Clustering Distributed Energy Resources for Large-Scale Demand Management

Elth Ogston
Vrije Universiteit Amsterdam
elth@cs.vu.nl

Astrid Zeman, Mikhail Prokopenko and Geoff James
CSIRO ICT Centre
{astrid.zeman, mikhail.prokopenko, geoff.james}@csiro.au

Abstract

Managing demand for electrical energy allows generation facilities to be run more efficiently. Current systems allow for management between large industrial consumers. There is, however, an increasing trend to decentralize energy resource management and push it to the level of individual households, or even appliances. In this work we investigate the suitability of using adaptive clustering to improve the scalability of decentralized energy resource management systems by appropriately partitioning resources. We review the area of distributed energy resource management and propose a simple yet realistic model to study the problem. Simulations using this model show that straightforward clustering and distributed planning methods allow systems to scale, but may be limited to only a few hundred-thousand appliances. Results indicate that there is an opportunity to apply adaptive clustering techniques in order to discover more advanced grouping criteria that would enable groups to change as appliances' behavior changes. The simulations further suggest that even an extremely limited exchange of information between clusters can greatly improve management solutions.

1 Introduction

The electricity industry in many countries is facing a number of pressures due to increasing demand, the increasing "peakiness" of demand, and the prospect of carbon accountability. Managing demand for power is an alternative to the traditional solution of investing in expanding centralized generation capacity and the associated infrastructure. Controlling distributed energy resources, customer loads and small generators located close to load centers, in response to price signals and network constraints can give relief from the volatility of wholesale electricity prices and can assist constrained distribution networks during summer and winter demand peaks. Such responses also provide a mechanism to reward customers who choose to manage their energy use thoughtfully.

The inherently distributed nature of energy resource management, and the large savings to be made by even a small gain in efficiency, make it a compelling application in which to apply decentralized planning and control algorithms. Software resource agents running on customer premises or embedded in appliances, can be used to plan future energy consumption and to shift loads according to constraints placed on the system. Australia's Commonwealth Scientific and Industrial Research Organisation (CSIRO) currently has a live multi-agent test-bed for monitoring and controlling industrial loads and generators in electricity distribution networks [6]. The physical deployment of the GridAgents Platform is located in the CSIRO Energy Center in Newcastle, New South Wales, Australia.

To demonstrate the feasibility of using such platforms in the domestic household area, it is necessary to address system scalability. Current agent-based energy resource management systems employ central auctioneers or aggregating agents to combine information on energy consumption from resource management agents and feed this information back to the agents. The use of central components, and more significantly their communications infrastructures, is undesirable in systems that contain hundreds of thousands or millions of devices. We therefore investigate the possibility of clustering devices into more manageable, relatively independent, groups. In particular, in this paper we address two questions: (1) *to what extent is the quality of global solutions reduced by partitioning agents into groups,* and (2) *is it sufficient to choose these groups at random or can agents be classified so as to create more effective groups?*

To avoid demand peaks, resource management agents must be coordinated. One alternative is to cluster agents into separately coordinated groups. Even more interesting is the possibility of creating dynamic adaptive clusters which instead of using fixed criteria can group agents according to unforeseen behaviors and current demand characteristics. Before developing such complex algorithms, however, it is necessary to assess the problem area to determine if they are likely to improve on simpler methods. In this paper we make the following contributions to the study of using cooperating agents to reduce short term peaks in electricity demand:

- We define a simple yet realistic scenario in which to study the problem and propose and test a minimal demand planning/coordination algorithm which can be used as baseline when developing more advanced planing and control mechanisms.
- We study decentralizing the coordination and planning problem by dividing agents into independent clusters. We show that for small numbers of clusters the quality of global plans is not significantly reduced by partitioning the problem in this manner.
- We show that the choice of clustering criteria can make a difference to the system's ability to stabilize demand.
- We demonstrate that for larger numbers of clusters an extremely limited exchange of information between clusters can greatly improve global plans.

The work presented in this paper considers an abstract model of the energy requirements of simple refrigeration devices, verified by additional experiments using measurements of the behavior of a real appliance. The results listed above justify further research into the area of adaptive decentralized coordination and clustering algorithms for distributed energy resource applications, involving larger scale test-beds and the collection of more varied measurement data.

2 Background and Related Work

This section provides background on electricity generation and distribution, in particular in the Australian market, and describes current work on using agents for distributed energy resource management. Adaptive decentralized clustering methods will be described in Section 5 in relation to our experimental results.

2.1 Demand Management

The management of energy resources can take place on a variety of time-scales. Seasonal peaks in demand, for instance caused by widespread use of air conditioning in hot weather, can place considerable strain on the power grid. When demand exceeds capacity in a grid sector, brownouts can occur, making power temporarily unavailable to customers in the area. Additionally, power draining from neighboring sectors can cause relays to malfunction, resulting in costly blackouts. Measures for managing seasonal demand peaks include intelligent load shedding, in which certain areas are strategically blacked out, or contract agreements are made with larger consumers for supply to be reduced when necessary, in return for lower prices.

On a weekly and daily time scale, energy demand also follows regular highs and lows, corresponding to business hours and common daily routines such as meal times. Many energy retailers employ a dual pricing system with peak and off-peak hours to encourage households to shift flexible tasks, such as washing, to times when there is spare capacity.

