

COORDINATION OF DISTRIBUTED ENERGY RESOURCE AGENTS

Jiaming Li, Geoffrey Poulton, and Geoffrey James CSIRO Information and Communications Technology Centre, Locked Bag 17, North Ryde NSW 1670, Australia

Email: jiaming.li@csiro.au, geoff.poulton@csiro.au, geoff.james@csiro.au

Abstract This paper describes our research in technologies for the management and control of distributed energy resources. An agent-based management and control system is being developed to enable large-scale deployment of distributed energy resources. Local intelligent agents will allow consumers who are connected at low levels in the distribution network to manage their energy requirements and participate in coordination responses to network stimuli. Such responses can be used to reduce the volatility of wholesale electricity prices and assist constrained networks during summer and winter demand peaks. In our system, the coordination of energy resources is decentralized. Energy resources coordinate each other to realize efficient autonomous matching of supply and demand in large power distribution networks. The information exchange is through indirect (or stigmergic) communications between agents. The coordination mechanism is asynchronous and adapts to change in an unsupervised manner, making it intrinsically scalable and robust.

1 INTRODUCTION

With the increasing gap between electricity supply and demand, the electricity industry in many countries is facing a number of new pressures. Distributed electricity generation technologies together with improved demand-side management techniques have been identified as a possible solution to this challenge [USDE 2000]. The idea of controlling the switching loads and generators to respond to price signals and network constraints is technically achievable and becoming more economically viable for businesses requiring greater supply reliability, flexibility, and lower cost to the consumers.

We have been developing multi-agent technology for the management and control of distributed energy resources [Guo 2005; Jones 2005; Li 2007; Li 2008], aimed at deployment in the Australian National Electricity Market within the next few years. A component of this work is the development of algorithms for coordinating distributed energy resources (DERs) comprising customer loads and generators. Coordination is mediated by local intelligent agents that control each DER, called resource agents, and additional agents that present an interface to the electricity

industry, called broker agents. DERs are coordinated to aggregate sufficient distributed capacity to be of strategic value to market participants such as retailers and network businesses. Such aggregation is a significant challenge, particularly for large numbers of DERs and when centralized control techniques are not feasible.

The main research focuses of improving distribution include: (i) how to coordinate all the agents' actions, (ii) how these agents communicate with one another during coordination, (iii) how the system can be made scalable so that the system can operate effectively even as number of resources increases arbitrarily, and (iv) how the system can be adaptable to non-accurate predictions and unexpected events. Four methods reported in recent literature are currently under trial or in active use.

1.1 Price-Based Control

Indirect control over customers' resources is achieved by asking human "owners" of each resource to respond to a varying, broker-determined price for power [Luh 2003]. Typically the price can be at several discrete levels, for example, low and medium rates that are less than the average retail price, and high and critical rates that are more, and the price is increased sharply at times when a reduction in demand is desired. Customers get advance notice of high prices by one or several means and can then choose which appliance settings to change, if any. The varying price may approximate the wholesale market price to some extent, or it may be based on local network loading, not reflected in wholesale pricing, according to the energy utility or business that is asking services of the broker.

Coordination is achieved through an iterative price-updating process carried out in a distributed and asynchronous manner without accessing others' private information or intruding on their decision-making authority. Although mitigation of peak-period power usage has been reported, there is considerable debate in the electricity industry about the long-term effectiveness of such programs [Hopper 2007], due to disadvantages such as the following.

- Human owners may not exist for some resources, or may not be able or willing to respond when asked, so there is no guaranteed level of system response.
- The process may lead to customer dissatisfaction since it requires effort from them and they are being asked to choose between cost and comfort or convenience.

1.2 Direct Load Control

Direct control over customers' resources can be achieved using one or several "circuits" that allow different categories of household appliances to be switched off by the utility at times of peak demand. Switching may be through a physical circuit, interrupting the flow of electricity, or through a broadcast communications method that activates a local switch at the appliance or circuit board. Often the broadcast signal is delivered by superimposing a communications signal on the energy-transporting

fields using "power-line carrier" technology. Customers receive a discount on their electricity bill, or another kind of reward, that may be based on which appliances are signed up to different circuits, how often those circuits are switched, or on the achieved system outcome which is what generates value for the utility. This kind of program has been extensively used for many years to control hot-water systems. More sophisticated forms of direct load control are now being trialled using a wider range of appliances [Energex, 2007]. The disadvantages are as follows.

- Insensitive intervention in the operation of some appliances can cause significant inconvenience to resource owners.
- There remains considerable uncertainty in the level of system response obtained, due to the lack of state information describing which appliances are on at the time of intervention.

1.3 Market-Based Control

Agent-based market-oriented algorithms [Carlsson 2007; Clearwater 1996; Dinias 2005; Kamphuis 2006; Kok 2005; Oyarzabal 2006; Ygge 1998; Ygge 2000], with real or virtual currency, have one or more broker agents to carry out a negotiation process with each resource agent to fix usage and price. Generally, market algorithms for solving flow-resource problems have two scalability problems: one regarding the number of participants in the market and the other regarding the inter-dependency in the participant's demand over time.

The first agent research applications and simulations carried out under the heading of market-based control were brought together in [Clearwater 1996]. Most early research was aimed at climate control in office buildings with many office rooms, where local control agents compete in the allocation of cool/hot air [Hudson 1999]. Then, a system-level theory of large scale intelligent and distributed control was formulated [Kok 2005]. This theory unifies microeconomics and control theory in a multi-agent theory. Kamphuis introduces the PowerMatcher algorithm [Kamphuis 2006], which is a market-based control concept for supply and demand matching in electricity networks with a high share of distributed generation. Real-time matching of supply and demand is crucial to the safe and reliable operation of electricity networks because electricity cannot be stored in sufficient quantity, and with sufficient speed, to absorb imbalance between production and consumption. The most successful agent-based market algorithm for power load management was published by Ygge [Ygge 1998; Ygge 2000]. Like other algorithms described above, Ygge only tried to solve the first scalability problem regarding the number of participants. In their solution to the problem, the demand functions of the individual agents are aggregated in a binary tree. This opens the possibility for running the optimization distributed over a series of computers in a network in a way that fits nicely to power systems architectures [Ygge 1998]. Although some performance advantages have been reported, this method also has disadvantages, which are as follows.

- There is a lack of simple scalability - existing market-based algorithms require hierarchies of brokers to negotiate with very

large numbers of resources, leading to potentially fragile structures.

- Market-based algorithms also require adaptation or replication to account for relationships between resource controlled actions at different times arising from their physical properties.
- Although the efficiency of market-based algorithms may be quantified there is no reported guarantee of an adequate level of service at resource or system level.

The second scalability problem, the one regarding the inter-dependency in the participant's demand over time, is harder to solve in such a way that the usability in the power field remains intact. One way of dealing with this problem is to ignore it and just suppose there is no inter-dependency between electricity used in different time periods. Then, a single-commodity market algorithm can be used, where the commodity is the amount of energy to consume in one time period. Then, the trading agents must totally rely on market price predictions in order to utilize flexibility in their demand over time. On the other end of the scale one could consider a multi-commodity market algorithm in which agents can formulate demand functions that are fully inter-dependent among the commodities, which are amounts of energy to consume in a series of consecutive time periods. This scalability problem was partly solved by Carlsson and Anderson who propose a market algorithm that can handle demand functions which are tree-structured in the time domain [Carlsson 2007]. Agents are able to express dependencies between bids in different time periods, but in a limited number of ways.

1.4 Planning Algorithms

Planning algorithms [Clement 2003; Clement 2000; Guo 2005; Müller 2001] for coordinating a group of distributed energy agents have been developed. An early method for a distributed energy management system based on offline planning was introduced in [Müller 2001]. This has a co-generation system with different generating units and energy storage mechanisms. It uses short term optimization with the aim of minimization of the operating costs based on forecast functions. It is actually a "top-down" centralized algorithm. Then, Clement and Barrett [Clement 2003] introduce the decentralized shared activity coordination (SHAC) algorithm, which negotiates the scheduling and parameters of shared activities until consensus is reached. Protocols are defined which determine when to communicate, what to communicate, and how to process received communication. Distributed energy resource agents coordinate their plans by establishing consensus on the parameters. Protocols are the mechanisms assigned to each agent that allow the agents to change constraints on the shared activity. Since the protocols are pre-defined the algorithm has difficulty adapting to some emergent system behaviours. As well, for a system with a large number of agents, it is hard to establish consensus among agents within a short time period. A coordination algorithm using summary information has been illustrated in [Clement 2000]. The summary information is used to guide the search for

a global plan that resolves conflicts and optimizes the total completion time of the agents' plan. It is shown that summary information can find solutions at higher levels exponentially more quickly than at lower levels. Even so, the algorithm still lacks scalability because the summary information grows exponentially with increasing numbers of agents. For very large numbers of agents the search is time constrained. Recently, Guo [Guo 2005] developed a planning algorithm for coordinating a group of distributed energy agents. The algorithm combines predicted environmental conditions, models for the constraints and behaviour of loads and generators, and a system goal to calculate plans for each resource for a period into the future. Each plan is a state sequence, for example, a set of switching actions and times that an agent will carry out in the future. A centralized genetic optimization algorithm was used in [Guo 2005] to simultaneously calculate the plans for each resource. Although the plan can coordinate distributed agents under ideal situations (e.g. accurately predicted environment conditions and no sudden change for any agent), the disadvantages are as follows.

