The Decentralised Control of Electricity Networks- Intelligent and Self-Healing Systems

Glenn Platt

CSIRO Energy Technology
10 Murray Dwyer Circuit, Mayfield West, NSW, 2304, Australia

glenn.platt@csiro.au

Keywords: agents, distributed energy, demand management

Abstract

This paper reviews the state of the art in distributed energy control systems- decentralised control techniques that coordinate the actions of devices such as electricity loads or generators. The paper reviews two recently proposed control techniques that bring significant advantages over the first-generation distributed energy or demand management systems currently being trialled. It introduces the basic operating principles of these systems, and reviews the challenges involved in realising these techniques in practical applications.

1. INTRODUCTION

There is a growing interest around the world in the benefits available from more involved control of the demand side of electricity networks. Essentially, by coordinating the responses of the many small generators or loads operating in the electricity network, system-wide gains can be realised. For business operators, the benefits here can include better network utilisation, more accurate control of loads, and improved response to system outages. These benefits, and the related costs, are now being explored by many companies in deployments around the world, mostly in so-called *demand management* trials, targeted at improving the control of loads and small-scale generation in the network.

Typically, these first-generation demand management deployments can be characterised by the method through which they elicit a response from the demand side resource- the load or small generator under control. Most techniques rely on one or both of the following mechanisms:

- Getting a person to change the operating state of a load or generator in response to a time-varying price- so for example, the customer may disable a load when the price is expensive, but enable it at lower prices. Such techniques can scale to very large systems- the network company generally only needs to communicate a price to the network. Yet these systems are limited by the reliability, or *firmness* of response they can offer the network company.
- The network company directly controls the operating state of a load or generator via a dedicated communications and control system. These systems can offer relatively high levels of firmness, yet can be difficult to scale, as the technical challenges of controlling many thousands of devices are not insignificant.

Recognising the limitations of these first generation techniques, there are now a number of research organisations working on more advanced demand side control systems. Such systems are intended to bring a variety of benefits, including consideration of local user preferences, scalability whilst also offering known firmness, and minimal requirements for expensive infrastructure. Whilst immediately applicable to demand management projects, such systems are also being considered as a way for local users to deal with network outage situations, for operating remote area power supplies, and for coordinating localised generation and control in a way that brings benefit to surrounding users. Such benefits, and the control systems they are based on, are the subject of this paper. We will review a variety of state of the art demand side control systems, discussing their benefits and challenges, including the steps necessary before these systems are ready for commercial scale deployments.

2. IMPROVED CONTROL OF LOADS AND GENERATORS

Before describing the most recent techniques being considered for the control of loads and generators in the electricity network, it is worthwhile first reviewing what the characteristics of an optimal control system are.

As introduced above, one of the first measures of success for a control system managing large numbers of small loads and generators is its scalability- how well a given technique can cope when the number of devices under control increases arbitrarily. Importantly, in parallel with any consideration of the system's scalability must be an awareness of the system's depth of control- whilst a simple broadcast based control system may be highly scalable, such shallow consideration of the implications of control will significantly limit uptake of such a simple system. For example, consider a simple demand management system that broadcasts a "turn off" command to thousands of air-conditioners. Without consideration of the operating parameters of those air-conditioners- for example, whether a homeowner is comfortable, there is likely to be a public backlash against this system. Additionally, without an awareness of how many air-conditioners were actually *on*, it is difficult to obtain any degree of firmness of response from such a system. Thus, not only is scalability important, but the control technique must have a reasonable depth of control- it should consider local device constraints such as temperature boundaries for loads such as air-conditioning or refrigeration, fuel costs for generators, and so on.

Whilst, as introduced earlier, a firmness of response is necessary in a well performing control system, this firmness should continue through changing system conditions- so the control system should be dynamic and responsive. Additionally, the optimal control system should be robust against attack or failure- there should be no single point whose failure will jeopardise the operation of the entire system.

