] A. Macias-Estrada, O.-I. Lepe-Aldama, and A. Garcia-Macias, "Implementing true multicast communications support for the jade/leap agent framework," in *Sixth Mexican International Conference on Computer Science on*, 2005, pp. 206–213.

[177] G. Hutchinson, R. Hair, and P. C. Taylor, "Leveraging location awareness for distributed energy resources: Self-organising energy communities," in *International Conference and Exhibition on Electricity Distribution (CIRED)*, 2013.

[178] European Commission. (2010) CORDIS: Community research and development information service. [Online]. Available: http://cordis.europa.eu

[179] INTEGRAL. (2011) INTEGRAL: Integrated ICT-platform based distributed control in electricity grids. EU Project. [Online]. Available: http://www.integral-eu.com

[180] IS-POWER. (2008) IS-POWER: Isolated power systems: knowledge and technology sharing: distributed generation, grid and demand management. EU Project. [Online]. Available: http://cordis.europa.eu/project/rcn/86579_en.html

[181] ADINE. (2010) ADINE: Active distribution network. EU Project. [Online]. Available: http://cordis.europa.eu/project/rcn/85682_en.html

[182] CRISTAL. (2008) CRISTAL: Control of renewable integrated systems targeting advanced landmarks. EU Project. [Online]. Available: http://cordis.europa.eu/project/rcn/86354_en.html

[183] VSYNC. (2010) VSYNC: Virtual synchronous machines for frequency stabilisation in future grids with a significant share of decentralized generation. EU Project. [Online]. Available: http://cordis.europa.eu/project/rcn/85687_en.html

[184] GROW-DERS. (2010) GROW-DERS: Grid reliability and operability with distributed generation using flexible storage. EU Project. [Online]. Available: http://cordis.europa.eu/project/rcn/86584_en.html

[185] DESIRE. (2008) DESIRE: Dissemination strategy on electricity balancing for large scale integration of renewable energy. EU Project. [Online]. Available: http://cordis.europa.eu/project/rcn/85711_en.html

[186] RELIANCE. (2009) RELIANCE: CooRd.perspectives of the European transm.network research activities to optimise the reLIAbility of power supply,usiNg a systemiC approach,inv.growing distrib.generation and renewable energy markEts. EU Project. [Online]. Available: http://cordis.europa.eu/project/rcn/75099_en.html

[187] CEERES. (2009) CEERES: Large-scale integration of RES-E and co-generation into energy supplies in Associated Candidate Countries. EU Project. [Online]. Available: http://cordis.europa.eu/project/rcn/75657_en.html

[188] EU-DEEP. (2011) EU-DEEP: The birth of a EUropean Distributed EnErgy Partnership that will help the largescale implementation of distributed energy resources in Europe. EU Project. [Online]. Available: http://cordis.europa.eu/project/rcn/73969_en.html

[189] G4V. (2011) G4V: Grid for vehicles. EU Project Website. [Online]. Available: http://www.g4v.eu

[190] WiseGRID. (2016) WiseGRID: Wide-scale demonstration of integrated solutions and business models for european smartGRID. EU Project. [Online]. Available: http://www.wisegrid.eu

Appendix A

Highlighted EU research projects

The EU FP6-FP7 research programme included research on microgrids and DG integration as does the current Horizon2020 programme; the most significant are described in detail in Section 2.2.5.2. The list is expanded here for completeness. Many of the EU project consortia hosted websites containing their project outputs, and these links have been referenced where possible. Some of these concluded project websites have subsequently closed. Relevant project outputs (journal articles, whitepapers, etc.) have been referenced directly in the main document where appropriate, but links to the EU CORDIS database descriptions have been included in Table A.1 where the web presence is unavailable. A further 11 H2020 projects into demand flexibility with a launch of November/December 2016 are not included but details will be available via the CORDIS service[178].

Table A.1: Other relevant EU research projects into microgrids and DER integration, programmes FP6-H2020

Project	Date	Abstract
INTEGRAL [179]	11/2007 - 10/2010	*INTEGRAted ICT-platform based Distributed Control (IIDC) in electricity grids with a Large share of distributed energy resources and renewable energy sources.* The INTEGRAL project aimed to build and demonstrate an industry-quality reference solution for DER aggregation-level control and coordination, based on commonly available ICT components, standards, and platforms.
IS-POWER [180]	10/2007 - 4/2010	Isolated Power Systems: Knowledge and technology sharing: Distributed generation, grid and demand management