- Lack of scalability to large numbers of resources - as the behaviour of all agents was optimized centrally for a particular set of events, the solution was not expected to scale well, particularly as genetic algorithms are used as the optimization tool and the assembly of agents must satisfy system global as well as local goals.
- No adaptability to changes in either local or global conditions - sudden changes in the situation of one or more agents are not anticipated or accounted for. For example, if large quantities are added to or removed from a cool room; the whole system would require re-optimization.

1.5 Challenges

In summary, in almost all distributed energy resource management and control algorithms machine learning technology has been used to optimize the plan to solve a given task. Two challenges exist for machine learning planning algorithms. One is scalability, which is a problem for many multi-agent learning techniques. The dimensionality of the search space grows rapidly with the complexity of possible agent behaviours, the number of agents involved, and the size of the network of interactions between them. This search space grows so rapidly that it seems clear that one cannot learn the entire joint behaviour of a large, heterogeneous, strongly intercommunicating multi-agent system. The other challenge is adaptability. Multi-agent systems are typically dynamic environments, with multiple learning agents competing for resources and tasks. Such dynamics present a unique challenge not normally found in single-agent learning: as the agents learn, their adaptation to one another changes the world state. How do agents learn in an environment where the goalposts are constantly and adaptively being moved? These dynamics also present the interesting problem of quality assessment. In a decentralized domain, such quality assessment is relative to or in the context of other agents in

the environment. Thus there may be no absolute quality measure that can be assigned to any one agent.

In this paper we introduce a distributed multi-agent algorithm which coordinates distributed energy resources by attempting to enforce a time-variable supply cap on the power drawn from the grid. The information exchange is through indirect (or stigmergic) communications between resource agents and one or more broker agents. The coordination mechanism is asynchronous and adapts to change in an unsupervised manner, making it intrinsically scalable and robust. In the system, individual agents are selfish and reasonably simple. However, the desired (complex) system response emerges out of low level agent coordination, which is in stark contrast to traditional centralized control systems. This work will bring potential solutions to the volatility of wholesale pool prices and an alternative way of dealing with network constraints during summer and winter peaks.

The present algorithm overcomes all of the difficulties mentioned above in regard to existing methods. Specifically, no human action or effort is required at the resource level; solutions implicitly include the satisfaction of the local constraints of resources and also offer system-level users a defined service, the reliability or "firmness" of which may be quantified; the system is scalable to both very large numbers of resources and inter-dependency in the participant's demand over time, even with a single broker; resource agents act autonomously, so no central adaptation or replication is required when conditions change.

The paper will be organized as follows. Section 2 suggests an approach to coordinating distributed energy resource agents using indirect communications mediated by a "stigspace". Section 3 describes the coordination algorithm. Section 4 demonstrates the performance of the algorithm through the results of simulation experiments, and Section 5 analyses the convergence of the algorithm through comparison with theoretical limits of performance.

2 INDIRECT COMMUNICATION BETWEEN DISTRIBUTED ENERGY RESOURCE AGENTS

2.1 Resource Agents

Local intelligent agents are a natural means to manage quite complex data and control action for individual DERs while providing a simple interface by which the DER interacts with the energy system at large. Considering the listed disadvantages of the approaches discussed above, and wishing to find an approach that offers both scalability of the number of resources and adaptability to possible sudden changes in the situation of one or more resources, we suggest that the following properties of multi-agent systems tend to promote both scalability and adaptability.

- Agents should be as simple as possible (naive) regarding their interaction with the agent system. Here we distinguish between an

agent's function as a member of a multi-agent system and an agent's local resource management function. The latter is generally not simple because it concerns engineering details of the load or generator being managed and the requirements and preferences of the customer who owns it.

- Agents should satisfy local goals preferentially (selfish). As well as being simpler to design than agents that must simultaneously satisfy goals relating to the system, selfish agents will always ensure that customer requirements are met as far as possible without violating local goals. This will assist in successful adoption of the technology. System goals can be met through design of agent interactions and responses.
- Agents should be identical (or of a few varieties only). Since only a few agent designs are needed the number of optimisation parameters is small and does not increase with the number of resources. This will lead to faster, scale-independent, system design. It should also be relatively simple to add resources, keeping the same agent design.
- There should be little or no inter-agent communication (non-communicating). Lack of inter-agent communication will lessen the likelihood of unexpected (unplanned emergent) behaviour.

We have developed, in simulation and also in hardware for demonstration, resource agents that satisfy these properties. Their purpose is to fulfil the requirements of the electricity customer who owns the resource under control. They are nevertheless able to respond to information about the multi-agent system and any system goals, provided that this response doesn't compromise the customer's requirements and thereby constitute a cost. The resource agents used in this study control refrigerators and the customer requirement is to maintain temperature within normal operating bounds. In steady-state operation, without any changes in contents, door openings, or multi-agent responses, this results in a square-wave power consumption as the compressor is turned on and off when the internal temperature reaches upper and lower temperature bounds respectively. These resource agents, therefore, take over the control of the compressor and have the ability to change switching times to effect a response to information about the multi-agent system. FIGURE 1 shows steady-state switching of a real refrigerator under agent control.

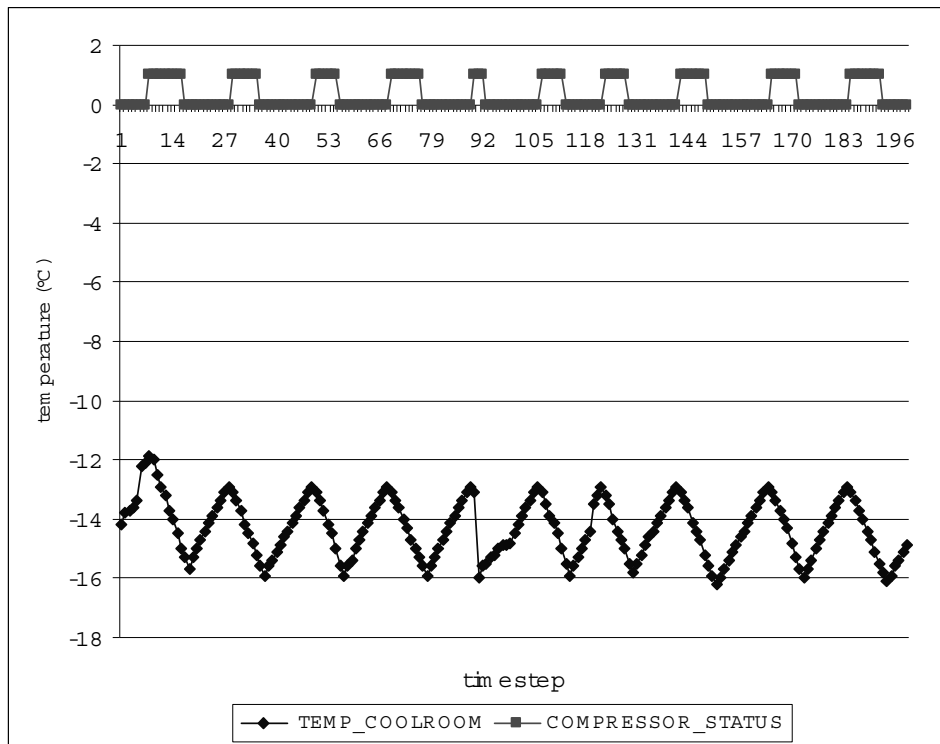


FIGURE 1: Steady-state behaviour of a real refrigerator under agent control. The upper graph shows when the compressor is on and off, the lower graph shows the corresponding internal temperature of refrigerator, and the horizontal scale is time in steps of 10 s.

The information communicated with the multi-agent system is chosen to address the second kind of scalability discussed in Section 1.3: it is necessary to account for the inter-dependency in the resource agent's energy demand over time, but in doing so we should avoid replication of algorithm function. We provide each resource agent with the ability to plan its energy consumption for a nominated period into the future. This requires it to model its own physical properties such as thermal mass, hysteresis, and compressor power. Our simulation and hardware demonstration agents have this ability, and significant complexity is required to fit models to measured behaviour, to use models to plan future switching actions, and to reconcile the execution of planned behaviour with real-time control of the refrigerator under varying conditions. This complexity is hidden from other agents, however, and the only information communicated in the multi-agent system is the planned power consumption of each agent for a planning time T_{plan} into the future. Since this paper concerns multi-agent behaviour the formulation and fitting of models will not be discussed in detail here.

These properties may be generalized to other kinds of physical resources including those that have continuous control rather than on/off switching. For example, heating, ventilation, and air conditioning in a building is more complex than operating a simple refrigerator, nevertheless, each zone of control may be modelled to allow the forecast electricity demand to be estimated according to control settings such as temperature set-points and fan speeds. It is also possible to incorporate

generating resources, even those that provide no opportunity for control such as renewable generators without the ability to dump or store power, by using weather forecasts and other data to predict output which is represented in a plan as negative demand. For demand management it is beneficial to consider heating and cooling loads because thermal inertia allows considerable flexibility in the control of these loads. They are also a very significant component of the total demand, for example, space heating/cooling, water heating, and refrigeration together comprise 75% of the energy use of an average Australian home and in most homes this is all electrical energy [AGO 2007].