Given these desires- a system that provides firmness, yet considers local user constraints, is scalable and can respond dynamically to network conditions, many researchers are trending away from the more traditional control techniques used in electricity systems. Such centralised control systems, where a large central controlling entity makes decisions and communicates these to the wider network, are increasingly being pushed to their limits [1]. The growing complexity of control needed, particularly when faced with the large, diverse range of devices operating at the demand side of the network, means that centralised control systems are facing significant challenges of reliability and scalability [1], [2]. Given these limitations, the research community is trending towards a decentralised approach to the control of electricity networks. Such techniques often employ agent-based technology, where the overall behaviour of the system emerges from the behaviour of individual *agents*- individual smart devices that manage particular network components, and communicate with each other to achieve given global goals.

In work such as [2], [3], these decentralised agent-based techniques are considered for the control of relatively large network assets, with a focus on applications such as network protection, system operation and restoration after outage. In this paper, we are more interested in the use of decentralised control techniques for managing loads and generators in the demand side of the system- initially for application in demand management programmes, but later as a way of intelligently managing low-level network behaviour. An example deployment of a system of agents being used to manage the consumption and generation of electricity in a residential situation is shown in figure 1.

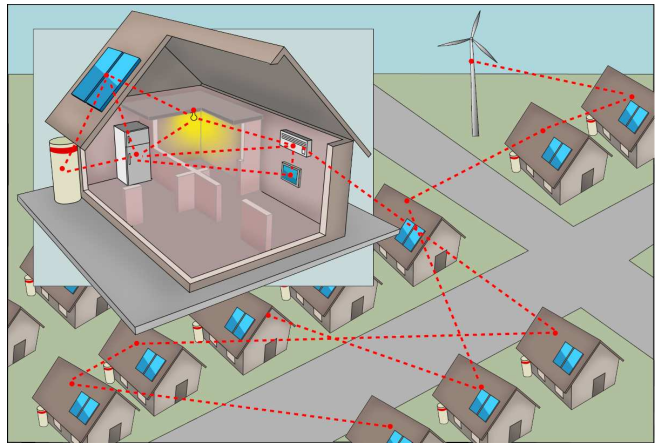

Figure 1. Agents used for controlling various loads and micro-generators in a residential setting

This is a relatively new application of this technology, and is quite different in approach to the first generation techniques currently operating in the demand side of electricity systems. In the following sections we review some of the most significant work in this area.

3. CONSUMERS AND SUPPLIERS- MARKET TYPE CONTROL SYSTEMS

One of the fundamental challenges when wanting to design a sophisticated, flexible control system is meeting the often conflicting requests of individual components of the system, whilst trying to steer the system to a common goal. As mentioned in the previous section, whilst centralised control systems may be able to find solutions to a given problem using powerful computational analysis, the complexity of modern electricity scenarios means that communications and computational overheads become a significant problem.

Decentralised agent based techniques are an ideal way to address this- they attempt to push much of the local computational load back on to the local agents, meaning local constraints can continue to be considered, whilst system goals are still achieved. To resolve the often conflicting requirements of multiple agents, one of the most common techniques used is to construct a "market", where a currency is introduced to the system, and local agents will negotiate with a broker to determine the cost of their desired action.

Running a successful market based system for controlling demand side energy devices is a challenging and novel concept, and thus there has recently been a significant amount of research dedicated to this area. Whilst limited to simulations, Ygge's work in [4] introduced the concept of a market for managing generation supply and demand. Further theoretical analysis considered features of both economics (for the market) as well as control theory, to prove the validity of this basic approach [5], [6]. Most recently, this work has resulted in an algorithm that has been trialled in practical deployment- the Powermatcher algorithm. As described in [7], the main goal of Powermatcher is to match the supply available from many small electricity generators operating in a minigrid, with a variety of small loads operating in the same minigrid.

In the market-based control paradigm, each load or generator is considered as a resource agent (RA), and there exists a broker agent (the "SD Matcher") whose aim is to fairly distribute the limited generation resources amongst the consumers. Resource agents issue bids to the broker agent, consisting of a proposed demand or supply at a given price. The broker evaluates all the bids, and adjusts the resource price in an attempt to make the total requested demand equal to the available supply. Thus, price becomes a signal of the relative scarcity of generation capacity at any given time- agents will continuously revise their bid, ensuring that the total amount of resource requested or offered (and thus its cost), matches the value (benefit) they will gain from the resource.