While seasonal, weekly, and daily trends are the most visible to energy consumers, changes in demand on a much shorter time scale are likely to profit most from control by agent technology. In Australia, energy retailers purchase capacity, in the form of hedge contracts, many months in advance, according to forecasts of future needs. Since this energy use is planned, it is relatively cheap, with a price near the average expected price, typically around $40 (AUD) per megawatt hour. However, if customer demand exceeds the pre-purchased quota, additional energy must be bought on a spot market. In this market, prices are set per five minute interval. In January 2007 spot market prices in New South Wales ranged between $10.80/MWh and $5091.95/MWh [10]. Managing resources so as to reduce short term peaks in demand can thus result in significant savings. Micro-managing a device by shifting its energy consumption by a few minutes can have very little impact on the functioning of some appliances, but requires constant monitoring and communication, making it an appropriate task for computer agents.

2.2 Energy Generation and Distribution Infrastructure

Large energy producers and consumers (including retailers) in eastern Australia participate in a single wholesale market, the National Electricity Market (NEM) [10]. Within this market energy production is centrally managed so as to meet demand, set fair prices, and ensure reliability of supply. Producers submit bids stating the amount of energy they can generate at what cost and consumers submit predictions for consumption. These are matched, 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. These dispatch and prediction intervals influence 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 NEM covers an area 4000 miles long, and serves eight-million end-use customers. Though it is possible for electricity to be traded over a distance of 4000 miles [13], transport of electricity between regions is limited by the capacity of the high-voltage transmission lines that interconnect them. Thus, in theory the whole of the NEM can be seen as a single very large-scale computational system in which tradeoffs in power use can be made between agents, regardless of location. In practice some restrictions apply and agent location may place some constraints on which agents may interact with each other. Currently, information on electricity demand is only exchanged between major players in the market. However, the installation of "smart-meters", and the ability to use transmission lines to carry data, means that in the future, even small consumers such as households could participate in energy management systems [2, 7].

2.3 Agents for Distributed Energy Resource Management

Previous research on multi-agent systems for electricity management has mainly focused on agent-based simulations of electronic market mechanisms [20, 18]. A couple recent efforts described below are building deployed agent systems.

The European CRISP project [17] carried out field studies in Sweden to assess the advantages of using intelligent agent and electronic market technologies. In the project's system model the electricity grid is divided into local areas, called cells. Coordination between agents and cells takes place through the Powermatcher electronic market [9]. In Powermatcher, agents within a cell submit bids for electricity production or use to a designated aggregator, called an SD-matcher, which manages the cell. Aggregated bid information is passed up a hierarchical tree of aggregators. The market equilibrium prices is calculated at the root node, and distributed back down the tree. This network configuration remains fixed over the duration of system operation.

Market mechanisms coordinate current power consumption and production based on a single price. Energy needs of devices at a particular time are, however, often dependent upon device behavior at other time intervals. Multi-commodity combinatorial markets can be used to coordinate prices over a number of time intervals, but quickly become overly complex as the number of intervals grows [4].

The GridAgents framework, developed for the Australian CSIRO testbed, explores a variety of alternatives to market-based control [6]. In particular, a method that uses genetic optimization to coordinate agent plans has been studied. A cap on energy use for a given period is set for a group of agents, which accordingly rearrange plans to minimize cost whilst considering local objectives [5].

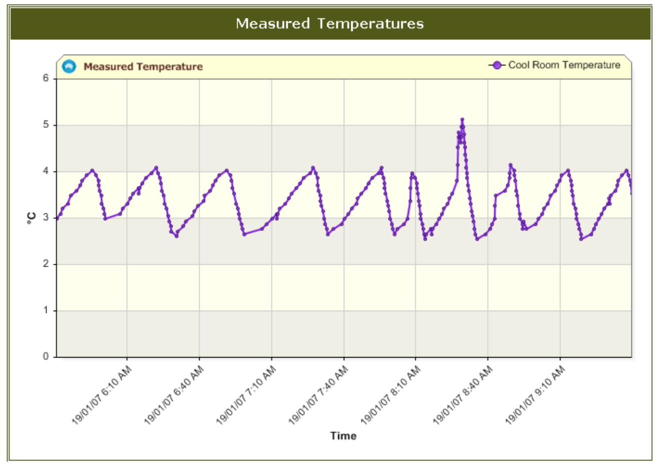

Figure 1. CSIRO cool-room agent operation.

Both the CRISP and GridAgents projects propose dividing agents into groups in order to meet system scalability objectives. In the CRISP project these groups take the form of cells, whose composition is based on the physical locations of devices. The project report, however, notes that group configurations which differ with the seasons are likely to improve on the system's ability to match supply to demand [17]. The GridAgents method of genetic optimization is also based on coordination within groups, of a size of around 500 agents. The project aims to further explore clustering methods by which these groups are defined [6]. In this paper we provide a detailed investigation of the factors involved in successfully creating improved clusterings for such systems.

3 Methodology and Scenario

The experiments presented in Section 4 evaluate the effectiveness of clustering distributed energy resources for the purposes of coordinating planned energy demand. Coordinating energy use locally within clusters of resources in place of centrally coordinating all resources has several advantages. First, it removes a dependance on a single central point of failure, improving system robustness. Second, it reduces the communication delay between resource agents and the coordinator, allowing planning to be done on a shorter time scale, or for additional planning/coordination rounds to take place. Third, it reduces the complexity of the coordination problem, allowing coordination to be done by a simpler, cheaper, controller. Fourth, it creates a system which is more scalable as adding new resource agents only requires creating additional clusters, rather than expanding the capacity of a central coordinator.