In following, we consider a form of active stigmergy [Stone 2000]. In our simulation and hardware demonstration systems we implement a simple model (stigspace) using a bulletin board hosted by a single computer. A full stigmergy-based scenario can be implemented, if agents leave planning data locally within a communication network and the broker agent is able to access the network at any node, and aggregate necessary information by utilising a service discovery protocol within the network. In other words, stigspace would cover the whole network, agents would communicate with the stigspace locally, and the aggregation mechanism (needed by the broker agent) would be implemented within the network layer (e.g, using a scheme similar to directed diffusion, or Gradient-based cost fields [Estrin 1999; Ye 2001]). The current implementation (stigspace as a bulletin board) is viable as long as there are no communication bottlenecks. More complex stigspaces, where agents place and search for messages within a distributed region, are also possible and may have advantages for hierarchical systems and systems in which time-varying agent clusters may form for increased performance.

2.2 Broker Agents and Summarizing Agents

The above properties refer to resource agents. There may also be one or more broker agents to manage the interface with the electricity network and market. A broker agent

- receives information on predicted market and network usage and prices,
- interacts with resource agents through stigspace, where it can read and place information,
- constructs global goals, such as grid supply "cap" for a certain period of time, using market and predicted usage data, and
- may also act as the stigspace manager.

Regarding the final point, the role of stigspace manager is to act on information in stigspace to produce derived information; for example, the predicted total resource agent demand as a function of time may be derived from the planned power consumption of all the resource agents that submit their plans to stigspace. This paper assumes the summarizing agent acts for the broker agent, but other options, such as an independent summarizing agent, are possible as well. Derived information is placed in stigspace for resource or broker agents to access.

To limit the amount of data that must be stored in stigspace we introduce time bins (or intervals) of duration T_{bin} . Resource agents' output power plans are averaged into time bins to produce step functions rather than continuous functions of time. All resource agents use the same time bins to make the production of derived information possible and efficient: agent plans may be compared directly within each time bin. This adds a requirement that the resource agents must have synchronized clocks, which may be achieved with acceptable accuracy using the clocks provided on standard computing hardware, provided that T_{bin} is not a small fraction of a second. In our simulation experiments the smallest time bin used had $T_{bin} = 1$ min. A consequence of binning agents' plans is that adherence to the desired total output power is only measured with resolution T_{bin} and any smaller-scale variations cannot have any influence on the algorithm. This is consistent with the operation of electricity markets that have an ancillary market to provide real-time balancing of supply against demand within each generation dispatch interval. In the Australian market the dispatch interval is 5 min so as long as $T_{bin} < 5$ min we can be reasonably confident that the ancillary market will deal with smaller-scale variations in total resource-agent demand.

3 COORDINATION ALGORITHM

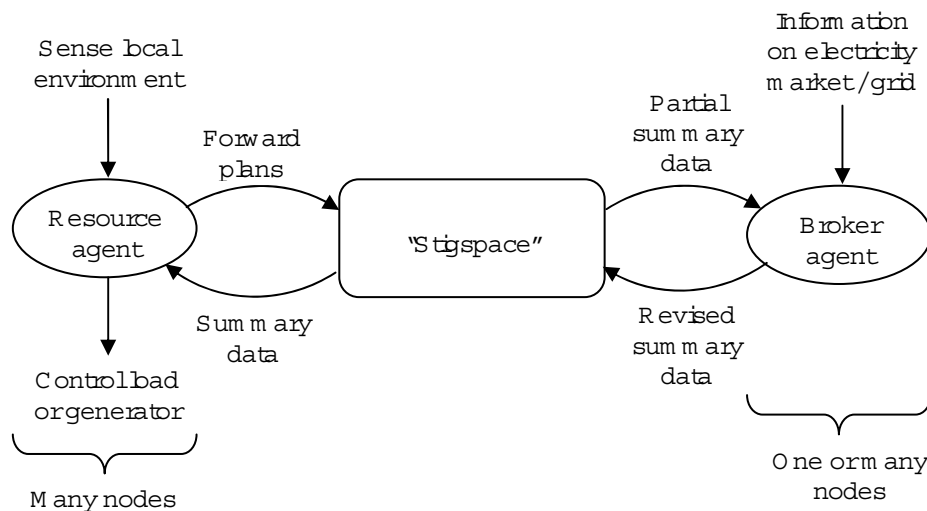


FIGURE 2: The coordination system. Each resource agent is the decision maker and controller of the resource, able to sense local conditions and plan its actions based on an internal model. All the agents' plans are sent to stigspace, a medium for indirect communication, in which summary data are computed from the plans and are available to all agents. In particular, the broker agent uses summary data and market or grid information to set a cap on the total demand for power drawn from the grid based. This supply cap is placed in the stigspace and is then available to resource agents, which can revise their plans to help satisfy the cap while continuing to adhere to their local constraints.

FIGURE 2 is the diagram of our coordination system. The resource agents have information about local constraints imposed by the electricity customer who owns the "distributed energy resource" (load or generator). At convenient intervals, timed to allow convergence before the next market cycle, an iterative process begins. Each resource agent applies their local constraints to a physical model of their resource to calculate a plan for electricity demand or supply for a period into the future. These plans are transferred into market cycles, and then sent to stigspace.

In stigspace the plans are summed to get the total predicted power demand in each interval. This is then available to any resource agents that wish to use it and also to the broker agent. The broker agent has knowledge of predicted electricity market price as well as information about the plans communicated by participating resource agents. The broker acts for electricity market participants, such as retailers and network operators, who provide additional information leading to a desired cap on the total demand and for power drawn from the grid. This supply cap is placed in the stigspace and is made available to any resource agents that wish to use it.

Although the broker agent has no direct control over resources, resource agents agree to satisfy any global goals as long as local goals remain satisfied. Therefore, the resource agents, when they retrieve the total demand and supply cap from stigspace, can revise their plans using our CordCap algorithm to help satisfy the cap while continuing to adhere to their local constraints. This process is iterated until it stabilizes. By submitting revised plans they participate in a real-time process which is asynchronous: no explicit coordination is needed between plan submission, plan summing, and broker action.

When the total demand is stable, the broker agent is in a position to buy power for the next time period. The process is repeated for every market cycle. The heart of the algorithm, which allows the coordination process to be scalable and adaptable, lies in the broker-derived features and the means of communication between broker and resource agents.

In our system, the constraints for resource agents are temperature bounds for a heating/cooling environment; the plan calculated for electricity demand or supply is for future half hour, i.e., $T_{plan} = 30$ min; these plans are transferred into average power demand or supply in each interval, e.g. $T_{bin} = 5$ min.

3.1 The Electricity Market in Australia

In Australia, the National Electricity Market Management Company [NEMMCO 2007] has been established to manage the operation of the wholesale electricity market and security of the power system. Our homes, businesses and industries depend on a reliable supply of electricity to function. NEMMCO plays a central role in ensuring South-Eastern Australia's electricity supply through its responsibilities as market and system operator of the National Electricity Market (NEM). Within NEM producers submit bids stating the amount of energy they can generate at

what cost and consumers submit predictions for consumption. These are matched centrally, the lowest cost producers are instructed to supply energy, and a single price is set for all participants. This planning process is based on short-term forecasts of the volume of energy required over the next 24 hour period. Generators are scheduled in 5 minute dispatch intervals. Prices are set for each dispatch interval and provide a signal by which consumers can manage their individual demand. Electricity is charged every half hour on the hour and half hour, e.g., 4:00 am and 4:30 am. These dispatch and prediction intervals dictate the time scale at which electricity management agents can operate. The large volumes of electricity used in the NEM make it impossible to store energy for future use. This means that the NEM is unable to respond quickly to significant unpredicted changes in demand. On the whole, the less oscillation there is in demand, the better. The unpredicted demand in NEM is reflected in NEM MCO peak electricity price.

Significant electricity demand and price information is available from NEM MCO. The main available electricity prices are:

- 5-minute pre-dispatch price. This contains 5-minute pre-dispatch (forecast) data by region, showing short term price and demand forecasts looking out one hour ahead and is updated every 5 minutes. It is usually published one minute before the time of the first prediction price. The 5-minute pre-dispatch price file is in CVS format on the NEM MCO website [NEM MCO 2007].
- 30-minute pre-dispatch price. This is the forecast 30-minute price to the end of the next market day.
- 30-minute trading price. This is the real time 30-minute price, at which retailers buy electricity from the market.
- 5-minute dispatch price. This is the real time 5-minute price, which is averaged to give the 30-minute trading price.

3.2 Using a Supply Cap in the Electricity Market

The broker agent buys electricity from the grid at the 30-minute trading price, which is variable as described above, and sells the electricity to consumers at a consumer price (set by the broker) which is more constant. To maximize its own profit the broker wants consumers to use less power in higher price periods, and more power in lower price periods. To accomplish this outcome, the broker reads 5-minute pre-dispatch prices from NEM MCO to get predicted prices for the next hour. It uses this together with the total predicted demand to set a variable supply cap, in 5-minute intervals, on the total power to be drawn during the next 30 minutes. This supply cap can be a percentage of total power demand.