Particular resource agents will always strive to optimise the economics of their operation (minimise cost for loads, maximise revenue for generators), but are constrained by local parameters such as temperature boundaries, fuel supply, etc. Thus, the local constraints of an agent are implicitly recognised in the market process- for example, in a refrigerator agent if turning off the load will cost too much due to goods spoilage, then the agent will bid a high price so it can consume electricity. This selfish behaviour of local agents causes, over time, electricity consumption to be moved into periods of low price, and electricity generation to be moved into periods of high price. As a result, a match between supply and demand gradually emerges at the global system level.

To deal with very large systems of loads and generators, Powermatcher uses a tree structure of brokers to group market functions, as shown in figure 2. Here, a relatively small group of agents communicates with one particular broker, and the functionality of these brokers is aggregated upwards. The broker at the root of the tree (who is not aware of whether the agents below it are other brokers, or actual resource agents) forms a price for the entire network, and this price then propagates through the other brokers down to the bottom of the tree.

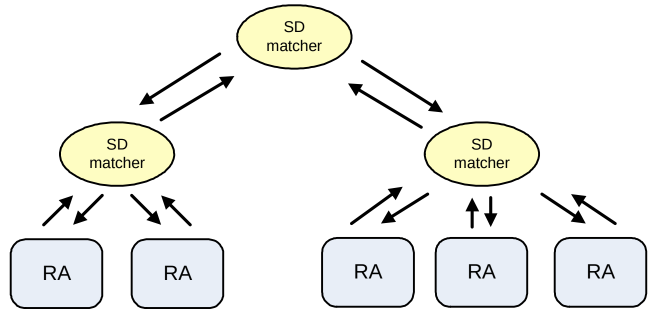

Figure 2. The hierarchy of the market-based coordination technique

Powermatcher has been tested in a variety of deployments. In one deployment Powermatcher was used to coordinate the power outputs of loads such as cool stores and residential properties, with a variety of distributed generation including residential combined heat and power (CHP) plant, diesel generator sets and wind farms. The aim of the coordination was to attempt to level the output of the combined set of loads and generators, relative to a situation where there was no coordination of the devices [8]. In another trial, the Powermatcher system was used to reduce the peak load on a residential sub-station by coordinating the output of many micro (1kW) CHP plant [9].

4. CAP BASED COORDINATION

In contrast to the market based work described in the previous section, CSIRO has been exploring an even less centralised way of coordinating the behaviour of a variety of agents controlling distributed energy resources.

CSIRO's coordination algorithm is based around four entities- a collective of resource agents, one or more brokers, an information repository (or "bulletin board"), and a summing agent. In the system, resource agents plan their local electricity demand for some period into the future, and then place these plans (which consist of simple statements of power consumption per interval of time) in to the information repository. The plans for all the agents are summed by the summing agent, to get the total predicted power demand for a particular time interval. This sum is then made available to the resource agents, as well as a demand cap figure, which indicates a desired *total* power consumption, for all agents, in the given time interval. The power cap is set by the broker agent, based on information such as prices from electricity market brokers, or status information from network operators.

Once a resource agent has observed the total power and cap figures, it will then try and modify its planned power consumption, to minimise consumption during intervals where total planned consumption is greater than the cap. In modifying its power consumption, a resource agent will

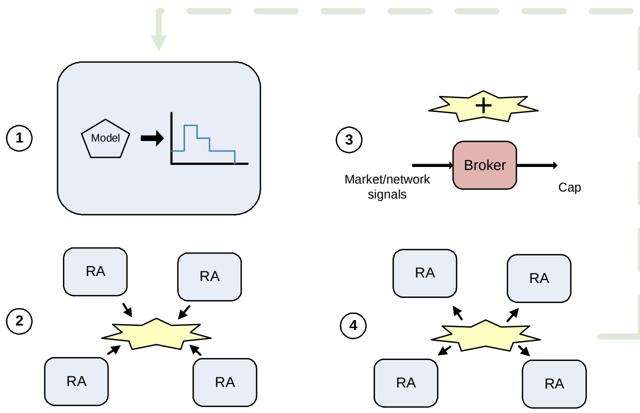

Figure 3. Operation of the cap based coordination method, including steps for resource agent (RA) future planning (1), submission of plans to the information repository (2), sum and cap setting (3), and retrieval of these from the information repository (4).

attempt to shuffle its power consumption into adjacent time intervals, creating a new planned power consumption profile. In forming this profile, resource agents will always respect their local constraints (such as temperature boundaries)- a resource agent will continue to consume energy in an interval that has excessive total power consumption, if it needs to due to local constraints.