There are, however, a number of questions involved with the relative quality of centralized coordination of plans versus distributed cluster-based coordination. Clusters only have partial information on the total energy use of the system, implying that distributed coordination should produce a global plan that is worse than found with a central coordinator. We would like to know exactly how much worse in order to be able to weigh the relative costs and benefits of clustering resources. Additionally, there are questions about how to create clusters: what size should clusters be, what composition function should be used to choose cluster members, should cluster membership be fixed or dynamic? Creating a fixed set of fixed size clusters with random membership would greatly simplify the coordination problem. We thus, in particular, would like to know how much of an advantage more advanced clustering methods would have over simpler settings.

The experiments in Section 4 are designed to examine the difference between joint plans coordinated only within clusters and those produced using centralized coordination, and to compare the quality of joint plans produced using random versus carefully selected clusterings. In this section we first present out experimental methodology, and in particular the energy resources, coordination algorithm, and clustering criteria we consider.

3.1 Scenario

We consider a scenario where the power consumption of individual refrigerators in a city must be coordinated so as to keep their total energy use as constant as possible for a given short period. Exact data on the preferable scale of this problem, the precise range of energy requirements of household refrigerators, or the possible planning behaviors of agents controlling refrigerators is unavailable. In the following sections we describe the existing data and the estimations we use to set parameters for our experiments.

Refrigerators and freezers typically make up over 20% of total residential electricity consumption in Australia [1]. We thus focus on modeling refrigerators as our resource agents, being the largest single end-use offender of household energy consumption. Nearly all households in Australia have at least one refrigerator and about 30% own two. Nearly 60% of households own a separate freezer [1]. While it may possibly be desirable to coordinate the behavior of all devices within the NEM, initial systems are likely to be more modest. We consider the case of coordinating refrigerators in the city of Newcastle NSW which has a single 600MW coal-fired power plant, serving approximately 200,000 households. We would thus like to ensure system scalability to around 400,000 resource agents.

Our model of energy resources and their consumption requirements focuses on representing variable flexibility of resource agents in adjusting initial plans. During a typical planning cycle in the multi-agent systems on which we base this work, resource agents send initial plans to a coordinator, which computes aggregate energy use then sends suggestions for changes to the agents. In general, this planning and coordination occurs in multiple cycles in which agents make readjustments, send their new plans to the coordinator, and receive additional suggestions. Plans are also readjusted and re-coordinated as predictions of future energy use changes. In this study we consider relatively simple agents. For this reason we simplify the planning process to a single round in which agents each send multiple possible plans for the coordinator to select among. We assume that once set, and agent is free to follow its chosen plan. By varying the number plans an agent submits we can model resources that are more or less flexible in their energy consumption profiles.

3.2 Energy Resources

In this study we consider a single type of energy resource, an abstract cool-room, of which refrigerators are an example. Cool-rooms consist of an insulated chamber which must be kept within a given temperature range, and a compressor that can be turned on to lower the temperature. Their basic behavior is straightforward, when the compressor is off the room's temperature slowly rises towards the outside temperature, until the maximum of the temperature range is reached. The compressor is then turned on until the chamber cools to the minimum value of the temperature range.

The simple cool-room agents studied in this paper are modeled after the basic operation of the more advanced CSIRO cool-room agent currently being used to control industrial cool-rooms as well as domestic-sized refrigerators in the Newcastle Energy Centre. Figure 1 shows the monitoring interface of the agent controlling the canteen cool-room[1]. The agent is given the goal of keeping the temperature of the cool-room within a certain range by turning on and off either the compressor or a fan. The agent receives data on current electricity prices and attempts to avoid using energy when prices spike. Predictions for future plans are made based on the current temperatures inside and outside the cool-room, and a model of heat-loss properties based on a history of temperature behavior.

Realistic models of cool-rooms show that their short-term power consumption follows a periodic square wave [20], as illustrated in Figure 2. In this study we model refrigerators with three parameters, a compressor cycle length (C) that states the length of one period of this square wave, an on-time (T), which states the amount of time in which the compressor is on during that period, and a power value (P) that gives the amount of power used by the compressor when on.

[1] Available online at http://demc.com.au/DEMC/Canteen/Canteen.aspx.

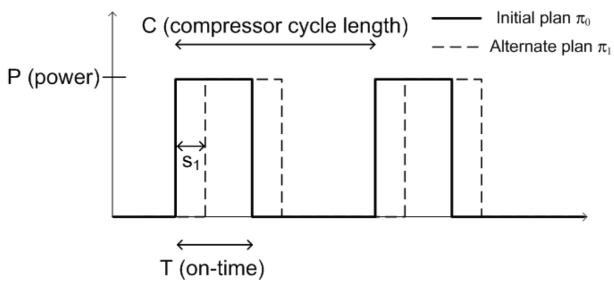

Figure 2. Plan Representation

We base the experimental values of the cool-room parameters on data giving the typical energy consumption per year of refrigerators and freezers currently sold in Australia [1]. The values of C and T for a refrigerator vary depending on use. From observations of refrigerators we estimate that cycle lengths vary between 20 and 60 minutes. Whirlpool's customer service website states that compressors in typical refrigerators are on for between 40% and 80% percent of the time. Based on these estimates we create a set or random refrigerator profiles for our experiments by choosing the yearly power consumption of an appliance at random from the list of Australian refrigerators, and choosing values of C and T uniformly at random from the ranges given above. The power used by the compressor, P, is calculated from the yearly power consumption, assuming that C and T are averages for the year. Plans are divided into one minute intervals in which the power consumption is either P when the compressor is on or 0 when the compressor is off. Each refrigerator agent is given an initial plans starting at a random point in its compressor cycle.