ADINE [181]	10/2007 - 9/2010	*Active Distribution Network.* The aim was to develop, demonstrate and validate a new Active Network Management (ANM) method of distribution network including distributed generation (DG) and enabling solutions to support it. The enabling solutions operate as active components in managing the network allowing easier interconnection of DG units. The solutions cover protection, voltage and reactive power control and planning and information systems of networks.
CRISTAL [182]	12/2007 - 12/2009	*"Control of renewable integrated systems targeting advanced landmarks".* The project focused on the development of enabling technologies for distributed smart energy networks, with high power quality and service security. The technical issues to be coordinated were concerned with solar, wind and micro-hydro systems control in conjunction with compensatory energy storage systems (fuel cells, hydrogen) and connection to the grid.
VSYNC [183]	10/2007 - 9/2010	*"Virtual synchronous machines for frequency stabilisation in future grids with a significant share of decentralized generation".* Demonstrated of the Virtual Synchronous Generator (VSG) concept are at the level of the individual owners of distributed generators and at the level of the network operator for groups of many distributed generators.
GROW-DERS [184]	9/2007 - 8/2010	*"Grid Reliability and Operability with Distributed Generation using Flexible Storage".* Aimed to demonstrate technical and economical possibilities of existing electricity storage technologies.
DESIRE [185]	6/2005 - 5/2007	*"Dissemination Strategy on Electricity Balancing for large Scale Integration of Renewable Energy".* Aimed to disseminate practices for integrating fluctuating renewable electricity supplies into electricity systems using combined heat and power.
RELIANCE [186]	10/2005 - 9/2007	*CooRd.perspectives of the European transm.network research activities to optimise the reLIAbility of power supply,usiNg a systemiC approach,inv.growing distrib.generation and renewable energy markEts* consisted of transmission network research involving increasing DG and the use of renewable energy markets.

CEERES [187]	04/2005 - 07/2006	*"Large-scale integration of RES-E and co-generation into energy supplies in Associated Candidate Countries"*. CEERES examined large-scale integration of renewable- and co-generation electricity into energy supplies with particular focus on the eight Central European Countries which joined the EU in May 2004.
EU-DEEP [188]	1/2004 - 6/2009	*"The birth of a EUropean Distributed EnErgy Partnership that will help the largescale implementation of distributed energy resources in Europe"*. A consortium of eight European energy utilities examined technical and non-technical barriers preventing large-scale deployment of DER in Europe.
ADDRESS [99]	6/2008-5/2013	*Active Distribution networks with full integration of Demand and distributed energy RESourceS.*. ADDRESS aimed to develop Active Distribution Networks architectures to balance power generation and demand in real time to increase flexibility of the entire system. It combined communications, automation and household technologies with novel market trading mechanisms and algorithms.
W2E [87]	1/2010 - 12/2012	*Web to Energy (W2E)* investigated standardisation of communications infrastructure and data management as an enabling technology for smart grids, predominantly by extending and existing control architectures. It examined the lack of standardisation of communication technologies valid throughout the entire electricity grid. Aggregation of SSEG was included in its project objectives of implementing communication down to consumers at the metering level.
G4V [189]	1/2010 - 6/2011	*Grid for Vehicles*: Analysis of the impact and possibilities of a mass introduction of electric and plug-in hybrid vehicles on the electricity networks in Europe, in particular for energy efficiency and operation of smart grids.

MIRABEL[94] (formerly MIRACLE)	1/2010 - 4/2013	*Micro-Request-Based Aggregation, Forecasting and Scheduling of Energy Demand, Supply and Distribution.*MIRABEL aimed to aggregate data from the "prosumer" (producer consumer, a consumer that also produces electricity for export), to balance supply and demand through market mechanisms[95], hence allowing a DNO to manage renewable energy sources in higher penetrations. It used the Experimental Microgrid at CRES, and applied the VPP-like technology to create low-emissions zones through optimising energy sources on a carbon basis with a trial in Germany.
P2P-SmartTest [100]	1/2015 - 6/2011	*Peer to Peer Smart Energy Distribution Networks(P2P-SmartTest)* will examine peer-to-peer operation for integration of demand side flexibility (i.e. provision of ancillary services through DER); this will, however, primarily focus on market participation of small-scale DER operators.
GOFLEX [101]	11/2016 - 10/2019	*Generalized Operational FLEXibility for integrating renewables in the Distribution Grid* - "will innovate, integrate, further develop and demonstrate a group of electricity smart-grid technologies, enabling the cost-effective use of demand response in distribution grids"[101]
WiseGRID [190]	11/2016 - 4/2020	*Wide scale demonstration of Integrated Solutions and business models for European smartGRID (WiseGRID)* will examine energy cooperatives as business models in systems with extensive DER and using demand response.