Resource agents submit the revised power consumption plan to the information repository, these are summed, a new total power sum made available, and so on. This process continues to iterate until the cap is met, or the number of system iterations exceeds a predetermined threshold, indicating the cap simply cannot be met for the given interval. It is important to note that the entire process here is asynchronous- no explicit coordination is needed between plan submission, summing and cap setting.

The various steps involved in the cap coordination control technique are shown in figure 3.

This cap based coordination approach has been tested in both simulation and practice, controlling real electricity loads such as refrigerators. We have analysed a variety of features, such as how long the system takes to converge to satisfactory consumption plans for different power reduction goals, or the amount of warning agents need before a cap will occur, in order to be able to shuffle their power consumption around to meet the given cap.

5. INTEROPERABILITY & IMPLEMENTATION ISSUES

The previous two sections discussed the state of the art in control systems for realising a common outcome from a group of distributed network resources. Both the techniques discussed have been implemented in real-world trials, and it is worthwhile discussing some of the common interoperability and implementation issues encountered in these trials.

5.1. Intelligent local devices- the ability to model and plan

One of the key components needed for operation of both the market and cap based coordination techniques is for local resource agents (say, loads or generators) to be able to model and plan their behaviour. In a market based scheme such a model is needed to evaluate the cost one is prepared to accept for a given action, whilst in the cap scheme a model is needed so the agent can submit a plan of its future consumption. Given the dynamics involved, formation of such a model may not be a trivial process. For example, consider the situation of a resource agent being associated with a large cool room. Refrigeration plant such as cool rooms makes up a very significant percentage of Australia’s electricity load, and is thus an ideal candidate for dynamic control. Most importantly, cool rooms have significant thermal mass, meaning that they are essentially a discretionary load- they can be turned off for a period of time, with little effect on the operation of the cool room, but potentially great benefit during times of network constraint. To participate in a market or cap based coordination technique as described in this paper, the resource agent controlling a cool room will need to plan operation of the cool room for some time in to the future. To do this, the resource agent will need a model of the cool room, so it can determine when the refrigeration plant will need to run in order to maintain the cool room’s temperature within given boundaries. Such a model must be dynamic- it should cope with different stocking conditions of the cool room, and will need to consider ambient weather conditions, heat loads and so on. We use so-called *machine learning* techniques to learn this model of the cool room, which are essentially a “black box” learning technique- we are able to form a model of the cool room’s behaviour with minimal understanding of the internal operation, or first-principles characteristics of refrigeration plant. More specifically, we use a support vector machine (SVM) based learning method, a technique which has been of significant interest in recent research publications- see for example [10].

Basically, the SVM model “watches” the cool room’s behaviour during normal operation, collecting several heating/cooling cycles worth of temperature and refrigeration plant (on/off) data. This data is used to train a learning model of the cool room, which essentially finds a characteristic system temperature curve. This model can then be time-stepped into the future, providing accurate temperature predictions of the system.