In these experiments we will consider planning based on expected energy use of a refrigerator for a term of two hours, to cover behavior over several compressor cycles. Because refrigerators only need to maintain a given temperature range, planning agents can be very flexible in choosing exactly when to turn compressors on or off. They thus have the option of creating a wide variety of plans for a given 120 minute interval. Since we are only interested in the effect of flexibility, we do not model exactly how agents can choose to change these plans. Instead we make the simplifying assumption that any plan that involves a time shift of the periodic power consumption square wave is acceptable to maximally flexible refrigerators. A refrigerator thus has a finite set of alternative energy-use plans it can follow, $\Pi = \{\pi_0, \ldots, \pi_{k-1}\}$. Given a refrigerator's initial plan, π_0, each alternative plan π_i is defined by a time shift s_i such that the power consumption at time t in π_i is that of plan π_0 at time $t - s_i$. In our experiments we consider time shifts s that are multiples of one minute intervals. Thus a refrigerator with a compressor cycle length of C minutes has C possible plans. Less flexible refrigerators are modeled as accepting only smaller subsets of plans.

3.3 The Desync Coordination Algorithm

We study a simple heuristic coordination algorithm that, given a fixed set of possible plans for each resource agent, attempts to choose the combination of plans that results in the best aggregate plan. This could be done with straightforward brute-force search of all the possible combinations. However, such a search quickly becomes prohibitively expensive as the number of resources considered grows. Instead we use a simplifying heuristic in which the resources are chosen sequentially, and their plans are selected according to a metric, without attempting to find the exact optimal combination. Plan coordination thus follows the sequence:

1. Each energy resource x_i sends to the coordinator its set Π_i of possible acceptable plans.
2. The coordinator goes through the resources sequentially, according to a certain order, and for each resource selects the plan that would best fit in the current aggregate. The selected plan is added to the current aggregate. The metric used to determine the "best fit" is described in Section 3.4.
3. The coordinator notifies each agent which of its possible plans was selected and is to be carried out.

The Desync coordination algorithm has the advantage that it is simple, requires minimal communication between the resource agents and the coordinator, and requires minimal computation by the coordinator. It has the disadvantage that there is no feedback to the agents which might have been able to be more flexible in their plans. A further disadvantage is that a complete search is not made to find the optimal combination of plans. For the purposes of analyzing the possible benefits of clustering, however, this simple algorithm is sufficient and provides a clear picture of the effects of clustering. The case for more complex algorithms is discussed in Section 5.

3.4 Plan Selection

The coordinator in our experiments aims to minimize the peak power usage in the aggregate plan. Aggregation of two plans, say π_j and π_k, denoted as $\pi_j + \pi_k$, simply involves adding the planned powers for each time interval. We denote by $f(\pi)$ the maximum (peak) power used in the plan π over its entire duration, in our case the time interval $[0, 120]$. For a given current aggregate plan, π_α, and a given set of possible plans for a resource agent, Π, the coordinator chooses the agent plan π_i for which $f(\pi_\alpha + \pi_i)$ is minimized. The full Dysnc algorithm is summarized in Algorithm 1.

Energy prices paid by energy distributors in Australia have two components, a price for a fixed amount of capacity, set in advance, and a spot price for energy use that exceeds that capacity. In this paper we do not consider the more complex planning problem of matching energy use to current costs, or a changing "cap" on energy demand [5]. We leave this problem to future work, discussed in section 5. Our main goal is to show that the underlying power consumption of agents can be coordinated so as to reduce variability in the "fixed" part of the energy use.

Algorithm 1: Desync algorithm.

Set the current aggregate plan π_α to zero usage
for *each resource agent* $x_i \in X$ **do**
  retrieve set of plans Π_i
  set the chosen plan for x_i as $\pi_i = \pi_{i,0}$
  set temporary maximum as $\mu = f(\pi_\alpha + \pi_i)$
  for *each subsequent plan* $\pi_{i,j} \in \Pi_i$ **do**
    if $f(\pi_\alpha + \pi_{i,j})$ *is less than* μ **then**
      set $\pi_i = \pi_{i,j}$
      set $\mu = f(\pi_\alpha + \pi_i)$
    end
  end
  update current aggregate as $\pi_\alpha = \pi_\alpha + \pi_i$
  send chosen plan π_i to the agent x_i
end

4 Experiments

In this section we present experiments to examine the following questions:

- How effective is Desync coordination? (Section 4.1)
- To what extent does reducing the flexibility of resources reduce the Desync algorithm's ability to coordinate energy use? (Section 4.2)
- In respect to planning, is actual cool-room behavior well represented by our abstract cool-room model? (Section 4.3)
- To what extent does dividing resources into clusters reduce the Desync algorithm's ability to discover improved global plans? (Section 4.4)
- Does clustering similar resources together differ from random clustering? (Section 4.5)
- How does system behavior change as the number of agents is increased? (Section 4.6)
- Does adding communication between clusters improve on non-interacting clusters? (Section 4.7)

4.1 Basic Desync Coordination Ability

Figure 3 compares two example aggregates of the plans of 640 maximally flexible refrigerators. The figure shows the aggregate power used at each minute interval when the refrigerators are left to choose their own plans, and the aggregate for the same refrigerators after Desync coordination. The figure shows that Desync coordination is able to effectively shift individual behaviors so as to flatten peaks in the overall power use.