3.3 The CordCap Algorithm

The CordCap algorithm is used by each resource agent to modify its power usage so as to help satisfy global goals – here the grid supply cap –

and has been designed for agents whose actions are limited to on/off load switching. Although each agent acts independently, the small but significant stigmergic communication drives the multi-agent system toward the global goal as an emergent property. Specifically, the response of an individual resource agent depends on its own local goals, total predicted power needs and the supply cap. As already mentioned, the latter two quantities, specified for each 5-minute interval in the 30-minute planning period, are obtained from stigspace where the data has been placed by broker and summarizing agents. If the supply cap is not satisfied for certain time intervals, each resource agent's switching strategy in those intervals will be updated as described below.

Each agent iterates the process until either (a) local and global goals are satisfied, (b) no further improvement is possible, or (c) a specified time limit is reached. Once this occurs, all agent actions are "locked in" for the next 5-minute interval, the 30-minute planning period is advanced by five minutes and the process begins again.

3.4 Resource Agent Switching Strategy

In the CordCap algorithm the resource agent modifies its predicted switching sequence to shift power consumption from each cap-violating interval into its left and right-hand neighbours. The process includes three steps, and is carried out for all offending intervals.

1. Locate a random point t_x in the interval.
2. Partly shift power usage in the interval on the left and right of t_x into the left and right-hand neighbouring intervals respectively.
3. Revise the resultant switching strategy to satisfy the resource agent's local constraints.

The following is a more detailed description of the algorithm for a resource agent controlling a simple cool-room resource. Agent actions are limited to switching the cooling on and off.

Suppose the agent plans its action for future period of time, e.g. half an hour, $T_{\text{plan}} = 30 \text{ min}$. The bin period is 5 minutes, $T_{\text{bin}} = 5 \text{ min}$, so there will be 6 five-minute bins in the plan period. At regular intervals each cool-room agent applies internal temperature constraints to its physical model to calculate plan actions, which are represented as a state sequence:

$$A = [A_1, A_2, A_3, A_4, A_5, A_6] = [a_{11}, a_{12}, \dots, a_{1N}, a_{21}, \dots, a_{2N}, \dots, a_{61}, a_{62}, \dots, a_{6N}],$$

where $N = T_{\text{bin}} / \Delta t$, Δt is the time resolution for planning, and a_{ij} is the resource's on/off state during the i^{th} time step in the j^{th} bin. The planned actions in each bin, $A_i = [a_{i1}, a_{i2}, \dots, a_{iN}]$, result in the consumption of power $[p_{i1}, p_{i2}, \dots, p_{iN}]$, and these may be summed to compute the average power demand in each bin, $P = [P_1, P_2, P_3, P_4, P_5, P_6]$ where

$P_i = \sum_{j=1}^N p_{ij} / N$, $i=1, \dots, 6$. The sequence of average powers is sent to stigspace. FIGURE 3 shows the time intervals and corresponding state and power sequences.

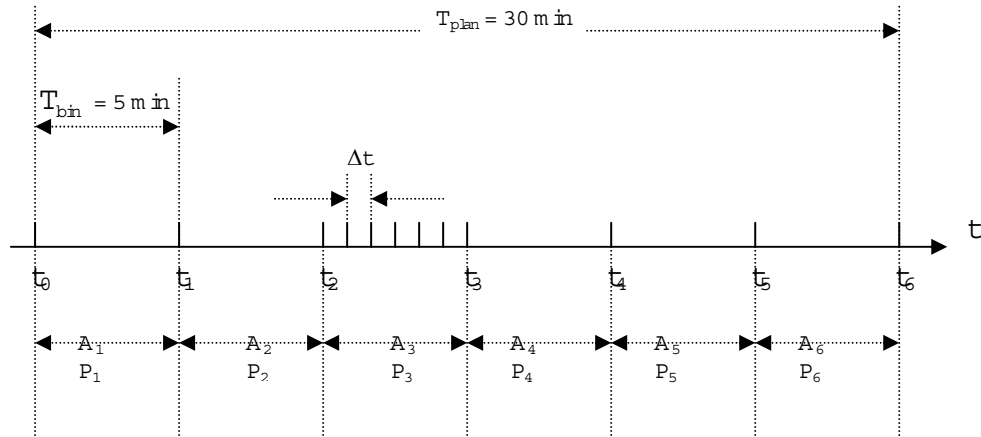


FIGURE 3: Example of a half-hour plan. There are 6 five-minute bins in 30 minutes. In the j^{th} bin the planned actions are state sequences $A_j = [a_{j1}, a_{j2}, \dots, a_{jN}]$, resulting in consumption of power $[p_{j1}, p_{j2}, \dots, p_{jN}]$ and thence an average power in each bin, $P = [P_1, P_2, P_3, P_4, P_5, P_6]$, which is communicated to stigspace.

Using its known resource model, the resource agent calculates its expected power usage during the planning period, taking care to satisfy its local temperature constraints. Each resource agent then examines the total predicted power and the grid cap (obtained from stigspace) and identifies time segments in which the cap is violated. It then modifies its planned power usage (if any) within these segments. The modification depends on whether one or more segments violate the cap, and these two scenarios are examined in some detail below.

3.5 Dem and Shifting Procedure

Suppose there are an unspecified number of resource agents running the algorithm, and there is only one violating interval, such as third interval $i = 3$. FIGURE 4 (a) illustrates the steps taken by a typical resource agent. The grey and black blocks show the time when the power being used. Ψ_{max} and Ψ_{min} are upper and lower temperature constraints of the cool-room and the curve is the predicted internal temperature based on the resource model and currently planned power usage. The bar code indicates an interval for which a cap violation is predicted. The agent follows the following process to help mitigate the problem.

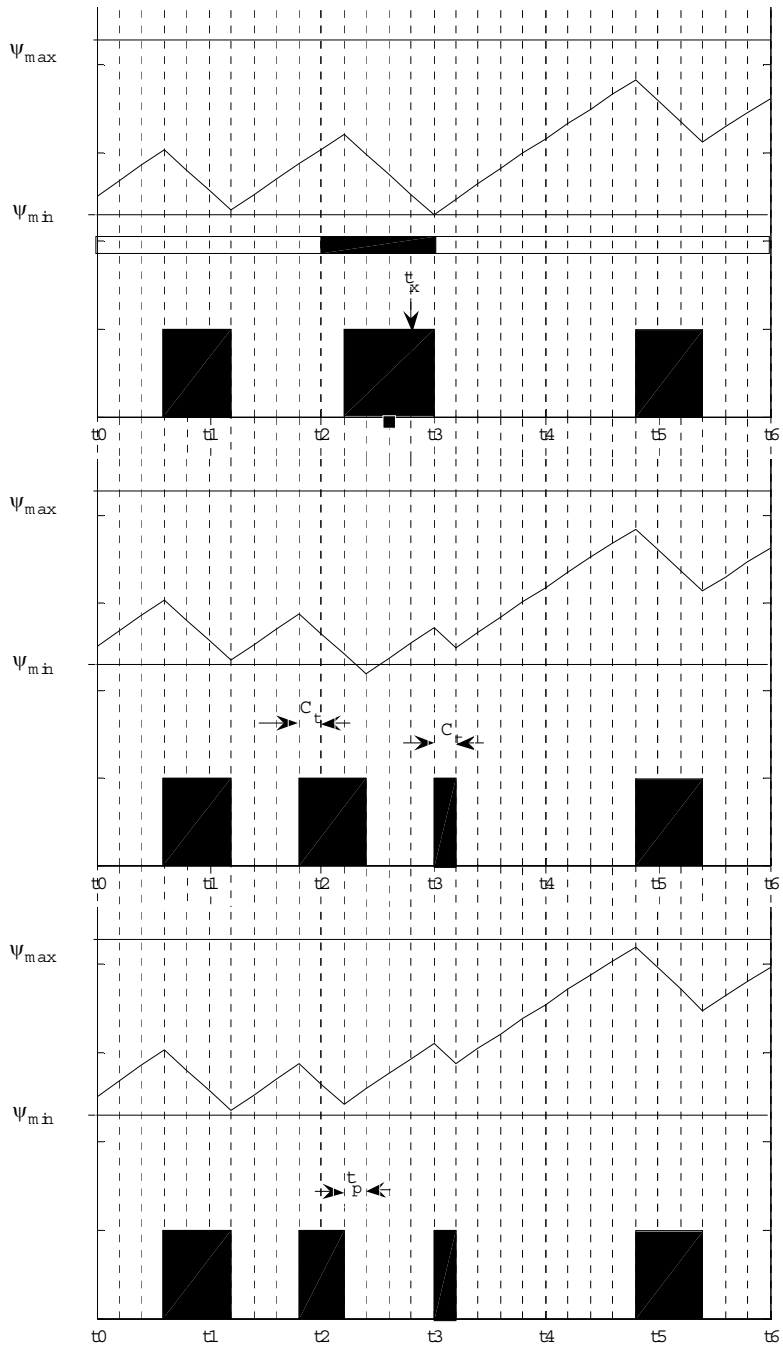
- Step1.1: A time t_x is randomly selected between t_2 and t_3 .
- Step1.2: Power in $[t_2, t_x]$ (if any) is shifted so that its left-hand edge is C_t minutes outside the left interval boundary. C_t defines how much energy in the violating interval can be shifted outside.