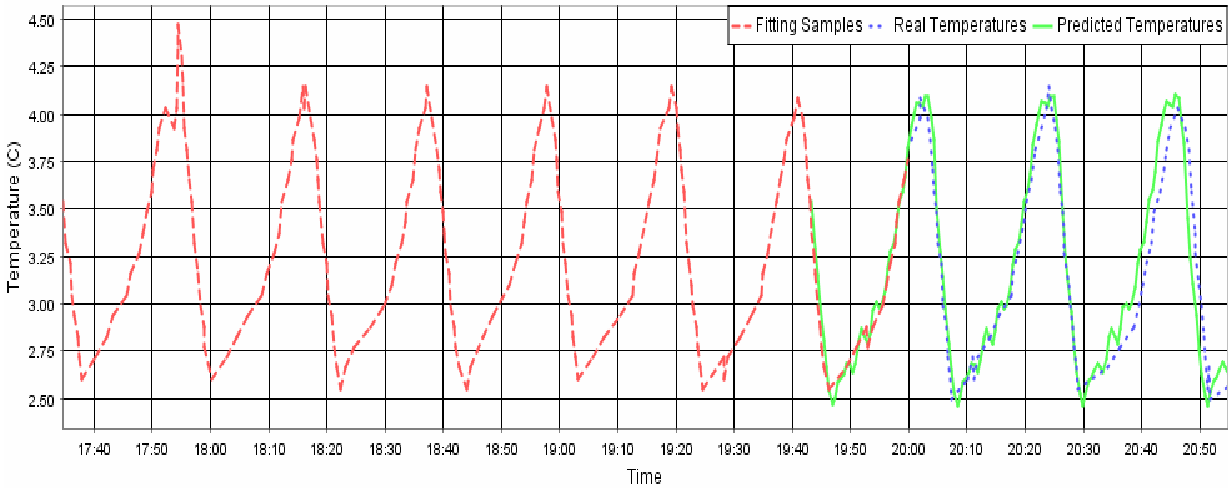

Figure 4: Operation of the machine learning technique for determining a thermal model of a cool room

This learning behaviour can be seen in figure 4, which shows the training samples, the predicted and actual temperatures for a period of cool room operation. It is interesting to observe the behaviour of the system at time 17.54, when the cool room door being opened causes a spike in the internal temperature. The learning model was able to identify this sample as having a minor effect, and so the contribution of this training data to the predictive model is minimal. This is a key advantage of SVMs- their ability to intelligently filter outlying data points, and model non-obvious system subtleties like overshoot and leakage which affect the fitting samples.

The types of models discussed in the preceding paragraphs are critical to the operation of an intelligent and dynamically reactive electricity control system. As another example, consider a renewable energy generator such as a wind or solar plant- for such a generator to participate in coordination systems such as those introduced in this paper, it will need to determine the electrical power it can contribute some time in the future. This is a challenging task for renewable energy systems with intermittent sun or wind availability, and CSIRO has spent a significant amount of time working on machine learning techniques that can autonomously form models of dynamic systems such as renewable generators or thermal loads. Details of these techniques can be found in [11].

5.2. Implementation Challenges

A key feature of the learning and modelling techniques described in the previous section is the need for a reasonable computational ability at each resource agent, so the agent can run these modelling algorithms. Our experience is that a variety of economical and reliable controllers are now available for associating with plant such as generators or cool rooms. We have experimented with a variety of computing platforms for running these models, from thin-client based devices, to personal digital assistant (PDA) type platforms.

Another challenge to the implementation of distributed systems such as discussed in this paper is the need for a communications network to link the various agents in the system. At face value such a requirement does not seem particularly arduous- reasonably reliable, high throughput communication networks are almost ubiquitous now. However, the practical implementation of such systems has proven challenging- we have encountered issues such as:

- Maintaining connectivity through different corporate firewall systems
- Given the plethora of communication platforms currently available, it is difficult for utilities to invest in a given technology with any confidence, particularly considering the long (10 year plus) investment cycles typical to electricity networks
- Ensuring the multi-agent system performs reliably and intuitively when faced with the brief but common communications outages typical to modern Internet protocol (IP) based communications systems

These challenges are gradually being mitigated by recent standardisation activities focussed on introducing reliable, ubiquitous and economical communication systems targeted at electricity network operation and control. For example, the recent IEC61850 standard is aimed at applying common IP based communications techniques to the control of electricity network infrastructure, but with the necessary reliability and robustness built in [12]. Another relevant standard to our work is the Standards Australia standard AS4755, targeted at creating a standardised communications system for the control of distributed energy devices [13]. In the scenarios envisaged by the IEC61850 and AS4755 activities, the resource agents as discussed in this paper might reside on a smart meter appliance, or home "gateway" product, thus addressing the communications and computation functionality requirements discussed above.