In general the alternative plans for a resource agent can have different total power consumptions over the total time

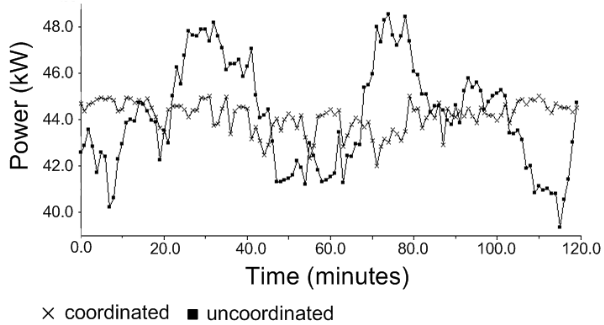

Figure 3. Aggregate power use, with and without coordination.

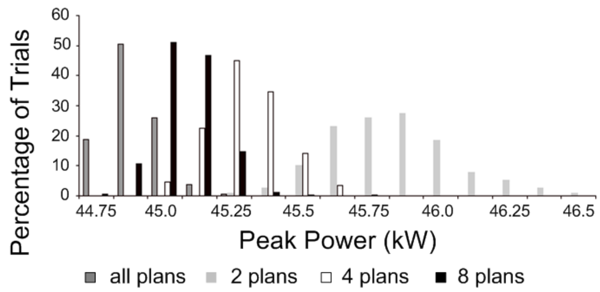

Figure 4. Distribution of peak power, less flexible agents.

interval. In order to keep the power consumption of all plans equal, and thus aggregate plans using different possible plans comparable, we artificially raise or lower the compressor power slightly when a shifted plan results in a refrigerator being on for a different period of time than in its initial plan. Thus a lower maximum power for a plan indicates a more constant plan. Without planning the maximum power used is 48.55 kW (out of a maximum possible 71.90 kW) while planning reduces this value to 45.01 kW.

The aggregate plans of refrigerators can vary, both with and without planning, because in the simulations each refrigerator chooses its initial plan at random for each trial and because agents can be aggregated in a different order in each trial. Without planning, the distribution for 1000 trials of the peak aggregate power varied between 46.00 kW and 51.51 kW. With planning, peak power consistently improves on these values, and is more predictable, varying from 44.70 kW to 45.27 kW.

4.2 Reducing Agent Flexibility

We next examine the effects of reducing the flexibility of resource agents in changing their consumption plans. We repeat the experiment from Section 4.1, but instead of allowing the coordinator to choose out of any possible plan for each refrigerator, each refrigerator submits only a small number of randomly chosen plans. Fully flexible refrigerators have between 20 and 60 possible plans. Figure 4

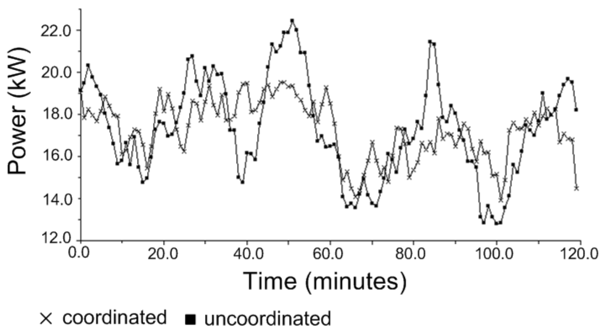

Figure 5. Aggregate power use, measured refrigerator data.

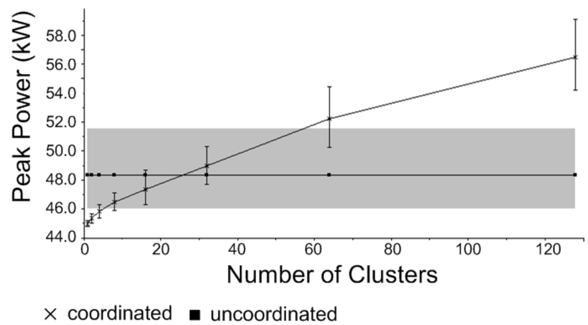

Figure 6. Peak power: random clusters.

shows the distribution of peak power for coordinated aggregate plans when refrigerators submit only 2, 4 or 8 possible plans. Even with very few plans to choose from, the Desync algorithm is able to find almost as good an aggregate plan as it does when given fully flexible agents.

4.3 Comparison to an Actual Refrigerator

Actual refrigerators have more complex behavior than the idealized cool-room model. We measured the behavior of a standard refrigerator over a period of several days. Without user intervention power usage of the appliance followed a regular square wave, in which the compressor was on for about 20 minutes then off for about 15 minutes. With user intervention however, spikes in power occurred when the door was opened, turning the light on. Opening the door also caused the compressor to stay on for longer. Every 24 hours a very high spike in power use occurred followed by a long cooling period, due to the operation of the automatic defroster.

The irregularity of behavior of actual appliances makes the coordination problem more difficult than for the model cool-rooms. Figure 5 shows example aggregates, with and without coordination, for 94 virtual measured refrigerators, created by breaking several days worth of measurements into two hour segments. Each virtual refrigerator is given 8 plans, made by shifting its initial plan by successive one minute intervals. Figure 5 shows that while Desync coordination becomes less effective in the non-ideal case, it still has a noticeable effect on lowering peaks in aggregate power. We will discuss possible effects of the behavior of real appliances further in Section 5.

4.4 Effect of Clustering

Figure 6 plots the peak power found in the global aggregate when dividing agents into 1, 2, 4, 8, 16, 32, 64, an 128 clusters, chosen at random for 640 agents. Values for the aggregate before and after coordination are plotted, marked at the average peak power found over 100 trials. Error bars, for the coordinated aggregate, and the shaded region, for the uncoordinated aggregate, indicate the minimum and maximum values seen. The graph shows that as agents are separated into more clusters, the aggregate plan formed by running the Desync algorithm independently in each cluster deteriorates. Interestingly, with more than 32 clusters the coordinated aggregate tends to be worse than the average aggregate obtained when no coordination was done at all and each refrigerator was simply left to follow its original randomly chosen plan. For less than 32 clusters however, clustering is an effective manner of distributing the coordination problem. With 128 clusters, coordination has a markedly detrimental effect on the systems. In the following experiments we shall investigate why distributed coordination performs so poorly with larger numbers of clusters.