Appendix B

JADE Behaviour Scheduler bugs

B.1 Round-Robin Bug

The Scheduler class in the JADE libraries determines execution of behaviours (plans, intentions and any other routines). However, it does not guarantee execution. If the final behaviour in the execution list (`readyBehaviours`) adds a new behaviour, that new behaviour is *not* executed in the current cycle around the list. But if the penultimate behaviour in `readyBehaviours` adds a new behaviour, that new behaviour *is* executed in this cycle. This is because the next behaviour is selected before the current one has been executed. We can see this with the following code:

```
public class BuggyAgent extends Agent
        {

        protected void setup()
                {
                System.out.println("Hello World!");
                addBehaviour(new CPUEater());
                }

        public class CPUEater extends CyclicBehaviour
                {
                int i=0;
                public void action()
                        {
                        System.out.println("Iteration "+ i);
```

```
            if (i==5)
                {
                addBehaviour(new Ender());
                }
            i++;
            }
        }

    public class Ender extends OneShotBehaviour
        {
        public void action()
            {
            System.out.println(”End!”);
            myAgent.doSuspend();
            }
        }
    }
```

A correctly-implemented scheduler should terminate on the 5th iteration; however, it does not. This can be solved by calculating `currentIndex` immediately before it is used:

```
currentIndex = (currentIndex + 1)
Behaviour b = (Behaviour)readyBehaviours.get(currentIndex);
return b;
```

instead of

```
Behaviour b = (Behaviour)readyBehaviours.get(currentIndex);
currentIndex = (currentIndex + 1)
return b;
```

Then `currentIndex` must be initialised to -1 throughout.

B.2 Skipping behaviours

`currentIndex` is not decremented if the behaviour removed is the current one, eg. in a call to block(). As a result, we end up skipping over behaviours. The following code misses out behaviour `Third`.

```
public class StepOver extends Agent{
        public First a = new First();
        public Second b = new Second();
        public Third c = new Third();
        public Fourth d = new Fourth();

        protected void setup() {
                addBehaviour(a);
                addBehaviour(b);
                addBehaviour(c);
                addBehaviour(d);
        }

        public class First extends CyclicBehaviour
        {
                public void action()
                {
                        System.out.println("1");
                }
        }
        public class Second extends CyclicBehaviour
        {
                public void action()
                {
                        System.out.println("2");
                        block();
                }
        }
        public class Third extends CyclicBehaviour
        {
```

```
                public void action()
                {
                        System.out.println("3");
                }
        }
        public class Fourth extends CyclicBehaviour
        {
                public void action()
                {
                        System.out.println("4");
                        myAgent.doSuspend();
                }
        }
}
```

As a result, we need

```
if (index <= currentIndex)
        {
        --currentIndex;
        }
```

Now, If the behaviour removed is the final one, then

```
if (index == currentIndex && currentIndex == readyBehaviours.size())
        {
        currentIndex = -1;
        }
```

causes us further problems, because the scheduler will skip over one more behaviour added at the end if we call

```
removeBehaviour(this);
addBehaviour(new endbehaviour());
```

We can safely remove this, because `currentIndex` has already been decremented, and then the fixed round robin algorithm correctly handles going past the end of the readyBehaviours list.

B.3 Patch

The following patch to the core libraries solves these problems.

```
--- SchedulerJADEv4-0.java
+++ StepOverfixedScheduler.java
@@ -81,7 +81,9 @@

        public Scheduler(Agent a) {
                owner = a;
-               currentIndex = 0;
+
+               //altered to solve Round Robin bug
+               currentIndex = -1;
        }

        // Add a behaviour at the end of the behaviours queue.
@@ -222,13 +224,15 @@
                        owner.idle();
                }

+               //Moved to solve Round Robin bug
+               currentIndex = (currentIndex + 1) % readyBehaviours.size();
                //#MIDP_EXCLUDE_BEGIN
                Behaviour b = (Behaviour)readyBehaviours.get(currentIndex);
                //#MIDP_EXCLUDE_END
                /*#MIDP_INCLUDE_BEGIN
                 Behaviour b = (Behaviour)readyBehaviours.elementAt(currentIndex);
                 #MIDP_INCLUDE_END*/
-               currentIndex = (currentIndex + 1) % readyBehaviours.size();
+
                return b;
        }

@@ -274,7 +278,9 @@
```

```
                }

                // The current index is not saved when persisting an agent
-               currentIndex = 0;
+
+               //Altered to solve Round Robin bug to -1 from 0
+               currentIndex = -1;
        }

        //#MIDP_EXCLUDE_END
@@ -303,12 +309,14 @@
        /*#MIDP_INCLUDE_BEGIN
        readyBehaviours.removeElement(b);
        #MIDP_INCLUDE_END*/
-       if(index < currentIndex)
-         --currentIndex;
+
        //if(currentIndex < 0)
        //  currentIndex = 0;
-       else if (index == currentIndex && currentIndex == readyBehaviours.size())
-               currentIndex = 0;
+
+       //Remove list overflow, dealt with in schedule()
+       //make <= to remove stepover bug
+       if(index <= currentIndex)
+               --currentIndex;
          }
        return index != -1;
        }
```