6. CONCLUSION

First-generation demand-side control systems are being rolled out in electricity networks across the world as a way of improving network reliability, managing operating cost and infrastructure investment. Whilst it is certainly encouraging to see these systems and the benefits they bring, such systems have a number of drawbacks related to flexibility, consideration of local user constraints, and available firmness.

This paper introduces two new techniques being studied by researchers for more optimal control of demand side resources such as electricity loads and distributed generation plant. These techniques are based around the decentralised control of such plant- there is no centralised hierarchical control system managing individual system devices. Rather, individual devices are controlled by agents, and a system of agents negotiates amongst themselves on how to achieve a desired outcome, with known firmness, and whilst considering local constraints.

The market and cap based coordination mechanisms introduced in this paper have both been trialled in real-world situations, with encouraging results. Importantly, such systems require relatively advanced computational ability at a local load or generator for forming predictive models of that resource's behaviour, and communication networks that can facilitate the inter-agent negotiations necessary to meet a system request. Recent standardisation work, and the ongoing growth of cheap, ubiquitous computing and communications networks means that these are not particularly difficult requirements; we thus look forward to growth in the uptake of these intelligent control systems in years to come.

References

[1] J. S. Thorp, X. Wang, K. H. Hopkinson, D. V. Coury, R. Giovanini, "Agent technology applied to the protection of power systems", Proceedings of Securing Critical Infrastructures conference, October 2004.

[2] J. M. Solanki, S. Khushalani, N. N. Schulz, "A multi-agent solution to distribution systems restoration", IEEE Trans. Power Systems, vol. 22, no. 3, August 2007.

[3] T. Nagata, Y. Tao, H. Sasaki, H. Fujita, "A multi-agent approach to distribution system restoration," in Proc. IEEE/Power Eng. Soc. General Meeting, July 2003.

[4] F. Ygge and H. Akkermans "Power Load Management as a Computational Market", Proceedings of ICMAS, 1996.

[5] J. M. Akkermans, J. F. schreinemakers, J. K. Kok "Emergence of Control in a Large-Scale Society of Economic Physical Agents", Proceedings of the AAMAS'04 conference, 2004.

[6] I. g. Kamphuis, J. K. Kok, C. J. Warmer, M. P. F. Hommelberg "Massive Coordination of Residential Embedded Electricity Generation and Demand Response Using the PowerMatcher Approach", The 4th International Conference on Energy Efficiency in Domestic Appliances and Lighting, June 2006.

[7] K. Kok, C. Warmer, R. Kamphuis, P. Mellstrand, R. Gustavsson, **"**Distributed Control in the Electricity Infrastructure", Proceedings of the 2005 Future Power Systems conference, November 2005.

[8] R. Kamphuis, F. Kujiper, C. Warmer, M Hommelberg, K. Kok, **"**Software agents for matching of power supply and demand: a field-test with a real-time automated imbalance reduction system", Proceedings of the 2005 Future Power Systems conference, November 2005.

[9] M. Hommelberg, C. Warmer, I. Kamphuis, K. Kok, G. Schaeffer, "Distributed Control Concepts using Multi-Agent technology and Automatic Markets: An indispensable feature of smart power grids", Proceedings IEEE Power Engineer Society General Meeting, June 2007.

[10] R. v. B. U. Thissena, A.P. de Weijerb, W.J. Melssena, L.M.C. Buydens, "Using support vector machines for time series prediction," Chemometrics and Intelligent Laboratory Systems, vol. 69, 2002.

[11] G. Platt, G. C. James, J. Wall, S. West, "The Intelligent Control of Distributed Energy Networks", Proceedings of the Energy, Informatics and Cybernetics Conference, 2006.

[12] R. E. Mackiewicz, "Overview of IEC61850 and benefits", Proceedings IEEE Power Engineer Society General Meeting, June 2006

[13] See AS4755-2007, available from http://www.saiglobal.com

Biography

Having completed a PhD in telecommunications and worked for a variety of companies from engineering consultancies to Nokia Mobile Phones, Glenn now works as a project leader for CSIRO's Division of Energy Technology. There, he runs a research programme focussing on applying innovative information and communications technology to improving the way we distribute and utilise energy.