4.5 Non-Random Clusters

The behavior seen in Section 4.4 should improve if the aggregate plans created in the individual clusters improve. The Desync algorithm's ability to coordinate plans should thus be enhanced by better clusterings in which refrigerators that have more compatible behaviors are grouped together. In Figure 7 clusters are created not at random, but so that clusters contain refrigerators with similar values of P, the power used by their compressor. This should make it easier for the Desync algorithm to make use of pairs of refrigerators whose plans can be shifted so that they complement each other. In fact, the figure shows that this *power-based clustering* criterion yields clusters with about the same behavior as random clusters.

In Figure 8 clusters are created in which refrigerators with similar compressor cycle times, C, are grouped together. Such *frequency-based clustering* does improve the aggregate plans. From this we conclude that clustering criteria used when dividing up agents does have an impact on the overall system behavior. The optimal criteria is likely to depend strongly on the actual behavior of the individual agents. We thus leave a further exploration into clustering methods for future work, discussed in Section 5.

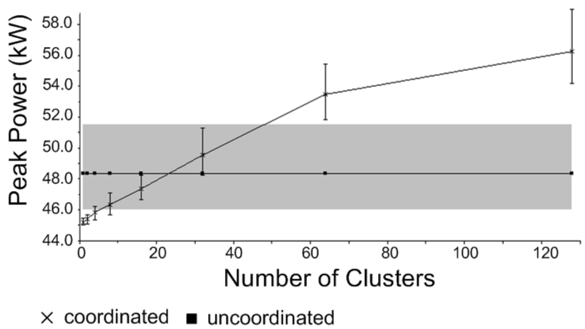

Figure 7. Peak power: P-based clusters.

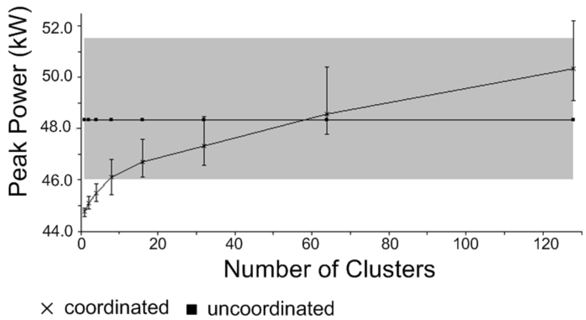

Figure 8. Peak power: C-based clusters.

4.6 Clustering and Scalability

A possible explanation for the poor coordination ability when agents are divided into more than 32 clusters is the small size of these clusters. Dividing 640 agents between 32 clusters produces clusters with only 20 agents each. It is possible that the Desync algorithm simply requires more agents to be able to produce acceptable aggregate plans.

We repeat the frequency-based clustering experiment from Section 4.5 using twice as many agents. We find, however, almost no difference in the influence of the number of clusters on coordination ability, in spite of the fact that the clusters in this experiment are twice as large. This indicates that there exists a limitation on the degree to which the energy resource coordination problem can be decentralized by dividing it into separate more manageable problems, independent of the overall size of the system.

Intuitively, this limitation can come from the nature of coordination itself. Looking back at the example aggregates in Figure 3 it can be seen that while coordination lowers the peak power in an aggregate, in doing so it increases the amount of time during which the aggregate is relatively high. When combining a two random aggregates, the chance of two high points overlapping thus increases after coordination. As the number of clusters increases the chance of undesired overlaps also increases. Hence, independent coordination in large numbers of clusters can result in worsening the global aggregate. Table 1 illustrates this. It compares coordinated and uncoordinated aggregates for 16 randomly chosen clusters, in an instance when the uncoordinated aggregate is slightly better than the coordinated one. The left side of the table shows the peak power in the individual plans of each of the clusters. It shows that within clusters coordination consistently improves over uncoordinated behavior, by between 8.5 and 21.5 percent. The right side of the table shows the result of aggregating the plans from the given clusters and those in the preceding lines. It shows that as more clusters are added to the aggregate, the global improvement made by coordination within clusters decreases, in spite of the local improvements made.

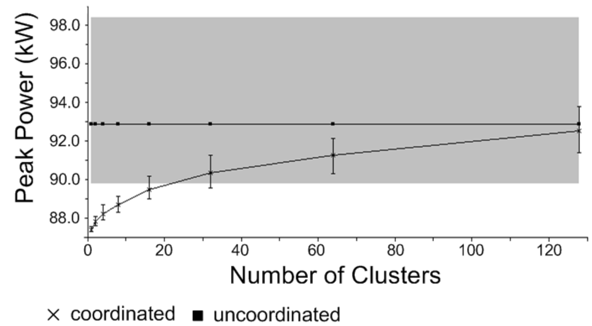

Figure 9. Frequency-based hierarchical clustering.

4.7 Hierarchical Desync Coordination

The experiment in Section 4.2 shows even a relatively small amount of flexibility on the part of individual resources allows Desync coordination to improve overall aggregates. This property may provide a solution to the limitations encountered when dividing the system into separate clusters. It indicates that by exchanging only a small amount of information, clusters may be able to sufficiently coordinate their plans. We test this by creating a simple hierarchical version of the Desync coordination algorithm. Resources are divided into separate clusters in which the Desync coordination algorithm is run twice, to produce two possible aggregate plans. These plans are then aggregated by a single higher level aggregator, again using Desync coordination. Figure 9 shows the result, when using frequency based clustering for 1280 agents. Compared to the results of frequency based clustering without the second level of coordination, this very low degree of synchronization between clusters does indeed significantly improve the overall aggregate plans. Better chosen information sharing between clusters could possibly have an even greater effect, and is a topic for future research, as discussed in Section 5.