The experimental value of C_t in our system is 1 minute. Similarly, power in $[t_x, t_3]$ (if any) is shifted so that its right-hand edge is C_t minutes outside the right interval boundary.

- If the violating interval is $i = 1$, i.e., the period $[t_0, t_1]$, the power outside the left interval boundary t_0 will vanish.
- If the violating interval is $i = 6$, i.e., the period $[t_5, t_6]$, the power outside the right interval boundary t_6 will vanish.
- Step 1.3: Revise the modified switching strategy to satisfy the temperature constraints of the cool-room, using the following rule:
 - If the predicted internal temperature is outside the temperature constraints $[\Psi_{min}, \Psi_{max}]$, leave the plan status t_p minutes (either increasing or decreasing power), then join the plan as soon as possible. t_p is the defined parameter representing how long the resource agent can leave the plan. The experimental value of t_p in our system is 1 minute.
- Step 1.4: Submit the revised forward plan to stigspace in readiness for the next iteration.

If two or more intervals violate the cap, each interval is treated independently as shown in FIGURE 4 (b) for the case $i = 3 \& 4$.

- Step 2.1: Times t_{x1} and t_{x2} are randomly selected in $[t_2, t_3]$ and $[t_3, t_4]$, respectively.
- Step 2.2: The process described in Step 1.2 is applied to $[t_2, t_3]$ and then $[t_3, t_4]$.
 - If 'ON' states overlap the result is 'ON'.
- Step 2.3: The modified switching strategy is revised to satisfy the cool-room's temperature constraints as in Step 1.3.
- Step 2.4: Submit revised forward plans to stigspace as in Step 1.4.



(a)

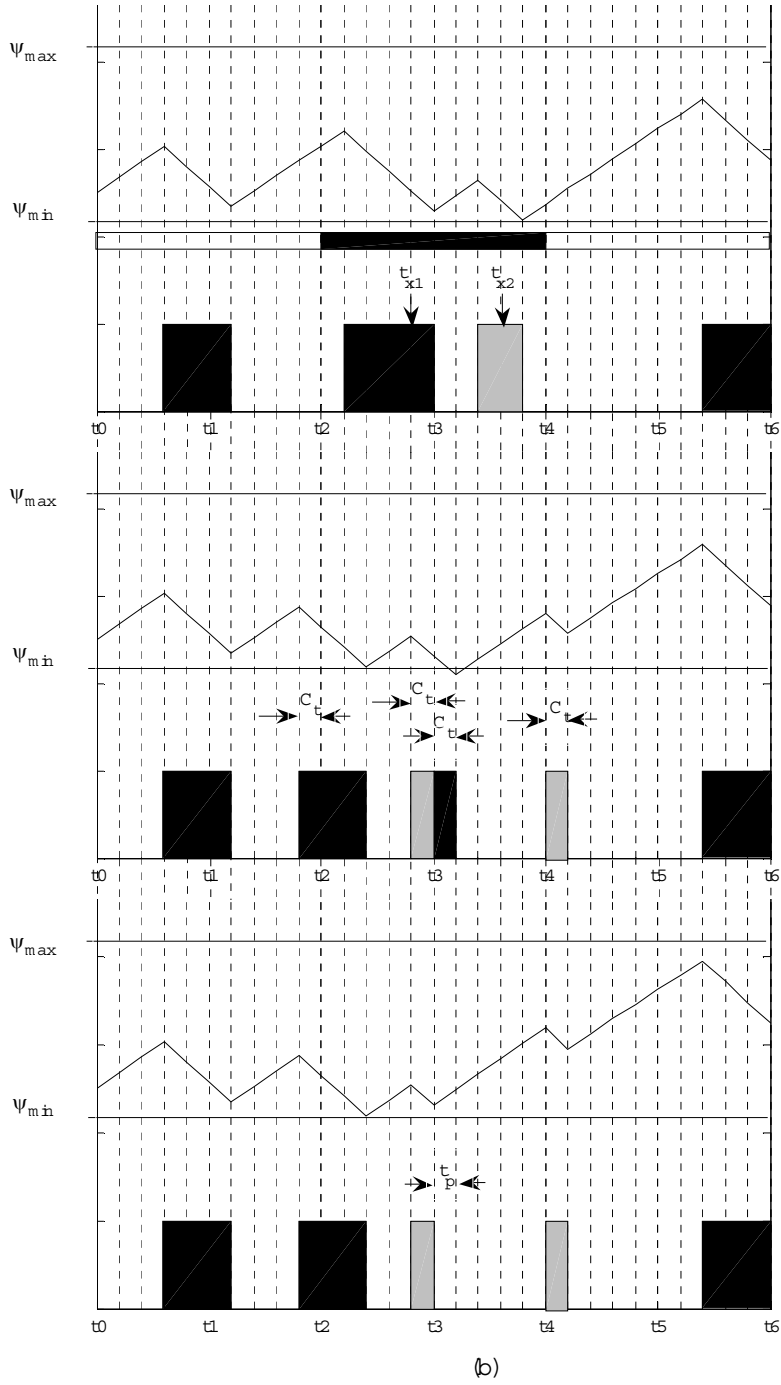


FIGURE 4: Coordinating with a single violating interval (a) and adjacent violating intervals (b). The grey and black blocks show the time when power being used. The bar code indicates the cap violation interval. Ψ_{max} and Ψ_{min} are upper and lower temperature boundaries. t_p is the defined parameter representing how long the resource agent can leave the plan. C_t defines how much energy in violating interval can be shifted outside.

4 SIMULATION RESULTS

We have implemented a system with one broker agent, one summarizing agent, one stigspace and a number of loads - all cool-room agents. The broker agent reads 5-minute pre-dispatch wholesale electricity price from NEM MCO, sets the grid supply cap accordingly and places it in stigspace. Each cool room agent calculates a plan of electricity demand for the next half hour which satisfies its internal temperature constraints. This plan is then transferred into mean power demand in 5 minute intervals and put in stigspace. The summarizing agent sums the mean power demand from all agents in each interval for the next half hour planning period, placing the total demand in stigspace. Cool-room agents then apply the CordCap algorithm to satisfy the system supply cap whilst continuing to adhere to their local constraints as described in above.

A series of experiments has been completed to investigate system coordination performance, which include coordination scalability, the effect of resource agent diversity on coordination performance, maximum system demand reduction for a short period supply cap and continuous coordination benefits for resource agents.

4.1 CoolRoom Model

Cool rooms have internal temperature constraints with boundaries 1° and 6° Celsius. The internal temperature is governed by the model developed in [Clement 2000] with most room features removed for the purposes of these experiments:

$$T_{ai}(t) = \frac{1}{C_a f + k} (C_a f T_{ai}(t-1) + Q_p(t) + k T_{ao}(t)) \quad (1)$$

where T_{ai} is internal air temperature ($^\circ\text{C}$), T_{ao} is outside air temperature ($^\circ\text{C}$), Q_p is the maximum cooling power (kW) of the cool room plant, C_a is the thermal capacity of air ($\text{kJ/g}^\circ\text{C}$) in the cool room, f is the sampling rate (Hz), $k = A_w U_w + U_v$, A_w is a unit-less coefficient, and U_w , U_v are thermal resistances (Ω) of the wall and ventilation path. We modelled non-identical resource agents by allowing different thermal capacities C_a in the room model. In our experiments, $f = 3000$ (Hz), $A_w = 1/58.9$, $U_w = 6$ (Ω), and $U_v = 0$ (Ω).

4.2 Coordination Scalability

We modelled a set of cool rooms with power capacity $Q_p = 3 \sim 6.67$ kW and thermal capacity of air varying within a range given by

$$C_a = C_{ao} (1 + 0.9^n) \quad (2)$$

where $C_{ao} = 8.8964 \times 10^{-4}$ (kJ/g°C) (and $n = \{1, \dots, N\}$ for a set of cool rooms. The thermal capacity is reflected in the typical period of a heating/cooling cycle.

FIGURE 5 shows initial room temperatures and power demands in each 5-minute interval for future half hour periods for a system with 3 cool rooms. The time constants of all cool rooms are very close to each other. A constant supply cap, 3kW, is applied to the system. Before coordination, the supply cap is not satisfied for the time 9:25 to 9:40. After several steps of coordination among resource agents, supply cap is satisfied as shown in FIGURE 6.

TABLE 1 lists coordination performance for the system with constant supply cap and different numbers of resource agents. From the table we can see that when agent numbers increase, the number of coordination steps does not increase accordingly. This indicates that the coordination is scalable for large numbers of agents.