5 Discussion and Future Work

The scenario in Section 3 set a goal of building a clustered energy planning system that can scale to at least

peak power of individual clusters				peak power in aggregate of current and preceding clusters			
	coordinated (kW)	uncoordinated (kW)	improvement		aggregate, coordinated (kW)	aggregate, uncoordinated (kW)	improvement
cluster 1	3.01	3.63	17.26%	cluster 1	3.01	3.63	17.26%
cluster 2	3.07	3.40	9.51%	clusters 1-2	5.94	6.57	9.70%
cluster 3	2.93	3.61	18.88%	clusters 1-3	8.71	9.24	5.65%
cluster 4	3.22	3.77	14.49%	clusters 1-4	11.58	12.51	7.45%
cluster 5	3.31	3.62	8.66%	clusters 1-5	14.80	15.47	4.38%
cluster 6	3.16	3.65	13.36%	clusters 1-6	17.72	18.93	6.39%
cluster 7	2.79	3.25	14.18%	clusters 1-7	20.31	21.27	4.50%
cluster 8	3.08	3.92	21.51%	clusters 1-8	23.08	24.64	6.32%
cluster 9	2.93	3.41	14.03%	clusters 1-9	25.88	27.70	6.56%
cluster 10	3.28	4.04	18.84%	clusters 1-10	29.04	30.93	6.11%
cluster 11	3.70	4.45	16.92%	clusters 1-11	32.58	33.56	2.92%
cluster 12	3.37	3.86	12.79%	clusters 1-12	35.92	37.38	3.91%
cluster 13	2.93	3.46	15.33%	clusters 1-13	38.57	40.42	4.58%
cluster 14	3.51	4.35	19.29%	clusters 1-14	41.64	42.62	2.30%
cluster 15	3.00	3.28	8.59%	clusters 1-15	44.57	44.76	0.41%
cluster 16	3.14	3.59	12.56%	clusters 1-16	47.59	47.46	-0.27%

Table 1. Effect of aggregating plans from random clusters with 40 agents each.

400,000 resource agents. Section 4 shows that using straightforward algorithms a system can easily be broken into 16 clusters in which coordinated energy use consistently improves on uncoordinated use, or up to 64 clusters in which coordinated use usually improves on uncoordinated. This gives clusters of between 25,000 and 6,250 agents. Given the simplicity of the Desync coordination algorithm, clusters of a few thousand agents are probably acceptable. However, there is a justifiable need to explore more sophisticated methods to both improve the number of clusters into which the system can be divided, and to increase the degree of improvement in power use created by coordination.

The experiments in Section 4 demonstrate three factors affecting global behavior of the system: the flexibility of individual agents in shifting power-use, the composition of clusters within which plans are coordinated between agents, and the degree of communication to coordinate plans between clusters. Each of these areas deserves further study.

Coordinating Resource Agents: This paper examined one difficulty in coordinating planned energy use: a lack of flexibility on the part of resource agents. Another main stumbling block is the unpredictability of the devices controlled by the agents. For this reason plans need to be constantly re-coordinated as predicted power-use changes. Refrigerators are amenable to planning because they can in effect store energy using the thermal mass of their contents. Other appliances, like pool pumps can be flexible because the exact time at which they operate is unimportant. There are, however, many devices, like light-bulbs that must act on demand. Given a fixed set of devices there is, thus, a limit on the degree to which the variability of their energy use can be controlled. On the other hand, the increasing prevalence of household devices with significant battery capacities, such as electric vehicles, is likely to make the planning and coordination problem easier, as their reserves can be called upon to make up for a lack of flexibility or predictability elsewhere.

Clustering: The presented experiments show that choosing more appropriate criteria for dividing agents into clusters can make coordination easier by matching up devices whose power consumption is complementary. The main problem that remains is how to discover the best clustering criteria. In addition, since device behavior can change over time, using the best criteria requires that clusters also change. A number of algorithms have shown that decentralized dynamic clustering is possible given predefined criteria [12, 16]. Clustering algorithms that discover the correct criteria based on current application requirements are still in early stages of research [11, 15, 19]. To study these further, more detailed models, test-beds and measurements are required.

Cluster Coordination: The experiments further indicate that some degree of coordination between clusters is needed. In current distributed energy resource systems this coordination is achieved through central auctioneers or aggregators. Inter-cluster coordination can also be imposed by means of energy caps: fixed limits on the amount of energy that should be used in a given period. A cap is usually expressed as a directive from an energy utility company to reduce electricity consumption for peak demand periods. Our experiments indicate that low quality information can also serve to coordinate clusters, suggesting that less structured aggregation algorithms may be also applicable [8].

Self-organization: The presented solution is hybrid — it combines distributed cluster-based coordination with a centralized method of determining clusters and hierarchical coordination between clusters. Ideally, all of these functions would be achieved in a completely self-organized way. Self-organization is typically defined as “a process in which pattern at the global level of a system emerges solely from numerous interactions among the lower-level components of the system. Moreover, the rules specifying interactions among the systems components are executed using only local information, without reference to the global pattern” [3]. There are, however, a few reasons why a hybrid solution is more suitable at this stage. First, the “numerous interactions among the lower-level components” that are essential for self-organization are often costly in terms of communication overhead. For example, many peer-to-peer approaches for distributed energy resource management are not feasible if the resource agents are limited to power-line communications. Second, self-organization results in nondeterministic outcomes. Often, this is one of its strengths, and one should not avoid far-from-equilibrium dynamics and symmetry-breaking behavior but rather exploit the opportunities for creating stable patterns out of fluctuations [14]. However, in order to be adopted by industry, the non-determinism of self-organizing patterns requires an appropriate verification process, and the search for most suitable verification methodology is still open. Finally, a complete self-organizing system for distributed energy resource management would depart too strongly from incremental advancements typically accepted by the industry. A more realistic approach suggests, instead, to deploy hybrid systems such as the one described in this work as an intermediate step on the path towards a completely self-organizing solution.

6 Conclusions

This paper examined two questions about the use of clustering in order to improve scalability of distributed energy management systems: (1) to what extent is the quality of global solutions reduced by partitioning resources into groups, and (2) is it sufficient to choose these groups at random or can resources be classified so as to create more effective groups? We have presented a simple model of energy management, and a simple coordination algorithm to study these questions. Experiments have shown that dividing resources into independently coordinated clusters improves the system's scalability, but the reduced quality of aggregate plans limits the number of potential clusters. We demonstrated that three aspects affect this limit: the flexibility of energy-use plans for individual resources, the criteria used to divide resources into clusters, and the information exchanged to coordinate plans between clusters.

References

[1] Australian Government, Department of the Environment and Heritage. www.energyrating.gov.au, January 2007.

[2] S. Borenstein, M. Jaske, and A. Rosenfeld. Dynamic pricing, advanced metering, and demand response in electricity markets. Technical Report CSEMWP-105, Center for the Study of Energy Markets, October 2002.

[3] S. Camazine, J.-L. Deneubourg, N. R. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-Organization in Biological Systems*. Princeton University Press, 2001.

[4] P. Carlsson and A. Andersson. A flexible model for tree-structured multi-commodity markets. *Electronic Commerce Research*, 7(1):69–88, 2007.

[5] Y. Guo, J. Li, and G. James. Evolutionary optimisation of distributed energy resources. In *AI 2005: Advances in Artificial Intelligence, 18th Australian Joint Conference on Artificial Intelligence*, volume 3809 of *Springer - LNCS*, pages 1086–1091, December 2005.

[6] G. James, D. Cohen, R. Dodier, G. Platt, and D. Palmer. A deployed multi-agent framework for distributed energy applications. In *AAMAS 2006: Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 676–678, 2006.

[7] I. G. Kamphuis, P. Carlsson, C. J. Warmer, J. C. P. Kester, and J. K. Kok. Distributed intelligence for supply/demand matching to improve embedding of distributed renewable energy sources. In *CRIS 2004: Proceedings of the 2nd International Conference on Critical Infrastructures*, October 2004.

[8] D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *FOCS 2003: Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, page 482, 2003.

[9] J. K. Kok, C. J. Warmer, and I. G. Kamphuis. Powermatcher: multiagent control in the electricity infrastructure. In *AAMAS 2005: Proceedings of the 4th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 75–82, 2005.

[10] NEMMCO, National Electricity Market Management Company Ltd. An introduction to australia's national electricity market, June 2005.

[11] E. Ogston, B. Overeinder, v. Steen, M., and F. Brazier. Group formation among peer-to-peer agents: Learning group characteristics. In *AP2PC 2003: Proceedings of the 2nd International Workshop on Agents and Peer-to-Peer Computing*, volume 2872 of *Springer - LNCS*, pages 59–70, 2004.

[12] E. Ogston, M. van Steen, and F. Brazier. Group formation among decentralized autonomous agents. *Applied Artificial Intelligence*, pages 953–970, October-December 2004.

[13] L. Paris, G. Zini, M. Valtorta, G. Manzoni, and N. D. Franco. Present limits of very long distance transmission systems. In *CIGRE: International Conference on Large High-Voltage Electric Systems*, 1984.

[14] M. Prokopenko. Design versus self-organization. In M. Prokopenko, editor, *Advances in Applied Self-organizing Systems*, page to appear. Springer-Verlag London, 2007.

[15] M. Prokopenko, P. M. rajah, and P. Wang. On convergence of dynamic cluster formation in multi-agent networks. In *ECAL 2005: Proceedings of the 8th European Conference on Artificial Life*, pages 884–894, 2005.

[16] L. Ramaswamy, B. Gedik, and L. Liu. A distributed approach to node clustering in decentralized peer-to-peer networks. *IEEE Transactions on Parallel Distributed Systems*, 16(9):814–829, 2005.

[17] G. J. Schaeffer and H. Akkermans. CRISP: Distributed intelligence in critical infrastructures for sustainable power. Technical Report ECN-FSP-WP53-002, Energy Research Center of the Netherlands, July 2006.

[18] L. Z. Varga, N. R. Jennings, and D. Cockburn. Integrating intelligent systems into a cooperating community for electricity distribution management. *International Journal of Expert Systems with Applications*, 7(4):563–579, 1994.

[19] R. Virrankoski and A. Savvides. TASC: topology adaptive spatial clustering for sensor networks. In *MASS 2005: Proceedings of the 2nd IEEE International Conference on Mobile Adhoc and Sensor Systems*, 2005.

[20] F. Ygge and H. Akkermans. Power load management as a computational market. In *ICMAS 1995: Proceedings of the 1st International Conference on Multiagent Systems*, 1995.

We kindly thank Alec Ogston for collecting the data used in Section 4.3.