In our simulation experiments, all resource agents were executed on one computer and used non-threaded calculations. In a deployed environment, each resource agent will have a dedicated machine and use threaded calculation. The total time for system coordination is less than 76 milliseconds for a system with 10,000 agents, which is extremely fast. Therefore real time, deployed coordination is certainly possible. The effects of communication speed are ignored in this paper. Preliminary studies indicate that the system will be resilient to this factor, which will be fully analysed in a future paper.

TABLE 1: Coordination times for different numbers of resource agents.

AgentNo.	Supply Cap (kW)	Coordination Steps	Total Time (sec.)	Time per Agent (sec.)
10	10	5	0.341	0.03
100	100	2	0.957	0.0096
1000	1000	1	7.268	0.0073
10000	10000	1	76.129	0.0076

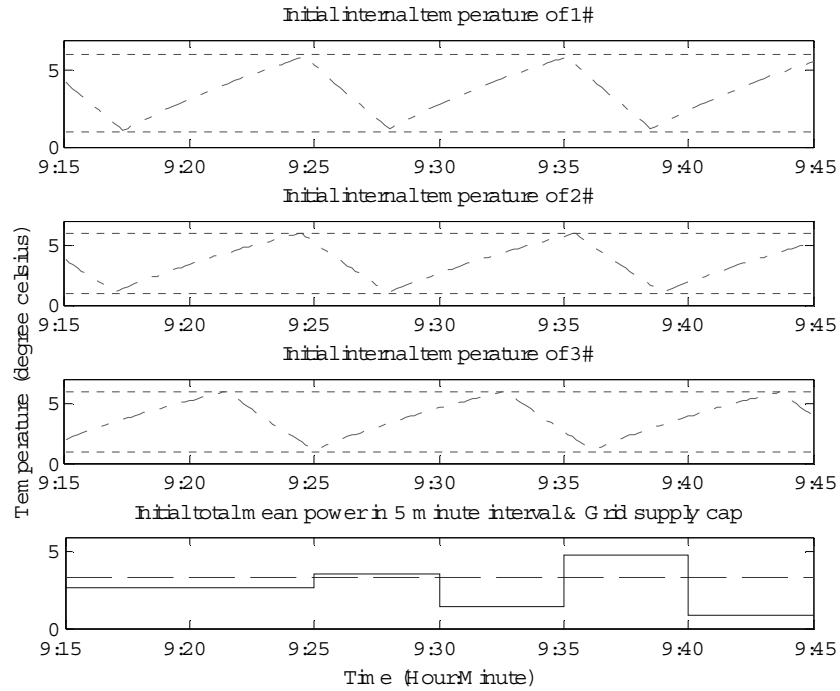


FIGURE 5: Room and system states before coordination (dotted line - cool room temperature constraints; dash-dot line - internal temperature of cool room ; dashed line - the system supply cap, solid line - system total demand).

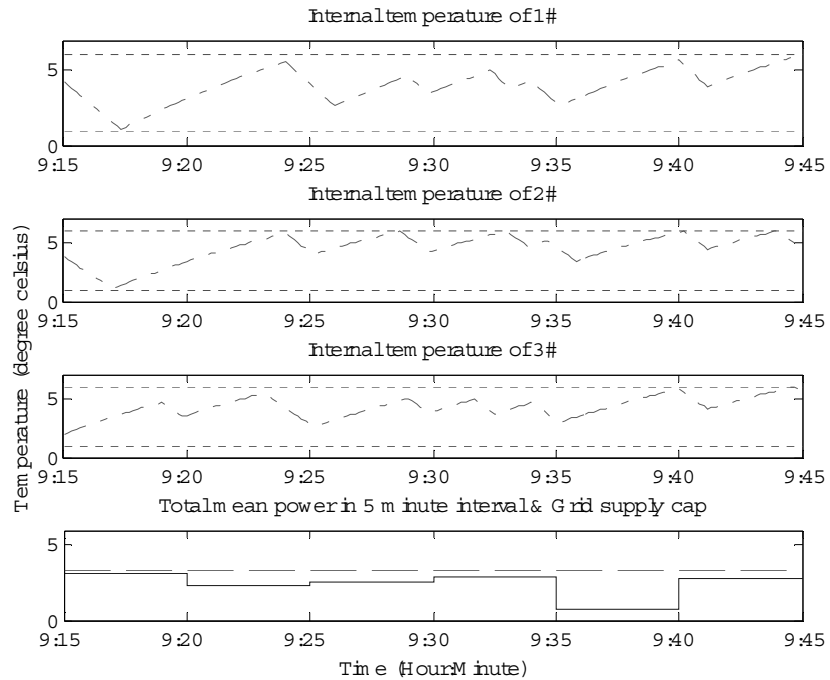


FIGURE 6: Room and system states after coordination. The resource agents revise their plans to help satisfy the system supply cap (solid line is below dashed line) whilst continuing to adhere to their local constraints (dash-dot lines are always between the dotted lines).

4.3 Effects of Diversity

Systems consisting of 10,000 cool rooms with different variation ranges in parameters have been set up to test the effects of diversity on coordination performance. The cool rooms have the same temperature constraint, but diverse power capacities and time constants, random starting internal temperatures and random initial switching states for the cooling plan.

TABLE 2: Caps achieved through coordination of 10,000 resource agents with average compressor power rating of 3 kW .

System	Compressor power (kW)	T_{on} (min)	T_{off} (min)	Cap (kW)	Steps
1	[3, 3]	[9.5, 17.8]	[8.8, 16.7]	10000	6
2	[1.12, 4.88]	[4.83, 10.3]	[4.5, 34.8]	4364	6
3	[0.35, 5.65]	[4.83, 37.3]	[4.5, 35]	3615	6

TABLE 2 lists minimum supply cap achievable for the systems comprising resource agents with different range of parameters, but the same system average power. T_{on} and T_{off} are respectively the turn on and turn off time constants of resource agents, 'Cap' is the minimum system supply cap which could be satisfied, and 'Steps' is the number of coordination steps for resource agents to satisfy supply cap. Bracketed entries are ranges of parameters over a number of cool rooms. From top to bottom, the systems have increasing diversity of resource agents. From the table, we can see that the system with more diverse resource agents will tolerate a smaller supply cap.

4.4 Maximum Demand Reduction for Short-Period Supply Cap

In a deployed environment, resource agents continuously coordinate their plans with each other every 5 minutes. To investigate how the system performs under continuous coordination we carried out a series of tests based on systems 2 and 3 in TABLE 2.

We expect that a supply cap would be useful in two circumstances: when the electricity price is high retailers could request a cap to reduce their expenditure on the wholesale market, and when the physical network is near capacity the network operator could request a cap to ensure a safe operating margin. Electricity price forecasts are published by the market operator, but network capacity may become a problem suddenly due to equipment failure, and then the broker may give only limited advance notice to resource agents before applying the cap. To test the system response to different cap notice times, the following tests have been executed for a supply cap of 15 minutes' duration.

Suppose a 15-minute cap occurs between 9:30 and 9:45. Cap notice times of 5, 10 and 15 minutes were investigated for two of the systems defined in TABLE 2. The demand of system 3 for 15 minutes advance cap

notice is shown in FIGURE 7, where supply cap, total dem and before and after using CordCap algorithm are shown in different style lines respectively. TABLE 3 gives the minimum system dem and that could be achieved for different cap notice times. We can see that different advance notice times give the same achievable minimum system dem and; but the greater the diversity of resource agents in the system, the more dem and reduction can be achieved. This tells us that achievable minimum system dem and for the CordCap algorithm depends largely on the diversity of resource agents. It also indicates a possible limitation of the algorithm for short cap durations. This point will be discussed further in section 5.2 when theoretical limits of dem and response are presented.

TABLE 3: Dem and reduction for different cap notice times.

System	Advance Notice (minutes)	Minimum System Dem and (kW)
2	5	368
	10	365
	15	365
3	5	310
	10	310
	15	310

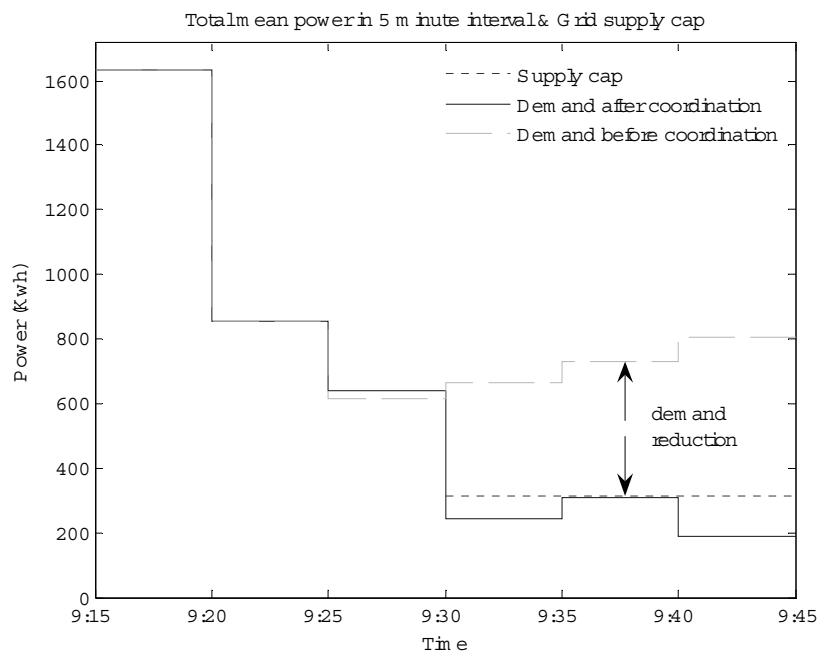


FIGURE 7: 15-minute advance notice for system 3 in TABLE 2. A 15-minute cap occurs between 9:30 and 9:45. Notification of the cap is given at 9:15. The dashed line represents initial (non-coordinated) planned power dem and, which is over the system supply cap (dotted line). The solid line represents coordinated planned power dem and, which is under the system supply cap. Dem and reduction is defined as the average reduction in power dem and between coordinated and non-coordinated situations.

5 CONVERGENCE

5.1 Firmness of Convergence

Electricity utilities will require reliability of convergence in any algorithmic technique to provide a demand-response service. The degree of reliability is called "firmness" in the electricity industry and depends on the reason for calling on demand response. Some electricity markets permit demand bids in the wholesale market, in which case there are well-defined specifications for quality of service that must be met [Nord Pool 2008]; presently the Australian market does not permit such bids. If demand response is an internal capability that the utility uses to manage the demand it presents to the market, the quality of service may be less well defined, and existing methods of demand management provide a variety of levels of reliability. Price-based control depends on voluntary switching of loads by customers and is inherently unfirm for this reason [Hopper 2007]. Direct load control provides a firm time of response by issuing a precise broadcast signal [Energex 2007], although the magnitude of the response depends on the statistical distribution of states of individual customer loads and is therefore less firm without additional data. These extremes present a wide range of useful levels of firmness that is nevertheless useful in assisting utilities to manage their overall demand.

We have studied reliability of convergence through numerical experimentation and comparison against a theoretical limit of coordination performance. FIGURE 8 shows a typical graph of convergence of system power demand for one of the experimental systems in TABLE 2. The demand that exceeds the cap reduces rapidly and in this case becomes zero after six steps of the algorithm.

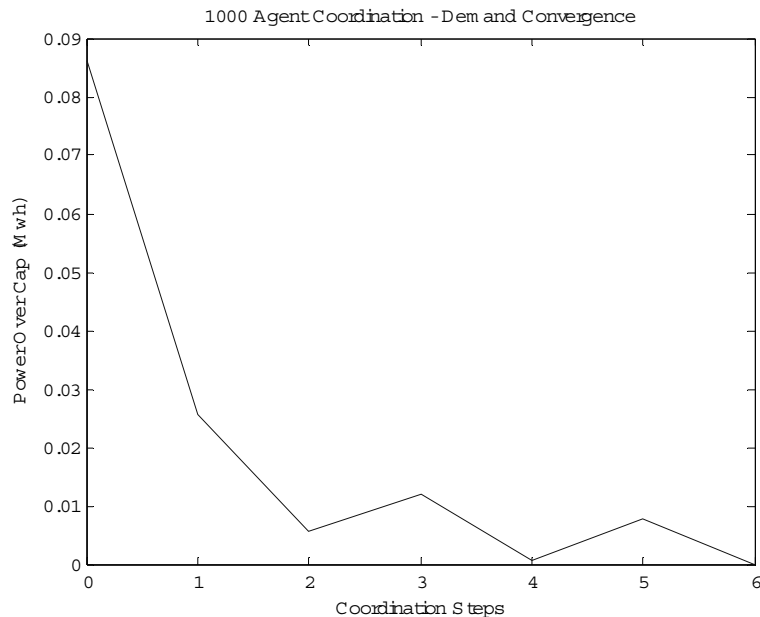


FIGURE 8: Coordination convergence, measured by the reduction of total power demand that exceeds supply cap during coordination process.

There is no single definition of firmness that is universal in the industry. What is required from an algorithmic perspective is assurance that the coordination process can achieve a requested cap level. This is a necessary part of the evaluation of firmness but not sufficient. In practice firmness will depend on a number of extraneous variables, which it is outside the scope of this paper to examine. We define firmness for a particular system as the probability that convergence has occurred after a given number of steps, and measure it by repeated simulations with identical caps but with different initial conditions on the resources. The experiment shown in FIGURE 7 was repeated 100 times for each of several different cap levels to generate the firmness results shown in FIGURE 9. Demand reduction is defined as average difference in demand between the coordinated and non-coordinated situations. It is noticeable that the probability of convergence is high and insensitive to the levels of demand reduction of 45% and less, and then changes rapidly as the demand reduction increases, indicating the range of cap levels that may be reliably achieved. The critical cap level at which the convergence behaviour changes gives a precise measure of the achievable algorithm performance.




FIGURE 9: Convergence firmness comparison for different levels of demand reduction for coordination of 1000 agents.

The cap level was fixed during each of these experiments. We also tested the effectiveness of progressively reducing the cap level during the simulation. Taking the case of 48% demand reduction in FIGURE 9, which showed marginal convergence, we applied the demand reduction in three steps:

- Firstly, apply the algorithm to obtain a set of resource agent plans that achieve a demand reduction of 38% .

- Starting from this set, apply the algorithm to obtain a revised set of plans that achieve a demand reduction of 44.5% .
- Finally, starting from this set, apply the algorithm to achieve a demand reduction of 48% .

It can be seen in FIGURE 10 that this improved the probability of convergence by about 10% , which is marginal compared to the dramatic deterioration of convergence as the demand reduction varies from 45% to 51% . This deterioration is an unambiguous and useful indication that the algorithm has reached its limit of performance. This limit is a characteristic of the algorithm and of the system of energy consumers, and a real system could be subjected to a series of convergence tests aimed at finding the point of deterioration under a range of conditions. Such a characterisation of performance would be useful information for a broker agent to use in setting cap levels.




FIGURE 10: Convergence firmness comparison between fixed cap and variable cap for coordination of 1000 agents. The demand reduction is 48% and the fixed-cap curve is taken from FIGURE 9.

5.2 Theoretical Limit of Demand and Response

We are fortunate that a system of simple refrigerators is amenable to theoretical analysis. For a given length of cap we may calculate the minimum achievable power consumption for a set of cool rooms subject to an "ideal" control regime. As will be discussed below this regime has slightly different constraints to those of our current algorithm. However, it provides a valuable insight into what might be achieved by such algorithms.

Assume that the temperature variation of cool rooms is linear between switching operations; this is approximately true for the real cool rooms in FIGURE 1 and the simulated cool rooms in FIGURE 5, and will only be violated significantly when the compressor is operating close to its capacity or when the allowed temperature range is particularly large, such cases requiring a more sophisticated analysis. Define c as the period or period in steady-state operation and r as the duty cycle or the fraction of each period that the compressor is switched on.

Assume also that the system has perfect coordination ability supported by infinite-bandwidth communications – clearly these are strong assumptions but it is useful to see what can be achieved in ideal circumstances. Consider at what level a supply cap of duration τ may be applied. For cap durations up to $\tau = (1-r)c$, which is the normal off-period for a cool room, a cap level of zero may be achieved by aligning the off-period of each cool room with the cap period. If τ increases and there are m any cool rooms, then each one must be switched on for a part of the cap period, and a minimum constant power consumption $P_{\min}(c,r)$ can in principle be achieved through coordination and multiple switching. For example, if $\tau = 2c - rc$ then each cool room must be switched on for a total time rc during τ and, with compressor power p_0 , the minimum constant power consumption is $P_{\min}(c,r) = p_0 rc / \tau$.

This consumption increases linearly with τ , due to linear temperature variation, so the minimum power consumption during a supply cap of duration τ for a system of $N(c,r)$ cool rooms having period c and duty cycle r is

$$P_{\min}(c,r) = N(c,r) p_0 \frac{[\tau - (1-r)c]r}{\tau} \quad (3)$$

when $\tau > c(1-r)$ and $P_{\min}(c,r) = 0$ otherwise. It is assumed that c and r do not depend on τ . Consider an experimental system of N_{tot} refrigerators with fixed duty cycle $r_0 = 1/3$ and periods uniformly distributed between $c_{\min} = 15$ min and $c_{\max} = 30$ min. Then

$$N(c,r) = \frac{N_{\text{tot}} \delta(r - r_0)}{c_{\max} - c_{\min}} \quad (4)$$

so that

$$N_{\text{tot}} = \int_{c_{\min}}^{c_{\max}} \int_0^1 N(c,r) dr dc. \quad (5)$$

Integrate the minimum power consumption across this set of cool rooms to calculate the minimum achievable level of a system supply cap of duration τ :

$$P_{\min} = \int_0^1 \int_L^U P_{\min}(c, r) dc dr = N_{\text{tot}} P_0 r_0 \left[\frac{U - L}{c_{\max} - c_{\min}} - \frac{(1 - r_0)(U^2 - L^2)}{2\tau(c_{\max} - c_{\min})} \right] \quad (6)$$

where $L = \min(c_{\min}, \tau / (1 - r))$ and $U = \min(c_{\max}, \tau / (1 - r))$.

In FIGURE 11 this curve is graphed together with an experimentally determined curve. For each cap duration the stigspace algorithm was applied to a simulated system of 100 cool rooms rated at 3 kW and with r_0 , c_{\min} , and c_{\max} as above. The cap level was reduced progressively as long as convergence succeeded; the criterion for success was that five successive experiments should converge successfully to satisfy the cap over its entire duration.

It is worth noting that the maximum peak power for this system when uncontrolled is 300 kW. The figure shows that the stigspace algorithm performs well over the entire range of cap duration, achieving a minimum cap of 50 kW for short caps, rising to 90 kW for longer caps. Comparison with the minimum cap limit shows two points of interest:

1. For short caps the best cap achieved is significantly higher than the minimum possible.
2. As cap duration increases performance approaches the theoretical minimum and even exceeds it for very long caps.

Both these points will now be discussed in detail.

1. Short caps: Good performance in this regime relies on agents shifting most of their ON times outside the cap interval. In section 4.4 it was shown that agents failed to do this for short caps even when advance notice of the cap was given. The main reason for this failure seems to be the limited control available to CordCap, which only shifts power from a cap-violating interval to both adjacent intervals, a process observed to be subject to local minima for short caps. Clearly, seeking an improvement to the algorithm in this regime will be a priority for future research.
2. Long caps: For caps much longer than the period c the optimum strategy is different. It is no longer possible for individual agents to shift their ON times outside the cap, so the best option is to reduce peak power by evenly distributing the ON times of all agents over the cap interval. The CordCap algorithm performs very well in this regime, approaching the theoretical limit and even exceeding it by a small amount in some cases. This apparent anomaly is due to simplifying assumptions made in calculating the theoretical limit. In particular, the assumption of constant duty cycle r does not hold exactly for systems modelled by equation (1).




FIGURE 11: Performance comparison against theoretical limit of achievable cap level as a function of cap duration.

6 CONCLUSION

A distributed multi-agent coordination system has been introduced in this paper, which coordinates distributed energy resources in an attempt to achieve a supply cap on the power drawn from the electricity grid while satisfying local constraints of the agents. This system features separation of the coordination mechanism from the information exchange mechanism by using indirect (or stigmergic) communications between resources and a broker. The coordination mechanism is asynchronous and adapts to change in an unsupervised manner, making it intrinsically scalable and robust. The inspiration for using indirect communications comes from the study of natural systems such as ant colonies. This system also features averaging (or more complex processing) of energy consumption plans over appropriate cycles, such as market cycles, before such information is communicated. This both reduces the message size and ensures that aggregated quantities of power created by coordination are aligned with the time intervals in which they are valued in the electricity market.

This system overcomes many of the difficulties of previously reported coordination systems. It should particularly be noted that the system remains robust under changing circumstances of resources, even for large resource numbers, and the system automatically includes different scales of temporal dependency through the amalgamation of energy consumption plans. Using cool rooms as representative loads under agent management, this paper has introduced the coordination approach in detail and demonstrated through simulation that it is scalable at least to 10,000 resource agents. It has examined the effect of the diversity of cool-room

parameters, as will be found in application, and shown that this improves performance. The reliability or “firmness” of convergence has been studied and the dependence of convergence on supply cap level gives an unambiguous and useful indication that the algorithm has reached its limit of performance. A theoretical performance limit was calculated for an “ideal” coordination system and allowed an instructive comparison against simulated performance, showing that for long periods of supply cap the new coordination system performs well and achieves near-optimum performance; for shorter periods of supply cap the system, while giving significant improvement, performs well below the theoretical limit. This seems to be a property of the agent switching strategy of modifying power in 5-minute cap-violating intervals by shifting power to adjacent intervals. This strategy works well for longer caps, but will need to be modified to give improved performance for shorter cap intervals. This will form an important part of the continuing research into agent-based coordination of distributed energy.

REFERENCES

- Australian Greenhouse Office (AGO). 2007. <http://www.greenhouse.gov.au/inventory/enduse>, accessed June 2007.
- Carlsson, P. and A. Anderson. 2007. A flexible model for tree-structured multi-commodity markets, *Electronic Commerce Research*, vol.7, no.1, pp.69-88.
- Clearwater, S.H. 1996. Market-based control – A paradigm for distributed resource allocation, book in World Scientific, Singapore.
- Clement, B. J. and A. C. Barrett. 2003. Continual coordination through shared activities, *Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems*, Melbourne, Australia, pp.57 – 64.
- Clement, B. J. and E. H. Duffee. 2000. Performance of coordinating concurrent hierarchical planning agents using summary information, *Proceedings, 7th International Workshop on Intelligent Agents VII: Agent Theories, Architectures, and Languages*, Boston, MA, USA, pp.202 – 216.
- Dineas, L. and N. D. Hatziaargyriou. 2005. Operation of a multi agent system for microgrid control, *IEEE Transactions on Power Systems*, vol.20, no.3, pp.1447-1455.
- Energex, 2007. http://www.energex.com.au/environment/cool_change/cool_change.html, accessed June 2007.
- Estrin, D., R. Govindan, J. Heidemann, and S. Kumar. 1999. Next century challenges: scalable coordination in sensor networks, *Proceedings of the Fifth Annual ACM/IEEE International*

- Conference on Mobile Computing and Networks, Seattle, ACM Press, New York, USA, pp.263-270.
- Guo, Y., J. Li, and G. James. 2005. Evolutionary optimisation of distributed electrical loads and generators, Proceedings, 18th Australian Joint Conference on Artificial Intelligence, Sydney, Australia, pp.1086-1091.
- Hopper, N., C. Goldman, R. Bharvirkar, and D. Engel. 2007. The Summer of 2006: A milestone in the ongoing maturation of demand response, The Electricity Journal, vol. 20, no. 5, pp. 62-75.
- Hudson, G. and C. P. Underwood. 1999. A simple building modelling procedure for MATLAB/Simulink, Proceedings, International Building Performance and Simulation Conference, Kyoto, Japan, vol. 2, pp. 777-783.
- Jones, T. E. and G. C. James. 2005. The management and control of distributed energy resources, Proceedings, CIGRE Symposium on Power Systems with Dispersed Generation, Athens.
- Kamphuis, I. G., J. K. Kok, C. J. Wamer, M. P. F. Hommelberg. 2006. Massive coordination of residential embedded electricity generation and demand response using the PowerMatcher approach, Proceedings, 4th International Conference on Energy Efficiency in Domestic Appliances and Lighting, London, pp. 1249-1259.
- Kok, J.K., C.J.Wamer, I.G.Kamphuis. 2005. PowerMatcher: multiagent control in the electricity infrastructure, Proceedings, 4th International Joint Conference on Autonomous Agents and Multiagent Systems, Utrecht, Netherlands, pp. 75-82.
- Li, J., G. Poulton and G. James. 2007. Agent-based distributed energy management, 20th Australian Joint Conference on Artificial Intelligence, Gold Coast, Queensland, Australia, pp. 569-578.
- Li, R., J. Li, G. Poulton and G. James. 2008. Agent-based optimisation systems for electrical load management, accepted by 1st International Workshop on Optimisation in Multi-Agent Systems, Estoril, Portugal.
- Luh, P.B., M. Ni, H. Chen, and L.S. Thakur. 2003. Price-based approach for activity coordination in a supply network, IEEE Transactions on Robotics and Automation, vol. 19, no. 2, pp. 335-346.
- Müller, H., A. Rudolf, and G. Aumayr. 2001. Studies of distributed energy supply systems using an innovative energy management system, Power Industry Computer Applications, pp. 87-90.
- National Electricity Market Management Company (NEMMCO), Australia, 2007. <http://www.nemco.com.au/>, accessed December 2007.
- Nord Pool, 2008. http://www.nordpoolspot.com/trading/The_Elspot_market/159/, accessed March 2008.
- Oyarzabal, J., J. Jimeno, J. Ruela, A. Engler, and C. Hardt. 2006. Agent based micro grid management system, International Journal of Distributed Energy Resources, vol. 2, no. 3, pp. 195-209.
- Regan, T., H. Sinnock, and A. Davis. 2003. Distributed energy neural network integration system: year one final report, Technical Report of National Renewable Energy Laboratory (NREL), Golden, CO., 1 June 2003, OSTI ID: 15003911, NREL/SR-560-34216.

- Stone, P. and M. Veloso. 2000. Multiagent systems: a survey from a machine learning perspective, *Autonomous Robots*, vol. 8, no. 3, pp. 345-383.
- United States Department of Energy (USDE). 2000. Report of the US department of energy's power outage study team. <http://certs.lbl.gov/certs-rtinakey-outage.html>, accessed April 2008.
- Ye, F., A. Chen, S. Lu, and L. Zhang. 2001. A scalable solution to minimum cost forwarding in large scale sensor networks, *Proceedings of the 10th International Conference on Computer Communications and Networks*, Dallas, Texas, pp. 304-309.
- Ygge, F. 1998. Market-Oriented Programming and its Application to Power Load Management, Ph.D. Thesis, ISBN 91-628-3055-4, Lund University, Sweden.
- Ygge, F. 2000. Resource-Oriented Multi-Commodity Market Algorithm, *Journal of Autonomous Agents and Multi-Agent Systems*, vol. 3, pp. 53-72.